

Level 3: Advanced

All Librarians

Module 18

Your AI strategy & next steps

Completing this curriculum is a beginning. This module helps you build the ongoing practice, community, and professional presence that turns a learning journey into a professional identity.

By Yulia Brusova

15 min read

WHAT YOU'LL BE ABLE TO DO

- 1 Draft a 12-month AI learning and implementation roadmap for yourself or your department
- 2 Identify two or three communities of practice for ongoing AI professional development
- 3 Outline a conference presentation proposal on a library AI topic
- 4 Develop a system for staying current with AI developments in library contexts
- 5 Define your own position and contribution as a practitioner voice in the AI conversation
- 6 Identify how to contribute library expertise to your institution's AI policy development, recognizing policy co-authorship as a professional contribution alongside presentations and community involvement

Completing fifteen modules of structured curriculum is a beginning, not an ending. The practitioner who has worked through this curriculum has developed a mental model of how AI systems work, practical experience with a range of tools, a prompt library in progress, and at least one workflow that operates differently than it did before. The question this module addresses is not what to learn next but what kind of practitioner to become in the professional conversation about AI that the library field is having right now, and how to build the ongoing practice that makes that contribution sustainable.

Building a twelve-month AI practice roadmap

A roadmap is useful not because AI develops slowly enough for plans to remain intact over a year, which it does not, but because deciding in advance what to accomplish forces prioritization in a way that reacting to the news cycle does not. Without a roadmap, the practitioner responds to whatever AI development receives the most coverage in a given week rather than building deliberately toward a professional goal that matters.

A practical roadmap structure divides the year into four phases of roughly equal length, each with a single primary objective. For example, the first three months are best devoted to deepening existing practice rather than expanding to new tools: selecting two or three workflows where AI has been used experimentally and making them systematic, writing

prompt library entries for them, and measuring the time difference over a consistent period. Such a phase builds the habits that make subsequent expansion sustainable rather than scattered.

The second phase, covering months four through six, is well suited to one visible professional contribution. For example, a short article for the library's internal newsletter, a presentation at a staff meeting, or a LibGuide describing AI tools relevant to a specific subject area all accomplish the same objective: translating personal practice into something that benefits colleagues and builds a visible record of engagement with the topic. Indeed, the act of explaining a practice to an audience forces a clarity of understanding that internal use alone does not.

The third phase is appropriate for extending AI practice to the department level in a deliberate and limited way. For example, identifying one colleague who has expressed curiosity about AI and showing that person one specific workflow, not attempting to train the whole department, not converting skeptics, is a more sustainable approach to departmental culture change than a formal training program introduced before the practitioner's own practice is fully developed. Such incremental extension builds peer support networks that prove more durable than top-down initiatives.

The fourth phase, months ten through twelve, is an honest evaluation of what the roadmap produced. What worked and what did not, what has changed in the AI landscape that alters the approach, and what the roadmap for the following year should prioritize are all questions that benefit from the full year of documented practice as evidence. Furthermore, this evaluation phase is the appropriate moment to distinguish between a personal roadmap and a departmental AI strategy, which is a separate and more complex undertaking requiring institutional authority and stakeholder involvement beyond the individual practitioner.

Staying current without being overwhelmed

The AI news cycle is structured to create urgency. Most of what receives coverage is either further from practical library application than the headline implies, a product announcement presented as a research development, or a capability that is technically real but not yet relevant to library workflows in any actionable way. In order to stay current without being overwhelmed, the practitioner needs a curated information diet rather than broad consumption.

A sustainable weekly practice requires one source: a newsletter specifically covering AI in library contexts, such as Library Tech Insights, scanned for headlines relevant to the

practitioner's specific role and read more closely only when the content is directly applicable. Such a practice takes fifteen minutes per week and is sufficient to remain aware of significant developments without the cognitive overhead of following general AI news.

A monthly practice adds one deeper read: a journal article, a conference proceeding, or a publication from ARL or ACRL. For example, Library Technology Reports, published by ALA TechSource, covers AI in library contexts regularly and at a level of depth appropriate for practitioner decision-making. Such sources are more reliable than general technology coverage because they are written by and for library professionals rather than technology generalists.

An annual practice adds the Pulse of the Library survey from Clarivate, which provides benchmark data on AI adoption, attitudes, and use across the profession. For example, the 2025 data showing seven percent of U.S. librarians reporting optimism about AI, compared to twenty-seven to thirty-one percent in other parts of the world, provides a calibration point for understanding where the field is and what work remains. Such benchmark data is more useful for professional positioning than day-to-day coverage because it reflects actual practice rather than aspiration.

It is evident that filtering is as important as selection. Vendor announcements, general AI news not specific to library contexts, and speculative coverage of capabilities that have not been in library practice for at least six months all consume attention without producing actionable insight. The ACRL AI competency framework identifies staying current with AI applications via reliable sources as a professional obligation; reliable in this context means curated, library-specific, and professionally vetted, not comprehensive.

Contributing to the professional conversation

The practitioner voice in AI conversations about libraries is underrepresented relative to its value. Most of what receives broad professional attention comes from vendors, researchers, or administrators; the perspective of the librarian who tests tools in actual reference workflows, instruction sessions, and cataloging practice and writes honestly about what those tests reveal is rarer and more useful than it is common.

There is no doubt that contributing to the professional conversation does not require seniority, a research appointment, or a national platform. For example, a thoughtful comment on a LinkedIn post by a library AI researcher, a response to a professional organization's practice survey, or a brief honest account of a pilot project shared with immediate colleagues are all contributions that move the collective understanding of the field forward. Such low-

investment contributions also build the professional visibility that makes subsequent contributions at higher levels more likely to find an audience.

At a medium investment level, a short piece for a library's internal newsletter or staff blog, a proposal for a session at a regional conference, or a response to a journal call for practitioner perspectives each accomplish more than their modest scale suggests. For example, a twelve-hundred-word account of what a specific AI tool did and did not accomplish in a specific library workflow is more useful to the profession than a general review, because it provides the contextual specificity that allows other practitioners to evaluate relevance to their own settings. Such specificity is the practitioner's distinctive contribution and the thing that vendor documentation and research literature most frequently lack.

Additionally, at higher levels of investment, proposals to national conferences such as ALA Annual or ACRL, submissions to peer-reviewed library journals, and workshops developed for delivery to other libraries all extend the reach of practitioner experience into the channels that shape professional norms and educational programs. Furthermore, each of these higher-investment contributions is made more accessible by the lower-investment contributions that preceded it: a practitioner with a record of regional presentations and published short pieces is better positioned to propose at national venues than one approaching those venues without a track record.

Policy co-authorship: from AI user to policy contributor

Scholarship published in the Journal of Library Administration in 2025 makes a finding that a practitioner completing this curriculum should take personally: institutional AI policies, where they exist at all, frequently have gaps in exactly the areas librarians understand best, including patron privacy, information literacy, and equitable access, and libraries are systematically under-represented among the offices and committees that write those policies. In order to understand why this matters for professional identity, it is useful to connect this finding back to Module 04's discussion of the ARL Guiding Principles and the ACRL framework: the expertise those documents describe as core to librarianship is precisely the expertise that institutional AI policy development is currently missing.

This module has framed the practitioner's development as a progression from AI user to contributor to the professional conversation. Policy co-authorship is the institutional expression of that same progression, and it belongs in the roadmap alongside conference presentations and community participation. For example, a practitioner who has spent six months building a documented, calibrated AI practice, the kind this module's first phase describes, brings something to an institutional AI policy committee that most committee

members do not have: direct, current, hands-on experience with what the tools actually do, where they fail, and what a workable acceptable use policy needs to specify. Such experience is a credential in this context, regardless of title or seniority.

The practical first step is finding out whether an AI policy committee, task force, or working group already exists at the institution, and if it does not, raising the question of who is responsible for AI policy with library administration. For example, a librarian who asks "who is writing our AI acceptable use policy, and has the library been asked to contribute?" is often the first person at the institution to ask that question directly, and asking it is itself a form of professional leadership that requires no formal authority. If a policy process is already underway without library representation, volunteering specific expertise, on patron privacy, on information literacy standards, on accessibility, is a concrete and welcome offer rather than an abstract request to be included.

None of this replaces the institutional decision-making that AI policy ultimately requires, and a librarian's contribution to a policy committee does not make the librarian responsible for decisions that remain the institution's to make. What it does is ensure that those decisions are made with the professional perspective this curriculum has built present in the room rather than absent from it. That presence, more than any individual recommendation a librarian might make, is the lasting contribution: a policy written with library expertise in the room differs from one written without it, and no amount of after-the-fact commentary on a finished policy substitutes for having been part of writing it.

Conference speaking on AI topics: what gets accepted

Conference proposals on AI in libraries are being submitted in large numbers, and the selection process has become more discerning as a result. The proposals that are accepted share a consistent set of characteristics: they are specific in their framing, grounded in actual practitioner experience, and honest about what did not work as well as what did. The proposals that are rejected are typically too broad, too general, or too closely resembling vendor presentations.

For example, a proposal titled "AI in libraries: an overview" or "The future of AI and library services" is unlikely to be accepted at any conference with competitive selection because it describes a subject rather than a finding. Such proposals could have been written without ever using an AI tool in a library setting, and reviewers recognize this. In contrast, a proposal titled "What six months of AI use in research consultations at a community college reveals" makes a specific claim about a specific context and implies findings that attendees could not encounter elsewhere. Such specificity is the quality that makes a proposal worth accepting.

The proposal structure that works most consistently follows four elements. The first element states the specific problem the practitioner was trying to solve. The second states what was tried, including which tools and which workflows. The third states what was found, including failures and unexpected results alongside successes. The fourth states what attendees will be able to do or decide differently as a result of attending. For example, a proposal describing a failed AI pilot for metadata enhancement, the specific reasons it failed, and the revised approach that followed is more compelling to reviewers than a description of a smooth success, because failures contain more transferable learning than successes and are less likely to duplicate what attendees have already heard.

Regional and state library conferences are the appropriate starting point for practitioners who have not yet presented at national venues. Such conferences have less competitive selection processes, audiences that value practitioner specificity as much as national audiences do, and the added advantage of producing a presentation record that strengthens subsequent proposals to ACRL, ALA Annual, ER&L, and similar venues with more competitive submission processes.

Communities of practice and what comes next

The most durable professional development in any rapidly changing field is peer-to-peer. Communities of practitioners who share what is working, what is not, and what has changed since the last conversation are more valuable over time than any static curriculum, including this one, because they update continuously as the tools and practices themselves evolve.

Several communities are worth joining for library professionals working with AI. ALA's Core division, which incorporates the former LITA membership and its technology-focused interest groups, provides listservs and discussion spaces where working practitioners address specific implementation challenges rather than general questions. For example, the interest groups covering library technology and metadata are where questions such as "has anyone built an AI-assisted workflow for institutional repository metadata using DSpace?" receive answers from practitioners who have actually tried it. Such communities reward specific questions with specific answers in ways that general professional channels do not.

AIRUS, the AI in Reference and User Services interest group, is the most directly relevant community for practicing reference librarians working with AI tools. Library Technology Reports readers and contributors tend toward thoughtful practitioner perspectives rather than advocacy, which makes that community particularly useful for calibrating assessments of new tools and approaches. LinkedIn, filtered by following specific practitioners rather than

general hashtags, surfaces a diverse and active library AI community that includes international perspectives underrepresented in domestic professional channels.

It is evident that the state of AI in library practice will continue to change at a pace that makes any fixed curriculum incomplete over time. The professional obligation is not to have mastered the current state but to have developed the habits of evaluation, honest practice, and peer exchange that allow the practitioner to adapt as the state changes. For example, a practitioner who has built a reliable system for testing new tools in low-stakes contexts, documenting findings honestly, and sharing those findings with colleagues is better positioned to navigate the next significant AI development than one who has memorized the current landscape without those habits. Such habits are what this curriculum has attempted to build. The landscape will change; the habits remain applicable.

PRACTITIONER NOTE

The shift from practitioner who uses AI to practitioner who contributes to the professional conversation about it happens incrementally: taking notes on what works and what does not, writing those notes into something shareable, saying yes to presenting at a staff meeting, then at a regional conference, then further. None of that requires being the most technically advanced person in the room. It requires being the most honest about what actually happened when the tool was used in a real library context. That combination of specificity and honesty is rarer in the professional literature than it should be, and it is entirely within the reach of any practitioner who has completed this curriculum.

KEY TAKEAWAYS

- A twelve-month roadmap structures AI practice development into four phases: deepening existing practice, making one visible professional contribution, extending practice to the department in a limited way, and evaluating honestly to reset for the following year.
- A sustainable information diet for staying current requires one weekly newsletter covering library AI, one monthly deeper read from a professional journal or ARL or ACRL publication, and the annual Pulse of the Library benchmark survey.
- The practitioner voice is underrepresented in professional AI literature; specific, honest accounts of what AI tools did and did not accomplish in real library workflows are more

useful to the field than general reviews or vendor-adjacent coverage.

- Conference proposals that are accepted are specific in framing, grounded in actual practice, and honest about failure; those that are rejected describe a subject rather than a finding and could have been written without using AI in a library.
- Journal of Library Administration (2025) scholarship finds that institutional AI policies often have gaps in patron privacy, information literacy, and equitable access, and that libraries are systematically under-represented in writing them; policy co-authorship is a professional contribution within reach of any practitioner who has built a documented AI practice.
- Communities of practice, including ALA Core interest groups, AIRUS, Library Technology Reports communities, and LinkedIn filtered by specific practitioners, provide peer exchange that updates continuously as tools and practices evolve.
- The professional obligation is not to have mastered the current state of AI but to have developed the habits of evaluation, honest practice, and peer exchange that allow adaptation as the state continues to change.

References

APA 7TH EDITION

1. Association of College and Research Libraries. (2025, October). *AI competencies for academic library workers*. American Library Association. <https://www.ala.org/acrl/standards/ai>
2. ALA TechSource. (n.d.). *Library technology reports*. American Library Association. <https://www.ala.org/pubs/periodicals/ltr>
3. Association of Research Libraries. (n.d.). *ALA Core: Library and information technology association*. American Library Association. <https://www.ala.org/core/>
4. Clarivate. (2025). *Pulse of the library 2025*. <https://clarivate.com/pulse-of-the-library/>
5. Michalak, R., Dawes, T. A., & Wallace, A. (2026, May). Envisioning AI's role in libraries: Perspectives from an LIS student, a library director, and a university librarian. *College & Research Libraries News*, 87(5). <https://crln.acrl.org/index.php/crlnews/article/view/27321>

ACRL AI Competencies covered

Ethical Considerations

Analysis & Evaluation

Use & Application

