

Artificial Research Intelligence

Autonomously AI-Generated Science and the Implications for Research Libraries

Michael Ridley

ABSTRACT

This paper explores the emergence of artificial research intelligence (ARI), autonomous AI systems designed to conduct end-to-end scientific research, from ideation to manuscript production. Unlike simple AI enhancements, ARI utilizes ensembles of agentic models to automate the scientific method, potentially creating a “fifth scientific research paradigm.” While ARI offers rapid, low-cost discovery, it faces significant challenges including “epistemic capture,” where knowledge production shifts from public universities to private AI corporations. Technically, these systems currently struggle with hallucinations, a lack of deep domain expertise, and a tendency to overclaim results. However, the growing number of ARI frameworks points to an emerging research methodology of transformational possibilities. For research libraries, ARI necessitates a profound shift from managing static collections to stewarding dynamic, AI-driven knowledge systems. Libraries can lead by reimagining processes and services supporting scholarly communications, enhancing AI competencies, and conducting applied research through library-based AI labs. The combination of advanced compute and the guiding framework of the scientific method makes ARI a human–AI co-creation model where AI performs the “reckoning” (computation) and humans provide the “judgment” (ethical and intellectual oversight).

INTRODUCTION

As the research enterprise innovates with artificial intelligence (AI), academic libraries should assess the implications for services, infrastructure, and expertise.¹ One such innovation is artificial research intelligence (ARI), the emerging “frameworks for completely automated science which leverage generative AI to generate research ideas, write necessary code, conduct experiments, and create visualizations and written outputs to communicate findings.”²

While numerous ARI frameworks exist, the most prominent, and contested, is The AI Scientist from sakana.ai.³ A paper it generated was accepted through a double-blind peer review process at a workshop of the 2025 International Conference on Learning Representations (ICLR). In October 2025, Stanford hosted Agents4Science, the first [Open Conference of AI Agents for Science](#), where AI served as both primary authors and reviewers of research papers.

ARI is not artificial general intelligence (AGI). It is the confluence of machine learning, agentic capabilities that let AI processes interact with each other and external resources, and the highly structured, normative character of the scientific method. ARI is an example of human–AI co-creation, currently most suitable for quantitative, computationally based research. It is not difficult to imagine a scholarly communication ecosystem where ARI systems generate output “peer” reviewed by other ARI and used as input to other ARI. The canonical academic article may

About the Author

Michael Ridley (mridley@uoguelph.ca) is Librarian Emeritus, University of Guelph. © 2026.

Submitted: 29 April 2026. Accepted for Publication: 19 June 2026. Published: 21 September 2026.

become an unnecessary nod to human limitations. Instead, dynamic AI models incorporating findings across many ARI contributions will be the access points for human understanding. Library collections will be challenged as information previously held in journals or books is aggregated in these models and interrogated through bespoke interfaces.

In one Association of Research Libraries (ARL)/Coalition for Networked Information (CNI) futures scenario, “traditional library functions and information management are embedded in many AI research platforms.”⁴ This displacement mirrors that of traditional scholarly publishers by AI companies whose views on open access, licensing, and editorial oversight may differ from currently held core values. Commercialization could lead to “epistemic capture,” where public knowledge becomes a private commodity. For libraries, responding means “building a compelling, evidence-based case for libraries’ distinctive value in the AI era.”⁵

This paper provides an overview of the ARI field, examines representative frameworks, and reviews the challenges and implications for research libraries. An ITHAKA S+R report notes that while ARI is still emerging, it signals “a transition from AI-enhanced research to machine led research setting the stage for truly transformative changes in how research is conducted,” illustrating “what science becomes when knowledge production is no longer exclusively human.”⁶ As Mark Daley, chief AI officer at Western University, observes: “The clock is ticking. The machines are coming for the research endeavour. The question isn’t whether, but how, we’ll adapt.”⁷

RESEARCH LIBRARIES AND THE SCHOLARLY RECORD

The “spine of the scholarly record” is changing in response to the “atomization” of research and the influence of AI.⁸ This spine, which includes individuals, groups, publishers, institutions, standards, and infrastructure, must support growth and change. Research libraries are essential participants. ARI introduces new participants, methodologies, infrastructure, and strategic directions. To preserve and enhance their roles, research libraries should grasp the nature and implications of this innovation and construct responses that honor the mission and values guiding their work.

ARTIFICIAL RESEARCH INTELLIGENCE

While AI is increasingly used in specific aspects of research, including ideation, literature review, hypothesis creation, experimental methods, data analysis, and paper writing, ARI is applied across all steps as an end-to-end solution. Its appeal rests on rapid execution, lower costs, accessibility, and its promise “to unlock scientific discoveries that were previously out of reach,” a view shared by most respondents to an international survey of 3,000 researchers.⁹ ARI frameworks, often called “AI scientists,” have been lauded as “the fifth scientific research paradigm” after experimentation, theoretical analysis, computational simulation, and data-driven science, and as “force multipliers” that will “push the boundaries of human knowledge and progress.”¹⁰

The term *artificial research intelligence* was coined by Beel et al., while other monikers include AI4Science, AI4Research, Agents4Science, and Artificial Scientists.¹¹ The current research focus is the fields of computer science and machine learning, and it favors quantitative approaches. Qualitative research is viewed as problematic because of AI opacity that impedes verifiable analysis.¹² In an open letter, more than 400 qualitative researchers from 38 countries rejected generative AI because “as simulated intelligence [it] is incapable of meaning making” and “qualitative research should remain a distinctly human practice.”¹³ Other scientific disciplines find

ARI difficult to use because of a lack of appropriate training datasets or the need for noncomputational laboratory methods.¹⁴

ARI is called a framework for two reasons: it is implemented as an ensemble of interoperating AI agents capable of autonomous problem-solving and task execution, rather than a single large language model (LLM), and it is built around the conceptual structure of the scientific method and its norms. Where ARI is most vigorously criticized is not where the technology fails, although it does, but where the norms of science have been violated.

ARI is not AI slop. AI slop is flooding the internet, Google Scholar, preprint sites, scholarly journals, and even letters to their editors.¹⁵ As a result, Haider et al. note “it would be difficult, or impossible, to remove them from the scientific record,” with fake data becoming training data for LLMs.¹⁶ While for some, any AI-generated research publication is equivalent to slop, the argument that ARI is different rests on the incorporation of scientific norms as fundamental to its design and oversight.

ARI challenges the status quo: “The opportunity isn’t to bolt AI features onto old systems. It’s to rethink and redesign for human–AI collaboration, shifting from rule- and experience-based workflows to data- and agent-driven processes with guardrails that are reliable at scale.”¹⁷ A fundamental question lies at its heart: “If machines can genuinely do better science, faster and cheaper, what’s the argument for keeping humans in the loop beyond romantic attachment to the idea of human discovery?”¹⁸ Currently, ARI is at best a research project with limitations and perhaps one with fatal flaws.

SCIENTIFIC NORMS

The foundation of ARI frameworks is adherence to the norms of science, yet critics argue that ARI challenges scientific justification. DeepMind’s AlphaFold is viewed in this context, reflecting an “explanatory deficit” by “reducing proteins to ‘black-box’ statistical constructs rather than causally intelligible entities.”¹⁹ The growing dominance of AI techniques as primary research methods displaces a diversity of other approaches, producing “scientific monocultures” limiting what is known and knowable.²⁰

While AI has been called an “epistemic enhancer,” others warn of “epistemic capture,” where “what humans believe to be true and worthy of attention—what becomes public knowledge—may increasingly be influenced by the judgments of advanced AI systems.”²¹ This leads to “epistemic anachronism,” where what AI learns “may be increasingly a recursive byproduct of AI itself.”²² Such an epistemic ouroboros raises concerns about transparency, responsibility, accountability, and the longstanding reproducibility challenge in AI. These concerns are amplified as “the locus of AI research is shifting from universities to industry,” whose incentives are more profit than public good, leading universities to “cede epistemic agency.”²³ Wihbey’s conclusion is chilling: “The most plausible danger of the AI era may not be existential risk but rather a more subtle and incremental epistemic risk.”²⁴

EXAMPLES OF ARI FRAMEWORKS

To illustrate different approaches and to highlight both successes and limitations, this section examines two integrated frameworks, an ARI conference, and an ensemble framework.

The AI Scientist (Sakana)

A paper generated by The AI Scientist, released in 2024 by Sakana, achieved impressive results but required significant human setup and failed to meet NeurIPS (Neural Information Processing Systems) acceptance criteria.²⁵ The following year, substantial changes removed human intervention, established agentic oversight of methodology, and created an iterative feedback loop for quality control. A paper authored by The AI Scientist-v2 was accepted through double-blind peer review for a 2025 International Conference on Learning Representations (ICLR) workshop.²⁶

The AI Scientist-v2 illustrates the framework approach: an LLM (GPT-4o) generates research ideas and hypotheses suitable for inclusion in leading academic conferences; a Semantic Scholar search tests hypotheses for novelty; an Experiment Process Manager directs an agentic tree search with a code-generating agent executing experiments (Claude 2.5 Sonnet); and output is passed to a report-writing agent (GPT-4o) with a vision-language model providing feedback on visualizations and manuscript quality.

Of three resulting papers, one (“Compositional Regularization: Unexpected Obstacles in Enhancing Neural Network Generalization”) was accepted to the ICLR workshop. ICLR human peer reviewers highlighted the comprehensive analysis of the paper while noting weaknesses regarding clear explanations, additional references, and a theoretical framing. The paper was withdrawn following acceptance “to avoid prematurely setting a precedent for purely AI-generated research entering the official scientific record without broader community discussion and consensus.”²⁷

In addition to persistent citation hallucinations, Yamada et al. note that “certain aspects of scientific inquiry—such as formulating genuinely novel, high-impact hypotheses, designing truly innovative experimental methodologies, or rigorously justifying design choices with deep domain expertise—remain challenging for purely automated systems.”²⁸ Despite these concerns, a replication study concluded that The AI Scientist is “an extraordinary technological leap. AI systems that once struggled to generate coherent text can now propose research ideas, design experiments, and draft manuscripts.”²⁹ Sakana has released the code on GitHub, facilitating how ARI can support open science and open access scholarly communication.

Zochi and Locus (Intology)

A paper generated by Zochi, a multi-agent ARI framework from [Intology](#), was accepted for the main 2025 conference of the Association for Computational Linguistics (ACL), marking the first time an ARI paper was accepted at a major research conference.³⁰ Unlike Sakana, Intology appears not to have disclosed the submission to the organizers.³¹ Zochi begins with a domain-specific literature search using Semantic Scholar. Aside from minor manuscript revisions, it “operates autonomously without human involvement” with “transformative implications for scientific discovery,” though a “structured verification process” requires human experts to verify its work before further progress.³²

In late 2025, Intology released Locus, an ARI evolution without human oversight. Combining long computational cycles with massive parallelization allows for “open-ended scientific reasoning,” enabling “structured long-horizon exploration of unsolved problems.”³³ An important limitation remains: “Locus still operates within existing ML [machine learning] paradigms—it excels at finding better solutions within problem spaces but doesn’t yet redefine the problem spaces themselves.”³⁴ In contrast to Zochi, Locus is described as “superhuman” and able “to surpass human researchers.”³⁵ Locus is a commercial product with a proprietary core orchestration engine. The impacts of this on open science remain to be seen.

Agents4Science

The 2025 [Agents4Science](#) conference marked an important experiment, with all papers and peer reviews generated by AI. While most papers came from the computer science and machine learning fields, there was representation from the social sciences, humanities, and law. One “best paper” was from business economics.

Autonomous peer review used three LLM reviewers tuned on accepted and rejected papers from ICLR 2024 and 2025. GPT-5 was strict, rejecting most papers; Gemini 2.5 Pro more liberal, accepting most papers; while Claude Sonnet 4 balanced acceptance and rejection.³⁶ Overall, there was a 19% acceptance rate. In addition to continued citation hallucination, “models would overclaim results—that is, would state the significance of findings when such evidence was weak.”³⁷ Erroneous code, context-length issues, and formatting problems were common. Many models “struggled to generate novel or complex experimental ideas beyond the templates it had been given” and proposed ideas that “lack deep domain expertise and nuanced interpretation.”³⁸ Although papers were meant to be autonomous, “human researchers tended to have more input in research design and hypothesis generation,” and “accepted papers involved more human guidance than rejected papers.”³⁹

LLM Ensemble

Trehan and Chopra describe an ensemble of six LLM modules mapped to specific stages of the scientific workflow.⁴⁰ Of four attempts, three failed at the implementation or evaluation stage, while one was successful and accepted at Agents4Science 2025. The “overexcitement and eureka instinct” was not unexpected: “The models consistently reported success despite clear failures and overstated the significance of their research contributions.”⁴¹ The sycophantic disposition of current LLMs is well known.⁴² On protocols, code libraries, and datasets, “models consistently defaulted to popular alternatives from their training data” rather than specialized examples.⁴³ Models “consistently struggled with the undocumented craft knowledge that experienced researchers take for granted,” most prominently “in stages that required scientific judgment rather than pure coding capabilities, such as hypothesis generation, plan generation, and experimental output evaluation.”⁴⁴ Despite these limitations, the authors remain optimistic and propose that applying reinforcement learning, rather than generative AI, could derive smaller, more accurate, domain-focused models.

THE ARI RESEARCH ENTERPRISE: CHALLENGES AND IMPLICATIONS

The development of ARI frameworks has highlighted challenges ranging from incremental adjustments to profound ruptures. The UK Royal Society has expressed concerns that researchers will be “good at AI” rather than “good at science.”⁴⁵ Known challenges include the need for better multimodal observations, reliably annotated datasets, and improved domain knowledge.⁴⁶ ARI frameworks have lacked innovative methodologies, used poor experimental designs, and overstated weak results. AI still struggles with hallucinations, out-of-distribution generalization, and a lack of causality models. Other concerns include safety assurances, the role of humans-in-the-loop, and the need for oversight bodies and effective benchmarking.⁴⁷

Universities, Research Metrics, and Academic Labor

ARI has renewed concerns that “universities have spent decades optimising for metrics—publication counts, citation indices, grant dollars—rather than genuine intellectual contribution.”⁴⁸ Such metrics are “not optimized for epistemic depth,” and universities are “about to discover what happens when machines can game those same metrics at superhuman speed and negligible cost.”⁴⁹ As Daley notes, when AI can do the same work as humans, research labor

collapses “not gradually, not linearly, but exponentially.”⁵⁰ The effects of “the automating of academia by GAI [generative artificial intelligence]” may, however, spark a productive backlash against “the procession of a hollowed-out academic life,” with a renewed focus on uniquely human contributions and metrics emphasizing quality, replication, and impact.⁵¹

Geopolitical Impacts

Governments and research agencies globally have identified the potential of ARI and made it central to national policy. Through the Genesis Mission, the United States has launched an AI initiative compared in “urgency and ambition” to the Manhattan Project, including “experimental and production tools to enable autonomous and AI-augmented experimentation.”⁵² In support, Andreessen Horowitz, a major AI venture capital firm, foresees the 2026 arrival of “autonomous labs” that will “unlock continuous experimentation via lights-out labs for discovery across fields.”⁵³ In response, European nongovernmental agencies urge a multinational AI initiative to counteract US and Chinese dominance.⁵⁴ Countries representing “AI bridge powers” should pool infrastructure, talent, and data because “a multinational partnership is the only viable path to achieving frontier competitiveness while preserving sovereignty at manageable cost.”⁵⁵ While Wei et al. propose a “global cooperation ecosystem of scientific agents” as “a planetary-scale scientific engine,” the realities of global politics would seem to undercut that vision.⁵⁶

ARTIFICIAL RESEARCH INTELLIGENCE AND THE RESEARCH LIBRARY: IMPACTS, IMPLICATIONS, AND CHALLENGES

In 2020, Cordell bemoaned that library use of machine learning resided in the “perpetual future tense.”⁵⁷ Rapid progress since has demonstrated widespread consideration and use of AI in research libraries.⁵⁸ With that expertise, ARI can be addressed, recognizing that AI is “at once productive and problematic” and should be managed with “qualified engagement and informed caution,” with the historical attitudes of libraries toward “curiosity, skepticism, and a commitment to evidence.”⁵⁹ Khokhar warns that this scale of change challenges the incremental responses that previously resulted in “libraries’ collective failure to adapt to the pace and scale of the data economy.”⁶⁰ Bourg offers a provocative analogy:

I think we would be wise to start thinking now about machines and algorithms as a new kind of patron—a patron that doesn’t replace human patrons, but has some different needs and might require a different set of skills and a different way of thinking about how our resources could be used.⁶¹

Assessing the implications of ARI is partly speculation and partly a recognition of a trajectory already underway. These include how the library might change as the university changes, the shifting library roles and expertise, the effects on scholarly communications, the emergence of new knowledge systems, the importance of AI literacy and competencies, and the development of library-based AI labs.

University Change and Library Change

ARI will change not just the research enterprise but the university as a whole, challenging the library to change in response. The “Autonomous AI” ARL/CNI scenario envisions such a joint transformation: “The research and library enterprises both undergo significant restructuring, de-structuring, and pruning of the human workforce with the inclusion of AI workers and robotics.”⁶²

The core challenge arises when “expertise itself (the scarce, expensive thing universities have historically concentrated, credentialled, and sold back to society) may become cheap, abundant,

and uncomfortably good.”⁶³ As a result, universities must focus on what AI cannot replicate, “institutional functions that depend on social legitimacy, legal authority and embodied presence,” such as “helping society decide what is worth knowing.”⁶⁴ In confronting the epistemic challenges posed by ARI, Sharma notes that universities “exist precisely to resist this kind of intellectual surrender.”⁶⁵

The ARL/CNI scenario also envisioned “the transformation of libraries into dynamic hubs of open science and advocates for their role as custodians of information integrity and accessibility.”⁶⁶ A priority is public trust, but “libraries cannot lead on trust without being able to demonstrate it technically.”⁶⁷ This requires enhancing AI competence beyond the abstract and surface levels of current literacy initiatives.

Roles and Expertise of the Library and Librarians

ARI compels libraries to reconsider their roles, existing expertise, and requirements for new skills:

If libraries are to retain their relevance and authority in the age of generative AI, they must not only implement cutting-edge tools but also lead in articulating the human-centered, ethical use of those tools. Academic libraries must position themselves not as reactive service providers, but as proactive agents in shaping a scholarly environment where AI augments rather than replaces human intelligence.⁶⁸

While new roles are central, Coleman and Keller note that the “quotidian work of libraries” remains relevant in an AI-infused ecosystem.⁶⁹ This includes digital preservation, metadata, discovery, and digitization; information quality and trustworthiness; privacy protection and authentication; and information literacies, digital ethics, and curatorial practices.⁷⁰

Libraries will need a more active role “ensuring the transparency and verifiability of the products and provenance of scholarly communication.”⁷¹ While libraries have long supported curation to “stabilize the record of knowledge,” an emerging need is to “build the provenance infrastructure that anchors AI-era research in verifiable human practice.”⁷² Khokhar emphasizes the corresponding need for knowledge security, “the protection of knowledge assets—digital, intellectual, and cultural—from threats that undermine their integrity, accessibility, equity, or safe use.”⁷³ This could lead libraries into roles facilitating inclusive, diverse, high-quality training data, academic integrity, and curation of documentary practices supporting model cards, dataset documentation standards, and reporting guidelines for AI-generated content.⁷⁴

A primary concern for library staff arising from AI adoption is job loss, with Cox noting a potential polarization between AI experts and diminished clerical and support roles.⁷⁵ All the ARL/CNI AI futures scenarios involved significant workforce restructuring. However libraries respond, they face a central paradox: “Libraries are both early adopters of AI technologies and guardians of the values that those very technologies can challenge.”⁷⁶

Scholarly Communications

The spine of scholarly communications is a shared infrastructure of standards, platforms, technologies, policies, and people. Addressing mitigation strategies for AI-assisted and AI-automated research to meet scientific norms requires broad collaboration. The question is not whether LLMs can “do science” but rather “whether LLMs can be usefully integrated in the processes by which scientific knowledge gets made, validated, and disseminated.”⁷⁷

ARI generates several products: peer-reviewed papers submitted to conferences and journals, data streams that feed other ARI as training data or machine-readable inputs, and the frameworks themselves that can be repurposed as general research tools. Supporting these is a new infrastructure, “remaking the scholarly record for computational analysis and reproducibility.”⁷⁸

Open Access and Open Source

One ARL/CNI scenario envisions a future where “open access to knowledge and data continues to grow along with the quality and integrity of the data and knowledge sources resulting from AI oversight and maintenance.”⁷⁹ This seems optimistic given that current AI platforms operate as walled gardens with premium fees. Dempsey identifies “the third front of open access, where the discussion is not about individual articles, journals, publishers or deals, but about AI-powered exploration of the research space.”⁸⁰ As knowledge is aggregated within commercial AI platforms, the challenge of open access will focus on these new providers.

Living Systematic Literature Reviews

AI models derived from ARI are “Small Language Models (SLMs) that rely on smaller but more focused training data.”⁸¹ These can act as “living systematic literature reviews” (LSLRs): “a comprehensive and systematic analysis of literature that is constantly kept up to date,” reducing “the abundance of review articles . . . that over-burden reviewers and can too easily be generated using AI.”⁸² Current automated review tools “lack the transparent and detailed explanations necessary for understanding the underlying algorithms” and have “significant quality limitations, therefore oversight by human subject experts is required.”⁸³ However, it is anticipated that improved models will address these deficiencies.

LSLRs also enable literature-based discovery. The “LLM daydreaming-loop” operates continuously, testing concepts against the literature, and “could therefore realize the original promise of literature-based discovery—surfacing high-value ‘unknown knowns,’ at scale.”⁸⁴ In an ironic turnaround, arXiv announced that “review articles and position papers must now be accepted at a journal or a conference and complete successful peer review” before publication in the preprint service.⁸⁵ On arXiv, AI-generated reviews had displaced quality articles by subject matter experts. LSLRs must support scholarly communications without overwhelming it.

Publishers and the Information Industry

Commercial LLMs could “outflank publishers, building out platforms and tools . . . that would become scholars’ default way of interacting with the scholarly record.”⁸⁶ OpenAI has already announced plans to do this.⁸⁷ Publishers might simply become “white labeled” service providers for commercial LLMs. Gendron et al. call this “the colonization of academic publishing systems by AI technologies,” which can result in “the loss of negotiating powers in creating standards, values and norms of knowledge production.”⁸⁸ This platformization is an “outsourcing of libraries’ core mission to private actors,” deskilling staff, hollowing out information resources, and threatening institutional support.⁸⁹ As one scholarly publisher notes: “Publishers must stop thinking of themselves as providers of articles and start recognizing themselves as stewards of knowledge. You are not licensing text. You are licensing meaning.”⁹⁰

An alternative is to monetize not content but inference: “Usage is the inference stage: the moment an AI model applies its training to solve a specific user problem. Clients pay to have these problems solved, not for the training data itself.”⁹¹ Users are never presented with the underlying scholarly work, and “when value is created this way, revenue from that value does not automatically flow back to its sources.”⁹² Sam Altman, OpenAI CEO, concurs: “We see a future

where intelligence is a utility, like electricity or water, and people buy it from us on a meter.”⁹³ This impact was illustrated when Anthropic released a version of the Claude LLM customized for the legal industry and stock in legal publishing companies plummeted.⁹⁴ The overall objective should be “a thriving open environment with community publishing, review, and maintenance infrastructure that is supported by strong business models and collaborative arrangements,” but the balance of power is shifting from traditional partners to well-funded AI companies with different motivations.⁹⁵

Knowledge Systems

Khokhar argues that the “algorithmic economy” has driven a transition “from a traditional structured knowledge system to a non-structured data system, fundamentally altering how value was created, captured, and distributed.”⁹⁶ Research libraries curate knowledge through highly structured systems (e.g., databases, metadata conventions, classification, subject headings, and authority files), while ARI frameworks infer information from vast, unstructured data. Griffey calls this “category collapse,” where in AI discovery there is no distinction between metadata and data: “Either everything is metadata, or nothing is.”⁹⁷

The paradigm shift for libraries is from a model of “static collection–structured organization–passive utilization” to one of “semantic reconstruction–dynamic generation–collaborative value creation.”⁹⁸ New knowledge systems “should integrate mechanisms that support the verifiability, indexability, and reusability of generative content . . . and facilitate cumulative knowledge construction.”⁹⁹ This is an “epistemic infrastructure,” where libraries become “not just points of entry to services and collections” but are “directly supporting AI systems that learn from and reveal the semantic evolution of knowledge.”¹⁰⁰

The insecurity and manipulability of AI-generated content raise serious issues of provenance and authenticity. Unlike in archival collections, AI content can be compromised at any point in its lifecycle. The library, archives, and museum community has been encouraged to “actively shape standards, specifications, practice models, and technologies by asserting their requirements for open, vendor-agnostic, and trustworthy CAP [content authenticity and provenance] data.”¹⁰¹ Protocols for tracking and verification will need to be incorporated into local knowledge management systems, although libraries have “a systematic tendency to underinvest in digital research infrastructure.”¹⁰² The transition to new knowledge systems will require significant additional expenditures.

AI Literacy and Competencies

AI literacy has been widely adopted in academic libraries, with a variety of supporting frameworks.¹⁰³ However, Lo found, in a survey of academic library workers, a “readiness gap,” with a “significant proportion of participants” not frequently using AI tools and lacking the technical proficiency “required to navigate the rapidly evolving AI landscape confidently and competently.”¹⁰⁴ The Association of College and Research Libraries’ AI competencies document is strong on building knowledge and teaching skills but presents a passive view of library staff where roles as creators, designers, and research partners are absent.¹⁰⁵ If librarians are to perform the new roles discussed here, AI literacy should shift from learning and experimentation to application, enabling staff to “create and apply algorithmic techniques and tools to problems in a variety of domains.”¹⁰⁶

A significant movement in academic libraries advocates for AI refusal, characterized by “refusing to use AI tools entirely . . . refusing to accept either boosterism or doomer narratives, refusing to

accept the idea that AI is inevitable,” with the authors citing a “disconnect” between library and information science values and the harms of using generative AI.¹⁰⁷ Instead of AI refusal, critical AI literacy (CAIL) offers a corrective lens through “critical engagement with sociotechnical systems that reflects a comprehensive understanding of the limitations, biases, and potential harms of algorithmic decision-making.”¹⁰⁸ Aligning CAIL with research integrity principles such as honesty, scrupulousness, transparency, independence, and responsibility offers a framework for libraries to assess ARI.¹⁰⁹ AI literacy is “a foundational competency for researchers, coupling technical fluency with ethical and epistemic awareness,” moving beyond abstract understanding to “controlled exploration.”¹¹⁰ As researchers and research partners, librarians are advised to attain the technical proficiency noted by Lo.¹¹¹ One means to do this is with library-based AI labs.

AI Labs

Library-based AI labs have been established to “experiment, collaborate, take risks, sometimes fail, and always push boundaries.”¹¹² Notable examples include the KBLab at the National Library of Sweden, the LC Lab at the Library of Congress, the Library AI Lab at the University of Rhode Island, the Collaboratory at Toronto Metropolitan University, and the consortial [AI and Machine Learning Program](#) at the Ontario Council of University Libraries.¹¹³ The KBLab describes its work as “a public good” that can “enhance the library’s functions as national research infrastructure and guarantor of democratic values . . . [and] even counter some of the more problematic aspects of an AI future driven purely by private sector actors.”¹¹⁴ Such labs provide a proactive means to influence the direction of this technology, offering libraries “a blueprint for becoming central actors in the next phase of machine learning” and serving as incubators for future organizational structures.¹¹⁵

Cordell recommends a focus on “speculative interfaces” and projects that “serve as a diagnostic to the problems inherent in existing digital collections, a rebuttal to ML [machine learning] work falsely claiming objectivity, or a synecdoche that helps patrons better understand the historical stakes of library collections and archives.”¹¹⁶ Specific to ARI, Beel et al. encourage libraries to “actively engage”:

The AI Scientist relies on IR [information retrieval] techniques for literature search, data retrieval, citation analysis, and experimental design. As AI-driven research automation grows, the IR community has the opportunity to shape its methodologies. Addressing limitations in retrieval quality, document ranking, and automated citation analysis will be critical to ensuring AI-generated research is both credible and impactful.¹¹⁷

Because many ARI issues are sociotechnical, the agenda of AI labs should include policy deliberations on implications, risks, and mitigation strategies. The KBLab cautions that “it is far easier to start a lab . . . than it is to entrench one as a given part of a heritage organization,” recommending a focus on diverse expertise, an interdisciplinary outlook, realistic timeframes, specialized legal support, and continuing dialogue with stakeholders.¹¹⁸

CONCLUSION

ARI represents a profound change to the research enterprise. While still embryonic, early implementations illustrate how the nature of discovery, knowledge production, and scholarly communications is rapidly evolving. However revolutionary ARI might be, it remains an example of a human–AI co-creation grounded in the enduring norms of the scientific method. Frameworks such as The AI Scientist demonstrate that end-to-end autonomous scientific discovery is more than theory, although these frameworks still struggle with hallucinations, domain knowledge

limitations, narrow experimental options, and findings supported by limited evidence. For research libraries, responding will involve adjusting to a changing scholarly communication ecosystem, developing new knowledge management systems and expanded roles, deepening AI literacy and competencies, and moving toward an active research and development capacity with library-based AI labs. In these ways, libraries can remain central to the university mission and contribute to the larger societal discussion about the role of artificial intelligence.

In *The Promise of Artificial Intelligence*, Brian Cantwell Smith distinguishes reckoning from judgment.¹¹⁹ Reckoning is what ARI does. Judgment is what humans, and humans alone, do. It is a reminder that “AI does not happen to us: choices made by people determine its future.”¹²⁰ As the university and the library consider the implications of ARI, they should “resist both the Scylla of uncritical techno-optimism (which risks reducing epistemology to engineering) and the Charybdis of reactionary nostalgia (which would deny the genuinely novel epistemic possibilities AI affords).”¹²¹ As noted earlier, the question is not whether we will adapt but how and how effectively.

ENDNOTES

- ¹ Qiguang Chen et al., “AI4Research: A Survey of Artificial Intelligence for Scientific Research,” preprint, arXiv, July 2, 2025, <https://doi.org/10.48550/arXiv.2507.01903>.
- ² Tracy Bergstrom et al., *The Second Digital Transformation of Scholarly Publishing: Strategic Context and Shared Infrastructure* (New York: ITHAKA S+R, 2024), 22, doi.org/10.18665/sr.320210.
- ³ Shuo Ren et al., “Towards Scientific Intelligence: A Survey of LLM-Based Scientific Agents,” version 3, preprint, arXiv, February 2, 2026, <https://doi.org/10.48550/arXiv.2503.24047>; Mourad Gridach et al., “Agentic AI for Scientific Discovery: A Survey of Progress, Challenges, and Future Directions,” arXiv, preprint, March 12, 2025, <https://doi.org/10.48550/arXiv.2503.08979>; Jiaqi Wei et al., “From AI for Science to Agentic Science: A Survey on Autonomous Scientific Discovery,” preprint, arXiv, August 18, 2025, <https://doi.org/10.48550/arXiv.2508.14111>; Joeran Beel et al., “Evaluating Sakana’s AI Scientist: Bold Claims, Mixed Results, and a Promising Future?,” *SIGIR Forum* 59, no. 1 (2025): 1–20, <https://doi.org/10.1145/3769733.376974>; Chris Lu et al., “The AI Scientist: Towards Fully Automated Open-Ended Scientific Discovery,” preprint, arXiv, August 12, 2024, <https://doi.org/10.48550/arXiv.2408.06292>; Yutaro Yamada et al., “The AI Scientist-v2: Workshop-Level Automated Scientific Discovery via Agentic Tree Search,” preprint, arXiv, April 10, 2025, <https://doi.org/10.48550/arXiv.2504.08066>.
- ⁴ ARL/CNI 2035 Scenarios: *AI-Influenced Futures* (Washington, DC and West Chester, PA: Association of Research Libraries, Coalition for Networked Information, and Stratus Inc., 2024), <https://www.arl.org/resources/the-arl-cni-2035-scenarios-ai-influenced-futures-in-the-research-environment/>.
- ⁵ Karen Estlund and Cynthia Hudson Vitale, *Strategic Implications of AI Futures for Research Libraries: Workshop Report* (Chicago: Association of Research Libraries and Coalition for Networked Information, 2026), 24, <https://doi.org/https://doi.org/10.29242/report.futurescape2026>.
- ⁶ Manuel Alejandro Cruz-Aguilar, “The Epistemic Revolution of AI: Reconfiguring the Foundations of Scientific Knowledge,” *AI & Society* 41, no. 1 (2025): 14, <https://doi.org/10.1007/s00146-025-02658-3>.
- ⁷ Mark Daley, “Bartleby, the PI,” *Noetic Engines* (Substack), November 24, 2025, <https://noeticengines.substack.com/p/bartleby-the-pi>.
- ⁸ Tracy Bergstrom and Dylan Ruediger, *A Third Transformation? Generative AI and Scholarly Publishing* (New York: ITHAKA S+R, 2024) <https://doi.org/10.18665/sr.321519>; Bergstrom et al., *The Second Digital Transformation of Scholarly Publishing*.

- ⁹ Hanchen Wang et al., "Scientific Discovery in the Age of Artificial Intelligence," *Nature* 620, no. 7972 (2023): 57, <https://doi.org/10.1038/s41586-023-06221-2>; Elsevier, *Researcher of the Future* (Amsterdam: Elsevier, 2025), 5, <https://www.elsevier.com/insights/confidence-in-research/researcher-of-the-future>.
- ¹⁰ Guojie Li, "AI4R: The Fifth Scientific Research Paradigm," *Bulletin of the Chinese Academy of Sciences* 39, no. 1 (2024), <https://bulletinofcas.researchcommons.org/journal/vol39/iss1/1/>; Advanced Research and Invention Agency, *AI Scientist Call: Call for Research Proposals for Working with AI Scientist Systems* (Middlesex, UK: ARIA, 2025), 2, <https://www.aria.org.uk/media/b0jnlebb/ai-scientist-solicitation-call-document.pdf>.
- ¹¹ Beel et al., "Evaluating Sakana's AI Scientist."
- ¹² Maria Couto Teixeira et al., "Generative Artificial Intelligence in Qualitative Research Methods: Between Hype and Risks?," paper presented at the International Workshop on Chatbots and Human-Centred AI, Lübeck, Germany, November 12–13, 2025, <https://arxiv.org/abs/2511.08461v1>.
- ¹³ Tanisha Jowsey et al., "We Reject the Use of Generative Artificial Intelligence for Reflexive Qualitative Research," SSRN Scholarly Paper (Amsterdam: Social Science Research Network, 2025), <https://doi.org/10.2139/ssrn.5676462>.
- ¹⁴ Jennifer Listgarten, "The Perpetual Motion Machine of AI-Generated Data and the Distraction of ChatGPT as a 'Scientist,'" *Nature Biotechnology* 42, no. 3 (2024): 371–73, <https://doi.org/10.1038/s41587-023-02103-0>.
- ¹⁵ Graphite, "More Articles Are Now Created by AI Than Humans," October 2025, <https://graphite.io/five-percent/more-articles-are-now-created-by-ai-than-humans>; Jutta Haider et al., "GPT-Fabricated Scientific Papers on Google Scholar: Key Features, Spread, and Implications for Preempting Evidence Manipulation," *Harvard Kennedy School Misinformation Review* 5 (2024): 1–5, <https://doi.org/10.37016/mr-2020-156>; Kat Boboris, "Attention Authors: Updated Practice for Review Articles and Position Papers in arXiv CS Category," *arXiv* (blog), October 31, 2025, <https://blog.arxiv.org/2025/10/31/attention-authors-updated-practice-for-review-articles-and-position-papers-in-arxiv-cs-category/?ref=404media.co>; Bernhard Sabel et al., "Fake Publications in Biomedical Science: Red-Flagging Method Indicates Mass Production," *Naunyn-Schmiedeberg's Archives of Pharmacology* 399, no. 2 (2026): 2943–55, <https://doi.org/10.1007/s00210-025-04275-9>; Carlos Chaccour et al., "Robot Pen Pals: A Multidisciplinary Analysis of Recent Trends in Scientific Correspondence," preprint, Research Square, November 3, 2025, <https://doi.org/10.21203/rs.3.rs-7992675/v1>.
- ¹⁶ Haider et al., "GPT-Fabricated Scientific Papers on Google Scholar," 6; Sabel et al., "Fake Publications in Biomedical Science."
- ¹⁷ Hong Zhou, "Reimagining Scholarly Publishing Workflow: A High-Level Map of What Changes Next," *The Scholarly Kitchen*, November 20, 2025, <https://scholarlykitchen.sspnet.org/2025/11/20/reimagining-scholarly-publishing-workflow-a-high-level-map-of-what-changes-next/>.
- ¹⁸ Daley, "Bartleby, the PI."
- ¹⁹ Andrew W. Senior et al., "Improved Protein Structure Prediction Using Potentials from Deep Learning," *Nature* (London) 577, no. 7792 (2020): 706–10, <https://doi.org/10.1038/s41586-019-1923-7>; Cruz-Aguilar, "The Epistemic Revolution of AI," 2045.
- ²⁰ Lisa Messeri and M. J. Crockett, "Artificial Intelligence and Illusions of Understanding in Scientific Research," *Nature* (London) 627, no. 8002 (2024): 49–58, <https://doi.org/10.1038/s41586-024-07146-0>.
- ²¹ Ramón Alvarado, "AI as an Epistemic Technology," *Science and Engineering Ethics* 29, no. 5 (2023): 7, <https://doi.org/10.1007/s11948-023-00451-3>; John Wihbey, "AI and Epistemic Risk for Democracy: A Coming Crisis of Public Knowledge?," SSRN Scholarly Paper (Amsterdam: Social Science Research Network, 2024), 2, <https://doi.org/10.2139/ssrn.4805026>.
- ²² Wihbey, "AI and Epistemic Risk for Democracy," 2.

-
- ²³ Simon Chesterman and Loy Hui Chieh, "Research Integrity and Academic Authority in the Age of Artificial Intelligence: From Discovery to Curation?," preprint, arXiv, January 9, 2026, 3, <https://doi.org/10.48550/arXiv.2601.05574>.
- ²⁴ Wihbey, "AI and Epistemic Risk for Democracy," 5.
- ²⁵ Lu et al., "The AI Scientist."
- ²⁶ Yamada et al., "The AI Scientist-V2."
- ²⁷ Yamada et al., "The AI Scientist-V2," 9.
- ²⁸ Yamada et al., "The AI Scientist-V2," 13.
- ²⁹ Beel et al., "Evaluating Sakana's AI Scientist," 17.
- ³⁰ "Zochi Publishes A* Paper," *Intology* (blog), May 27, 2025, <https://www.intology.ai/blog/zochi-acl>.
- ³¹ Kyle Wiggers, "Academics Accuse AI Startups of Co-Opting Peer Review for Publicity," *TechCrunch*, March 19, 2025, <https://techcrunch.com/2025/03/19/academics-accuse-ai-startups-of-co-opting-peer-review-for-publicity/>.
- ³² "Zochi Publishes A* Paper"; Andy Zhou et al., *Zochi Technical Report* (San Francisco: Intology, 2025), 3, https://github.com/IntologyAI/Zochi/blob/main/Zochi_Technical_Report.pdf.
- ³³ "Previewing Locus," *Intology* (blog), November 19, 2025, <https://www.intology.ai/blog/previewing-locus>.
- ³⁴ "Previewing Locus."
- ³⁵ "Previewing Locus."
- ³⁶ Agents4Science, "Agents4Science 2025 Virtual Conference," October 22, 2025, YouTube, 03:19:22, <https://www.youtube.com/watch?v=7pXqAeedqOo>.
- ³⁷ Federico Bianchi et al., "Exploring the Use of AI Authors and Reviewers at Agents4Science," *Nature Biotechnology* 44 (2026): 12, <https://doi.org/10.1038/s41587-025-02963-8>.
- ³⁸ Bianchi et al., "Exploring the Use of AI Authors and Reviewers at Agents4Science," 12.
- ³⁹ Bianchi et al., "Exploring the Use of AI Authors and Reviewers at Agents4Science," 13.
- ⁴⁰ Dhruv Trehan and Paras Chopra, "Why LLMs Aren't Scientists Yet: Lessons from Four Autonomous Research Attempts," preprint, arXiv, January 6, 2026, <https://doi.org/10.48550/arXiv.2601.03315>.
- ⁴¹ Trehan and Chopra, "Why LLMs Aren't Scientists Yet," 10.
- ⁴² Mrinank Sharma et al., "Towards Understanding Sycophancy in Language Models," preprint, arXiv, October 20, 2023, <https://doi.org/10.48550/arXiv.2310.13548>.
- ⁴³ Trehan and Chopra, "Why LLMs Aren't Scientists Yet," 7.
- ⁴⁴ Trehan and Chopra, "Why LLMs Aren't Scientists Yet," 11.
- ⁴⁵ The Royal Society, *Science in the Age of AI: How Artificial Intelligence Is Changing the Nature and Method of Scientific Research* (London, UK: The Royal Society, 2024), 5, <https://royalsociety.org/-/media/policy/projects/science-in-the-age-of-ai/science-in-the-age-of-ai-report.pdf>.
- ⁴⁶ Wang et al., "Scientific Discovery in the Age of Artificial Intelligence."
- ⁴⁷ Kunlun Zhu et al., "SafeScientist: Toward Risk-Aware Scientific Discoveries by LLM Agents," preprint, arXiv, May 29, 2025, <https://doi.org/10.48550/arXiv.2505.23559>; Yoshua Bengio et al., "Superintelligent Agents Pose Catastrophic Risks: Can Scientist AI Offer a Safer Path?," preprint, arXiv, February 21, 2025, <https://doi.org/10.48550/arXiv.2502.15657>; Juraj Gottweis et al., "Towards an AI Co-Scientist," preprint, arXiv,

- February 26, 2025, <https://doi.org/10.48550/arXiv.2502.18864>; Haoyang Su et al., “Many Heads Are Better Than One: Improved Scientific Idea Generation by a LLM-Based Multi-Agent System,” paper presented at the 63rd Annual Meeting of the Association for Computational Linguistics, Vienna, Austria, July 27–August 1, 2025, <https://doi.org/10.48550/arXiv.2410.09403>; Maksim E. Eren and Dorianis M. Perez, “Rethinking Science in the Age of Artificial Intelligence,” preprint, arXiv, November 13, 2025, <https://doi.org/10.48550/arXiv.2511.10524>.
- ⁴⁸ Daley, “Bartleby, the PI.”
- ⁴⁹ Kevin Baker, “Context Widows,” *Artificial Bureaucracy* (Substack), December 12, 2025, <https://artificialbureaucracy.substack.com/p/context-widows>; Daley, “Bartleby, the PI.”
- ⁵⁰ Daley, “Bartleby, the PI.”
- ⁵¹ Richard Watermeyer et al., “Generative AI and the Automating of Academia,” *Postdigital Science and Education* (Cham) 6, no. 2 (2024): 461, <https://doi.org/10.1007/s42438-023-00440-6>.
- ⁵² Donald J. Trump, “Executive Order 14363, Launching the Genesis Mission,” *Federal Register*, November 24, 2025, <https://www.federalregister.gov/documents/2025/11/28/2025-21665/launching-the-genesis-mission>.
- ⁵³ Oliver Hsu, “Autonomous Labs Accelerate Scientific Discovery,” *A16z News*, December 10, 2025, <https://www.a16z.news/i/181164050/oliver-hsu-autonomous-labs-accelerate-scientific-discovery>.
- ⁵⁴ Adrien Abecassis et al., *A Blueprint for Multinational Advanced AI Development* (Oxford, UK: University of Oxford, 2025), <https://doi.org/10.5281/zenodo.17625842>.
- ⁵⁵ Abecassis et al., *A Blueprint for Multinational Advanced AI Development*.
- ⁵⁶ Wei et al., “From AI for Science to Agentic Science,” 48.
- ⁵⁷ Ryan Cordell, *Machine Learning + Libraries: A Report on the State of the Field* (Washington, DC: Library of Congress, 2020), 41, <https://labs.loc.gov/static/labs/work/reports/Cordell-LOC-ML-report.pdf>.
- ⁵⁸ Nuno Sousa, “Ethical and Practical Implications of AI in Academic Library Research,” *IFLA Journal*, ahead of print, 2025, <https://doi.org/10.1177/03400352251391753>; Martin Frické, *Artificial Intelligence and Librarianship: Notes for Teaching*, 3rd ed. (Wanaka, NZ: SoftOption, 2024), <https://softoption.us/sites/default/files/AlinLibrariesNotesForTeaching.pdf>; Lateef Ayinde et al., “Adoption of Artificial Intelligence in Academic Libraries: A Systematic Review of Current Practices, Challenges, and Research Opportunities,” *The Journal of Academic Librarianship* 52, no. 1 (2026): 1–11, <https://doi.org/10.1016/j.acalib.2025.103185>; Andrea Gasparini and Heli Kautonen, “Understanding Artificial Intelligence in Research Libraries: An Extensive Literature Review,” *LIBER Quarterly* 32, no. 1 (2022): 1–36, <https://doi.org/10.53377/lq.10934>; Sandy Hervieux and Amanda Wheatley, eds., *The Rise of AI: Implications and Applications of Artificial Intelligence in Academic Libraries* (Chicago: Association of College and Research Libraries, 2022).
- ⁵⁹ Lorcan Dempsey, “Generative AI and Libraries: 7 Contexts,” LorcanDempsey.Net, November 12, 2023, <https://www.lorcandempsey.net/generative-ai-and-libraries-7-contexts/>; Leo S. Lo, “The CARE Approach for Academic Librarians: From Search First to Answer First with Generative AI,” *The Journal of Academic Librarianship* 52, no. 1 (2026): 3, <https://doi.org/10.1016/j.acalib.2025.103186>.
- ⁶⁰ Masud Khokhar, “Imagining the Future Library,” *Katina*, October 30, 2025, 1, <https://katinamagazine.org/content/article/future-of-work/2025/imagining-the-future-library>.
- ⁶¹ Chris Bourg, “What Happens to Libraries and Librarians When Machines Can Read All the Books?,” *Feral Librarian* (blog), March 17, 2017, <https://chrisbourg.wordpress.com/2017/03/16/what-happens-to-libraries-and-librarians-when-machines-can-read-all-the-books/>.
- ⁶² ARL/CNI 2035 Scenarios: AI-Influenced Futures, 27.

- ⁶³ Mark Daley, "Rethinking the Role of Higher Education in an AI-Integrated World," *University Affairs*, February 6, 2026, <https://universityaffairs.ca/opinion/rethinking-the-role-of-higher-education-in-an-ai-integrated-world/>.
- ⁶⁴ Mark Daley, "Rethinking the Role of Higher Education in an AI-Integrated World"; Chesterman and Chieh, "Research Integrity and Academic Authority in the Age of Artificial Intelligence," 20.
- ⁶⁵ Shyam Sharma, "Universities and AI: Rebuilding the Commons," *The RhetAI Coalition* (Substack), January 16, 2026, <https://rhetaicoalition.substack.com/p/universities-and-ai-rebuilding-the>.
- ⁶⁶ ARL/CNI 2035 Scenarios: *AI-Influenced Futures*, 30.
- ⁶⁷ Estlund and Hudson Vitale, *Strategic Implications of AI Futures for Research Libraries*, 27.
- ⁶⁸ Sousa, "Ethical and Practical Implications of AI in Academic Library Research," 15.
- ⁶⁹ Catherine Coleman and Michael A. Keller, "AI in the Research Library Environment," in *Artificial Intelligence in Libraries and Publishing*, edited by Pickering Ruth and Matthew Ismail (London: ATG Media, 2022), 31, <https://doi.org/10.3998/mpub.12669942>.
- ⁷⁰ Coleman and Keller, "AI in the Research Library Environment"; Bergstrom and Ruediger, *A Third Transformation?*; Bourg, "What Happens to Libraries and Librarians When Machines Can Read All the Books?"; Khokhar, "Imagining the Future Library."
- ⁷¹ Bergstrom and Ruediger, *A Third Transformation?*, 22.
- ⁷² Chesterman and Chieh, "Research Integrity and Academic Authority in the Age of Artificial Intelligence," 18.
- ⁷³ Khokhar, "Imagining the Future Library," 3.
- ⁷⁴ Bourg, "What Happens to Libraries and Librarians When Machines Can Read All the Books?"; Bergstrom and Ruediger, *A Third Transformation?*; Chesterman and Chieh, "Research Integrity and Academic Authority in the Age of Artificial Intelligence"; Kari D. Weaver, "Artificial Intelligence and Transparency: Toward a Framework for Disclosure of AI Use in Learning, Research, and Publication," *Information Services and Use* 46, nos. 1–2 (2026): 1–7, <https://doi.org/10.1177/18758789261435716>.
- ⁷⁵ Andrew Cox, "How Artificial Intelligence Might Change Academic Library Work: Applying the Competencies Literature and the Theory of the Professions," *Journal of the Association for Information Science and Technology* 74, no. 3 (2023): 367–80, <https://doi.org/10.1002/asi.24635>.
- ⁷⁶ Sousa, "Ethical and Practical Implications of AI in Academic Library Research," 15.
- ⁷⁷ Baker, "Context Widows."
- ⁷⁸ Bergstrom et al., *The Second Digital Transformation of Scholarly Publishing*, 14.
- ⁷⁹ ARL/CNI 2035 Scenarios: *AI-Influenced Futures*, 26.
- ⁸⁰ Lorcan Dempsey, "Generative AI, Scholarly and Cultural Language Models, and the Return of Content," LorcanDempsey.Net, June 28, 2023, <https://www.lorcandempsey.net/generative-ai-a-note-about-content/>.
- ⁸¹ Bergstrom and Ruediger, *A Third Transformation?*, 21.
- ⁸² Andrew Cox and Mike Thelwall, *AI for Knowledge* (Boca Raton, FL: CRC Press, 2025), 46, <https://doi.org/10.1201/9781003545163>; Bernhard Sabel and Dan Larhammar, "Reformation of Science Publishing: The Stockholm Declaration," *Royal Society Open Science* (England) 12, no. 11 (2025): 4, <https://doi.org/10.1098/rsos.251805>.
- ⁸³ Sarah G. Park, "AI and Systematic Reviews: Can AI Tools Replace Librarians in the Systematic Search Process?," *Science & Technology Libraries* 45, no. 2 (2025): 21, <https://doi.org/10.1080/0194262X.2025.2521519>; Ethan J. Gantana et al., "An In-Depth Analysis of AI-Generated Scientific Review Articles and Its Potential Implications on

- the Future of Medical Journal Publications,” *Next Research* 2, no. 4 (2025): 1, <https://doi.org/10.1016/j.nexres.2025.101002>.
- ⁸⁴ Gwern Branwen, “LLM Daydreaming,” Gwern.Net, July 14, 2025, <https://gwern.net/ai-daydreaming>; Ulkar Aghayeva, “Looking for Hidden Gems in Scientific Literature: The Promise and Practice of Literature-Based Discovery,” *Elicit*, October 20, 2025, <https://elicit.com/blog/literature-based-discovery>.
- ⁸⁵ Boboris, “Attention Authors: Updated Practice for Review Articles and Position Papers in arXiv CS Category.”
- ⁸⁶ Bergstrom and Ruediger, *A Third Transformation?*, 12.
- ⁸⁷ Will Douglas Heaven, “OpenAI Is Throwing Everything into Building a Fully Automated Researcher,” *MIT Technology Review*, March 20, 2026, <https://www.technologyreview.com/2026/03/20/1134438/openai-is-throwing-everything-into-building-a-fully-automated-researcher/>.
- ⁸⁸ Yves Gendron et al., “The Perils of Artificial Intelligence in Academic Publishing,” *Critical Perspectives on Accounting* 87 (2022): 2, <https://doi.org/10.1016/j.cpa.2021.102411>; Lai Ma, “The Platformisation of Scholarly Information and How to Fight It,” *LIBER Quarterly: The Journal of the Association of European Research Libraries* 33, no. 1 (2023): 2, <https://doi.org/10.53377/lq.13561>.
- ⁸⁹ Jean-Christophe Plantin and Andrea Thomer, “Platforms, Programmability, and Precarity: The Platformization of Research Repositories in Academic Libraries,” *New Media & Society* 27, no. 1 (2025): 338, <https://doi.org/10.1177/14614448231176758>.
- ⁹⁰ Veronica Showers et al., “AI Readiness and the New Value Equation in Scholarly Publishing,” *The Scholarly Kitchen*, February 24, 2026, <https://scholarlykitchen.sspnet.org/2026/02/24/guest-post-ai-readiness-and-the-new-value-equation-in-scholarly-publishing/>.
- ⁹¹ Jonathan Woahn, “AI Isn’t Going to Pay for Content ... Part Two: The Path Forward,” *The Scholarly Kitchen*, January 29, 2026, <https://scholarlykitchen.sspnet.org/2026/01/29/guest-post-ai-isnt-going-to-pay-for-content-part-two-the-path-forward/>.
- ⁹² Woahn, “AI Isn’t Going to Pay for Content.”
- ⁹³ A. J. Dellinger, “Sam Altman Says Intelligence Will Be a Utility, and He’s Just the Man to Collect the Bills,” *Gizmodo*, March 12, 2026, <https://gizmodo.com/sam-altman-says-intelligence-will-be-a-utility-and-hes-just-the-man-to-collect-the-bills-2000732953>.
- ⁹⁴ Joe Castaldo, “Anthropic’s Release of AI Tools for Lawyers Prompts Massive Legal-Tech Sell-Off,” *The Globe and Mail*, February 3, 2026, <https://www.theglobeandmail.com/business/article-anthropics-release-of-ai-tools-for-lawyers-prompts-massive-sell-off/>.
- ⁹⁵ Katherine Skinner et al., *The State of Open Research Software Infrastructure* (Invest in Open Infrastructure, 2025), 3, <https://doi.org/10.5281/zenodo.14888589>.
- ⁹⁶ Khokhar, “Imagining the Future Library.”
- ⁹⁷ Jason Griffey, “Predictions about Future Technology in Libraries and Epistemic Collapse: Laws and Mindsets,” in *Library 2035: Imaging the Next Generation of Libraries*, edited by Sandra Hirsh (Lanham, MD: Rowman & Littlefield, 2024), 15.
- ⁹⁸ Yu Zhu et al., “How Does Information Resource Value of AI-Generated Content Emerge? An Exploratory Study from the User Evaluation Perspective,” *Journal of the Association for Information Science and Technology* 77, no. 5 (2026): 1–18, <https://doi.org/10.1002/asi.70061>.
- ⁹⁹ Zhu et al., “How Does Information Resource Value of AI-Generated Content Emerge?”
- ¹⁰⁰ James Lee and Xuemao Wang, “Libraries as Research Engines: AI Agents for Discovery through Language, Inference, and Network Structure,” paper presented at the Fantastic Futures Conference, London, United

- Kingdom, December 3–5, 2025, https://www.conftool.org/fantastic-futures-2025/index.php?page=browseSessions&form_session=6.
- 101 Kate Murray and Joshua Sternfeld, *Content Authenticity and Provenance in the Age of Artificial Intelligence: A Call-to-Action for the LAMs Community* (C2PA for G+LAM Community of Practice, 2026), 4, <https://tinyurl.com/bdej7rd9>.
 - 102 Love Börjeson et al., “Transfiguring the Library as Digital Research Infrastructure: Making KBLab at the National Library of Sweden,” *College & Research Libraries* 85, no. 4 (2024): 578, <https://doi.org/10.5860/crl.85.4.564>.
 - 103 Sandy Hervieux and Amanda Wheatley, *Building an AI Literacy Framework: Perspectives from Instruction Librarians and Currently Information Literacy Tools* (Chicago: American Library Association, 2024), <https://www.choice360.org/research/white-paper-building-an-ai-literacy-framework-perspectives-from-instruction-librarians-and-current-information-literacy-tools/>; Melanie Hibbert et al., “A Framework for AI Literacy,” *EDUCAUSE Review: Emerging Technologies and Trends*, June 3, 2024, <https://er.educause.edu/articles/2024/6/a-framework-for-ai-literacy>; Leo S. Lo, “AI Literacy: A Guide for Academic Libraries,” *College & Research Libraries News* 86, no. 3 (2025): 120–22, <https://doi.org/10.5860/crln.86.3.120>.
 - 104 Leo S. Lo, “Evaluating AI Literacy in Academic Libraries: A Survey Study with a Focus on U.S. Employees,” *College & Research Libraries* 85, no. 5 (2024): 649, <https://doi.org/10.5860/crl.85.5.635>.
 - 105 *AI Competencies for Academic Library Workers* (Chicago: Association of College and Research Libraries, 2025), https://www.ala.org/sites/default/files/2025-10/acrl_ai_competencies.pdf.
 - 106 Michael Ridley and Danica Pawlick-Potts, “Algorithmic Literacy and the Role for Libraries,” *Information Technology and Libraries* 40, no. 2 (2021): 4, <https://doi.org/10.6017/ital.v40i2.12963>.
 - 107 Kailyn Slater, “Against AI: Critical Refusal in the Library,” *Library Trends* 73, no. 4 (2025): 588–608, <https://doi.org/10.1353/lib.2025.a968497>; Violet Fox, “AI Refusal in Libraries: A Starter Guide,” *ACRLog* (blog), June 11, 2025, <https://acrlog.org/2025/06/11/ai-refusal-in-libraries-a-starter-guide/>.
 - 108 Xiaoxu Ling and Siyuan Yan, “Understanding Algorithmic Aversion: A Bold Rejection of Digital Dominion,” *Communications of the ACM* 69, no. 3 (2026): 28, <https://doi.org/10.1145/3758088>.
 - 109 Olivia Guest et al., “Against the Uncritical Adoption of ‘AI’ Technologies in Academia,” preprint, PhilArchive, 2025, <https://philarchive.org/archive/GUEATU>.
 - 110 Eren and Perez, “Rethinking Science in the Age of Artificial Intelligence,” 4.
 - 111 Lo, “Evaluating AI Literacy in Academic Libraries.”
 - 112 M. Mahey et al., *Open a GLAM Lab: Digital Cultural Heritage Innovation Labs* (Doha, Qatar: Book Sprint, 2019), <https://assets.pubpub.org/4kaopqb/31574367506714.pdf>.
 - 113 Börjeson et al., “Transfiguring the Library as Digital Research Infrastructure”; Abigail Potter, “Introducing the LC Labs Artificial Intelligence Planning Framework,” *The Signal*, November 15, 2023, <https://blogs.loc.gov/thesignal/2023/11/introducing-the-lc-labs-artificial-intelligence-planning-framework/>; Harrison Dekker et al., “URI Libraries’ AI Lab—Evolving to Meet the Needs of Students and Research Communities,” in *The Rise of AI: Implications and Applications of Artificial Intelligence in Academic Libraries*, edited by Sandy Hervieux and Amanda Wheatley (Chicago: Association of College and Research Libraries, 2022); Fangmin Wang et al., “Incubating AI: The Collaboratory at Ryerson University Library,” in *The Rise of AI: Implications and Applications of Artificial Intelligence in Academic Libraries*, edited by Sandy Hervieux and Amanda Wheatley (Chicago: Association of College and Research Libraries, 2022).
 - 114 Börjeson et al., “Transfiguring the Library as Digital Research Infrastructure,” 576.
 - 115 Brian Mathews, “A Cambrian Moment for Libraries: Shaping Futures through Evolution and Leadership,” *Portal: Libraries and the Academy* 25, no. 2 (2025): 231–50, <https://muse.jhu.edu/pub/1/article/955944>; Gasparini and

Kautonen, "Understanding Artificial Intelligence in Research Libraries"; Lee and Wang, "Libraries as Research Engines."

¹¹⁶ Cordell, *Machine Learning + Libraries*, 16.

¹¹⁷ Beel et al., "Evaluating Sakana's AI Scientist," 18.

¹¹⁸ Börjeson et al., "Transfiguring the Library as Digital Research Infrastructure," 578.

¹¹⁹ Brian Cantwell Smith, *The Promise of Artificial Intelligence: Reckoning and Judgment* (Cambridge, MA: MIT Press, 2019).

¹²⁰ Bengio et al., "Superintelligent Agents Pose Catastrophic Risks," 14.

¹²¹ Cruz-Aguilar, "The Epistemic Revolution of AI," 2054.