

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 · Reviewed October 2026

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 eighteen modules of structured curriculum is a beginning, not an ending. A practitioner who has worked through this curriculum has 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 runs differently than it did before. This module does not ask what to learn next. It asks what kind of practitioner to become in the professional conversation about AI that the library field is having right now. It also covers how to build the ongoing practice that makes that contribution sustainable.

Building a twelve-month AI practice roadmap

A roadmap is not useful because AI develops slowly enough for plans to survive a year intact; it does not. It is useful 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 chases whatever AI development gets the most coverage in a given week, instead of building deliberately toward a professional goal that matters.

Before drawing up the roadmap, establish a baseline. Leo S. Lo's AI Literacy for All: A Universal Framework, released in 2025 as a preprint that has not yet been peer reviewed, organizes AI literacy into five components: Technical Knowledge, Ethical Awareness, Critical Thinking, Practical Use, and Societal Impact. Each has four progressive levels, and the

framework includes a self-assessment rubric with indicators for every level (Lo, 2025). For example, a librarian who can define basic AI terms and recognize common applications sits at Level 1 of Technical Knowledge. One who can compare AI approaches and identify appropriate use cases has reached Level 3. Rating yourself honestly against all five components takes about fifteen minutes and turns a general intention to improve into a specific gap. A practitioner at Level 3 in Practical Use but Level 1 in Societal Impact knows where the year's reading and experimentation should go. The rubric is a planning aid rather than a credential. Which gap matters most for a particular role and institution remains the practitioner's own judgment.

A practical roadmap divides the year into four phases of roughly equal length, each with a single primary objective. The first three months are best spent deepening existing practice rather than adding new tools. Select two or three workflows where you have used AI experimentally and make them systematic. Write prompt library entries for them, and measure the time difference over a consistent period. This phase builds the habits that make later expansion sustainable rather than scattered.

The second phase, months four through six, suits one visible professional contribution. A short article for the library's internal newsletter, a presentation at a staff meeting, or a LibGuide on AI tools for a specific subject area all serve the same purpose. Each turns personal practice into something that benefits colleagues and builds a visible record of engagement with the topic. Explaining a practice to an audience also forces a clarity of understanding that private use does not.

The third phase is for extending AI practice to the department in a deliberate, limited way. For example, find one colleague who has expressed curiosity about AI and show them one specific workflow. That is more sustainable than trying to train the whole department or convert skeptics, and more sustainable than a formal training program launched before your own practice is mature. Incremental extension of this kind 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? What should next year's roadmap prioritize? A full year of documented practice provides the evidence for these questions. Repeating the self-assessment from the start of the year makes progress, or its absence, visible across all five components. This evaluation phase is also the right moment to separate a personal roadmap from a departmental AI strategy. The latter is a separate, more complex undertaking that requires 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 gets coverage falls into one of three categories. It is further from practical library application than the headline implies, or it is a product announcement presented as a research development. Or it is a capability that is technically real but not yet relevant to library workflows in any actionable way. To stay current without being overwhelmed, the practitioner needs a curated information diet, not broad consumption.

A sustainable weekly practice requires one source: a newsletter covering AI in library contexts, such as Library Tech Insights. Scan it for headlines relevant to your role, and read more closely only when the content applies directly. This takes fifteen minutes a week and is enough to stay 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, at a depth suited to practitioner decision-making. These 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 Clarivate's Pulse of the Library survey, which provides benchmark data on AI adoption, attitudes, and use across the profession. For example, the 2026 data shows that average confidence in AI concepts has stalled at 3.2 out of 5, unchanged from the year before, and that U.S. libraries trail Asia and Europe in implementation. That gives a calibration point for where the field is and what work remains. Benchmark data like this is more useful for professional positioning than day-to-day coverage, because it reflects actual practice rather than aspiration.

Filtering matters as much as selection. Vendor announcements, general AI news not specific to libraries, 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 names 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. Librarians who test tools in actual reference workflows, instruction sessions, and cataloging practice, and write honestly about what those tests reveal, are rarer than they should be. Their perspective is also more useful.

Contributing to the professional conversation does not require seniority, a research appointment, or a national platform. A thoughtful comment on a library AI researcher's LinkedIn post, a response to a professional organization's practice survey, or a brief honest account of a pilot project shared with colleagues all move the field's collective understanding forward. These low-investment contributions also build the professional visibility that helps later, more ambitious contributions find an audience.

At a medium level of investment, consider a short piece for a library's internal newsletter or staff blog, a session proposal for a regional conference, or a response to a journal call for practitioner perspectives. Each accomplishes more than its 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. It provides the contextual detail that lets other practitioners judge its relevance to their own settings. That specificity is the practitioner's distinctive contribution, and it is what vendor documentation and research literature most often lack.

At higher levels of investment, proposals to national conferences such as ALA Annual or ACRL, submissions to peer-reviewed library journals, and workshops for other libraries extend practitioner experience into the channels that shape professional norms and educational programs. The lower-investment contributions that come first make each of these more accessible. A practitioner with a record of regional presentations and published short pieces is better positioned to propose at national venues than one approaching them 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, often have gaps in exactly the areas librarians understand best: patron privacy, information literacy, and equitable access. Libraries are also systematically underrepresented among the offices and committees that write those policies. To see why

this matters for professional identity, connect the 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 what 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. Consider a practitioner who has spent six months building a documented, calibrated AI practice, the kind this module's first phase describes. They bring something to an institutional AI policy committee that most members lack: direct, current, hands-on experience with what the tools actually do, where they fail, and what a workable acceptable use policy needs to specify. In this context, that experience is a credential, regardless of title or seniority.

The practical first step is to find out whether an AI policy committee, task force, or working group already exists at the institution. If none does, raise 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 directly. Asking is itself a form of professional leadership that requires no formal authority. If a policy process is already underway without library representation, volunteer specific expertise, on patron privacy, information literacy standards, or accessibility. That is a concrete and welcome offer, not an abstract request to be included.

None of this replaces the institutional decision-making that AI policy ultimately requires. 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 those decisions are made with the professional perspective this curriculum has built present in the room. That presence, more than any individual recommendation, is the lasting contribution. A policy written with library expertise in the room differs from one written without it. No amount of after-the-fact commentary on a finished policy substitutes for having helped write it.

Conference speaking on AI topics: what gets accepted

Conference proposals on AI in libraries are arriving in large numbers, and selection has become more discerning as a result. Accepted proposals 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. Rejected proposals are typically too broad, too general, or too much like 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, 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. It implies findings attendees could not encounter elsewhere, and that specificity is what makes a proposal worth accepting.

The proposal structure that works most consistently has four elements. The first 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. 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 an account of a smooth success. Failures contain more transferable learning than successes and are less likely to repeat what attendees have already heard.

Regional and state library conferences are the right starting point for practitioners who have not yet presented at national venues. Their selection is less competitive, and their audiences value practitioner specificity as much as national audiences do. They also produce a presentation record that strengthens later 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 become more valuable over time than any static curriculum, including this one. They update continuously as the tools and practices evolve.

Several communities are worth joining for library professionals working with AI. ALA's Core division, which absorbed the former LITA membership and its technology-focused interest groups, offers listservs and discussion spaces where working practitioners tackle specific implementation challenges rather than general questions. For example, the interest groups covering library technology and metadata are where a question such as "has anyone built an

AI-assisted workflow for institutional repository metadata using DSpace?" gets answers from practitioners who have actually tried it. These communities reward specific questions with specific answers in ways 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. That makes the community particularly useful for calibrating your assessments of new tools and approaches. On LinkedIn, following specific practitioners rather than general hashtags surfaces a diverse, active library AI community. It includes international perspectives that are underrepresented in domestic professional channels.

The state of AI in library practice will keep changing at a pace that leaves any fixed curriculum incomplete over time. The professional obligation is not to have mastered the current state. It is to have developed habits of evaluation, honest practice, and peer exchange that let the practitioner adapt as things change. For example, a practitioner with a reliable system for testing new tools in low-stakes contexts, documenting findings honestly, and sharing them with colleagues is better prepared for the next significant AI development than one who has memorized the current landscape without those habits. Those habits are what this curriculum has tried 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 happens in increments. You take notes on what works and what does not, then turn those notes into something shareable. You say 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. It is entirely within 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 underrepresented 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* [Archive]. American Library Association. <https://journals.ala.org/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. (2026, September). *Pulse of the library 2026*. <https://doi.org/10.14322/pulse.of.the.library.2026>
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>
6. Lo, L. S. (2025). *AI literacy for all: A universal framework* [Preprint]. University of New Mexico Digital Repository. https://digitalrepository.unm.edu/ulls_fsp/213
Preprint, not yet peer reviewed. Licensed CC BY-NC 4.0.

ACRL AI Competencies covered

Ethical Considerations

Knowledge & Understanding

Use & Application

Sub-competencies: 2.3, 1.5, 4.2 · ACRL AI Competencies (2025)