Do Artifacts Still Have Politics? Technological Determinism and Professional Agency in Academic Libraries' AI Transformation

Stephen Zweibel

abstract: Academic libraries face a generative artificial intelligence (AI) transformation that differs fundamentally from earlier database automation. While 1980s systems offered transparent logic and local control, contemporary AI operates through opaque neural networks and concentrated corporate power. Drawing on Andrew Abbott's professional jurisdiction theory, Langdon Winner's technological politics framework, and the Social Construction of Technology, this article argues that generative AI's architecture embeds political arrangements that threaten library values, yet libraries retain agency at the implementation level. The proposed critical implementation framework emphasizes transparency, collective action, mission-alignment, and flexibility as strategic pillars for selective AI adoption that preserves professional jurisdiction and epistemological clarity about AI's limits.

Introduction

Academic libraries are grappling with rapid transformation as generative artificial intelligence (AI) reshapes information work. Large language models and related systems now produce fluent text, images, and code from statistical patterns, often without exposing a stable evidentiary trail. What reads to users like a conversational search is, in practice, probabilistic generation rather than information retrieval.

portal: Libraries and the Academy, Vol. 26, No. 3 (2026), pp. 495–512.

There is a striking mismatch in the pace of individual versus institutional generative AI adoption. Although 81 percent of researchers report having incorporated large language models into their workflows, institutional adoption in libraries is nascent. Clarivate's 2024 Pulse Library survey (vendor-conducted) finds only seven percent currently implementing generative AI features, even as 43 percent name AI tools their top priority for the next year and over 60 percent are evaluating or planning integration.[1] This divergence between individual uptake and institutional deployment signals a widening service gap that librarians need to address.

This gap is evident in the university's research and teaching infrastructure, where academic libraries reside. Libraries license and organize the scholarly record, maintain discovery and metadata systems that make scholarship findable, teach information literacy and academic integrity, and participate in campus debates over publishing, data governance, and intellectual property. Generative AI is already pressing on each of these domains. It shifts how students search and write, how researchers review literature and draft arguments, how vendors redesign discovery around conversational assistants, and how institutions negotiate ownership over scholarly data and outputs. Library decisions about generative AI therefore shape the knowledge infrastructure of the academy.

Langdon Winner asked whether technologies can embody specific forms of power and authority; that is, whether artifacts have politics.[2] Generative AI brings that problem into sharp relief for libraries. Contemporary systems are built on opaque models, developed within concentrated vendor ecosystems, and deployed through platforms that reach users directly. Training data, model behavior, and interface defaults are set upstream, while local institutions encounter the technology after many consequential decisions have already been made.

The technical pressures are immediate. Language models exhibit systematic reliability problems that challenge core library functions, including frequent citation fabrication.[3] Fabrication at this scale destabilizes reference work and instruction built around verifiable sources.

Libraries face a strategic choice about how to meet the moment. They can become passive consumers of commercial AI services or active participants in shaping implementation around professional values.

The organizational pressures are equally tangible. Leo Lo's national survey of 760 academic library employees finds that most feel unprepared for generative AI, with respondents emphasizing the need for comprehensive training and ethical guidelines.[4] Clarivate's survey confirms these barriers: lack of expertise is identified as the greatest concern related to AI adoption, with 32 percent of respondents reporting that institutional training is unavailable.[5] Professional capacity is developing on the fly, in the middle of campus-wide shifts already underway.

Libraries face a strategic choice about how to meet the moment. They can become passive consumers of commercial AI services or active participants in shaping implementation around professional values. This choice will determine whether technological constraints become deterministic or remain spaces for professional agency.

Some advocate complete withdrawal. Kailyn Slater argues for "active refusal of generative AI" as necessary resistance to embedded corporate control.[6] Yet refusal risks jurisdictional surrender. If librarians refuse to engage with AI implementation, other professional groups will claim authority over implementation, or no one will, and professional values simply will not be represented. Students already rely on ChatGPT for schoolwork, while instruction adapts unevenly.[7] And vendors offer campus-wide deployment options (for example, ChatGPT Edu), with little to no input from librarians.[8]

This article develops "critical implementation" as a framework for navigating this tension. The approach emphasizes selective adoption, transparency requirements, labor protections, and epistemological clarity about generative AI systems' limits.

Theoretical Framework: Understanding Technology and Professional Power

Debates about technology typically swing between technological determinism and social construction. Determinist perspectives treat technology as an autonomous force that imposes its logic on society. The Social Construction of Technology (SCOT) counters that technologies are shaped through negotiation among relevant social groups, exhibiting interpretative flexibility until dominant designs stabilize. Neither approach alone captures contemporary AI adoption in libraries, which involves both powerful technical drivers and institutional choices about implementation. Three frameworks help us locate where professional agency persists for librarians.

Abbott's Professional Jurisdiction

Andrew Abbott frames professional control as jurisdiction claimed in three arenas: legal (that is, formal, structural, and rule-based authority), public, and workplace. Abbott's framework treats jurisdiction as something that must be actively claimed and defended. It does not inhere in the work itself. Two professions may contest the same territory, and the outcome depends on organization, timing, and the ability to frame the work as belonging to a particular kind of expertise. When a profession fails to claim jurisdiction, the vacuum is filled—by another profession, by non-professionals, or by automated systems that render the question moot.[9]

Technologies become jurisdictional contests when they reorganize who does intellectual work and who receives credit for it.[10] The introduction of paralegals and document automation reshaped legal practice. The rise of nurse practitioners and diagnostic software reshaped medicine. Libraries have faced this phenomenon before. What remains uncertain is whether they will claim jurisdiction over AI implementation or cede it to vendors, IT departments, or no one in particular.

Winner's Politics: Foundation and Implementation

Langdon Winner asks whether technologies can embody specific forms of power and authority—whether artifacts have politics. His famous example is Robert Moses's Long Island parkway bridges, designed too low for public buses and thus effectively excluding poor and minority residents who depended on public transit from Jones Beach. The politics there were built into the concrete.[11]

Winner distinguishes between two ways artifacts can have politics. Some technologies are compatible with many social arrangements and can be shaped through use—their politics emerge from deployment decisions. Others effectively require specific political arrangements—centralized control, particular distributions of power, and certain forms of expertise. Nuclear power, in Winner's analysis, tends toward the latter; it requires security regimes, concentrated authority, and expert management regardless of who owns the plant.

The diagnostic task for any given artifact is identifying which decisions remain open and which have been foreclosed. For libraries, this means locating where intervention remains possible when so much has been decided upstream.

Most technologies fall somewhere between these poles. The diagnostic task for any given artifact is identifying which decisions remain open and which have been foreclosed. For libraries, this means locating where intervention remains possible when so much has been decided upstream. Deployment decisions may be constrained by upstream choices about training data, model behavior, and platform governance. Or they may meaningfully shape how the technology operates in practice. Identifying which is which is the diagnostic task.

SCOT and Interpretive Flexibility

The Social Construction of Technology (SCOT) argues that technologies possess interpretive flexibility. Their meaning and use remain open to negotiation among relevant social groups until closure occurs and dominant interpretations stabilize.[12] The bicycle, in the canonical SCOT example, meant different things to different groups. For instance, it was a sporting vehicle for young men, a danger to pedestrians, and a tool for women's mobility, until the "safety bicycle" design stabilized and closed off alternative developmental paths.[13]

Whether a group can shape a technology depends on whether it constitutes a "relevant social group" with sufficient power to influence development or implementation. Relevance requires organization, resources, and the ability to articulate interests in ways that matter to other actors. Consumer groups, professional associations, regulatory bodies, and technical communities all compete to shape how technologies develop. Groups that fail to organize find the technology shaped without their input.[14]

SCOT directs attention to the spaces where interpretation remains possible—where technologies can still have different meanings and be configured for different purposes. When closure happens at one layer, flexibility may persist at others. A locked-down operating system may still permit application-level customization. A proprietary database may still allow local configuration. The question is whether the remaining flexibility permits meaningful intervention.

Application

These three frameworks ask related but distinct questions. Abbott asks what is at stake: jurisdiction, professional survival, or the authority to define how work gets done. If libraries do not claim implementation decisions as professional territory, someone else will, or the decisions will be made by default, without professional input. Winner asks where agency lives. In other words, at what point in the technology's development and deployment can values be embedded or contested? Some choices are foreclosed by the time libraries encounter the technology. Others remain open. SCOT asks how agency operates, through interpretive flexibility exercised by groups with sufficient power to matter. Libraries must constitute themselves as a relevant social group through collective action, shared standards, and organized intervention, or find the technology shaped without them.

. . . at what point in the technology's development and deployment can values be embedded or contested?

Together, these frameworks suggest that professional agency depends on conditions: the characteristics of the technology, the organization of the profession, and the timing of intervention. The historical analysis that follows applies these frameworks to two periods marked by great technological change in libraries: database automation (1980–1995) and the current AI transformation (2022–present). The comparison reveals how different technological characteristics create different conditions for professional agency, and what that means for libraries navigating generative AI.

Historical Analysis Part I: The Database Automation Era, 1980-1995

Technological Characteristics and Professional Control

Libraries' navigation of database automation was neither straightforward nor inevitable. Three characteristics of that era favored professional control, but each required effort to secure. These were the transparency of underlying systems, the competitiveness of vendor markets, and the feasibility of building local technical capacity.

Hard-won Transparency

Online catalogs used Boolean logic and structured query languages, which made search logic inspectable and teachable. Librarians could examine how "author AND subject" queries worked, modify indexing rules, and explain search strategies to users. In practice, achieving this transparency required sustained professional effort. Contemporary reports document confusion with vendor documentation, heterogeneous standards, and significant retraining demands as librarians learned to administer integrated systems.[15] But crucially, the underlying logic remained accessible to professional analysis. Communities of practice formed around SQL debugging, Z39.50 troubleshooting, and MARC field interpretation.[16] By the mid-1990s, systems librarians could audit vendor claims about search algorithms and modify local implementations accordingly.[17]

Contested Market Competition

Multiple commercial and cooperative vendors served libraries, creating meaningful choice and competitive pressure. Librarians could evaluate and select among offerings for database access, often using competition to secure better terms.[18] While consolidation pressures and proprietary formats created messes for librarians to clean up, the competitive landscape gave libraries leverage that is absent in today's concentrated generative AI market. Clifford A. Lynch's analysis documents how consortia, including OCLC, Research Libraries Group, and the Library of Congress, collaboratively shaped the Z39.50 standard from 1988 through 1995, creating interoperability testbeds where libraries could verify vendor claims about standards compliance in real-time demonstrations.[19] This empowered professionals to hold vendors accountable for technical specifications.

Gradual Local Capacity

By the early 1990s, many research libraries had established systems departments and employed systems librarians who could modify and maintain local implementations; evidence of staffing growth over the period underscores this institutional capacity-building.[20] Building this capacity required sustained institutional investment, but the technical architecture made it feasible for individual libraries to develop expertise. Battin's pioneering work at Columbia, where she combined library administration with information technology services in 1978, demonstrated how libraries could claim technical authority rather than depending entirely on vendors and other institutional forces.[21] This organizational model spread as libraries recognized that local technical capacity provided leverage in vendor negotiations and control over system behavior.

Labor Market Transformation

Automation reshaped work content and staffing patterns in ways that suggest deskilling pressures. Contemporary accounts warned that "one of the great dangers of automation is that we will become technicians," as jobs became more routine, with increased monitoring and reduced professional discretion.[22] Employment data bear this out: paraprofessional roles expanded more quickly than professional positions, with library technician and assistant ranks growing by 40 percent and 27 percent, respectively, between 1988 and 2000, compared to just 12 percent growth in number of librarians.[23] This pattern suggests work was redistributed to lower-skill classifications as automation standardized previously professional tasks.

Despite institutional change, libraries preserved core knowledge practices around information retrieval. Database systems operated within what Michael Buckland termed the "information-as-thing" paradigm, organizing existing objects for retrieval rather than generating new content, as modern generative AI does.[24] This preserved the epistemological foundations of librarianship even as interfaces transformed.

Historical Analysis Part II: The Current AI Transformation, 2022-Present

Structural Constraints Limiting Control

The theoretical frameworks outlined above help illuminate three interrelated constraints that distinguish the current moment from earlier automation. These constraints create what might be called a jurisdictional emergency. Unlike the 1980s, when libraries could contest Abbott's three arenas sequentially — establishing legal authority, then public legitimacy, then workplace control — AI threatens all three simultaneously. In the legal/formal control arena, opaque models prevent professionals from auditing algorithmic logic.[25] In the public arena, users bypass librarians by querying general-purpose chatbots directly.[26] In the workplace, vendors embed AI features into library platforms before professional evaluation can occur. The compressed timeframe leaves little opportunity for gradual adaptation.

Algorithmic Opacity

Modern language models rely on multi-layer neural architectures with billions of parameters that resist professional analysis. Unlike Boolean queries that librarians could debug step by step, neural networks produce outputs through statistical associations across high-dimensional spaces that defy human comprehension.[27] Even with access to weights and code, individual outputs cannot be traced to specific training inputs or decision rules. This represents a fundamental shift from the 1980s, when systems librarians could examine SQL execution plans and modify indexing algorithms. Contemporary AI offers no equivalent to Z39.50 testbeds, where professionals could verify vendor claims about system behavior.

Market Concentration

The capital intensity of foundation model development creates structural barriers that eliminate the competitive leverage libraries once possessed. Unlike the 1980s landscape where libraries could negotiate between DIALOG, BRS, and OCLC, contemporary markets offer few meaningful alternatives at the foundation model level. Vertical integration compounds this problem by allowing model developers to control downstream applications, reducing libraries to passive consumers rather than negotiating partners.[28]

External Development and Misaligned Incentives

Database systems emerged from library and information science contexts, incorporating professional norms around accuracy and provenance. Contemporary AI development often emphasizes engagement and growth metrics rather than evidentiary practices.[29] As noted earlier, systemic citation fabrication rates make current models unsuitable for reference work without significant professional mediation. Generative AI vendor priorities have shifted from competing on precision toward optimizing fluency and user satisfaction.

That said, deployment architectures create intervention points that the foundation models themselves foreclose. LLM systems separate the base model from data access layers, parameter controls, and system integrations.[30] A library cannot change how a

model was trained, but it can decide what documents to feed into a retrieval-augmented generation system, what temperature settings to use, what guardrails to implement, and how to present outputs to users. This is implementation flexibility, interpretive space that reopens at deployment what was closed at the foundation.

Contemporary Implementation Patterns and Institutional Responses

Libraries are experimenting with a range of approaches within these constraints, from institution-wide platform negotiations to targeted local pilots, organic chatbot development, and coordinated professional governance.[31]

Campus-wide Platforms

Some institutions negotiate managed access to generative AI while preserving institutional control. The University of Michigan built campus-specific AI tools rather than accepting vendor terms, addressing "privacy, accessibility, and affordability questions that a public institution of higher education has to account for."[32] This represents selective adoption based on institutional values rather than vendor convenience. While OpenAI's ChatGPT Edu provides some administrative controls, Michigan's approach shows how libraries can work with IT to shape deployment conditions rather than passively consuming commercial offerings.[33]

Some institutions negotiate managed access to generative AI while preserving institutional control.

Local Collaborations and Pilots

Other libraries pursue targeted projects that put epistemological clarity about information sources into practice. Northwestern University Libraries received IMLS funding to build semantic discovery over verified digitized collections, using collection metadata to ground AI responses in verified institutional holdings.[34] Yale University Library developed systematic evaluation protocols, testing multiple LLMs against the same digitized texts to compare outputs and assess reliability before broader deployment.[35] Both projects prioritize collection integrity over engagement metrics, reshaping AI systems around professional values rather than accepting generic implementations.

Organic Growth through Chatbots

San José State University's approach shows how transparency requirements can evolve incrementally. Their KingbotGPT evolved from providing generic responses to grounding in a "local dataset curated by SJSU librarians and library staff," gradually establishing professional control over content and sourcing decisions.[36] This evolution demonstrates how existing systems can be reshaped toward library values even within vendor frameworks. Some institutions are creating new roles that explicitly claim jurisdiction over AI strategy. For example, San José State University appointed its first dedicated AI Librarian "to lead the integration and development of AI technologies for an academic

library."[37] This is Abbott's jurisdiction in action: claiming authority by creating positions that assert it.

Professional Governance and Ethics

Library associations are establishing frameworks that put critical implementation into practice at scale. The Association of Research Libraries adopted Guiding Principles for Artificial Intelligence that emphasize privacy, equity, transparency, and human oversight, providing member institutions with concrete guidance for evaluating generative AI proposals.[38] The International Coalition of Library Consortia articulated principles for AI clauses in licensing agreements, giving consortia negotiating language to preserve institutional control over vendor AI features.[39] International bodies like the International Federation of Library Associations and Institutions (IFLA) have issued frameworks for responsible AI adoption that reinforce professional commitments to equitable access and information integrity. These coordinated efforts demonstrate how professional associations can contest vendor control through collective action, establishing alternative evaluation criteria that prioritize library values over commercial metrics. There are signs of professional readiness for this work. Many surveyed librarians self-identify as innovators or early adopters with respect to AI.[40] Now it remains to be seen whether individual inclination can be channeled into collective capacity.

These coordinated efforts demonstrate how professional associations can contest vendor control through collective action, establishing alternative evaluation criteria that prioritize library values over commercial metrics.

Epistemological Disruption: From Retrieval to Generation

The shift from information retrieval to generation is epistemologically significant. As Birger Hjørland argues, retrieving known items and generating new text are distinct knowledge activities.[41] Libraries promise accurate, unbiased, and authentic information.[42] Generative systems, by contrast, produce fluent text based on statistical associations. Practitioners report that such outputs often combine plausible phrasing with subtle inaccuracies that are difficult to detect without domain expertise.[43]

Rather than information-as-thing, AI operates through what may be termed "information-as-probability," generating new content through statistical pattern matching. This represents what Thomas Kuhn would recognize as a paradigm shift, a fundamentally different conceptual framework, one that is less constrained by evidence.

Empirical studies corroborate these concerns. Katie Lai evaluated ChatGPT-3.5 on academic reference questions and, using a three-point rubric, found lowest accuracy on the most complex questions, which scored 1.50 out of 3 on average, and higher accuracy on simple ones, which scored 2.67 out of 3.[44] A survey of hallucinations in large language models reports that outputs can be fluent yet factually unsupported, with inaccuracies often difficult to detect.[45] William H. Walters and Tyler Wilder's bibliography experiment showed fabricated citations even in the most advanced general model tested.[46] These

results indicate that generative systems require new forms of verification and instruction to prevent the unintentional spread of error.

Labor Market Implications and Collective Response

Generative AI is already altering library work. Task-exposure analyses suggest heavy automation pressure on routine roles and significant redesign pressure on professional roles.[47] Lo's survey revealed uneven access to training, with support staff less likely to participate in training programs than senior staff, a pattern that risks widening internal organization inequities.[48] At the same moment, vendors are rapidly integrating generative AI into their systems. Ex Libris introduced an AI metadata assistant feature within the Alma discovery system. Other discovery providers have launched conversational research assistants that they claim to have grounded in trusted indexes.[49]

Labor organizations and library unions are responding. The American Federation of Teachers launched a national AI academy to expand AI fluency among educators.[50] In public libraries, contract language increasingly addresses new technologies. The Boston Public Library Professional Staff Association negotiated provisions requiring notice and bargaining over major technology changes.[51] These examples suggest a path for academic libraries to address training, consultation, and job redesign more proactively. What they lack is a unifying framework, a way to connect individual responses to professional strategy.

Part III: The Critical Implementation Framework

Defining Critical Implementation

Critical implementation, as developed here, is the strategic response that synthesizes the historical and theoretical analysis above, including the conditions under which professional agency persists, the structural constraints that distinguish contemporary AI from earlier automation, and the implementation patterns already emerging across institutions, into practice. It acknowledges genuine structural constraints while identifying spaces for professional agency within them. Professional power under technological constraint operates through intervention at accessible points—deployment, architecture, governance—where libraries retain meaningful choice.

The critical implementation framework builds on existing guidance from national associations. The Library of Congress's AI planning framework and the Association of Research Libraries (ARL)'s principles emphasize privacy, transparency, equity, human oversight, and mission alignment.[52] Critical implementation extends those commitments by insisting that AI adoption decisions be governed by professional judgment at every stage. Librarians should assess each proposed use case against values, local capacity, and verifiable benefits, and resist pressure to adopt tools simply because they are available.

Four Strategic Pillars

Drawn from ARL's Guiding Principles, the International Coalition of Library Consortia (ICOLC)'s licensing framework, and IFLA's AI recommendations, the author advocates four strategic pillars to serve as a foundation for critical implementation in practice. To-

gether, these pillars address the structural constraints identified above: opacity, market concentration, misaligned incentives, and the erosion of professional jurisdiction.

Transparency

Librarians must treat algorithmic opacity as a procurement problem, not an inevitable condition.[53] When systems cannot explain their outputs, professional oversight becomes impossible, and libraries lose the ability to verify that tools serve their values. It is essential to require documented provenance and model information in the Requests for Proposals (RFPs) that libraries issue when procuring AI-enabled tools and services. Questions to ask of vendors offering AI include: What training data was used? What content policies govern outputs? What logging is available? Libraries should label generated content in user interfaces so users can distinguish retrieval from generation. Libraries should log queries and outputs to create audit trails; when a generated answer fails, librarians (and users) need to be able to trace what happened.[54] The boundary between retrieval and generation should remain visible, whatever the interface. These measures expand interpretive space and enable professional oversight of system behavior.[55]

The boundary between retrieval and generation should remain visible, whatever the interface.

Collective Action

Libraries acting alone have little leverage against vendors whose models serve global markets; libraries acting together can set terms. Consortial procurement with AI clauses can require transparency provisions, data use restrictions, and exit rights that individual institutions could not negotiate. ICOLC's principles for AI in licensing agreements offer model language: provisions addressing how library data may be used to train models, requirements for disclosure of AI features in licensed products, and protections against lock-in.[56] Libraries can also establish shared evaluation protocols and testbeds that verify vendor claims independently. When a discovery vendor claims its AI assistant is "grounded in trusted sources," libraries should be able to test that claim against their own collections, much as Z39.50 testbeds once verified interoperability.[57] Pooled infrastructure, including shared prompt libraries, evaluation datasets, and governance templates, can further strengthen collective capacity. Open-source or collectively governed models that permit inspection and modification shift power away from single vendors and toward community governance. Open-source frameworks like Hugging Face Transformers provide industrial-grade AI tools freely available for local implementation.[61] FOLIO's community-governed model, where libraries collectively control development priorities, suggests collaborative AI projects can scale.[62]

Open-source or collectively governed models that permit inspection and modification shift power away from single vendors and toward community governance.

Mission-alignment

Librarians have the opportunity to embed library values in technology deployment. When evaluating AI features, librarians can ask whether the system makes provenance visible or obscures it, whether it encourages verification or discourages it, whether it treats user queries as data to be harvested or interactions to be protected. Privacy is a central concern: does the system log user queries? Are those logs used to train models? Can users opt out? Academic integrity matters equally: does the system make it easy for students to pass off generated text as their own work, or does it support proper attribution? Decision-makers can consider institutionally bounded options—controlled deployments with privacy and accessibility assurances—when they better protect communities than public cloud defaults. Michigan's decision to build its own tools rather than adopt ChatGPT Edu reflects this logic; accepting some additional cost and complexity in exchange for control over privacy, accessibility, and institutional values.[58]

Flexibility

The fourth pillar calls for maximizing implementation flexibility within constraints. Modular architectures that preserve configurability allow systems built from separable components to be reconfigured as needs change, while monolithic platforms create dependencies. Human override for consequential decisions should be maintained — when AI suggests a subject heading, a librarian should approve it; when AI generates a research summary, a user should be able to see the sources and reject the synthesis. Libraries should evaluate systems against local data and scenarios rather than vendor benchmarks, since a model that performs well on generic reference questions may fail on specialized collections or local user populations. Adopting architectures that make the distinction between retrieval and generation legible to users preserves professional judgment over how information gets presented. These practices enlarge the space for professional discretion and maintain the ability to reconfigure systems as evidence evolves.

Libraries should evaluate systems against local data and scenarios rather than vendor benchmarks, since a model that performs well on generic reference questions may fail on specialized collections or local user populations.

Technical Possibilities and Professional Opportunities

Retrieval augmentation and local grounding

Libraries can ground AI responses in their own collections rather than relying on whatever the base model absorbed during training. Retrieval-augmented systems retrieve relevant documents from curated holdings before generating responses, which means a query about local history can draw on digitized primary sources the library controls.[59] Northwestern's project demonstrates this approach with a metadata augmentation tool-

kit built on verified institutional holdings.[60] These implementations prioritize collection integrity over generic fluency.

Evaluation Frameworks and Independent Assessment

Vendors claim their AI discovery tools are grounded in trusted sources, that their assistants reduce hallucinations, that their systems respect privacy. Libraries need systematic ways to test these claims. Developing evaluation frameworks—standardized queries, known-answer tests, comparison matrices across platforms—allows libraries to verify vendor assertions independently rather than accepting marketing language. This work does not require programming expertise; it requires the evaluative skills librarians already possess, applied systematically to new tools. When a library can demonstrate that one discovery assistant fabricates citations at twice the rate of another, procurement decisions become evidence based.

Workflow Automation for Professional Tasks

AI-assisted coding lowers barriers for librarians who need to automate repetitive work. A cataloger extracting ISBNs from a spreadsheet and checking them against WorldCat can describe the task, receive working starter code, and review the results, a process that once required formal programming training or a request to systems staff. This redistribution of technical capacity is uneven; not every librarian will use it, and code still requires review and testing. But the barrier between "I do this manually" and "I automate this" becomes more permeable, expanding what individual librarians can accomplish without depending on scarce technical staff.[63]

Emerging Jurisdictions and Professional Evolution

Algorithm Auditing and Bias Detection

Librarians' commitments to intellectual freedom and equitable access extend naturally to auditing AI outputs for accuracy and bias. When a discovery assistant consistently underrepresents certain perspectives, or when a metadata tool systematically miscategorizes works from particular traditions, someone needs to notice and document the pattern. Professional training in source evaluation and critical assessment provides foundational skills for this work. Reference librarians already evaluate sources for authority and accuracy; extending that practice to AI outputs is a natural jurisdictional claim. Framing auditing as an extension of reference and metadata practice positions librarians to claim this work before others do.

Librarians' commitments to intellectual freedom and equitable access extend naturally to auditing AI outputs for accuracy and bias.

AI Literacy as Information Literacy Extension

Teen and student use of chatbots for schoolwork has grown rapidly, while instruction has struggled to keep pace. Students need to understand how language models differ from search engines—that they generate plausible text rather than retrieve verified sources, that confidence in tone does not indicate reliability, and that citations may be fabricated. They need practical skills, including how to prompt effectively, how to recognize hallucination patterns, and how to cross-verify claims against authoritative sources. This is information literacy for a generative environment. Libraries already teach students to evaluate sources, trace claims, and distinguish evidence from assertion. Extending that instruction to AI-generated content reaffirms libraries' educational mission rather than abandoning it. The alternative, leaving AI literacy to vendors, to faculty who may not have time, or to no one, means students learn by trial and error, often after the damage is done.

Ethical implementation leadership

Someone on campus needs to ask the hard questions about AI deployment:

- What happens to user data?
- Who is accountable when generated content causes harm?
- What populations are excluded by accessibility failures?
- What recourse exists when systems malfunction?

Libraries have the institutional standing to raise these questions and professional commitments that require doing so. Libraries can build on established association guidance to adopt concrete policies that require clear provenance for AI-assisted outputs, enable opt-outs for users who prefer non-AI alternatives, and keep a human in the loop for decisions with academic or ethical consequences.[64] Where appropriate, libraries should align licensing agreements with consortial AI principles, so that values are embedded in contracts rather than dependent on vendor goodwill.[65]

Why Libraries

Vendors build systems to maximize engagement and renewal. Their incentives align with ease of use, speed of response, and institutional stickiness rather than source transparency or epistemic caution. Product documentation from major discovery platforms emphasizes seamless experience and reduced friction; it does not emphasize the difference between evidence and assertion or the importance of teaching users to trace claims to their origins. This is not criticism. Vendors are doing what vendors do. But it means that the values articulated in the ALA Code of Ethics, the ARL Guiding Principles, and the IFLA framework will not appear in these systems unless someone insists on them.

Faculty care about their own research. They are users of library systems. When discovery tools make research faster, faculty benefit; when those tools obscure provenance or encourage overconfidence in generated answers, faculty may not notice until the damage surfaces in student work or their own citations. IT departments care about security, uptime, and integration with campus systems. These are genuine concerns, but they are not epistemological concerns.

Students want answers. They are not wrong to want answers. But the difference between an answer and a sourced, traceable, evaluable answer is precisely what information literacy instruction has always tried to teach. If the systems students use do not support that distinction, if the interface makes it easy to accept a confident assertion and hard to find where it came from, then the systems are working against us.

Librarians are the professional group whose core commitments (source verification, equitable access, user privacy, intellectual freedom) align directly with the governance needs of AI-mediated information systems.

Librarians are the professional group whose core commitments (source verification, equitable access, user privacy, intellectual freedom) align directly with the governance needs of AI-mediated information systems. If librarians do not bring these values to implementation decisions, no other campus constituency will supply them. Jurisdiction over AI implementation, in Abbott's framework, belongs to whoever claims it. Ceding that ground to vendors and IT departments means the epistemological commitments at the heart of librarianship simply go unrepresented.

Conclusion

Without strategic engagement, the risks are clear. Generative AI deployments will prioritize engagement over accuracy when library norms are not embedded. Professional authority over information literacy will drift elsewhere if librarians are absent from implementation decisions. Lock-in deepens as dependencies grow: the longer institutions rely on a single vendor's AI infrastructure, the harder it becomes to switch providers, renegotiate terms, or maintain independent evaluation capacity.

Librarians' professional codes commit them to transparency about sources, accuracy, and provenance for a reason. If these commitments are not embedded in the systems students and scholars use, research will be faster, carry a greater appearance of authority, and grow less connected to the evidentiary foundations that make scholarly inquiry trustworthy.

Author's Note

In preparing this manuscript, I used generative AI for exploratory discussions of theoretical frameworks and for editorial assistance in tightening prose. All ideas, arguments, analysis, and conclusions are my own.

Stephen Zweibel is digital scholarship librarian at the Graduate Center, CUNY, email: szweibel@gc.cuny.edu, ORCID: 0009-0009-1877-6978.

Notes

1. Zhehui Liao, Maria Antoniak, Inyoung Cheong, Evie Yu-Yen Cheng, Ai-Heng Lee, Kyle Lo, Joseph Chee Chang, and Amy X. Zhang, "LLMs as Research Tools: A Large Scale Survey of Researchers' Usage and Perceptions," in *Proceedings of the Conference on Language Modeling (COLM)* (2025), https://openreview.net/forum?id=p0BwJk3R1p; Clarivate, *Pulse of the Library 2024* (2024), https://doi.org/10.14322/pulse.of.the.library.2024.
2. Langdon Winner, "Do Artifacts Have Politics?" *Daedalus* 109, no. 1 (1980): 121–36.
3. William H. Walters and Tyler Wilder, "ChatGPT and Fabricated Bibliographic Citations," *Scientific Reports* 13 (2023), https://doi.org/10.1038/s41598-023-41032-5.
4. 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), https://doi.org/10.5860/crl.85.5.635.
5. Clarivate, *Pulse of the Library 2024*.
6. 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.
7. Emily A. Vogels, *Teens and ChatGPT: Technology and Education* (2025), https://www.pewresearch.org/short-reads/2025/01/15/about-a-quarter-of-us-teens-have-used-chatgpt-for-schoolwork-double-the-share-in-2023/; Liao et al., LLMs as Research Tools.
8. "Introducing ChatGPT Edu," OpenAI, May 30, 2024, https://openai.com/index/introducing-chatgpt-edu/.
9. Andrew Abbott, *The System of Professions: An Essay on the Division of Expert Labor* (University of Chicago Press, 1988).
10. Abbott, *System of Professions*.
11. Winner, "Do Artifacts Have Politics?"
12. Trevor J. Pinch and Wiebe E. Bijker, "The Social Construction of Facts and Artefacts: Or How the Sociology of Science and the Sociology of Technology Might Benefit Each Other," *Social Studies of Science* 14, no. 3 (1984): 399–441.
13. Ibid.; Wiebe E. Bijker, *Of Bicycles, Bakelites, and Bulbs: Toward a Theory of Sociotechnical Change* (MIT Press, 1995).
14. Ibid..
15. Clifford A. Lynch, "The Z39.50 Information Retrieval Standard: Part I: A Strategic View of Its Past, Present and Future," *D-Lib Magazine* 3, no. 4 (1997).
16. Ibid.
17. Ibid.
18. Karen Markey, "Twenty-Five Years of End-User Searching, Part 1: Research Findings," *Journal of the American Society for Information Science and Technology* 58, no. 8 (2007): 1071–1081, https://doi.org/10.1002/asi.20462.
19. Lynch, "The Z39.50 Information Retrieval Standard."
20. Ibid.
21. "Memorial: Patricia Battin, 1929–2019," Association of Research Libraries, https://www.arl.org/news/memorial-patricia-battin-1929-2019/.
22. Jo Bell Whitlatch, "Automation and Job Satisfaction Among Reference Librarians," *Computers in Libraries* 11, no. 8 (1991).
23. U.S. Bureau of Labor Statistics, *Employment Trends for Library Professions (1980s–1990s)*.
24. Michael K. Buckland, "Information as Thing," *Journal of the American Society for Information Science* 42, no. 5 (1991): 351–60.
25. Jenna Burrell, "How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms," *Big Data & Society* (2016), https://doi.org/10.1177/2053951715622512.
26. Vogels, *Teens and ChatGPT*; Liao et al., "LLMs as Research Tools."
27. Burrell, "How the Machine 'Thinks'."
28. Anton Korinek and Jai Vipra, "Market Concentration Implications of Foundation Models: The Invisible Hand of ChatGPT" (Brookings Institution, September 2023), https://www.

brookings.edu/articles/market-concentration-implications-of-foundation-models-the-invisible-hand-of-chatgpt/.
29. International Federation of Library Associations (IFLA), *AI Entry Point for Libraries and AI, 2025*, https://repository.ifla.org/items/f197f327-dc49-4743-bb57-0a373505da8b.
30. Patrick Lewis et al., "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," in *Advances in Neural Information Processing Systems* 33 (2020): 9459–74; Matthew Renze and Erhan Guven, "The Effect of Sampling Temperature on Problem Solving in Large Language Models," in *Findings of the Association for Computational Linguistics: EMNLP 2024* (Association for Computational Linguistics, 2024), 7346–56; Sehoon Kim, Suhong Moon, Ryan Tabrizi, Nicholas Lee, Michael W. Mahoney, Kurt Keutzer, and Amir Gholami, "An LLM Compiler for Parallel Function Calling," in *Proceedings of the 41st International Conference on Machine Learning*, PMLR 235 (2024): 24370–91, https://proceedings.mlr.press/v235/kim24y.html.
31. Geoffrey Schafer and Ravi Pendse, "How (and Why) the University of Michigan Built Its Own Closed Generative AI Tools," *EDUCAUSE Review* (2024), https://er.educause.edu/articles/2024/2/how-and-why-the-university-of-michigan-built-its-own-closed-generative-ai-tools; "Collections & AI," Northwestern University Libraries, 2024, https://collections-and-ai.library.northwestern.edu/; "Yale Library Developing AI Application Could Transform Research in Digitized Collections," Yale University Library, 2024, https://library.yale.edu/news/yale-library-developing-ai-application-could-transform-research-digitized-collections; Lesley Seacrist, "Commitment to AI Innovation: SJSU's First AI Librarian," SJSU NewsCenter, February 11, 2025, https://blogs.sjsu.edu/newsroom/2025/commitment-to-ai-innovation-sjsus-first-ai-librarian/.
32. Schafer and Pendse, "How (and Why) the University of Michigan Built Its Own Closed Generative AI Tools."
33. "Introducing ChatGPT Edu," OpenAI.
34. "Collections & AI," Northwestern University Libraries.
35. "Yale Library Developing AI Application Could Transform Research in Digitized Collections," Yale University Library.
36. Seacrist, "Commitment to AI Innovation."
37. Ibid.
38. Association of Research Libraries (ARL), *Research Libraries Guiding Principles for Artificial Intelligence*, April 2024, https://www.arl.org/news/association-of-research-libraries-releases-guiding-principles-for-artificial-intelligence/.
39. International Coalition of Library Consortia (ICOLC), *Statement on AI in Licensing*, March 22, 2024, https://icolc.net/statements/icolc-statement-ai-licensing.
40. Brady D. Lund, Isaiah Omame, Solomon Tijani, and Daniel Agbaji, "Perceptions Toward Artificial Intelligence Among Academic Library Employees and Alignment with the Diffusion of Innovations' Adopter Categories," *College & Research Libraries* 81, no. 5 (2020): 865–882, https://doi.org/10.5860/crl.81.5.865.
41. Birger Hjørland, "Information Retrieval and Knowledge Organization: A Perspective from the Philosophy of Science," *Information* 12 (2021), https://doi.org/10.3390/info12030135.
42. American Library Association (ALA), *Code of Ethics of the American Library Association*, 2021, https://www.ala.org/tools/ethics.
43. Matthew Pierce, "Academic Librarians, Information Literacy, and ChatGPT: Sounding the Alarm on a New Type of Misinformation," *C&RL News* 86, no. 2 (2025): 68–71, https://crln.acrl.org/index.php/crlnews/article/view/26659.
44. Katie Lai, "How Well Does ChatGPT Handle Reference Inquiries?" *College & Research Libraries* 84, no. 6 (2023): 974–995, https://doi.org/10.5860/crl.84.6.974.
45. Lei Huang, et al. "A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions," *ACM Transactions on Information Systems* 43, no. 2 (2024), https://doi.org/10.1145/3703155.
46. Walters and Wilder, "ChatGPT and Fabricated Bibliographic Citations."

47. Urban Libraries Council, *Leadership Brief: Explorations of Generative AI for Library Systems*, October 2023, https://www.urbanlibraries.org/resources/explorations-of-generative-ai-for-library-systems.
48. Lo, "Evaluating AI Literacy in Academic Libraries."
49. "Getting Started with Primo Research Assistant," Ex Libris, 2025, https://knowledge.exlibrisgroup.com/Primo/Product_Documentation/020Primo_VE/Primo_VE_(English)/015_Getting_Started_with_Primo_Research_Assistant; Katie Means and Karen Coombs, "Understanding AI and Its Role in Shaping the Future of Library Discovery" (webinar, OCLC, May 1, 2025), https://community.oclc.org/t5/devconnect-agenda/understanding-ai-and-its-role-in-shaping-the-future-of-library/ev-p/57220.
50. "AFT to Launch National Academy for AI Instruction with Microsoft, OpenAI, Anthropic and United Federation of Teachers," American Federation of Teachers, 2025, https://www.aft.org/press-release/aft-launch-national-academy-ai-instruction-microsoft-openai-anthropic-and-united.
51. Claire Kelley, "Solidarity in Action: Labor, AI, and the Future of Libraries at ALA 2025," *Library Journal*, July 16, 2025, https://www.libraryjournal.com/story/solidarity-in-action-labor-ai-and-the-future-of-libraries-ala-annual-2025.
52. ARL, *Research Libraries Guiding Principles*.
53. Burrell, "How the Machine 'Thinks'."
54. Ibid.
55. ARL, *Research Libraries Guiding Principles*.
56. ICOLC, *Statement on AI in Licensing*.
57. Lynch, "The Z39.50 Information Retrieval Standard."
58. Schafer and Pendse, "How (and Why) the University of Michigan Built Its Own Closed Generative AI Tools."
59. Lewis et al., "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks."
60. "Collections & AI," Northwestern University Libraries.
61. Nestor Maslej et al., The AI Index 2025 Annual Report (Stanford, CA: Institute for Human-Centered AI, Stanford University, April 2025), https://hai.stanford.edu/ai-index/2025-ai-index-report.
62. FOLIO Project, "FOLIO Governance Model," FOLIO Wiki, 2021, https://wiki.folio.org/display/COMMUNITY/FOLIO+Governance+Model.
63. Maslej, *AI Index 2025*.
64. ARL, *Research Libraries Guiding Principles*; IFLA, *AI Entry Point*; ALA, *Code of Ethics*.
65. ICOLC, *Statement on AI in Licensing*.