

Submitted on: 02.07.2026

Curating Trust and Connections: Transforming the Information Ecosystem with Instruction, Grounded Curation, and Open Pedagogy

Cyril Oberlander

Library, Cal Poly Humboldt, Arcata, United States

E-mail address: Cyril.Oberlander@Humboldt.edu



Copyright © 2026 by **Cyril Oberlander**. This work is made available under the terms of the Creative Commons Attribution 4.0 International License:

<https://creativecommons.org/licenses/by/4.0/>

Abstract

In an era of context collapse and a disrupted information landscape, where the origins and authority of information are increasingly obscured, this presentation demonstrates how Cal Poly Humboldt Library is advancing AI literacy and skills, while also developing innovative trusted digital environments. By leveraging the expertise of library professionals in curation and information literacy, we can transform AI from a source of misinformation into a powerful tool for discovery and learning.

This session presents a multifaceted case study from Cal Poly Humboldt, highlighting a spectrum of AI integration, from AI Worldcafés, to AI workshops, to the creation of custom, source-grounded AI applications within an open pedagogy framework. The instructional approach ranges from Slice of AI & Pie, to AI SkillShops, designed to help students, staff, and faculty move from using general AI chatbots to utilizing grounded systems. Unlike general models that rely on broad internet datasets, the instruction focuses on finding and uploading trusted sources in AI tools. These AI Skillshops address a need to share techniques to build trust and AI literacy using grounded tools informed by user-uploaded trusted sources that improve AI citing directly to high quality works, enhancing critical thinking, and reducing opportunities for hallucinations and unsupported claims.

The presentation further details the development of Ask Alex Humboldt, a collaborative research project that utilizes a Retrieval-Augmented Generation (RAG) pipeline. This application leverages the library's Digital Commons institutional repository as knowledge source to a vector database, connecting

academic repositories to large language models. Our goal is to ground AI responses in peer-reviewed scholarship and specialized open science repositories, our Library seeks to add meaningful connections of academic research with community science and observations.

Crucially, these projects employ an open pedagogy approach, involving students as active co-creators in the design and testing of agentic AI. By working under the constraints of open-source and proprietary models, students learn the mechanics and realities of AI; such as document parsing with Docling and parallel processing with Ray clusters, while grappling with the ethical implications of algorithmic bias and data privacy. This hands-on collaboration empowers students to move from passive consumers of AI to informed architects of trusted information systems.

Ultimately, this session demonstrates that the role of libraries in the age of AI is not just to teach literacy, but to lead the development of scholarly-informed AI agents. By combining expert curation with agentic AI design, libraries can strengthen authority and transparency within the digital information environment, ensuring that the future of research remains both innovative and trustworthy.

Keywords: AI literacy; grounded curation; retrieval-augmented generation; open pedagogy; information trust

Introduction: The Disrupted Information Landscape

Generative AI, recommendation algorithms, deepfakes, and post-truth media are reshaping how people create, encounter, and trust information. As provenance cues disappear, researchers need more than retrieval skills; they need strategies for understanding and shaping the systems that mediate information. How do we architect spaces that people can trust and want to use? This presentation paper advocates that libraries can do so by teaching users to curate the source contexts that AI systems rely on, building source-grounded discovery systems, and by inviting students as co-creators of transparent AI tools. Through Cal Poly Humboldt Library's Slice of AI & Pie, AI SkillShops, and the Ask Alex Humboldt project, the paper proposes grounded curation as a solid foundational strategy to combine crucial information literacy instruction with trustworthy AI system use and design.

The phenomenon of *context collapse* obscures shared understanding of information, and generative AI can obscure the origins, transparency, and authority of digital information, thereby disrupting the information landscape and reinforcing the need for trust. Historically, viewers often treated photographs as physical evidence. However, in modern digital environments,

information lacks a structural “container” such as a peer-reviewed journal or a printed newspaper, which serves as a critical framework for viewers to evaluate information provenance. For example, when an AI-generated image of Pope Francis wearing a puffer jacket went viral across social media in 2023, observers later identified the image as a fabrication. (Dolan, 2023) Kevin Sebastian describes generative AI as a “significant epistemic rupture: one in which the logic of information production is increasingly obscured, potentially leading to a mystification of knowledge systems. In this environment, the risk of misinformation, algorithmic bias, and information validity dependency escalate, demanding new literacy frameworks that extend beyond retrieval-based competencies.” (Sebastian, 2025) The rupture of the factual information landscape fundamentally destabilizes the ACRL frame Authority is Constructed and Contextual. For information literacy instruction, this reality suggests moving beyond traditional fact-checking checklists toward a deeper misinformation literacy that critiques how generative systems synthesize, or outright fabricate information, as well as inappropriate uses of AI. When metadata and traditional containers vanish from the information ecosystem, libraries must foster AI literacies that effectively meet researchers’ current and rapidly changing conditions. Those AI skills needed include finding trusted sources to architect quality context of AI systems. To understand how these concepts manifest in practice, we can look to the library workflow of metadata application.

This paper demonstrates examples of how a library can strengthen trust in disrupted information ecosystems by teaching users how AI systems mediate information, grounding AI outputs in curated sources, and involving students as co-creators of transparent and accountable AI systems. Metadata provides one practical site where the promise and risk of AI become visible. AI can assist description, discovery, and classification, and it can also misclassify, overgeneralize, or fabricate. A clear illustration for the need and value of human verification in AI systems, and for source-grounded architectures are found in Kyle Morgan’s 2025 article “Using AI to Auto-Tag Graduate Theses” at Cal Poly Humboldt provides a perceptive analysis of a practical approach of using AI for applying metadata to graduate theses in an institutional repository. Morgan documented the methods for identifying UN Sustainable Development Goals (SDG), and evaluated ChatGPT’s accuracy in applying 121 SDG tags, and found 22 errors, yielding an accuracy rate of about 82%. Morgan’s insightful comparison of human error and AI errors is astute: attributing human error to a “lack of careful reading of the metadata”, while AI errors were more “consistent choices that reflected a different interpretation of SDG relevance.” Morgan’s analysis of the AI results and testing prompt changes demonstrates the importance of human review, as well as grounded architectures. In an alarming finding, Morgan warns that when the AI chatbot could no longer access the past threads, it fabricated a plausible summary instead of disclosing the limitation. (Morgan, 2025)

When AI tools are not transparent and confident in misinformation, libraries must build new trusted architectures. When we provide students with real-world AI projects, we develop their literacy, systems knowledge, and technical skills. For instance, Morgan's postscript mentions working with Courtney Rowe, a Computer Science Library Scholar Intern to fine-tune the model and develop a weighted score and metadata tagging using ModernBERT model, and their approach was the foundation for expanding and fine-tuning the accuracy of applying UN SDG tags. (Morgan, 2025)

This approach aligns well with UNESCO's *AI Competency Framework for Students*, which prepares learners to become responsible users, including appropriate use of AI and co-creators of AI systems. The framework includes twelve competencies across four dimensions: AI competency, human-centered mindset, ethics of AI, AI techniques and applications, and AI system design. These dimensions provide a useful lens for library-based AI literacy instruction, which often include learning how AI systems work, evaluating their ethical and social implication, applying AI tools to authentic information needs, and participating in the design and evaluation of source-grounded systems, making them more transparent, explainable, and provenance and human verification more visible. (Miao 2024) In addition to AI literacy instruction with existing systems, libraries benefit greatly from designing internships or student research that evaluates AI systems, and furthermore, design new AI systems.

While AI tools can instantly synthesize essays or images, offering information seekers a single platform that effortlessly fabricates plausible information based on prompts, disrupts the information seeking and evaluating landscape, reducing the diversity of information sources into a single presentation. Libraries curate diverse, vetted information ecosystems, albeit less frequently used by information seekers; the declining use trends we see with print books are starting to appear in article download data from electronic databases at our and many other libraries. For instance, when this author compares the 2014 ACRL electronic usage data with 2024 across the 23 campuses of the California State University Libraries, 21 campuses saw a significant decrease of full text article use, collectively a 39.6% decrease from 20.6M to 12.4M downloads. This represents significant changes in research behaviors, requiring more evaluation to determine how much is this change driven by factors including: AI discovery, assignment design, open access, and/or enrollment changes.

This decline in traditional database usage underscores a deeper crisis in how users find and evaluate information, a reality explored in Amber Willenborg and Robert Detmering's insightful article "I Don't Think Librarians Can Save Us" and study of librarian instructional strategies with misinformation, participants expressed the CRAAP Test was not suited for today's information environment, many instead choosing lateral reading of websites. The article discusses the need for "changing the material conditions of instruction librarianship that limit and constrain librarians' role. (Willenborg, 2025) In fact,

these disrupted conditions present an ideal opportunity to scale information literacy from individual behavior to system-level architecture. In web-based environments, lateral reading requires a researcher to leave a webpage to cross-reference claims across independent, trusted domains. In a generative AI environment, where a chatbot synthesizes a single answer stripped of sources, traditional evaluation becomes nearly impossible.

To address this, we propose that grounded curation functions as the algorithmic and system-level equivalent of lateral reading. Instead of relying on the user to fact-check a generated unsupported output, grounded curation embeds lateral defense directly into the technical architecture of the system. By utilizing RAG or uploading vetted sources into AI tools, we constrain the large language model's retrieval context. Essentially, the system is forced to look laterally at a pre-vetted, library-curated corpus of trusted data before generating a response, effectively improves provenance and reduces unsupported generation, by anchoring outputs explicitly to selected sources, all cited, retaining provenance and attributions. Teaching students to understand the mechanics of the information architecture is a key component to advancing AI literacy, and a familiar role for library instruction, ironically, grounded AI helps us reinforce many of the learning objectives, and ease the researcher's workload. Furthermore, applying open pedagogy techniques of fostering student agency with co-design and co-creating knowledge provides an opportunity for developing critical thinking and team skills with real-world sensemaking at a time of rapid transformations.

The Shifting Paradigm of the Library

As AI-based search systems support repurposing of documents into a holistic answer, and reduce the human effort in retrieving, reading and analyzing of information, curating sources becomes an effective process to teach research skills today.

In 2014, Trevor Owens' vision of the archivist using digital tools as exo-suits, empowering armor; "We need tools that let us quickly work across massive amounts of items and objects by extending and amplify the seasoned judgement, ethics, wisdom, and expertise of the archivist-in-the-machine." (Owens, 2014) Today's AI tools are very much the exo-suits Owens envisions, and libraries remain architects of research environments that extend researchers' capabilities, wherever possible connecting people and ideas. Researchers are developing new 'reading' behaviors and research workflows with a variety of tools; skimming assistants, automated summaries and overviews, immediate analysis across large sets of documents, and infographic visualization of information, to name only a few.

As traditional "containers" of information vanish, the fundamental role of the academic library must adapt from traditional information curators and

collection managers to active curators of community engagement and connections that are supported by a range of digital literacies.

In the 2025 book *Social Intelligence*, researchers Wang, Shang, and Zhang provide a detailed understanding of a variety of human-AI collaborative learning activities that leverage the strengths of each intelligence to mitigate their individual limitations. Powerful examples of crowdsourcing and AI are especially of interest to addressing scale of learning, addressing immediacy with emergency firefighting operations, and even bias and privacy in social spaces. *Social Intelligence* may provide us an interesting way to reconceptualize keeping the ‘human in the loop’, and in fact, inviting humans to collaborate with AI and each other.

In a disrupted information ecosystem, academic libraries can strengthen trust by teaching users how AI systems mediate information, grounding AI outputs in curated sources, and involve students as co-creators of transparent, accountable AI systems. By reimagining the library as a living laboratory and an AI architect, institutions can empower trust through curated connections. Blending advanced information literacy instruction with grounded AI curation and open pedagogy frameworks allows libraries to build a scalable learning infrastructure, one that rebuilds trust through empowering students with real-world information and career-ready skills. To operationalize this framework, Cal Poly Humboldt Library developed a scaffolded, iterative approach to AI Literacy, inviting campus community members to informal conversational spaces to targeted tools, and use specific workshops.

Iterative AI Literacy: From Conversational Spaces to Targeted SkillShops

The rapid emergence of useful generative AI agents created an opportunity to examine information literacy in the age of AI everywhere. At Cal Poly Humboldt Library, in addition to adding AI literacy instruction to library instruction, we organized nine campus events entitled [Slice of AI & Pie](#), as a cooperative workshop series with the Center for Teaching, facilitating campus conversations and learning about AI. The series began on November 1, 2024, and became a periodic conversational campus gathering to learn and share out about rapidly evolving AI technologies. The Slice of AI & Pie was an informal combination of discussion and workshop, and included pizza or sweet pies, and it included topics ranging from “reading critically with AI and reading AI critically” to teaching specific tools, such as Copilot, Gemini, NotebookLM, as well as use cases for AI technologies. Presenters varied, including faculty, staff, students, and vendor representatives.

By 2025, the Library began providing specific AI workshops as [SkillShops](#), 50-minute co-curricular workshop series, and these AI SkillShops were either platform or use case specific. Kyle Morgan, Scholarly Communications Librarian offered *Using ChatGPT in Writing for Publication*. The author

provided 12 SkillShops in 2025-2026, several workshops on Gemini and NotebookLM for research grant writing, resume refining, and marketing; these SkillShops (50-minute workshops at the library) were tool specific and presented a variety of use cases, for illustrative purposes.

The instructional design of Oberlander's AI SkillShops balances technical skills with critical information literacy. Every session begins with a collaborative co-design exercise, asking attendees what tasks they hope to accomplish using AI tools. This immediately gauges the audience's prior knowledge and aligns the workshop's goals with their immediate needs. The instruction frames prompt engineering not as a technical trick, but as an extension of critical thinking. The facilitator introduces prompt anatomy by guiding participants to integrate clear context, explicit goals, precise constraints, and user-selected source material. To ground these techniques, the session explains how Large Language Models (LLMs) and Retrieval-Augmented Generation (RAG) function. This technical overview highlights the crucial distinction between individual data limitations and the robust privacy and security benefits of enterprise-level AI accounts. Finally, these practical AI literacy exercises are paired with ethical discussions covering user privacy, intellectual property, and copyright. Ultimately, attendees learn to treat AI interaction as an iterative loop testing, refining, and sharing successful strategies with colleagues, while ensuring their authentic human voice always drives the final output, and their verification as essential. One of the key questions raised for discussion is "To what extent is AI OK?" and the range of possibilities include what the syllabus states, personal preferences for starting with brainstorming or only asking for revisions. These conversations serve as a sharing of the variety of perspectives and thoughts of appropriate AI use.

One of the important learning outcomes is students appreciating the value of finding and utilizing vetted information. This notion of grounded curation of AI parallels the concept of lateral reading, akin to a researcher investigating provenance cues, the grounded AI researcher acts as information architect curating the knowledge base that AI output is structurally tethered to for verification. Students learn about information architecture using library databases and AI to improve the quality of results. By introducing source-grounded AI platforms alongside general multimodal chatbots, the author demonstrates how uploading curated datasets effectively restricts a large language model's retrieval context. This conceptual framework introduces a powerful shift; instead of relying solely on the end user to cross-reference and fact-check a synthesized, source-less output, the researcher acts as an information architect and curator. They proactively construct the trusted knowledge base to which the AI's generation is structurally tethered.

This hands-on methodology is reinforced through several practical, everyday use cases that demonstrate the value of constrained machine retrieval:

Ideation & Professional Development: We demonstrate how these tools can streamline administrative workflows, such as generating targeted email drafts while maintaining an authentic professional voice. For career readiness, students learn how to upload their resumes alongside specific job descriptions, prompting the system to perform a gap analysis and generate probing, reflective questions that help them uncover overlooked skills.

Institutional Communication & Strategic Analysis: In marketing and outreach contexts, the workflow showcases how a single AI response created with source documents can be dynamically adapted across platforms, from generating a press release based on a website, to tailoring that press release to a tailored donor email invitation. On an analytical level, the workshops guide attendees through AI-assisted policy reviews and website search engine optimization audits, using the technology to compare web interfaces, extract key metadata for search engine optimization, and run multi-document pattern recognition.

Visual Literacy & Research Inquiry: Even visual generation is critiqued with AI generated infographics and presentation slides, revealing the tremendous benefits of data visualization, and also exposes glaring text and typographical errors that require careful human verification.

Ultimately, these sessions conclude by demonstrating custom AI agents that can permanently store customized instructions and curated sources, eliminating the need to repeatedly upload files or re-type complex prompts. Whether analyzing data through smartphone-based computer vision features or mapping connections across 50 sources in a research notebook, the overarching objective remains unchanged. The library leverages these emerging, agentic AI tools to supercharge the traditional research frame of *Research as Inquiry*, transforming pattern recognition into an human driven adventure in curiosity and verification. During the AI SkillShop attendees were encouraged to follow-along, and after the SkillShop, attendees are encouraged to practice and ask for assistance. Those hands-on practice sessions help expand our understanding of what users are interested in doing with AI tools.

The NotebookLM AI SkillShop is similar, although the focus is creating larger curated projects that ground the LLM with up to 50 sources - PDF, videos, and many other formats, as well as, utilizing the agentic AI tools for creating podcasts, quizzes, slides, etc. Among the many practical uses of a notebook such as reviewing policies and regulations, identifying gaps and grounding recommendations in trusted sources, creating infographics or presentations slides has been an invaluable way of discovering and exploring crucial connections across a set of research articles. Although having read the articles and identified many of the connections across disciplines, AI sometimes surfaces unexpected connections that warrant human review, this is especially true in AI generated infographics and presentations, as well as prompts that

ask to explore connections across disparate sources. A very practical example of this in the AI SkillShop or one on one consultation with students focuses on mapping out career interest to review a resume aligned to a position announcement or type of career. In the SkillShop, the author demonstrates using their CV, position announcement, with curated sources helpful for seeking jobs. The example prompt asks to identify gaps in the CV, and to provide questions to help the author remember other experiences and skills that may align with the position. The relevance of probing AI questions was extremely useful to reflect and surprisingly rediscover details the author overlooked. Surfacing unexpected connections is very much a core part of the ACRL Frame 3: Research as Inquiry. AI tools provide gap analysis and pattern recognition that powerfully support research as inquiry, and enhance researcher's curiosity.

Demonstrations of Agentic AI tools to create customizable podcasts, video overviews, slides, interactive quizzes and flashcards to learn and practice with source topics, concept maps to navigate source materials, and infographics, broaden the potential uses for AI in research, teaching and learning, as well as practical work. Explaining how to customize the focus, style, and specifics is an important element to improving the alignment of the results with the needs of the attendees. Similarly, demonstrating Agentic AI tools as conceptually a Large Language Model composed of many AI agents with specific roles, and curated sources in a Retrieval-Augmented Generation (RAG) constrains the context of the generalized LLM, helps us understand the architecture of generative AI and Agentic AI. Essentially, Agentic AI includes conversational AI chatbots that can invoke external tools to perform actions, such as querying a library system API or enabling a Multi-Agent Orchestration to run a recursive series of agent-driven activities (analyze intent, formulate search strategy, execute searches, filter relevant results, and write). Navigating AI possibilities is challenging, given the plethora of emerging platforms and features. One of our current library exhibits shows how various AI models generate infographics although each had the same prompt and source document. The exhibit gamifies verification, inviting viewers to spot AI errors and receive chocolate at the library office. Highlighting AI literacy skills can be a simple engagement board that poses a question, providing examples, information, and certainly a treat.

In addition to teaching an AI tool with various examples, the author also integrated AI literacy as a component to specific topic workshops. For example, during a career & resume strategies workshop, the author introduces and demonstrates appropriate use of AI tools, showing that prompt design showcases AI's strength in identifying gaps, alignment with position descriptions, and most importantly, generating questions to help discover and determine whether additional experiences and skills were overlooked. Similarly, AI literacy was an important element of a website marketing and search engine optimization (SEO) SkillShop. The primary audience for these

SkillShops consisted of college administrative specialists, however, everyone was invited to build outreach and recruitment skills that included Google Maps to improve Google Search results, and use AI for extensive webpage comparisons, gathering key metadata for SEO, as well as, providing identifying revision recommendations. Each SkillShop was a crucial and timely opportunity for open discussion concerning appropriate use of AI, bias, environmental concerns, and misinformation, and because AI tools rapidly changed, attendees were encouraged to share their practices and tips. Given the perceptions of AI as powerfully transformative, environmentally destructive, biased and dangerous, it has been crucial to design the workshops with conversation in mind. In fact, having facilitated campus World Café (<https://theworldcafe.com/>) on various important topics, the author was asked by Associated Students to organize their World Café on AI for students. This structured conversation provided an opportunity to explore insights and questions students have with the use of AI. Students identified opportunities and challenges they see using AI, and discussed how to ensure equitable access, clear and fair expectations for AI use so students know what is ethical in their courses, and discussed issues of environmental impact, privacy, safety, and mental-health disclosures. One of their identified priorities was the need for training on appropriate uses of AI, and the Library responded by developing and facilitating more AI SkillShops. Setting time at the start of each workshop to discuss the interest in what they want AI to do, and at the end, providing some time with hands-on assistance has proven invaluable with collaboratively exploring and lowering barriers and risks for using AI. While our AI SkillShops teach individual show to ground AI tools with vetted sources, the Ask Alex Humboldt AI project extends this same principle from individual practice to institutional experimentation with a source-grounded system.

Grounded Curation: Co-Creating Agentic AI as Project Based Learning

One of the most important learning outcomes of our AI workshops is to explain how uploading sources to generative AI can reduce Large Language Model (LLM) unsupported outputs by narrowing the systems retrieval context, although these methods reduce errors, it also reaffirms that human verification remains essential. When the sources are intentionally selected peer-reviewed articles or curated within digital repositories, the curation can significantly enhance the quality of AI results and support of research, teaching, and learning. Although this can be done within generative AI chats by uploading files, AI platforms often limit the number of uploads to 10 articles, and others 50, however, techniques such as Retrieval-Augmented Generation (RAG) enable thousands of articles to be included as the knowledge base for the LLM to draw context from. In 2024, a Library team decided that we would create Ask Alex Humboldt AI, a RAG custom agentic AI to ingest content across our institutional repositories to support an integrated search with a natural language conversational chat. The SkillShops teach individuals how to ground AI tools with vetted sources. The Ask Alex Humboldt AI project extends the

same principle from individual practice to institutional experimentation with a source-grounded system built from curated institutional repository content.

The Ask Alex Humboldt AI Project

The Ask Alex Humboldt AI project brought together three goals, one for a natural language search across our various institutional repositories and digital libraries, another a metadata generation tool, and lastly, a longer-term vision to combine open science projects such as eBird and iNaturalist with academic content with immersive interfaces. The Ask Alex Humboldt AI project was a collaboration of the Library and the Information Technology Services (ITS), and included Brianne Hagen, the Metadata and Discovery Librarian, Cyril Oberlander, Library Dean, Ravi Chalasani and Brian Campbell, ITS Academic Technology Specialists, and John Gerving, a Computer Science Student Assistant, an ITS hire mostly funded by the Library.

Why build a specialized, source-grounded conversational search engine rather than rely on traditional library discovery layers? Publisher databases and library systems often have a rigid keyword matching that cannot infer, or worse, are weighted to discover content in their portfolio. Our prototype agentic AI experiment demonstrated inference by searching for “Bigfoot”; traditional searches were unable to match records labeled with Sasquatch, and we sought to enhance the search results across disparate repositories seamlessly. In addition, this project was about collaborative research and learning about RAG based agentic AI systems, and while the process included frequently changing rapidly evolving large language models, we recognized the learning curve would parallel changes in AI technologies. Essentially, the project ingests repository PDFs, extracts text, chunks the text, creates embeddings, stores the embeddings in a vector database, retrieves relevant passages in response to user questions, and prompts a language model to generate cited answers from those passages. The AI architecture lessons learned included:

Data Ingestion and Processing: The project relied on building the knowledge base from the library’s Digital Commons institutional repositories. Because we selected a text based large language model, we limited the source to an API call and extraction of over 6,500 PDF documents in the repository. From an extraction, the process required recognizing all the texts in the PDFs. These documents were parsed into machine-readable text using Docling, an advanced OCR tool capable of recognizing multi-column academic paper layout structures, and in order to process the scale of this endeavor, we used Ray Data in a Ray Cluster, a tool to handle computing at scale by splitting the processing workloads of thousands of PDFs across multiple GPUs in parallel.

Vector Database: The parsed text was segmented into small semantic chunks of (250 words with text overlap), and these chunks were converted into dense

vector representations using a Hugging Face Transformer embedding model, and the final content text chunks and their corresponding embeddings indexed into a PostgreSQL database with vector extensions.

Agentic Orchestration: When a user queries the interface, e.g. “Where does Bigfoot live?” the pipeline moves into an agentic RAG workflow, deploying an autonomous loop using a retrieval agent to execute precise vector, metadata and full-text searches over the PostgreSQL knowledge base. We started with the OLMo2 large language model, an open weight and open source LLM, however, due to its limitation in instruction and citation formatting, we selected Gemini 2.0 Flash. The writer agent synthesizes the top vector contexts into a natural language response, with every assertion made by the agent, it was required that a direct clickable URL citation lead back to the original source PDF document. In order to improve the accuracy of the results, 35 prompt test questions were developed alongside ground truths determined to evaluate the accuracy of the results, although we learned that some errors were the result of misconfigured questions.

We have never found Bigfoot, however, Ask Alex Humboldt effective response succeeded in citing the repository appropriately: “Bigfoot lives in the minds and in the collective consciousness of culture (1). However, if you’re looking for a place to explore the legend of Bigfoot, rural Humboldt County is a great place to start. (2) (5)...”

Ravi Varma Kumar Bevara et al in “Prospects of Retrieval-Augmented Generation (RAG) for Academic Library Search and Retrieval” outlines important beneficial features and ethical issues in designing RAG systems for libraries, concluding that RAG integration “holds promise for revolutionizing academic library services” and “Future research should focus on refining RAG’s technical architecture to address library-specific needs, exploring cost-effective adaptations for resource-constrained environments, and developing comprehensive frameworks to ensure ethical compliance.” (Bevara, 2025)

AI RAG projects are strongly aligned with the ACRL Frame 1: Authority is Constructed and Contextual, our understanding of AI advanced by enabling creative opportunities to learn and construct contextual authority using vetted scholarship, and we do see tremendous value in engaging students in the process. This project was well aligned with ACRL Frame 2: Information Creation as Process: Students are learning the process behind AI information generation, the engineered pipeline with structural choices, vulnerabilities, and biases.

Engaging students in the exploration of AI technology development is an open pedagogy framework that advances our collective understanding of AI, and for students, provides meaningful and transformational experiences as architects of trusted information environments. This project invited the team into critical

conversations and presentations regarding agentic AI, algorithmic bias, intellectual property, privacy and security, representation, search accuracy, sustainability, and more. That said, there are, however, many significant challenges and limitations with developing custom AI. The rapid changes in models, methods, and infrastructure was difficult to keep up with, and for sustainability, we needed a larger team of students that had varying graduation dates. In addition, RAG systems inherit the limitations and biases of the selected corpus, the trusted-vetted sources are socially and institutionally constructed, and any proprietary AI tools raise cost, licensing, privacy, sustainability questions, or external control policies may have a significant impact on equitable access. Building research prototypes are a significant step to advance our understanding of AI design, and aligned with our goal of open models and open-source AI, however, we would have to significantly expand the organizational infrastructure needed to support any open models for production level services.

While the AI environment continues to change rapidly, combining open pedagogy project-based learning with students to curate high-quality sources of AI agents provides a foundational transparency for rebuilding trust and digital authority for academic and community users alike. This practice can be done with a workshop demonstration of uploading sources to an AI tool, or building custom AI agents. In addition to our Ask Alex Humboldt project, we worked with a team of Computer Science students in a software engineering course to experiment with AI and Geographical Information System (GIS) georeferencing automation of aerial photos as a class project. That team was able to create a machine learning application that reasonably identified some potential georeferences for aerial photos of Humboldt County, however, the accuracy depended on clear landmarks, and did not leverage flight plan data to optimize the accuracy of the results to automate referencing our collection of over 40K aerial images. More recently, a team of Library Scholar Interns worked on the Terraform project, a new version of the augmented reality sandbox to provide realistic satellite imagery based on integrating an AI diffusion application with a Kinect camera.

Each project-based learning approach advances our understanding and application of AI, providing an invaluable opportunity for students to develop their analytical, communication, teamwork, and technology skills that prepare their launch into their future, and we continue to iterate how to address real-world challenges in transformative ways.

Curating Connections Curricular and co-curricular transformational foundations
Open Pedagogy is “a site of praxis, a place where theories about learning, teaching, technology, and social justice enter into a conversation with each other and inform the development of educational practices and structures... an access-oriented commitment to learner-driven education AND as a process of designing architectures and using

tools for learning that enable students to shape the public knowledge commons of which they are a part.” Robin DeRosa and Rajiv Jhangiani, [Open Pedagogy Notebook](#)

The principles of Open Pedagogy include a practice of developing meaningful connections by engaging students that contribute to public knowledge and transformation by advancing their portfolios, understanding, and launching their careers. Open pedagogy turns AI literacy from a lesson about tools into a shared practice of building, testing, questioning, and improving information systems. The practice involves learning together, co-mentoring of implementing disruptive technologies is more akin to an adventurous hike than some may realize. Real-world projects involve risk, but they also amplify opportunities to explore and foster professional identity as well as enhance a sense of belonging for teams and the library. In essence, the Library as a living laboratory cultivates interdisciplinary adventures with collaborative meaningful connections that positively transform lives. During the spring of 2026, we tested methods to increase the number of open pedagogy initiatives, including sponsoring 14 internships, 9 student teams in Engineering 205 building interactive exhibits, and sponsored the Hackathon. The varying strategies enabled us to better understand a variety of engagement opportunities, capabilities and limitations of curricular and co-curricular tactics. This year, we plan to offer more AI challenges, including inviting English Library Scholar Interns to reimagine literature book stacks with immersive AI experiences. Providing a range of opportunities to work on real-world projects is of great interest to students, and through an iterative design of our infrastructure to launch our students, we are building strong foundations for student agency and trust in a disrupted information landscape.

Implications for Libraries

Reframing the global library conversation requires shifting our focus from what a library has, in terms of its physical and electronic collections and spaces, to what a library inspires, its community connections and experiential opportunities. In a disrupted ecosystem where information containers are often stripped from content, a loss of contextual provenance, our traditional keyword-matching discovery systems are no longer sufficient. Nor are general AI chatbots.

The future of global libraries relies on building open and scalable learning infrastructures that democratize advanced technical and analytical skills, and empower people skills. By integrating real-world challenges that include AI into the university curricular and co-curricular programs, the library establishes itself as an essential partner in student retention, belonging, workforce readiness, and student recruitment. Combining AI literacy with information curation workshops to agentic AI design, we are ensuring that the future of

knowledge remains both deeply transformative and trustworthy. The library is an active architect and partner of a trustworthy future.

Conclusion

As academic ecosystems confront the misinformation environments, the examples of human-AI collaboration becomes the baseline for algorithmic accountability. By applying human curation to define the boundaries of AI retrieval, libraries effectively synthesize the computational scale of machine learning with the ethical oversight, verification, and critical thinking of human intelligence. Looking back to 1996, when Apple Researchers Bonnie Nardi and Vicki O'Day published "Intelligent Agents: What We Learned at the Library" the astute study of what reference librarians do, and the design of intelligent agents, recognizes a hybrid and balanced collaboration of human and machine agents:

"Our argument for computer agents is that, if designed correctly, they can personalize technology use, creating a revolution in the user experience... In turn, we ask librarians to consider and speak out about their critical role in the design of a diverse information ecology. Our vision of a diverse information ecology is that of a richly patterned collaborative system of users, human agents and software agents. There is a larger emerging vision of a global information infrastructure that is populist, egalitarian, inclusive." (Nardi, 1996)

Building trust as community impact is co-created when libraries step into the intersections of community, curriculum, and technology. Cal Poly Humboldt Library's case study demonstrates that through AI SkillShops, source-grounded curation, and student co-creation, libraries teach users to curate the contexts that AI systems use, and support verification of the outputs those systems. Ask Alex Humboldt AI illustrates both the promise and the limits of RAG; retrieval can improve transparency, however, human judgment remains essential. For libraries, the transferable lesson is to demonstrate the value of trusted collections, teach verification as a design practice, and invite students into accountable AI development and testing. In this way, libraries power transformation by building learning infrastructures that make AI more transparent, participatory, and trustworthy.

By blending information literacy, grounded AI curation, and open pedagogy frameworks, libraries can build a scalable learning infrastructure that restores trust in information, while equipping students with real-world, career-ready skills. This adaptable and transformational strategy can include:

Instructional engagement, from *Slice of AI & Pie* to platform and use specific AI SkillShops. These programs explicitly train students, staff, and faculty to transition from general AI chatbots to specialized systems grounded with user-

selected resources. By demonstrating tools, users learn to constrain Large Language Models to verified datasets that reduces the risk of hallucinations, and rebuild trust.

Building a customized agentic AI *Ask Alex Humboldt AI*, enabled us to learn how to engineer an agentic AI RAG that anchored the chatbot to the information in our institutional repository. Research efforts such as these align well with IFLA's publication in 2025 *Entry point to Libraries and AI* call for developing AI Retrieval Augmented Generation (RAG) application as an opportunity to improve library services and to foster an open approach. (Cox, 2025)

Inviting our students to be active co-creators in the design, testing, and implementation of AI will continue to be an engagement strategy to create meaningful connections and real-world experiences. This hands-on collaboration empowers students to be informed architects and researchers of trusted information. Libraries are rebuilding trust in the age of AI through connecting people and ideas.

Acknowledgments

Sincere gratitude and recognition to the time and talent of so many students, student assistants, and interns, and all the staff and faculty that make learning so inspiring, to the Ask Alex Humboldt Team, and to the support of Mary, Kate, and Julia Oberlander.

References

ACRL, (2016) *Framework for Information Literacy for Higher Education*, 2016. <https://www.ala.org/acrl/standards/ilframework>

Bevara, Ravi Varma Kumar et al., (2025) "Prospects of Retrieval-Augmented Generation (RAG) for Academic Library Search and Retrieval), *Information Technologies and Libraries*, June 2025. <https://doi.org/10.5860/ital.v44i2.17361>

CSU, (2026) *AI Commons*, <https://genai.calstate.edu/> viewed 6/16/2026

Cox, A.M., De Brasdefer, M. (2025), *IFLA AI Entry Point for Libraries and AI*, June 5, 2025. <https://repository.ifla.org/handle/20.500.14598/4034>

DeRosa, Robin and Rajiv Jhangiani, (2017) "Open Pedagogy and Social Justice," Digital Pedagogy Lab, <http://www.digitalpedagogylib.com/open-pedagogy-social-justice> Blog [Open Pedagogy Notebook](#) viewed 6/16/2016

Dolan, Leah, (2023) "Look of the Week: What Pope Francis' AI puffer coat says about the future of fashion", *CNN Style*, March 29, 2023.
<https://www.cnn.com/style/article/pope-francis-puffer-coat-ai-fashion-lotw>

Miao, Fengchun et al., (2024) "AI competency framework for students",
Unesco, 2024. <https://doi.org/10.54675/JKJB9835>

Morgan, Kyle, (2025) "Using AI to Auto-Tag Graduate Theses," *Information Technologies and Libraries*, December 2025.
<https://doi.org/10.5860/ital.v44i4.17381>

Nardi, Bonnie and Vicki O'Day, (1996) "Intelligent Agents: What We Learned at the Library," *Libri*, 1996, p. 59-88. <https://doi.org/10.1515/libr.1996.46.2.59>

Owens, Trevor, (2014), *Mecha-Archivists: Envisioning the Role of Software in the Future of Archives*, <https://trevorowens.org/2014/05/27/mecha-archivists-envisioning-the-role-of-software-in-the-future-of-archives/>

Sebastian, J. (2025). Reframing Information Seeking in the Age of Generative AI: A Critical and Humanistic Approach. *ACRL 2025 Conference* 527-539. Minneapolis, MN: The Association of College and Research Libraries (ACRL).
https://oasis.library.unlv.edu/lib_articles/771

Wang, D., Shang, L., & Zhang, Y. (2025). *Social intelligence : the new frontier of integrating human intelligence and artificial intelligence in social space*. Springer Nature Switzerland.

Willenborg, A., & Detmering, R. (2025). "I Don't Think Librarians Can Save Us": The Material Conditions of Information Literacy Instruction in the Misinformation Age. *College & Research Libraries*, 86(4), 535. doi:
<https://doi.org/10.5860/crl.86.4.535>

World Cafe (2026) <https://theworldcafe.com/> viewed 6/16/2026