

Level 2: Applied

Practicing Librarian

Module 07

Reference & instruction

Reference and instruction are where librarians spend the most time, and where AI provides the most immediate daily utility. This module is entirely practical: real tasks, real prompts, real time savings.

By Yulia Brusova

25 min read · Reviewed June 2026

AI for Academic Libraries · ai-in-academic-libraries.vercel.app

© 2026 · Licensed under CC BY-NC-SA 4.0

WHAT YOU'LL BE ABLE TO DO

- 1 Use AI to draft or improve a subject guide for a specific course or discipline
- 2 Develop a library instruction lesson plan using AI as a drafting partner
- 3 Create patron facing AI use guidance appropriate for a specific population
- 4 Draft a faculty consultation framework for AI policy questions
- 5 Apply AI to at least three patron communication tasks in your daily workflow

AI fits into instruction work in ways that produce genuine and measurable time savings, particularly for drafting, refining, brainstorming, and adapting content for different audiences. Not for everything; there are tasks where it would be slower than just doing the work directly. The functions where librarians spend the most time, including drafting LibGuides, developing lesson plans, responding to patron email, and collaborating with faculty on policy questions, are precisely the functions where AI's drafting and revision capabilities are most useful. In order to benefit from those capabilities without making the professional errors that come from uncritical adoption, a librarian needs to understand both how to use AI effectively for instruction work and where the professional judgment that AI cannot provide remains essential.

Subject guides: drafting, improving, and adapting with AI assistance

Creating a new subject guide from scratch is one of the most time consuming routine tasks in reference and instruction work. The structural thinking, the description writing, the introduction copy, the organization of resources into logical tabs: each component is individually manageable, but together they constitute a substantial time investment for a task that may need to be completed quickly when a faculty member makes a last minute instruction request. AI is well suited to accelerate the structural and editorial components of guide creation, while the librarian's professional knowledge of the actual resources remains essential and irreplaceable.

The most efficient workflow for guide creation proceeds in stages that the librarian can learn to execute quickly once the pattern is established. First, provide AI with the context it needs to generate a useful structure: the subject area, the course level and institution type, the

patron population, and the assignment the guide will support. For example: "I am creating a LibGuide for a sociology research methods course at a community college. Students will be writing an original research proposal with a literature review section. Suggest a logical tab structure for the guide and what each tab should include." Such a prompt typically produces a reasonable starting structure in seconds, though not necessarily the final structure the librarian will use, but a scaffolding to react to and revise. Reacting to a draft is consistently faster than generating structure from scratch.

Second, use AI for description writing once the structure is set. Database descriptions, section introductions, and annotated resource entries are repetitive writing tasks that consume significant time and benefit from AI drafting. For example: "Write a 60-word description of JSTOR for undergraduate sociology students who are searching for peer reviewed articles for the first time. Plain language, no jargon, emphasize what types of content students will find and how to access it from off campus." Such a prompt produces usable descriptive text that the librarian reviews and adjusts for accuracy, rather than starting from a blank field. Furthermore, the same prompt structure can be applied across multiple resources once the pattern is established, converting what was a forty five minute description writing session into a ten minute review and edit process.

Third, AI is equally useful for improving existing guides rather than only creating new ones. A librarian with a five year old guide that has accumulated outdated language, unclear descriptions, and inconsistent tone can paste individual sections into AI and ask for revision: "Here is a section of my library guide on evaluating sources. Rewrite it for first generation college students who may be unfamiliar with library terminology. Plain language, no jargon, welcoming tone, under 100 words." Such revision passes are considerably faster than rewriting from scratch, and they can be applied selectively to the sections most in need of updating while leaving well functioning sections unchanged.

The professional boundary that matters is the one between editorial and substantive content. AI can write the introduction to a guide, draft descriptions of resources, and suggest organizational structures. It cannot identify which databases are actually appropriate for a given subject area, assess whether a database's coverage is current and relevant to the patron population, or verify that resources it names are available through the institution's subscriptions. Such professional judgment is the librarian's contribution, and it remains the central, irreplaceable component of guide quality even when AI handles the writing.

Lesson plan development: AI as an instruction design partner

Library instruction sessions follow recognizable structural patterns (learning objectives, content sequence, active learning activities, closing assessment) that AI can scaffold reliably for any new context. The challenge is not generating a lesson plan structure; it is generating one that is appropriate for the specific course, assignment, audience, and time constraints a librarian faces. AI is substantially better at this contextual adaptation when the librarian provides rich context in the prompt rather than a generic request.

The most useful prompting approach involves specifying four variables: the patron population and course context, the assignment the session is supporting, the available time, and the specific learning outcome the librarian wants to prioritize. For example: "I have 50 minutes with a first year English composition class whose students are starting their first academic research paper. They will be using Academic Search Complete and the library catalog. The instructor has asked me to focus on finding and evaluating peer reviewed sources. Generate a lesson plan with a hook, two or three content segments, an active learning activity, and a closing. Active learning should require students to do something, not just watch." Such a prompt produces a structured session plan that the librarian can review, adjust, and execute, rather than spending preparation time on structural decisions that could be delegated.

AI is also useful for generating the variety of activity options and assessment questions that make instruction planning genuinely creative. When a librarian asks for five different ways to demonstrate Boolean searching rather than one, the list almost always contains an approach she would not have arrived at independently. For example, a prompt such as "Give me five different active learning activities for teaching database Boolean searching to first year students, ranging from low tech to digital, each taking no more than ten minutes" may produce a mix of familiar approaches and genuinely novel ideas. Such variety is particularly useful when planning for different course contexts that require different pedagogical approaches, given that the same content delivered to an honors seminar versus a developmental writing course requires very different activity design, and generating options quickly allows the librarian to select appropriately.

The single shot instruction challenge (designing a session that achieves meaningful learning outcomes in a single fifty minute encounter) is well served by AI's ability to help prioritize. A librarian uncertain whether to cover database searching, source evaluation, or citation management in a constrained time slot can ask AI to reason through the tradeoffs: "I have 30 minutes with a criminal justice seminar. The students need to know how to find peer reviewed sources and how to evaluate them. I cannot cover both well in this time. Which should I

prioritize given this population and assignment type, and what is the most efficient way to briefly address the other?" Such a structured prompt helps the librarian make a principled decision rather than defaulting to whatever was covered last time. In order to get the most useful response, the prompt should include any relevant information about what students already know and what the assignment actually requires.

Additionally, AI is valuable for generating formative assessment questions, discussion prompts, and session evaluations: the components of instruction design that require creativity but not subject expertise. For example: "Write three exit ticket questions I could use at the end of a single shot session on database searching to assess whether students can identify the difference between a popular and a peer reviewed source. Questions should take under two minutes to answer and be usable as a quick paper slip." Such assessment tools strengthen instruction design without requiring additional preparation time, and the resulting questions can be saved and adapted across multiple sessions.

Adapting instruction content for different audiences and contexts

One of the most time consuming dimensions of library instruction work is not creating new content but adapting existing content for new audiences and contexts. A lesson plan developed for a nursing seminar requires substantial revision to serve a history capstone course. A LibGuide written for graduate researchers needs to be reframed for community college first year students. A patron FAQ appropriate for a research university population may be entirely unsuitable for a public library community. AI is exceptionally well suited to this adaptation work, and developing a reliable adaptation workflow is among the highest return AI practices available to instruction librarians.

The adaptation workflow proceeds from a clear prompt structure: provide AI with the existing content, specify the original context, describe the new audience, and give format guidance. For example: "Here is a handout I developed for a graduate research seminar on database searching. Rewrite it for first generation community college students who have never used an academic database and may be anxious about library research. Same information, accessible language, welcoming tone, no library jargon, under two pages." Such a prompt produces a genuinely reworked document rather than a superficially revised version, because the audience specification gives AI the information it needs to make meaningful changes to vocabulary, examples, assumed knowledge, and tone.

Reading level adjustment is one of the most practically useful adaptations AI can perform. A librarian working with a patron population that includes English language learners, students with reading disabilities, or community members unfamiliar with academic vocabulary can use AI to adjust existing library content to a specified reading level consistently and quickly. For example: "Rewrite this database search tutorial for a sixth grade reading level. Keep all the factual content but simplify vocabulary and sentence structure. Add brief explanations of any library terms that cannot be avoided." Such a revision preserves the professional content of the original while making it genuinely accessible to a wider patron population: a significant service expansion that would otherwise require hours of careful rewriting.

Furthermore, AI is useful for repackaging the same content in different formats for different delivery contexts. A detailed lesson plan can be condensed into a single page handout; a handout can be expanded into a slide deck outline; a slide deck can be revised into a self paced LibGuide tutorial. For example, a librarian who has developed a strong single shot lesson plan can ask AI to generate a condensed version for a fifteen minute embedded session, a self paced version for asynchronous delivery through the course management system, and a single page quick reference card for students to keep. Such format multiplication from a single base document is one of the most efficient uses of AI in instruction work; the core content is developed once, and AI handles the structural adaptation. Indeed, for librarians whose instruction programs serve many course formats and delivery contexts, this single workflow can reduce content development time substantially across an entire semester.

Teaching information literacy in an AI world: updating the framework

The ACRL Framework for Information Literacy was developed before generative AI became a mainstream tool, and most of its six frames remain as relevant as ever, but the examples, entry points, and instruction strategies that bring those frames to life require updating for the current information environment. AI has not replaced the framework; it has created new occasions for each frame to become visible to students in ways that earlier information environments did not provide.

The **Authority is Constructed and Contextual** frame is the most directly implicated. AI has no authority in the framework's sense; it does not derive credibility from institutional position, expertise, peer review, or community recognition. It synthesizes patterns from sources that may or may not carry authority, without preserving or communicating the authority of those sources. For instruction purposes, this frame provides the vocabulary for a productive

conversation: rather than telling students that AI is untrustworthy, the librarian can ask students to articulate what makes a source authoritative and then apply that framework to AI output. For example, an instruction activity might ask students to evaluate an AI generated paragraph using the same authority criteria they would apply to a news article or a database result, and to identify what is different about a source that has no author, no institution, no publication history, and no peer review.

The **Information Creation as a Process** frame connects directly to how AI generates text. Students who understand, even conceptually at the level covered in Module 01, that AI text is produced through statistical prediction rather than research, expert judgment, or evidence evaluation will evaluate AI output more accurately than students who treat it as a search result. In order to teach this frame in an AI context, librarians can design activities that make the creation process visible: asking students to compare how a Wikipedia article, a peer reviewed journal article, and an AI response are each produced, and identifying what quality controls apply to each. Such comparison makes the production process, which is invisible in the final text, becomes a subject of critical attention.

Searching as Strategic Exploration is the frame most frequently violated by student AI use. Students who use AI as a search tool, asking it to find sources rather than using it to develop search strategies, are applying a search logic to a generation tool. In order to teach this frame in an AI context, instruction should make the distinction between generation and retrieval experiential rather than explanatory: show students what happens when they ask AI for sources (hallucinated citations), then show what happens when they use that same topic in a database search (actual retrievable documents). Such a demonstration, taking five to ten minutes in a single shot session, produces more durable conceptual change than any amount of verbal explanation.

The **Research as Inquiry** frame provides useful language for the librarian's role in an AI environment. Inquiry involves forming questions, evaluating evidence, reaching conclusions, and revising them: an iterative process of intellectual engagement that AI can support at specific points but cannot conduct on the patron's behalf. For instruction purposes, this frame is most useful when teaching students that AI assistance is appropriate when it supports their inquiry and inappropriate when it substitutes for it. Such a distinction (AI as inquiry support versus AI as inquiry replacement) is the conceptual foundation of appropriate research AI use, and the ACRL Framework provides the professional vocabulary to teach it clearly. Additionally, for librarians who teach embedded instruction and have ongoing relationships with courses, this framing gives students a durable evaluative principle that applies beyond any single session or assignment.

Working with faculty on AI syllabus policies

Faculty are consulting librarians about AI syllabus policies at an increasing rate, not always because librarians have actively positioned themselves as the appropriate resource, but because the intersection of academic integrity, information literacy, disciplinary norms, and assignment design that AI policy questions involve is genuinely library territory. Librarians who develop a clear consultation framework for these conversations will find themselves in a professional role that strengthens the library's relationship with faculty across disciplines and that directly serves the institution's need for coherent, pedagogically grounded AI guidance.

The consultation approach that works most consistently begins with the assignment's learning objectives rather than with the AI question. Faculty who arrive asking "Should I allow AI in this assignment?" are asking a policy question; what they need is a learning design conversation. In order to redirect productively, the librarian can open with: "Let's start with what this assignment is meant to develop in your students. Once we understand what learning the assignment is supposed to produce, we can identify what forms of AI use support that learning and what forms undermine it." Such a reframe shifts the conversation from prohibition versus permission (a false binary) to purposeful design, which is where the librarian's information literacy expertise is directly applicable.

AI can assist with the consultation itself. A librarian preparing for a faculty meeting on AI policy can ask Claude to generate a range of policy examples organized by restriction level (from fully restricted to fully permitted with disclosure required), drawn from the growing body of publicly available higher education policy language. For example: "Generate five examples of AI syllabus policy language representing different levels of restriction. Include a single sentence rationale for each level and note what assignment types each level is most appropriate for." Such a prepared overview gives the faculty member concrete language to react to, which is considerably more productive than a blank policy drafting session. Furthermore, the librarian's role in the consultation is not to advocate a particular policy position but to help the faculty member think clearly about what the assignment's learning goals require and what policy language accurately expresses those goals.

Additionally, librarians can help faculty think through the practical dimensions of AI policy that are often overlooked in initial drafting: how students will be expected to document their AI use, what constitutes required disclosure, what the consequences of policy violation are, and how the policy aligns with the institution's broader acceptable use standards. For example, a faculty member who writes "AI is not permitted" without specifying what counts as AI use (grammar checkers, autocomplete, translation tools) is creating an ambiguous policy that generates more confusion than clarity. Such specificity is the contribution the

librarian's information and policy background enables, and it improves the policy's practical effectiveness in ways that faculty who are expert in their disciplines but not in information policy may not anticipate.

There is no doubt that the library's position in faculty AI policy conversations is strongest when the librarian arrives with preparation: examples from peer institutions, a working knowledge of the ACRL and ARL guidance documents, and familiarity with what the institution's own AI acceptable use policy requires. Such preparation signals that the library is a substantive resource for this conversation, not merely a referral destination, and it establishes the kind of faculty partnership that supports instruction programs and collection decisions far beyond the immediate AI policy question.

Patron communication: drafting for every audience and channel

Patron facing communication (emails, FAQs, website content, signage, newsletter items, social media posts, handouts) represents one of the highest return, lowest risk applications of AI in library work. The tasks are repetitive and structurally familiar, the outputs are reviewed before use, and the efficiency gains are immediately visible. Such conditions make patron communication the natural starting point for librarians experimenting with AI for the first time, and the application where the professional case for AI use is easiest to demonstrate to skeptical colleagues.

Reference email response is the most immediately useful application for most reference and instruction librarians. A librarian who spends twenty minutes drafting a careful, patron appropriate response to a common research question can reduce that time substantially by asking AI to draft the response and then editing for accuracy and personalization. For example: "Draft a friendly, helpful email response to the following reference question from a community college student: [paste question]. The student should be directed to Academic Search Complete for this topic. Warm and accessible tone, under 200 words, offer to follow up if the student needs help." The resulting draft typically requires editing and accuracy review but is structurally sound and covers the essential elements. Such drafting is particularly useful for questions that require explaining database search to a patron who may be unfamiliar with library resources; the structural explanation is delegated to AI, and the librarian's time is spent on the elements requiring professional judgment.

FAQ content for library websites and LibGuides is well suited to AI drafting because the format is standardized and the patron population is knowable. For example, a librarian developing FAQ content for a new interlibrary loan service page can ask AI to draft answers to common ILL questions in plain language appropriate for undergraduate students, then

review and adjust the content for accuracy against the library's actual ILL procedures. Such a workflow produces a complete FAQ draft in considerably less time than writing from scratch, with the important professional step of verifying every procedural claim, since AI may generate plausible ILL procedures that do not match how the specific institution's service actually operates.

Signage, handouts, and print communication benefit from AI's ability to condense complex information into brief, patron accessible language. For example: "Write a 50-word sign explaining how to access library databases from off campus, for students who do not know what a proxy server is. Plain language, no jargon, focus on what the student needs to do, not how the technology works." Such a prompt produces immediately usable sign text that the librarian reviews for accuracy and adjusts for institutional branding. In order to produce patron communication that reflects the library's voice consistently across channels, it is useful to provide AI with a brief institutional description and sample existing communications; the resulting output will more closely match the established tone and register. Indeed, once a librarian has developed a prompt that produces patron communication in the library's voice reliably, that prompt becomes a reusable template for every subsequent communication task of the same type.

Building a sustainable AI assisted instruction practice

The efficiency gains available from AI in reference and instruction work are real, but they accumulate most reliably for librarians who develop deliberate, systematic practices rather than using AI ad hoc for individual tasks. A librarian who restates her institution and patron population context to the AI at the start of every AI session is losing half the value that custom instructions or persistent project contexts provide. A librarian who drafts each LibGuide introduction independently is not building on the successful patterns she has already developed. In order to realize the full efficiency available, a sustainable practice requires investment in three dimensions: persistent context, accumulated resources, and shared team knowledge.

Persistent context means setting up AI tools so they know, across every conversation, who the librarian is and what professional context shapes her work. As described in Module 02, custom instructions in ChatGPT, persistent project contexts in Claude, and similar features in other tools allow the librarian to specify institution type, patron population, preferred tone, and professional norms once rather than in every session. For example, an instruction librarian at a community college might set up custom instructions that include: "I work at a community college with a large first generation student population. I am an instruction and

reference librarian. My content should always use plain language, avoid academic jargon, and assume students may be anxious about library research." Such persistent context changes the quality of every subsequent output without requiring the librarian to reestablish it each time, and it represents a one time investment that returns value across every future session.

Accumulated resources refers to building a repository of AI drafted content that has been reviewed, approved, and used: the LibGuide introductions that tested well, the email templates that patrons found helpful, the lesson plan elements that worked in specific course contexts. Such a repository serves two purposes: it provides ready made content for adaptation to new but similar situations, and it preserves the successful prompts that produced it. For example, a LibGuide introduction that AI drafted and the librarian approved becomes both a usable document and a model prompt; the librarian can return to the original prompt, adjust the course and discipline specifics, and rapidly produce a new introduction for a different guide. Module 10 covers building and maintaining a prompt library systematically; the principle here is that every successful AI output is also a prompt pattern worth saving.

Shared team knowledge represents the practical dimension of sustainable AI practice that individual effort cannot replicate. Such knowledge is most valuable when shared with library colleagues, who are otherwise required to discover the same effective approaches independently. For example, a team of three instruction librarians who each develop their own AI practices independently may arrive at different approaches, with different levels of efficiency, without ever comparing results. A brief team discussion (what prompts have worked for lesson plan development? what context specifications produce the most consistent output?) develops a shared vocabulary for AI assisted instruction work that benefits the entire team. There is no doubt that the most sustainable institutional AI practice is one that is collective rather than individual, documented rather than implicit, and revisited regularly as tools and institutional needs evolve.

PRACTITIONER NOTE

The task that consistently converts skeptical colleagues is demonstrating AI on a recognizable, unglamorous task: drafting a LibGuide introduction in three minutes and editing it in two more. Once a colleague sees AI save time on a task they recognize as genuinely tedious, the conversation shifts from 'should we use this' to 'what else can we use it for.' The most effective demonstration task is not AI for a complex instruction challenge but AI for the kind of routine, time consuming writing task every librarian has on their list. The efficiency is visible immediately, and visible efficiency is the most

persuasive argument available. The professional questions about appropriate use and verification follow naturally once the efficiency case has been made, and those are conversations worth having, because they build the kind of deliberate, calibrated practice that produces real institutional value.

KEY TAKEAWAYS

- AI accelerates the editorial components of subject guide creation (structure, descriptions, introductions) while the librarian's knowledge of actual resources and database coverage remains irreplaceable.
- Effective lesson plan prompts specify four variables: patron population, the assignment being supported, available time, and the specific learning outcome to prioritize.
- Adapting existing content for new audiences (adjusting reading level, reformatting for different delivery contexts, rewriting for different patron populations) is among the highest return AI practices available to instruction librarians.
- The ACRL Information Literacy Framework remains the right conceptual foundation for AI instruction; the frames of Authority, Creation, Searching, and Inquiry each have direct application to AI evaluation.
- Faculty AI syllabus policy consultations are most productive when the librarian reframes the question from 'allow or prohibit' to 'what does this assignment's learning objective require.'
- A sustainable AI instruction practice requires three investments: persistent context that eliminates repetitive setup, accumulated resources that build on successful prompt patterns, and shared team knowledge that multiplies individual efficiency.

References

APA 7TH EDITION

1. Association of College and Research Libraries. (2016). *Framework for information literacy for higher education*. American Library Association. <https://www.ala.org/acrl/standards/ilframework>
2. Association of College and Research Libraries. (2025, October). *AI competencies for academic library workers*. American Library Association. <https://www.ala.org/acrl/standards/ai>

ACRL AI Competencies covered

Use & Application

Sub-competencies: 4.1, 4.2, 4.3 · ACRL AI Competencies (2025)