ARL/CNI AI Scenarios Strategic Context

June 2024

https://doi.org/10.29242/report.aiscenarioscontext2024

ASSOCIATION OF RESEARCH LIBRARIES

cni
Coalition for Networked Information

ARL/CNI AI Scenarios: Strategic Context

June 2024

Association of Research Libraries (ARL)

Coalition for Networked Information (CNI)

Stratus Inc.

Suggested citation: *ARL/CNI AI Scenarios: Strategic Context*. Washington, DC, and West Chester, PA: Association of Research Libraries, Coalition for Networked Information, and Stratus Inc., June 2024. https://doi.org/10.29242/report.aiscenarioscontext2024.

Table of Contents

Foreword

The Association of Research Libraries (ARL) and the Coalition for Networked Information (CNI) are undertaking a scenario planning process to imagine the future and explore a wide range of relevant dynamics associated with artificial intelligence (AI) that will reshape the landscape of the research and knowledge ecosystem. Scenario planning is a strategy methodology employed by many organizations and communities to test their long-term assumptions and investigate the uncertain landscape of their future environment.

The result of a series of interviews and focus groups of ARL and CNI members and stakeholders, this document gives voice to the collective wisdom and range of perspectives within the ARL and CNI communities on the uncertain future AI landscape as it relates to research and knowledge production, with additional focus on research libraries within that domain.

When the future is highly uncertain and the dynamics of change and progress are highly complex—such as AI in the research and knowledge ecosystem—scenario planning is an ideal methodology. Through scenario planning, a group creates a framework to span the relevant uncertainties.

The material presented in this document was a resource for the AI scenario planning process from start to finish. Please note there are no hard and fast answers in this document, but a tapestry of questions, perspectives, and insights that will allow the ARL and CNI communities to frame the core strategic question(s) it wishes to answer through the process. The themes explored by participants included strategic opportunities and challenges, core strategic questions, visions of success, and erroneous assumptions and missteps to avoid.

This document was an important pre-read for the scenario development workshop and sessions. Direct quotes from individuals are presented in quotation marks. Editing of direct quotes was minimized to stay true to the respondent's meaning.

Executive Summary

AI's future potential and impact is unclear and difficult to navigate. If the research and research library communities were to wait for certainty, it would not only be a false goal, but would lead to them losing the opportunity to determine how best they can advocate for and position further development of tools, methodology and guidelines to enable a thriving future research and knowledge ecosystem. In this context, **thriving means expanding, responsible and safe, effective, equitable, inclusive, diverse, and accessible**.

The ARL and CNI communities are in some ways challenged to leverage the thoughtfulness that comes from their cultural "conservatism" and still be proactive in an ecosystem that is currently transforming around AI. Put another way, how can research libraries more fully apply and share their expertise in data and knowledge collecting, organizing, curating, preserving, communicating, and sharing to ensure responsible and ethical AI is at the core of the transforming ecosystem?

As noted by one individual, within the ARL and CNI communities there are "so many erroneous assumptions! Assumptions about what AI is or isn't able to do in the next 15 to 25 years. How it works. Rush to regulate, assumptions about AI 'bias' versus AI reflecting human bias and agency." This scenario planning process is designed to allow the communities to gain insights on how to optimally enable the future research and knowledge ecosystem to achieve the aspirational vision of "Reshape and vastly expand the corpus of scholarly knowledge."

Strategic Focus

Based on strategic themes that emerged through the data gathering process, the following strategic focus was developed:

How do *we enable* the full potential of AI in the *research and knowledge ecosystem*?

- **Ensure responsible AI with data integrity, provenance, and persistence**
- **Achieve equitable and inclusive practices**
- **Optimally position the research and knowledge ecosystem for learning**
- **Clarify strategic role(s) for libraries that add *value***

The strategic focus is the starting point for ARL and CNI's communities to begin their strategic planning around AI and is the question this process is designed to inform. This question was discussed and further refined at the start of the scenario development workshop. The words were carefully chosen. The following are operational definitions and/or context associated with several terms/phrases:

- **We**—Allows the user of the AI Scenario Set to apply collectively, as an individual institute, organization, or department (research, higher ed, or other), as librarians, as researchers, or adjacent stakeholders
- **Research and Knowledge Ecosystem**—focused on knowledge creation and research; role of the library; knowledge for research, learning, and use
- **Enable**—proactive positioning (versus in support of)
- **Value**—fills a critical need, requires active engagement

Strategic Themes

Research Libraries Demonstrating Value in an AI Future

For research libraries, demonstrating value requires advocating for responsible, ethical, and trustworthy modeling and structuring of the research and knowledge ecosystem. The primary roles of collaborating on, collecting, organizing, preserving, communicating, and sharing the record of knowledge will continue to be of value if libraries embrace AI as a tool to enhance and enrich research and learning. Research libraries can be valued partners in developing AI research and learning models. To gain a seat at the table, the research library needs to be a proactive advocate for data integrity and trustworthy research. To avoid obsolescence and the devaluing of scholarship and libraries, the community needs to overcome its reluctance to change. Indeed, the conservatism of the library DNA allows for thoughtful planning, but in a time of accelerating change, the libraries must become comfortable with taking calculated risks.

> "How do we support academia being part of the conversation that business and governments and militaries have been having the last 5+ years about AI?"
>
> "Being part of THE assessment, evaluation life cycle of models and results."
>
> "The research library is the only actor in the scholarly research ecosystem who is involved in every stage of the research lifecycle."
>
> "The librarian needs to be the liaison—the trusted person between researchers/users and AI. Many flagship universities are having strong issues with literacy of the researchers and users. It goes beyond understanding to really influencing how it is being used."
>
> "Libraries can leverage their convening power on campus—linking it to open scholarship."

Proactively Advocating—Gaining a Seat at the Table

> "Silence is a huge problem [around copyright]. If we are not proactive now to fight for rights while it is still malleable, we will fail."

> "What is our role in all of this? What are the pieces we are OK with being out of our control, and where do we absolutely need to be 'at the table'?"

> "Do we really want to go with the flow or do we want to actively advocate? At what point do we, as a profession, start really trying to put ourselves into the conversation?"

> "I could imagine a misstep would be to not be an active partner in developing AI tools and falling behind. Not engaging would be the biggest misstep. I would also suggest that it is important for the academic library to be a partner in campus conversations around AI."

> "Libraries can assume a more expansive role as information science hubs and ethicists."

> "How can we be the leaders in deliberately thinking and taking action to benefit our individual parent institutions and the academic enterprise?"

> "If we want to be a follower—that is a very risky and dangerous position to take. My greatest concern is we remain a follower rather than a leader. We need to be a leader embracing AI, in ethics, in discovery, and in experience."

Avoiding Obsolescence—The Risk of the Devaluing of Scholarship and Libraries

> "You will find those who are in leadership in the libraries will demonstrate tremendous resistance and anxiety around a scenario that challenges the library environment. If I were to take myself 10 years back, that is the type of risk future leaders in our profession

need to take and apply and learn from these scenarios—the next generation of leaders today. There is little hope of current leaders working much beyond applying AI to improve the existing systems."

"Many library leaders are wedded to how libraries have operated. We are no longer the go-to place for information on campus. Truly planned obsolescence is the way to go. Universities don't need a big building full of books and journals. Planned obsolescence requires us to use the term librarian loosely and not refer to libraries. Our faculty had to have a degree in library school. I rarely hire people with this anymore. Instead, we are hiring people in a discipline or a postdoc that are good at working with data information."

"We are already horribly behind as libraries and academics in most places. How do we not get caught up in the details of generative AI right now? How do we get into the edges of the conversations instead of the mess in the middle, and get ahead of the conversations in knowledge creation instead of always catching up?"

"Why don't we seem to want to make decisions about AI? It is in our DNA—we are conservative."

"Through this work be sure to emphasize risk-taking—we are so risk-adverse that people don't dare try new things. Fear of failure is big as a librarian. We think in boxes—but life is fluid. We need to take some calculated risks. If not, we are gone. I have a son. What is his future? We owe it to them to prepare them for the future. To prepare systems to enable a positive future."

"We need to break the paradigm to 'change agents' from 'reluctant victims.'"

"The End of the Library as an Intellectual Center is the worst case scenario. It becomes just a place to get stuff because AI poorly replaces reference and cataloging but its good enough that budget folks can't justify human labor."

> "Ease of use of gen[erative] AI will specifically undermine critical assessment/thinking skills for library impacts all of us. Example: How Google's first page transformed information discovery. Can't always combat this, fully mitigate."

Addressing Issues of Trust in AI

Issues of trust in AI was one of the greatest challenges that surfaced through our information-gathering process. To overcome these issues, the assumptions of AI as being an evil that must be fought and AI being significantly different from human actors must be addressed. Indeed, the distrust and dangers of AI cannot be overcome unless scholars and the general public reach a threshold of AI/digital literacy knowledge and issues of disparity, inequity, and bias are effectively addressed.

> "A dangerous assumption is that AI is evil and we need to fight against it. We need people with an understanding of the issues to advocate for AI, clear away the paranoia. This is not about technology taking away your jobs. Ensure we successfully address ethics and transparency."

> "I'll lean hard into the most extreme version of my position on this, which is, I think, we over essentialize humanity relative to other tools. I'm trained in science and technology studies. I'm trained to think about socio-technical systems that include human and nonhuman actors sort of interacting with each other. I think we make the mistake of thinking there is something essentially distinct, that will matter in the deployment of these systems whether or not humans are in charge of them. I think that we're going to spend a lot of time trying to understand. Like, what are the motives of AI? What are the underlying causal motivations of an AI model? How do we trust what it is going to do? And, in fact, what's going to happen when I work with a human being? It is actually not that different. When I'm collaborating with a human, I have mechanisms for establishing trust. But even with that, I can't fully predict what someone's going to do. And I think we're looking for that from AI

systems as opposed to approaching it as we approach working with humans. Instead, with AI we want to open the black box of the AI system and understand it from the inside, as opposed to just being willing to treat them as black boxes."

Reaching AI/Digital Literacy Threshold

"Not learning enough about risks and opportunities to adequately teach AI literacy."

"Rhetoric on OpenAI—we don't understand what it means. How to meaningfully preserve these systems and maintain openness at the same time."

"The need for AI literacy—users need to understand how it works (in general) and that they need to evaluate the results for veracity. I have heard from a lot of people that they are afraid of generative AI and I think the position should be let's understand it."

"I think sometimes people think AI is evaluative—it is not, it is statistical and probabilistic, the answers still need review."

"Misapplication of the tools so that scientific integrity is undermined. Misunderstanding what the tools can and cannot do. Not knowing what is happening in the box of the AI tools and the implications of the lack of knowledge and understanding. The **perception** that AI is automagical and can solve problems without any human intervention."

Tackling Disparities, Inequities, and Issues of Bias

"That the bias inherent in the tools is comparable or similar to the bias that exists within people, and thus the solutions for debiasing datasets resemble the ways that we've attempted to manage or counteract bias within our institutions. Both are valid issues, but the root causes and ways they manifest are different and AI will need new solutions to continue to provide the quality of information that we currently do."

"That we will use AI to recreate existing practices and ways of knowing that are themselves problematic and underexplored—amplifying and scaling dangerous, insular, monolithic ways of thinking—versus the promise of multi-perspectives, multiple ways of knowing. I am not confident we can overcome the dominant cultures and systems."

"We need to create and maintain a healthy knowledge commons."

"If AI algorithms are trained on data that does not include lived experience and social determinants of health and we fail at developing a diverse workforce to be working with AI. The end result will be faulty results and conclusions based on biases and flawed data."

"AI conversation at a research university is different than the AI conversation at a small community college or small liberal arts college. Yet another example of creating haves and have nots. AI research versus others with more limited access to AI. Talk about egalitarianism—perfect for libraries. 'AI for everybody!'"

"Our systems that have the greatest funding tend to sustain and perpetuate their models, rules, and biases. How with this can we achieve improved accessibility, equity and inclusiveness around AI? Higher education in research, teaching, and resources is stratified."

"AI could deepen the digital divide, sometimes called the "capabilities overhang," growing gaps between users and non-users, especially as AI is integrated into our work products (e.g., Microsoft Office suite, Windows 11)."

Effective Guidelines for Responsible and Trustworthy Use of AI

Establishing guidelines for responsible and trustworthy AI use is a crucial endeavor. The research library community, with its knowledge of responsible data stewardship and open software development, is well positioned to lead this effort. These guidelines should tackle

risks related to the integrity, privacy, transparency, authenticity, and provenance of data and knowledge resources. Given the rapid pace of technological advancement, regulations may struggle to keep pace and could potentially be counterproductive. Therefore, emphasizing guidelines over regulations offers a more effective approach.

> "Lack of auditability and accountability of information sources leads to a reduction of new primary sources and newly created raw data."

> "My concerns center around the 'proper' implementation of the tools. Who builds them? Who manages them? Who supports them? Who checks for bias or errors in the models? What are they trained on?"

> "How will we ensure safety and prevention of harm to vulnerable communities?"

> "Will AI be properly regulated to avoid large-scale social and economic disruption?"

> "The technology is moving so quickly that no regulatory framework can keep up, and regulation may be harmful."

> "Much of our digital archives are made up of heterogeneous and often obsolete file formats with software dependencies. Without administrative metadata about what can and can't be rendered, will AI silently omit or hallucinate based on mis-readings of a corpus?"

> "Nature article on reproducibility crisis in AI. What we are finding with ChatGPT and stable diffusion is they have emergent behaviors. Science already can be pretty sloppy. The data that is used to train is critical. Garbage in means garbage out. Trustworthy AI needs to be the focus."

> "CISA (Cybersecurity and Infrastructure Security Agency), collaborated with the Australian Signals Directorate's Australian Cyber Security Centre (ASD's ACSC) on Engaging with Artificial Intelligence. In that work they compiled their list of AI-related threats:

- Data poisoning
- Input manipulation
- Generative AI hallucinations
- Privacy and intellectual property threats
- Model stealing and training data exfiltration
- Re-identification of anonymized data

This work also provides a good example of guidelines versus regulations."

Addressing Privacy and Transparency

"We need to focus on privacy. There is a huge opportunity to create AI governance for the campus....This then will enable AI research, teaching, and transform the whole business model of higher ed."

"The inability for higher education to afford appropriate infrastructure to ensure privacy and allowing them to compete with institutions who have the money to build and maintain better research infrastructures."

"How will we ensure the products/platforms used are network safe, accessible, and comply with privacy policies?"

"Scholars are either over cautious or they don't want to talk about their methodologies so research is somewhat opaque."

Addressing Authenticity, Provenance, and Copyright

"Formalize provenance through a series of digital signatures."

"Our ability to assess quality, reliability, trustworthiness of data and resources is critical. If we fail at this, we are creating knowledge chaos."

"Creators of data are upset with the use of their data to train AI. Could be data on artists, historical figures, or family history. What will it mean when we reanimate historical figures? To reanimate grandparents for future generations? Things that reanimate or

catalyze conversation? Today we need ground rules on how to approach this."

"Administrative metadata is critical in all of this, and is so boring that it may get missed as an important building block." [Note: Administrative metadata provides information that is useful in managing resources including information related to governance, access controls, security, technical data on copyright, rights management, and license agreements. Source: "What is Metadata? Definition, Types, Uses, and Examples," www.spiceworks.com]

"There are legal support that libraries need and guidance that searchers need. Libraries will need the most help with the way we acquire and license content for use by scholars—what publishers can lawfully compile and terms to grant licenses for use. Text and data mining with AI to detect patterns and relationships. The tool first needs to be trained. The data to train is not the same as the data for the research itself. AI with machine learning can do this on its own. This freaks publishers out—a large volume of work that they cannot control."

Identifying Strategic Partners across the Research and Knowledge Ecosystem

The ARL and CNI communities need to identify and partner with stakeholders beyond their traditional reach and silos. Research libraries need to collaborate effectively with researchers and academics across their institutions. Working with industry, venture capitalists, government agencies, and other sectors beyond academia and public institutes will be necessary and require the community to find common ground and interests and challenge its beliefs about these organizations.

"What partnerships are essential to ensuring that the library remains the heart of the academic enterprise that is reshaped by AI?"

"Opportunity to integrate new partners into the research enterprise."

"The key will be attaining the ability to collaborate with each other and AI to create new knowledge to progress humanity."

"Outsourcing the learning curve. Bending to industry practices even if it flies in the face of our values re: privacy, copyright, equity."

Upskilling and Becoming Adaptive, Continuous Learners

The AI workforce must develop new technical skills as well as the capacity to be adaptive, continuous learners given the accelerated pace of change. AI will most certainly take on simple and redundant tasks thus freeing up researchers and librarians for tasks requiring their deeper expertise. AI can and will augment the ability of humans to adapt and continuously learn new abilities and knowledge.

Although generative AI offers opportunities to assist individuals to learn and communicate in their native language and increase their ability to write in a non-native language, the process of writing is, in fact, important in both critical thinking and the research process itself. As a result, AI should be shaped to enhance, but not replace, the process of writing for the human researcher.

"Tasks and skill development have gone through or are going through a fundamental shift. Twenty years ago, professional knowledge and skills were at the core for a successful professional. Around this core, transferable skills were acquired. The person then operated in the organizational context. With the continuous advancement of technology and dynamics of the external environment, now the transferable skills are at the heart, the ability to learn this is the core. Around this core, professional knowledge and skills can be developed. The content is almost incidental. The core is how to work with people of different backgrounds, multidisciplinary team efforts."

"How do we need to educate current and future library workers to be effective in this emerging/evolving environment?"

"My focus is on the research enterprise. I do think there's a feedback connection into training and education around AI. We are going to see the obsolescence of a lot of things. I actually think we are going to see very rapid obsolescence of a lot of skills that are core to the curriculum right now. I worry that not enough people are thinking about what skills are going to be necessary to teach. The skills will change dramatically and very quickly. And I don't think the education system will be able to keep pace with the changing demands of students even year to year."

"The potential AI has for doing simple and redundant tasks that would free up researchers and librarians for tasks that require deep expertise"

"The work going on applying gen AI large language models is exciting. There is a peer-to-peer mentoring program with students in which AI is supporting the peer mentors. It is offering them real-time suggestions on how to share feedback and apply empathy with the peer they are mentoring. It may offer what are some conversational empathy keywords or sentences that you can use when you're mentoring peers. Suggestions on how you can better communicate and mentor."

"I am concerned about writing. I'm a huge fan of writing well. Developing an individual style. And I do think that there's a threat to that practice. I have two kids in college. I definitely hammer home to them that the best thing they can do in college is learn how to write well. Which, after all, is learning how to think well but also in whatever career they go into they will be able to advocate and assemble knowledge and points in a way that will be persuasive. You create your own style out of the bits and pieces of things you like from other writers, and I think if people rely too much on the generative AI stuff then they won't be able to do that. I'm worried about the loss of the ability to structure your thinking, and then

express it with the right words, and to edit yourself. I see this loss as a bigger fear than fabrication."

"There is a great potential for assisting with writing in scholarly research, particularly writing and grammar for people who might not have a good handle on academic writing."

Broaden Engagement in Research and Use, Engaging Community

Can AI upend the existing scholarly communication model and players? Will AI enable broader communication than ever before with instantaneous translation and communication in multiple languages? Broadened engagement of individuals who were unable to be a part of the research process or conversation on research findings will lead to an enriched dialogue and level of creativity.

"Impact on scholarly publishing by imagining a world where the writing ability doesn't supersede other contributions and AI helps to communicate ideas and discoveries without that barrier. And upends scholarly communication publishing oligarchs. Changing of the publishing business models."

"Enabling broader communication (and in turn, broader understanding) of scholarly discussions to audiences not as well engaged."

"I'm also very excited about the improvements in automated translation: both to help people consume information in other languages but also to decrease the needs for internationalization efforts in interface development—the same is true with accessibility: I expect transcription quality to be high enough in the future that it's seen as adequate for ADA compliance."

"If done right, it will improve communication across and within fields of scholarly research, especially with people who don't speak English as a first language. And it goes beyond non-English speakers

as well. Those who speak English, these tools can help to improve people's communication, efficiency, and productivity."

Creativity and Reasoning—Enhancing in Humans, Emerging in AI

Will AI augment and enhance human creativity and reasoning or might creativity and reasoning emerge in AI? What early indicators would we be looking for to foresee the first signs of creativity and reasoning in AI? Is AI's current work in medicine to aid in diagnosis and decision-making at the threshold?

Enhancing in Humans

"It would be a misstep to not fully embrace AI-augmented humans—to assume that this is just another hype."

"Can we use AI to inspire more creativity?"

"Faster ways to do things. It will be all about the role of the humans. A key question for us in the future is what will be the role of humans in scholarly creativity and activity? I like to think about humans and machines working together."

Emerging in AI

"I'm curious to see if AI can get beyond returning, compiling, and summarizing information, to reasoning. It might depend upon how we define reasoning. Just making connections, I think it can do. Coming up with new ideas? I'm not sure. "

"There is so much more coming. They don't reason yet, but there is no reason why that couldn't happen. Currently, AI helps with designing and conducting lab experiments, and hypothesis building. All automated by AI and robotic tools."

"We need to recalibrate the roles and rights of creators. Legislate for what AI creates. What it means to create."

"Eroded trust with doctors may develop with AI involved in diagnostics and making decisions."

Achieving Scalability in Research and Knowledge Processes

AI provides the possibility that human beings can leverage technology to solve complex problems at a scale not previously possible. Achieving an immensely greater scale in knowledge use and knowledge creation would accelerate the potential to successfully begin to learn about and address the grand challenges facing humanity. What does achieving scale then mean to the processes of vetting research, including peer review? Does it need to be completely rethought?

"Efficiencies at scale with AI are exciting and with AI—at scale may come to pass this time."

"I'm excited about the possibility of what AI methods could bring to humanities content at scale (beyond what has been done with machine learning and NLP)—new questions, programs, and services that could emerge as a result."

"If we're going to have a huge proliferation of the production of new knowledge, and there's still going to be a presumption that all has to be vetted by humans. That's not going to work. There's going to be a need for a huge system where you've got AI producing and then reviewing, and then everything becomes a purely automated system. And if humans are basically slower, what then is the competitive value of the human in the process? What is the relative value or the competitive advantage of humans relative to those algorithms? That's something that's going to have to be dealt with in the next 10 years."

"The potential to use AI to more efficiently describe and make information available to diverse users at scale is exciting. Applied to metadata generation."

"Achieving scale will allow humanity to make progress on grand challenges like climate change."

"There are important questions about how much we are going to be willing to allow systems in deployment to move at faster velocity than humans can manage to work at greater scale than humans can? I mean, what's tricky is that this is at odds with the logic of the technology which is, build whatever you can, and then see what it can do. So how can we reconcile maintaining a sense of intention on how this tech is deployed?"

Merging Technologies with AI—Greatest Hope, Biggest Disruption

AI is now merging with mobile devices, office software, and virtual and augmented reality equipment. What will happen when AI and quantum computing converge? As technologies emerge, disruptive breakthroughs have the potential to redefine what it means to be human and what a human being is capable of learning and accomplishing. Such mergers could be the greatest hope of the future or the biggest disruption.

"What will happen when AI and quantum computing converge...it feels like we are in truly uncharted waters even as much of the computational technology driving current AI is statistical and predictive."

"We can't separate AI from other technologies like sensor technology. It's a synergy between other technologies and AI. You see now fully automated labs that are clean spaces with automated cameras and sensors to look and observe with AI looking at specific characteristics of say mouse behavior. It blows you away. Multimodality and no humans in the lab."

"What excites me is the sense of possibility and unknowingness of it [merging of AI with other technologies]."

"The work of a researcher is oriented much more around experimental design. And you know, evaluating questions rather than the actual minute work of doing the analysis. One of the things

I've been paying attention to is the results of AI being employed in the physical sciences with lab automation tools. I think the iterative feedback loop of robotic labs combined with AI is going to mean a host of things to materials, discovery, chemistry, chemistry analysis. There is going be a lot more opportunity for researchers to just do more science."

Need for Thoughtful Planning and Collaborating

Thoughtful planning and collaboration is needed to enable a thriving research and knowledge ecosystem. Such a process of planning will require critical thinking, consideration of context, assessment, and learning, and a deepened knowledge and understanding of AI. Slowing the process down has benefits in the rigor of analysis and the decision-making process. However, slowing the process down could mean failure to keep pace or act in time. With the accelerating pace, how can individuals and institutions maintain flexibility in both their understanding and adoption of AI?

"Critical thinking and considerations of context in determining the integrity of a source; the value of assessment and learning (outcomes-driven learning, determination of best practices, definitions of success and effectiveness); patron privacy as a value; appreciation for needs-based explorations, user communities, and user experience."

"How can we maintain flexibility in our understanding and adoption as fundamental AI capabilities continue to develop at an accelerated pace."

"I feel like sooner or later, we're going to have the equivalent of the financial crash that happened in sort of algorithmic trading where some real world system gets taken down because of the unanticipated consequence of a not fully vetted AI system."

"Worried about the hype. Buying more stuff to do more stuff without a thoughtful understanding of what and how best to do it. Higher ed

likes to do things right, thoughtfully. With a new tech there are always many risks and new potential. The process may be too slow for some, but it is a thoughtful process."

"We are failing to invest in foundational research on AI necessary to uncover paths forward."

"At what point will users start accepting or indeed rejecting 'fast' over 'good'?"

"The institution's reaction to AI has been ad hoc and scattered. Nobody's quite sure where to focus."

"Secondary concern here is putting all our eggs into one basket when the technology isn't even really settled, who knows if ChatGPT is the final winner vs. Claude?"

Addressing Who Is in Control

Who is in control and what does this community do about it? How might our community's assumptions and biases about intentions be barriers to collaboration and progress? Are the tech companies in control? Or is it the government regulators? Or someone else? Tech companies that are run by individuals who have come out of the academic space tend to be better partners for collaboration, understanding the concerns and value the scholarly research community brings to the table. Are we looking at a race to market or a race to quality with a pause for ethics? Vendors and the research community must come together around the importance of ethics and responsible AI.

"We need to discuss how to be the counterbalance with legislators, lobbyists, big tech to influence decisions—so we can have an ethical and responsible AI—making AI essential, not peripheral. Influencing bias, ethics, information exchange, evaluating fake data, watermarking to know AI origination."

"My greatest concern is that these systems are run by technocrats and corporations that control their content based on their bottom line. We are effectively handing them complete control over censorship and we're not even aware of it. What happens when they decide that the climate change discussion negatively affects their bottom line and they do post-training to filter some of it out?"

"Power and transparency—trust in name and not deliver on it. Little incentive in the commercial sector for openness and transparency. AI companies want to gain access to the data, they aren't normally looking to partner with the research and library community. However, those companies run by individuals who came out of scholarly research seem to understand and see the benefit from partnering in the development and design, seeing the value we bring to the table."

"What will happen around the vendor community? Will this be a race to market with quick wins and fast money or will there be a pause for ethics so it instead becomes a race to quality? A misstep would be if the vendors and research community (researchers and libraries) don't come together around the same set of expectations."

"Monopolization of AI exacerbates the digital divide."

Minimizing Unintended Consequences

To minimize unintended consequences, the research and research library community need to be able to discern what is hype and what is not. It is dangerous to assume it is all hype or to believe the hype. Bad actors in the research and knowledge space have the potential to further misinformation or misuse of AI technology. To date, we have not figured out effective countermeasures for "poisoning" datasets with misleading or malicious results. How does this community minimize unintended consequences?

What's Hype, What's Not?

"We need to be careful of the erroneous assumption that it is all hype."

"I worry that people will expect AI to deliver more than it can, believing the hype, jumping to the potential for cost savings...it's like digitization all over again."

"What does the AI discussion look like when we move past the hype?"

"I wonder if AI will truly be more 'revolutionary' than other technological innovations that libraries and research have incorporated in the past."

Addressing Bad Actors

"That the bad actors who already control/enclose too much critical information will be further strengthened, enriched by their ability to develop new weapons and thereby deepen inequality."

"AI could be used to mass-produce content, enhance the quality of predatory journals/publishers, create more sophisticated fraudulent data/research, et cetera."

"Bad uses outside research could trigger disabling and indiscriminate regulation, preventing us from achieving the promise."

"I am concerned that we haven't figured out effective countermeasures for 'poisoning' datasets with misleading or malicious results."

Transformative Power of AI

AI has the potential to transform the research and knowledge ecosystem profoundly. Is the community underestimating either the extent or speed of the transformation? This theme is organized around

transformation of the research libraries, research workflows and outputs, higher education, learning, and scholarly communication.

> "We are thinking too small and too immediate; we are underestimating the transformation."

> "Phoenix rising from the ashes. How can we totally reconceive of scholarly publishing, higher education, and research impact?"

Transformation of Research Libraries

> "In the future there will be new jobs to manage ethics, privacy, biases. For these the university or research library will be the best entity to manage AI. Spaces will be completely different in 10 to 20 years. There will be a nostalgia for old-fashioned spaces that only rich libraries will be able to afford to maintain. Currently Chinese libraries have created a fake sense of nostalgia in their buildings by introducing fake shelves of books into the space."

> "Enhancing the discovery/connection-making process that researchers already use."

> "With focus of work on standards and creating well-formed data, we have set the stage for helping to drive the development of customizable responses to user queries. This is a 'win' and shows the value of the work of libraries."

> "I think that there are some great potential tools to help out with scholarly research. New kinds of searches, new kinds of analyses, better approaches not only to text, which, after all, we've been operating on computationally for decades, but now images can also be analyzed with the rise of better computer vision. Potentially also the rise of better OCR (optical character recognition) for hard problems like handwriting."

> "Ability to describe and make accessible material we could not afford the manual effort for in the past."

Transformation of Research Workflows and Outputs

"The profound change that comes to every aspect of research. People look at it myopically. The fundamental change in how we think about research together with every aspect of research—it's not just about predictive models. This means change in how we think about experiments, write grant proposals, whether to propose a hypothesis or not. With AI currently we can get the right answer but don't yet understand why we know it is the right answer. For instance, AI created a 3D model of a gene. We don't know why it worked, but it did. It means the whole philosophy of science has changed."

"AI is not just generative AI. How to leverage vast transdisciplinary power of AI that does not rebuild disciplinary silos, but enables grand-challenge collaborative research."

"The potential for divergent discovery and new fields of research and new methodological approaches to existing research is compelling—this certainly will be a big impact for research libraries."

"I'm excited about the potential for improved personalization of discovery results—that AI will be able to do a better job of contextualizing search results and allow people to navigate increasingly dense and complex networks of information."

"The relationship between the human mind and the space the AI algorithm uses. So far research has been based on acts of reductionism for human minds to manage the amount of data, we found ways to simplify. Simplification does not mean we are working with what actually is. With AI we can work with vastly larger data sets and remove this act of simplification."

Transformation of Higher Ed

"So far on our campus, the questions around teaching and learning have been most prominent—so thinking about how/if students are using it, and what are the parameters, etc."

"Arts and humanities are not separate and left out of the AI world. AI is deeply involved in arts and drama. There is much innovation happening in the arts with the introduction of AI."

"AI can be democratizing, but how to ensure we don't create AI haves and have-nots between STEM and arts & humanities?"

"How will we uphold our commitments to equity, diversity, inclusion, and accessibility while also integrating AI on our campus?"

Transformation of Learning

"To experiment, to model ways of leveraging assistive tech, to continue to be involved as the ways of learning evolve."

"Mutually collaborative partnerships across the research/learning ecosystems."

"Critical thinking and considerations of context in determining the integrity of a source; the value of assessment and learning (outcomes-driven learning, determination of best practices, definitions of success and effectiveness)."

"I also think that AI will make in-person experiences (especially in-person learning) more valuable."

Transformation of Scholarly Communication

"The issue of how to address bottlenecks is an interesting question. A lot of what makes open-source software work is the community dynamics. It's the trust. What does it mean if you have nonhuman contributors to a software project? And in the same sense, what does

it mean if you have algorithmic or nonhuman reviewers of an article? And when will those be accepted as legitimate?"

"In terms of scholarly research and communication, AI seems like it might offer a level of computer automation/integration that fully upends traditional scholarly outputs."

Where Do We (ARL and CNI) Collectively Focus?

Respondents shared a wide range of areas that would benefit from collective learning, planning, and action. These areas include advocating around areas where the community has a distinctive voice, developing AI tools for libraries, developing a system of trust markers and standards, developing guardrails or guidelines. One respondent noted it will be important to think about this globally, recognizing the limitation that this community is representative primarily of the US and Canada.

> "For ways we can support the larger field, focus on areas where we can have a distinctive voice. Many of the concerns of AI are shared across many disciplines, AI ethics for instance, and we can each support, but are unlikely to lead individually in those areas. We have distinct viewpoints and the ability to lead topics in this discussion given our unique skill set, and we should provide that voice into the conversation."

> "We are not recognizing the need for a coordinated approach to lobbying for federal guardrails."

> "More broadly speaking, I would like to see libraries working together to develop AI tools for libraries rather than everyone custom developing for themselves."

> "Are there trust markers that we can jointly develop with standards organizations? What guardrails should we advocate for? And how does this play out globally given our advocacy is limited to the US?"

> "I wonder if we need to pivot from looking to individual institutions and their 'success' in using AI to thinking more consortially and creating a national infrastructure."

> "Representing and participating in a uniquely collaborative community of practice, allowing knowledge and skills to be shared across institutions with an ethos of openness and collaboration."

"We have the ability to shape how AI is adopted and implemented in various tools, services, and ways in which we have expertise. We need to engage our patrons and stakeholders."

Aspirational Visions of the Future

Several respondents shared aspirational visions of an AI-influenced future:

> "Our current work and way to do business will be elevated tremendously by AI. It will be the co-piloting model—making our work faster, higher volume, deeper, and broader. I am 99% certain this will happen. The second dimension will transform totally what we do and be a new paradigm that is hard to imagine and describe. On this one, I would say the probability of it happening is 51%. AGI (artificial general intelligence) where we are no longer co-piloting. AI has autonomy and is a self-driving model. Implications and impact on society are profound. Knowledge creation then moves from the people we support (researchers, users) to the people **and** the AGI scientists and users. At a recent conference, AI scientists discussed that they felt, in 25 to 30 years, AI can create a Nobel prize–winning scientist for its work."

> "Looks like a much faster rate of reliable knowledge communication, much faster time cycle for resolution of open questions. Used to take three years. In this scenario it will take three days **and** we are asking better questions. Asking more meaty and important questions. In terms of process to ensure reliability and transference of scholarly output —> fake data and AI tools will be used to quickly assess credibility—allow for comprehending use, analyze at much more volume than possible with human beings. Looking at large swatches of human knowledge. For humanity—rate of progress much faster on critical questions and things we want to know more about. Human and AI are in knowledge-development together. It could go down the path of AI agents working together, but I much prefer AI working with teams of humans."

> "The research enterprise has been transformed by AI augmenting human research skills. AI has taken away many tasks from researchers that they didn't need to be doing. Research faculty focus their talent on identifying and answering big input questions, like

the UN sustainable development goals as an example. The librarian's part is to help in the data piece and in the space of skill development. Our librarian team helps in training the AI platform. We are discipline agnostic, so we are not skewed by any one discipline—this is a vital part of our credibility."

"Imagine you are a history major and want to study a historically significant individual. With a VR headset and a timeline to follow with digital content that tracks that individual's life, social network, and significant accomplishments, you can now travel to any point in that history and experience a conversation with the individual or interact with the individual and others in their network perhaps associated with a key event or decision. It is a simulated academic village housing the history in the past for a fully immersive experience. You are simulating history."

"Higher ed came together around core requirements for any AI software or product = it was the cost of entry. This made vendors work to meet those requirements and wishes. Higher ed rethought curriculum—writing across curriculum. Ethics around these technologies is built across every college in the university. AI will disrupt and teach students to be 'robot-proof' (Joseph Aoun book) with skills derived from the humanities and building up creativity. Higher education supports students developing themselves to participate successfully in a digital world."

"I have a really credible AI collaborator who I can get in a conversation with about interesting or novel potential experiments or research ideas. I think a successful world is where that collaborator is as generative as, if not even more than, I am as a researcher. I'm in a 10-year horizon with this, but I'm a little wary of the idea that research gets fully turned over to the AI systems. I think there's a role there no matter what for humans. Humans to continually be injecting novelty, for continually bringing different perspectives to bear and sort of destabilizing those AI systems that may be trained on all of the scientific literature."