

Level 3: Advanced

All Librarians

Module 15

Agentic AI: what it means

AI agents take actions rather than just answering questions. Understanding what they are, and what library workflows they could handle, is the next frontier for digital librarians.

By Yulia Brusova

15 min read · Reviewed June 2026

WHAT YOU'LL BE ABLE TO DO

- 1 Explain what an AI agent is and how it differs from a standard AI chatbot
- 2 Set up a Claude Project with custom instructions and persistent context for a library workflow
- 3 Identify at least two library workflows where agentic AI would reduce manual work
- 4 Understand the risks of agentic AI and when human oversight is required
- 5 Evaluate whether an agentic AI approach is appropriate for a specific library task

The term 'agent' is being applied to a wide range of AI systems right now, from simple chatbots rebranded under a more impressive name to genuinely autonomous systems capable of taking sequences of actions in the world with minimal human direction. For library professionals, the technical debates about what qualifies as a true agent are less important than a practical understanding of what is actually available today, what it can accomplish in library settings, and where human oversight remains non-negotiable. This module addresses each of those questions directly.

Agents versus chatbots: the operational difference

The distinction between a chatbot and an agent is not a matter of marketing terminology. It reflects a genuine difference in how the system operates. A chatbot responds to a single prompt and produces a single output; the interaction is complete when the response is delivered. An agent, by contrast, pursues a goal across multiple steps, taking actions and making decisions based on what each step produces before proceeding to the next.

For example, when a practitioner asks a standard chatbot to draft a reference follow-up email, the chatbot produces the email based on the prompt alone. An agent given the same goal might first retrieve the patron's original question from a connected system, examine the practitioner's previous correspondence to match the established tone, verify that the resource it intends to recommend is still available, draft the email, and present it for review before sending. Such a process involves multiple steps, tool use, and intermediate judgment, none of which characterizes a standard chatbot interaction.

The characteristics that define agentic AI are worth understanding in concrete terms. Multistep execution means the system breaks a goal into component tasks and completes

them in sequence rather than in a single response. Tool use means the system can take actions beyond generating text, including searching the web, reading and writing files, calling external APIs, and sending communications. Memory means the system can retain information across steps within a task, using earlier results to inform later ones. Judgment means the system makes decisions about what to do next based on what it has found, rather than following a fixed script.

It is evident that most AI systems library professionals encounter today sit at various points on a spectrum between simple chatbot and fully autonomous agent. Claude Projects and ChatGPT Projects occupy the lighter end of this spectrum; they are not truly autonomous, but they persist context across conversations and support more complex assistance than a single-turn chatbot. Fully autonomous agents capable of executing multistep library tasks without human oversight are emerging but are not yet common in library practice, which means the current moment is well suited for building familiarity with the lighter agentic tools before the more autonomous ones arrive.

Claude Projects: the practical starting point for library work

Claude Projects is the most accessible agentic-adjacent tool available for library work at present. It does not take autonomous actions without practitioner input, but it functions as a persistent, context-aware workspace that carries institutional knowledge across every conversation within the project. In order to understand what makes this practically significant, it is useful to compare it directly with a standard AI conversation.

In a standard AI conversation, every session begins without any knowledge of previous sessions, the institution, the patron population, or the practitioner's preferences. Each session requires the practitioner to reestablish context before the work can begin. A Claude Project eliminates this overhead entirely. Custom instructions set for a project persist across every conversation within it, uploaded files remain accessible for reference throughout, and the project functions as a standing briefing that never needs to be repeated.

For example, a reference desk project might be configured with custom instructions that read as follows: this assistant supports a reference librarian at a community college library serving primarily first-generation college students and working adults; always recommend specific databases from the attached list rather than generic suggestions; maintain a warm, clear, jargon-free tone in all patron-facing drafts. With those instructions in place and the library's database list and FAQ document uploaded, any subsequent conversation in that project operates with full institutional context. Such a setup means that drafting a patron response requires only pasting the patron's question, not explaining the library's context first.

Furthermore, a second project can be configured independently for a different library function. For example, an instruction design project might carry custom instructions describing the institution's course levels, typical session length, and ACRL Framework alignment, with existing lesson plans uploaded as templates. Such a project allows the practitioner to request a lesson plan draft for a specific course section without re-explaining institutional parameters each time. Additionally, projects can be created for collection development support, assessment drafting, or any other recurring library function where persistent context would reduce setup time at the start of each working session.

What agentic AI can realistically accomplish in libraries today

There is a meaningful gap between what vendors claim agentic AI systems can do and what is practical in library settings at present. An honest assessment of current capability is more useful than an optimistic one, because overestimating reliability leads to implementations that fail in consequential ways.

Several applications are practical today and have been demonstrated in library contexts. Multistep research assistance is one of the most reliable: an agent can search a topic across multiple sources, read and synthesize the results, identify gaps, generate follow-up search terms, and produce a structured report for the practitioner's review. For example, a practitioner preparing for a research consultation can submit the patron's topic and assignment description and receive a pre-consultation briefing document that includes relevant databases, preliminary search terms, and suggested clarifying questions, all in the time it would previously take to open a single database. Such a workflow does not replace the practitioner's judgment in the consultation itself; it reduces the preparation time required to enter that consultation well-prepared.

Document processing pipelines represent another practical application. For example, an agent can ingest a set of assessment survey responses, extract recurring themes, count references to specific library services, and compile the results into a structured summary ready for the practitioner's review and interpretation. Such pipelines are most reliable when the documents follow a consistent structure and the extraction task is well-defined.

Conversational FAQ interfaces, which allow patrons to ask questions about library services and receive responses drawn from a knowledge base the library provides, are also practical today, though they require careful configuration and ongoing review of the responses being generated.

Undoubtedly, several capabilities that receive significant vendor attention are not yet reliable enough for library use without substantial caution. Fully autonomous catalog record creation

without human review, autonomous management of patron communications, and multisystem workflows that take action in an ILS, send email, and update spreadsheets without human approval steps at each stage are all emerging but not yet sufficiently reliable for production library use. Any agentic system that takes action in a patron-facing context requires particularly careful evaluation, because the speed at which agentic AI executes steps means that errors propagate quickly before a human has the opportunity to intervene.

Custom instructions and memory: building a persistent institutional assistant

The most practically useful agentic capability available to library professionals today is not autonomous task execution but persistent memory and context. A well-configured set of custom instructions is, in effect, a standing briefing that transforms every AI interaction from a cold start into a contextualized one. There is no doubt that the thirty to forty-five minutes required to write good custom instructions once represents one of the highest-return investments available in AI-assisted library work.

In order to write effective custom instructions for library use, the practitioner should include several categories of institutional and professional context. The institution type and size establish the operating environment. For example, "community college library serving approximately eight thousand students" tells the AI system something materially different from "research university library." The patron population characteristics matter equally: first-generation college students require different communication register and different assumptions about prior research experience than graduate students in a specialized research program. Such differences should be stated explicitly rather than left for the AI to infer.

Role and primary responsibilities provide the system with the information it needs to calibrate the relevance of its suggestions. Preferred tone for different output types, patron-facing versus internal, formal report versus quick email, prevents the system from defaulting to a generic register that does not match the library's established voice. Additionally, specifying particular databases or resources to reference, and explicitly stating what the system should not do, such as recommending Wikipedia as a primary source or suggesting resources not available through the library's subscriptions, prevents outputs that require correction before use.

Memory features, where available in a given tool, extend the value of custom instructions beyond what can be stated in advance. For example, Claude's memory feature and

ChatGPT's memory allow the system to retain information the practitioner specifies across conversations over time, including preferences, ongoing project context, and institutional details that would otherwise need to be restated. Such features are best treated as a convenience rather than a guarantee: important context should always be verified as present in the memory before it is relied upon for consequential work, because memory features in current AI tools are not fully reliable and can fail to retrieve stored information under some conditions.

Library use cases for agentic AI that are practical now

The use cases described here are practical in current library settings, not aspirational scenarios dependent on capabilities that do not yet exist. Each requires practitioner review of the output; none operates without human judgment at the point of use.

Research consultation preparation is one of the highest-value applications. For example, before a scheduled consultation, the practitioner submits the patron's stated topic and assignment description to a configured Claude Project. The system identifies the most relevant databases from the library's subscribed list, generates a set of preliminary search terms organized by concept, notes any subject guides relevant to the topic, and drafts three clarifying questions to ask the patron at the start of the session. Such a workflow compresses fifteen minutes of preparation into approximately three, and produces a more systematic starting point than preparation done under time pressure without AI assistance.

Collection development scanning is a second practical application. In order to review a vendor's new title list against the library's collection criteria, the practitioner uploads the list to a configured project along with the collection development policy and, if available, a current holdings list. The system flags titles that meet specified criteria, notes any that appear to duplicate existing holdings, and produces a prioritized recommendation list for the practitioner's review. Such a process does not replace the practitioner's selection judgment; it performs the initial filtering that would otherwise require checking each title manually against multiple criteria.

Assessment work is a third practical application with significant time return. For example, a practitioner with raw assessment data from a patron satisfaction survey can submit the responses to an agent and request a summary of the three most significant trends, a two-paragraph narrative suitable for the annual report, and a list of areas that warrant follow-up investigation. Such a summary does not replace the practitioner's interpretation of what the findings mean for the library's services; it compresses the initial synthesis from hours to minutes, leaving the practitioner's time available for the interpretive work that requires

professional judgment. Additionally, instruction session follow-up drafts, collection use reports, and grant narrative sections describing library programs are all tasks where agentic AI can produce a competent draft from structured input, which the practitioner then reviews, corrects, and approves before use.

PRACTITIONER NOTE

A Claude Project configured for instruction design returns its setup time within the first week of use. The most significant change is not speed but consistency: sessions planned with access to previous lesson plans as templates start from a more substantive baseline than those planned without that context, and the quality of the starting draft improves the quality of the final session design.

KEY TAKEAWAYS

- An agent pursues a goal across multiple steps using tool use, memory, and intermediate judgment; a chatbot responds to a single prompt and produces a single output.
- Claude Projects is the most accessible agentic-adjacent tool for library work: it persists custom instructions, uploaded files, and context across every conversation within the project, eliminating the cold-start problem.
- Multistep research assistance, document processing pipelines, and conversational FAQ interfaces are practical in library settings today; fully autonomous catalog work and patron-facing communications without human review steps are not yet reliable.
- Custom instructions are the highest-return investment in AI-assisted library work: thirty to forty-five minutes spent writing them once produces compounding time savings across every subsequent interaction.
- Memory features in current AI tools extend the value of custom instructions but must be treated as a convenience rather than a guarantee; important context should be verified as present before being relied upon for consequential work.
- In every practical agentic library application, the practitioner's review and judgment shape the final output; the agent handles synthesis and initial drafting, not the decisions that require professional accountability.

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>

ACRL AI Competencies covered

Knowledge & Understanding

Use & Application

Sub-competencies: 2.1, 4.4 · ACRL AI Competencies (2025)