# Library Freedom Project's Critical AI in Higher Education Toolkit

LIBRARY FREEDOM PROJECT

June 2026

# Table of contents

Introduction

What this toolkit is and who it is for

Goals for this toolkit

What are we talking about when we talk about AI?

AI in higher education

Preliminary organizing with your colleagues

Assessing the landscape at your institution

When AI becomes a required job duty or deliverable

AI in library vendor tools

Talking points against AI

Resources

Credits

# Introduction

Aligned with library worker core values advocating for the free flow of information (access), equity, intellectual freedom and privacy, for the public good, and sustainability, this guide addresses the ways that the drive toward uncritical artificial intelligence (AI) use forecloses academic ideals of critical inquiry, ethical knowledge production, and maintaining the space for academic freedom.

The following guide and resources respond to the overwhelming unease academic library workers have expressed in the Library Freedom Project's survey on library workers and sentiments about artificial intelligence in libraries.[1] Academic library workers are at the forefront of innovation assisting researchers, faculty, and students in exploring, building, and using emerging technologies.

The American Association of University Professionals' ad hoc Committee on Artificial Intelligence and Academic Professions summarizes potential harms from haphazardly injecting AI into higher education. Those harms include, but are not limited to, "potential work intensification and job losses and through its implications for intellectual property, economic security, and the faculty working conditions that affect student learning conditions."[2]

[1] Library Freedom Project, "AI In Libraries Survey 2025" (forthcoming); National Information Standards Organization, "Survey of Library Professionals Reveals Perceptions of AI," NISO Insights, January 2021.

[2] American Association of University Professors. "Artificial Intelligence and Academic Freedom." AAUP. Accessed May 11, 2026.

# What this toolkit is and who it is for

These days, AI is seemingly everywhere in higher education. Administrators, vendors, and even some library workers laud the perceived benefits of outsourcing tasks to computer-assisted technologies. Universities and colleges launch new partnerships with large tech companies, announce AI implementation strategies, and promise a bright future where students are better prepared for the workforce with these technologies, while AI tech company CEOs brag about the number of jobs that will become obsolete and disappear. Educational tech (ed-tech) companies continue to scoop up student data and, with little-to-no warning, add AI tools to existing platforms and software products.

At the same time, AI skepticism grows. There are numerous concerns about academic dishonesty, privacy, environmental costs, copyright, lack of attribution, and more. Alongside all of these ethical issues, AI tools perform poorly at many tasks and rarely deliver on all of their creators' and/or providers' promises. Voices in libraries against AI are growing louder, despite AI proponents having the most institutional power over licensed software.

Library Freedom Project created this toolkit for fellow AI skeptics and refusers, as well as library workers who want to better understand the position of those refusing AI. We offer guidance for building our power against uncritical AI adoption. This toolkit outlines some of the main areas of struggle against AI, how AI is showing up in higher ed contexts, and ways to organize as workers against this normalization that outsources our thinking and creativity.

Learn more about our work at Library Freedom Project: libraryfreedom.org.

# What are we talking about when we talk about AI?

AI, or artificial intelligence, refers to a number of different technologies and tools, some of which have been in existence since the mid-20th century.[3] Generally, AI refers to technologies or tools which have been programmed to perform pattern recognition and/or optimize performance over time. Common examples of AI tools include chatbots, recommendation engines, and virtual assistants. Often what is meant when someone refers to AI is Generative AI, a form of this technology which uses pattern recognition and optimization, such as with large language models (LLMs), to generate text, videos, or images.

In addition to describing a collection of technologies, AI is a marketing term, as well as a sociopolitical project pushed by the largest tech companies in the world. At the time of this writing, it is also the dominant economic engine of the US economy. A chorus of economists are worried about what the inevitable AI bubble burst will mean for the global economy.[4]

[3] Ghost in the Machine. Directed by Valerie Veatch. 2026.
[4] "Here's Why Concerns about an AI Bubble Are Bigger Than Ever." NPR, November 23, 2025. https://www.npr.org/2025/11/23/nx-s1-5615410/ai-bubble-nvidia-openai-revenue-bust-data-centers.

# AI in higher education

Framing AI's incursion onto campuses as inevitable is a shortcut to justifying the uncritical spread of artificial intelligence tools throughout higher education. The range of higher education functions offloading human thought and care–including, but not limited to student services, assessment, teaching, and AI-bot use by Human Resources–many higher education administrators are eager to join the AI bandwagon.

**The Student Experience**

Students are being sold a frictionless college or university experience that promises convenience and job market appeal to justify skyrocketing tuition costs. However, authentic learning opportunities are put in jeopardy by administrators advancing an AI agenda. This is done under the guise of student demand for AI, disingenuous appeals to democratizing access to resources, and continuing commodification of education are prevalent rationales for encouraging widespread AI adoption on campuses.

Administrators eager to implement AI tools without thoughtful review and discussion with library workers are alleging student demand as one force behind their efforts. Students are, in effect, being framed as insisting on using AI. On the other side of the coin, they are being told that they are putting their education and future job prospects at risk by not keeping up with demands for AI-qualified workers.

Opening up access to University-branded AI tools for low-income students is another red herring administrators use. They create partnerships with Big Tech AI platforms or rebrand their models with University logos with claims of democratizing access when, in fact, what's being democratized is unfettered pillaging of students' data.

**The Instructor/Adjunct/Faculty Experience**

Pedagogical centers for teaching and learning are campus partners in developing and deploying institutional, school, department, and course-specific AI policies. These centers may offer workshops that introduce AI tools and explain how to use them in teaching and grading. They may also attempt to insert instructional staff into the role of policing AI use. Because of limited staffing, AI tool instruction or literacy sessions may be offered by the large commercial vendors that provide the tools themselves.

Because the goal of these companies is to capture and retain users, there may be little attention given to educating students about how the tools work from a technical perspective, their appropriate function, their source materials, or about the costs and potential pitfalls of using these tools. AI literacy is framed in terms of using the tool instead of critical understanding of the technology and its limitations.

AI products' active role in undermining learning, understanding, and synthesis of course content are overlooked or intentionally ignored. Students are presumed to be digital natives already enthusiastically using AI without ethical questions or qualms. This assumption robs educators of the opportunity to teach students about research skills, authorship, and citational ethics. It also reduces the emphasis on how to build arguments through deep research and the development of critical thinking skills.

Instructors may be encouraged to create policies outlining if and when students can or can't use AI. While on the surface, clear written policy is a positive aspect of open governance, students may be challenged by different schools, departments or professors having radically different policies around AI use in the classroom.

**The Campus Governance Experience**

Pushing AI onto campus communities without the explicit, structured inclusion of faculty, staff, and students in the planning and policy-making phases represents a fundamental violation of shared governance, effectively treating the university's core missions of teaching and research as technical problems to be solved by administration rather than intellectual endeavors to be stewarded by the community. When institutional policies are drafted without the direct input of the workers who will bear the burden of implementation—whether through altered pedagogical practices, compromised research integrity, or new surveillance mechanisms—the resulting initiatives lack legitimacy and fail to address the specific ethical and practical realities of the campus, reducing shared governance to a hollow formality while centralizing power in the hands of those who prioritize efficiency over academic autonomy.

**The Library Worker Experience**

Academic library workers may be subjected to library vendors inserting AI chatbots or assistants into popular databases or software products. These implementations may or may not be made in discussion with library middle and upper management, and it may or may not be possible to turn these AI tools off. Library administrators may also ask library workers to include these tools in their information literacy instruction, or demonstrate to library users how to implement AI tools in online search engines and leverage their summarization of scholarly articles. Despite austerity and budgetary concerns over rising software subscriptions costs, administrators may not move quickly enough to question companies hastily adding artificial intelligence into their platforms. Library workers should vocally object to the addition of any embedded AI tools unless there is a clear - and agreed-upon - method to review and assess these tools to determine what if any value they provide to library users.

There's a similar disregard for library worker expertise or interest in developing in-house AI tools. At one University, for example, a director of digital services sent non-unionized humanities subject librarians a set of ill-formed questions and prompts to enlist their labor in training a research assistant AI model. Subject librarians weren't consulted about the objectives of this labor, nor were their supervisors. As with previous knee-jerk initiatives, the librarians asked for clarification about the purpose of the exercise and for the opportunity to provide collaborative input, but received no response. Eventually, a graduate researcher walked the librarians through a beta AI assistant research platform that summarized materials for users, but, to date, the tool has not been released to the University community. This initiative and other vague remits for units to "come up with AI-focused" projects at the behest of the Provost's office that oversees the Libraries' budgets, in the end, seems merely an exercise devised to look engaged with the AI hype cycle.

# Preliminary organizing with your colleagues

Many of us have campus groups or committees that we're already involved with and have formed for information-sharing. Use these existing groups to begin your work around approaching AI tools and use them critically.

Who do you already organize with? Who do you know and trust?

A crucial aspect of accomplishing change is to work towards a concrete common goal. Some areas of critically assessing AI are easier to organize around than others. While wholesale AI refusal may be your ideal, consider in conversation with others the concrete steps required to work towards that ideal. Can you first build momentum around turning off embedded AI assistants in vendor-provided databases? Many library workers are critical of large, commercial library vendors and may be willing to contribute to a campaign that clearly identifies these commercial actors as the villain attempting to undermine professional expertise.

Having small, concrete goals also allows for, first, asking for smaller individual contributions of time and labor while reducing the potential for burnout and second, celebrating wins more frequently which helps maintain group morale and momentum.

Here are some general considerations for organizing on campus:

- The people you gather into a group should be independent of institutional entities to create meaningful actions outside bureaucratic confinements. Organizing outside official library channels, committees, academic departments, etc. affords your group flexibility in organizing.
- Is your workplace unionized? Your union may be a good place to start amplifying your work. Hosting a discussion at a membership meeting or putting out a survey to members is a great starting place. Also review the collective bargaining agreement–are there areas where AI may be affecting your working conditions or rights as protected under the contract?
- Instead of institutional email and messaging tools, consider using non-institutional communication tools, such as Signal, Discord, Slack, or personal email addresses.

- Depending on the public or private status of your organization, bear in mind communications that might be subject to state-level Freedom of Information Laws (FOIL) and/or the federal Freedom of Information Act (FOIA) requests. Treat these communications as if they were public or could be read aloud in a court proceeding.
- Survey your group to get a sense of the diversity of skills and identify external connections who can aid your organizing.
    - Who has technical expertise?
    - Who is knowledgeable about institutional bureaucracy?
    - Are there noticeable absences (e.g. types of workers, departments, etc)? You may need to make face-to-face, one-to-one conversations to connect people across your workplace who may not typically be in contact.
    - Your group may want to do a needs assessment to figure out which skills and connections are most valuable in this fight.
    - Who makes decisions about AI, and how could you influence them? Power-mapping helps identify potential allies and barriers and think about how to sway them.
    - How big of a push is there for AI at your institution? Threat modeling is a way to think about risks and what we can realistically do about them. This might help you recognize that there are huge pushes for AI in some departments, but not others.
    - You may reassess threats and map power structures as your group considers the campus landscape and strategizes about actions to take.
    - If your organization hosts pro-AI events, consider hosting your own AI-critical events to connect with others who are concerned. With this issue, we often hear from people who feel like the only one in the room who has concerns about AI. A book group, speaker event, or brown bag lunch discussion are ways to find one another.
    - Similarly, if you are in a meeting or at an event and someone asks a good question that is critical of AI, introduce yourself to them! They may be looking for like-minded colleagues to connect with beyond the event.
    - Finally, organizing can be a slog, especially when it feels like the forces of power and capital are stacked against you. Make sure you celebrate every win in your group, and find ways to build social connection and fun into your work.

# Assessing the landscape at your institution

Once you have found a trustworthy community of colleagues to organize with, the next step is to assess how AI shows up at your institution and find opportunities to complicate or challenge the pro-AI narrative locally.

To start, gather information about where and how AI is being promoted at your institution. Are there particular people in power who are pushing for AI on your campus? Where are those with power situated–which departments or units? These key players could range from high-level administrators, like provosts and university presidents, to middle managers with influence over smaller groups.

As you research the powerbrokers who are promoting AI at your institution, make sure to also research who they are collaborating with. You may find collaborations between individuals or large-scale partnerships with businesses invested in AI. Are there specific investors or private companies partnering with your institution to promote AI products? Are there schools, departments, or centers within your institution that are leading the push towards AI? Has your institution created new centers, departments, or other units that revolve around AI?

The State University of New York (SUNY) at Binghamton and private investor partnerships is one example of a large-scale collaboration to promote AI. In January 2026, the New York State Governor's Press Office announced that a group of donors had given $30 million to SUNY, which was combined with a $25 million investment from New York State, to create the Center for AI Responsibility and Research at SUNY Binghamton. According to the SUNY Binghamton newspaper, this $55 million investment was the largest gift in the history of the university.[5] SUNY is a member of the Empire AI consortium, a collaboration between ten New York higher education institutions and private funders to promote AI in the state.[6] The Empire AI consortium and the Center for AI Responsibility and Research exemplifies how colleges and universities have invested significant financial and human resources to implement AI-focused initiatives. Although the name of the center and mission statement of the Empire AI consortium both refer to responsible AI use, it is unclear what steps they are taking to address the harmful effects of AI, if any.

[5] Brhel, John. "Binghamton Receives Largest Academic Gift in University History to Establish AI Center." BingUNews, January 21, 2026.
[6] Empire AI. "About."

Another important question at the information-gathering stage is whether your institution has an AI policy in place or is considering creating an AI policy. At this writing, most institutions have formulated some version of an AI use policy for students and faculty. Depending on the size of your institution, multiple AI policies might be in place that guide how students use AI in their coursework, how instructors use AI in their teaching and research, and how administrators and staff use AI in their work.

There might also be library-specific AI policies that determine how AI is used in collection development or in electronic resources, e.g. the Jasper-Dubois County Public Library's "Collection Development Policy." For other examples of AI policies in place at colleges and universities, browse the "Collection of Guidelines and Policies for AI Use in Higher Education" by Angelo Belardi and search the websites of other academic institutions.

Once you have surveyed existing AI policies on your campus, discuss whether those policies are sufficient in terms of guardrails and library worker input. If they are not, a next step is to conduct power research (also known as power-mapping) to figure out who is in charge of creating and implementing such policies and how you can get them on your side. Power-mapping stakeholders, as described in the "Map the Power Toolkit," can help you understand who holds power, how they view AI, and what kind of AI policy they are likely to support. Start by gathering basic information about each stakeholder using a search engine and your institution's website. Then, expand your search to determine what connections each stakeholder has and how those networks may affect their views on AI. Finally, develop an action plan for how to use power-mapping to facilitate your efforts to create an AI policy or to create space for critical conversations about AI at your institution.

# When AI becomes a required job duty or deliverable

Many institutions of higher education are now requiring faculty or staff to make AI a part of the job. Academic labor unions and advocacy groups are working to resist these new job duties and implement strong worker protections. Some unions have adopted contract language to resist overbroad AI adoption and implement guardrails.

When AI-for-AI's-sake becomes a non-negotiable job duty, resistance can come in the form of creating work alongside AI that's informed and intentional. In our everyday tasks and tools where AI is appearing, we can offer to beta test new tools with vendors with a critical eye. We can also perform user or accessibility testing of AI tools. In publishing, we can create reports assessing the state of AI in your specialized area of higher ed or librarianship (e.g. teaching, info literacy, systems, archives, etc). We can also submit reviews of AI tools in databases and other systems for journals, such as the Journal of Electronic Resources Librarianship or Information Technology & Libraries.

We can create programming and events that bring AI-critical speakers to instruction sessions or campus. Contact organizations such as the AI Now Institute, the Distributed AI Research Institute (DAIR), or the Electronic Frontier Foundation (EFF), for ideas. AI-critical authors are keen to discuss their ideas and cross-disciplinary resistance; some include Karen Hao, Alex Hanna, and Emily Bender. We can also create book displays, booklists, or subject guides that are though-provoking about AI, aimed at the many audiences we reach on campus.

Curriculum development offers numerous opportunities, such as incorporating AI skepticism into information literacy curriculum, using LFP's AI curriculum materials or flyers, critiquing or drafting AI use statements on syllabi, ensuring that AI skeptics and critics are included on any faculty senate committees that determine policies for instructional faculty, creating workshops normalizing critical inquiry about AI, and creating spaces for participants (students, faculty, staff, etc) to practice their talking points based upon their values and what drives their position on AI use.

Starting or joining cross-community campus groups are another option for critically engaging AI hype. You can start or join a Lamp Club or Luddite Club, or your campus AI policy committee. Think creatively about which campus groups might be on your

side about AI – for example, many academic institutions have climate groups or committees, some with official institutional goals around climate and energy usage. These groups may be interested in partnership.

# AI in library vendor tools

Many vendors within the higher education ecosystem are integrating AI into products and services. In most cases, they use commercially available models such as ChatGPT or Claude rather than developing their own. Some vendors include AI in existing products for an additional charge; some develop entirely new products requiring separate subscription fees; and some include AI in existing products for no additional charge, though, considering how vendors tend to operate, it is likely that these additions will be reflected in future rising prices.

In an attempt to keep up with the onslaught of AI integration by library vendors, LFP has compiled an ongoing list of tools and functions that use AI. We intend for this to be a living document, and invite others to contribute additions, edits, and feedback using this form.

This database includes a variety of context about each tool or function, including:

- The vendor offering the product
- Product status: Is this a new product or an update to an existing product?
- Purported product function
- Types of data stored
- AI mode(s) used
- Ability or process for turning AI integration off

If you are working to resist AI integration from contracted vendors, the following tools may be helpful in your advocacy:

- LFP's AI questions to ask vendors
- Existing documents about privacy, vendor relations, collections priorities, or principles for contracts. Here are some examples:
  - University of Massachusetts Amherst's Framework for Provider Agreements
  - Columbia University's Principles for the Assessment of Text and Data Mining Acquisitions
  - Cornell University's Licensing for Privacy
  - Jasper-Dubois County Public Library's Collection Development Policy
  - Library Futures' Understanding Confidentiality and Risk Assessment, see also their resources on reading contracts
- Ithaka S+R's Generative AI Product Tracker

# Talking points against AI

Here are some fast facts about why AI hype requires library workers' response and action:

- Academic integrity is threatened.
- AI use is not, in fact, inevitable and there's still time to listen to library workers' expertise and voices in establishing ethical guardrails.
- Labor intensification and job losses due to AI adoption will impact library workers disproportionately.

**Speaking with specific audiences**

When preparing talking points, take some time to anticipate specific audiences' arguments and how you'll respond. In developing your talking points consider: typical tactics and arguments in favor of widespread AI buy-in, methods for gaining traction and buy-in from specific audiences, and pushing for direct and concrete action to empower communities.

**For Speaking with Administrators/Boards**

*Funders and the board are asking for AI use and innovation.*

Critical responses:

- Not unlike previous economic bubbles and crashes, the current AI investment landscape shows signs of overvaluation and overpromising on AI's capabilities.;
- Academic freedom and shared governance should always protect the curriculum from being shaped directly by funders or the board of trustees. The faculty have the primary responsibility for teaching and research.

*Students are already using AI.*

Critical responses:

- Are they, though? For what? Seek out institution-specific information about usage, or review national data like the Pew Internet reports on AI, which are data-informed checks on "everyone is using it!" arguments.
- Use of AI tools is demonstrably harming students' learning. It's the educational institutions' responsibility to make sure they are actively engaging in critical thinking and learning.
- University provosts responsible for AI integration, such as New York University's Clay Shirky, understand that generative AI is harming student's ability to connect with one another and create a sense of self in a collective world.

# Talking points against AI

*We will fall behind our peer institutions/organizations if we don't get ahead or keep up with what's happening in AI.*

Critical responses:

- For institutions that go to great lengths to avoid risk, uncritical adoption of commercial AI tools rooted in fear-of-missing-out is out of alignment with responsible, evidence-based decision-making and fiscal practices.
- Technology companies are increasingly providing disclaimers about false information gathered by their products as their product testing and implementation isn't comprehensive enough and safety measures revealed to be insufficient. Individuals and NGOs are suing these companies for devestarting harms aided by their technologies, including user suicide and AI-induced pychosis.

*AI is a university- / system-wide initiative.*

Critical response:

- If a university launches a campus-wide push for generative AI, it must simultaneously champion the academic integrity that such tools often undermine, ensuring that these initiatives do not devolve into unpaid marketing campaigns for commercial vendors seeking to monetize students and cultivate lifelong brand dependence rather than serve genuine educational interests.

*AI is good for accessibility.*

Critical responses:

- Despite claims that AI enhances accessibility, it currently performs worse than established, purpose-built screen readers, meaning that relying on it could actually degrade the experience for users who depend on these critical tools.
- Relying on AI to solve accessibility challenges violates the core disability justice principle of, "nothing about us without us," as it prioritizes algorithmic efficiency over the lived experiences and direct input of disabled people, often resulting in tools that reinforce ableist assumptions rather than delivering solutions that truly serve the community.

# Talking points against AI

**For Speaking with Students and Faculty**

*AI helps with efficiency and/or is time-saving.*

Critical responses:

- Does it, though? So far it looks like they are not. Data & Society's policy brief argues that the narrative of AI as a "waste-slasher" relies on a dangerous technosolutionist myth. The brief suggests that prioritizing AI for speed and cost reduction without robust safeguards often sacrifices the very quality and fairness that public services are meant to uphold.
- Commercial AI tools create a critical FERPA compliance gap, as these third-party platforms often lack the necessary data privacy safeguards to ensure that students' protected educational records and personal learning needs remain confidential and secure.
- Relying on AI to manage academic work risks eroding student resilience by outsourcing the cognitive struggle essential for deep learning, fostering a dependency that heightens anxiety when the technology fails and diminishing the sense of agency and accomplishment that comes from authentic intellectual effort.
- Research indicates that AI reliance induces "cognitive debt" and "metacognitive laziness," where the convenience of offloading thinking tasks erodes the deep processing and critical engagement necessary for genuine learning, ultimately leaving students with a fragile understanding that collapses when the tool is removed.

**For Speaking with Library Worker Colleagues or IT**

*AI will be so much better so soon and all these current problems will be fixed*

- Relying on the promise that "AI will fix itself soon" is a dangerous form of technological deferment that ignores the reality that current failures—from hallucinations and bias to massive energy consumption and labor exploitation—are often inherent to the architecture of these systems, not just temporary bugs waiting for a patch.
- History shows that scaling flawed models tends to amplify their harms rather than erase them.
- Treating these systemic issues as mere "growing pains" allows stakeholders to avoid the hard work of regulation, transparency, and ethical redesign, effectively betting public trust on a speculative future while accepting tangible damage in the present.

# Talking points against AI

*AI will make research easier and more efficient for students and faculty under pressurized grant-writing and publication deadlines.*

- Relying on AI to generate core academic work:
  - violates the principle of intellectual honesty, by substituting algorithmic output for the authentic cognitive struggle and original thought that constitute the very purpose of education.
  - undermines the core purpose of scholarly communication: the rigorous, transparent, and reproducible generation of new knowledge.
  - risks eroding the intellectual integrity and methodological accountability that define academic work, turning research into a "black box" exercise where the provenance of ideas becomes unclear and the ability to verify findings is compromised.
- True efficiency in academia isn't just about speed; it's about the reliability of the record, and substituting human critical judgment with generative shortcuts threatens to flood the ecosystem with plausible but unverified content, degrading the very trust and precision that grants and publications are meant to secure.
- We cannot allow a lifetime of intellectual and creative labor to be silently harvested by Big Tech to train proprietary models that enrich shareholders while devaluing the very human ingenuity that made those systems possible in the first place.
- AI-generated misinformation that fabricates citations in syllabi, ILL requests, and publications poison the scholarly record, forcing researchers to chase phantom sources and eroding the foundational trust essential for academic integrity.
- Replacing transparent curation with opaque, metric-driven AI in library databases transforms these neutral repositories into black boxes that prioritize popularity over accuracy, burying vital research behind un-auditable algorithms that compromise equitable access.
- AI-driven discovery in libraries creates a feedback loop that amplifies established topics while systematically burying emerging research, effectively siloing cutting-edge scholarship and cementing existing paradigms at the expense of intellectual diversity.

# Talking points against AI

*Including AI in research tools will diversify perspectives in research topics, projects, and outcomes.*

- In reality, vendor-controlled algorithms may silently censor or shadowblock topics that conflict with corporate partnerships or business interests, effectively narrowing the scope of inquiry to only what is commercially safe rather than intellectually necessary.
    - There is, for example, evidence of ongoing "content filtering" with Ex Libris blaming Azure OpenAI for this behavior in its research assistant tool.
- The academic apparatus of writing, submissions, peer review, and publishing are already under strain due to maxed out workloads for journal staff, editors, peer reviewers, and copyeditors. AI-induced noise creating verification bottlenecks and misinformation merely exacerbates an already broken system.

# Resources

## Library Freedom Project

- AI ethics
- What's the deal with AI?: an AI Skeptics FAQ
- A series of for removing AI LFP's resource page
- Critical AI talking points
- Questions to ask vendors about AI
- AI in Library Vendor Tools Database


## Research Guides from Librarians

- Brown University's Research Guide on Generative Artificial Intelligence
- Zotero AI Resources List Libguide
- Chemekta Community College's Research Guide on Generative AI and Ethics
- Oregon State University's Generative Artificial Intelligence (AI) and Ethics
- Syracuse University Not Neutral: A Critical Guide to AI

## Teaching

- AI and Digital Accessibility: What's Up?
- Anna Kornbluh, Krista Muratore, Eric Hayot's Against AI teaching resources
- Anchor Insights Consulting's Syllabi Policies for Generative AI Repository
  - Spanish version of these policies by Tatiana Torres Zapata
- Aaron Tay's Classifying the Ways LLMs Summarise in Academic Search: Understanding AI Summaries in EBSCO, ProQuest, and More
- Creating AI Resistant Assignment
  - A Stoplight Model for Guiding Student AI Usage
  - The Importance of Clarity in AI Usage -Guidelines for Classroom Activities
- Civics of Technology curriculum to help students critically engage with AI
- Montana State University's Viewfinder toolkit to help with ethical reflection on AI through different stakeholder scenarios.
- New York University, "Which databases can provide AI-enhanced search in the library?"

# Resources


- Elizabeth Palumbo, "A Student's Right to Refuse Generative AI"
- Refusing GenAI in Writing Studies
- Yizhou Fan et. al. "Beware of Metacognitive Laziness: Effects of Generative Artificial Intelligence on Learning Motivation, Processes, and Performance." British Journal of Educational Technology, 56, 489–530.
- Nataliya Kosmyna, et. el. "Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task." arXiv preprint arXiv:2506.08872.
- Gerlich, Michael. 2025. "AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking" Societies 15, no. 1: 6. https://doi.org/10.3390/soc15010006.


## Organizing

- AI Now Institute
  - Anatomy of an AI system
  - Roadmap for action against AI
- "AI" Worker's Inquiry Zine: Provides questions for workers to reflect on their own work, AI, and organization support.
- The Foundation Model Transparency Index (May 2024): comprehensive assessment of the transparency of foundation model developers
- Kairos Fellowship's guide on fighting data centers
- North Star Data Center Policy Toolkit: Policy Interventions to Stop Rampant AI Data Center Expansion (AI Now)
- Toolkit for the Field: An Organizer's Guide to Stop Data Centers (Media Justice)
- Little Sis powermapping toolkit for taking action against powerful forces in your community
- Public First Action - "a bipartisan organization designed to educate Americans on key AI issues and advance an AI policy agenda supporting safeguards."
- Richard Mathenge on Data Workers Organizing - The African Content Moderators Union
- Techworker Community Africa (TCA)
- Workersdecide.tech for workers to resist AI in their workplaces
- The AI Resist List - highlights communities taking actions to counter AI and different approaches from around the world

# Resources

Professional Associations

- Committee on Artificial Intelligence and Academic Professions, American Association of University Professions (AAUP), "Artificial Intelligence and Academic Professions," July 2025,
- International Federation of Library Associations and Institutions (IFLA)'s " AI Entry Point for Libraries and AI"
- Association for Information Science and Technology (ASIS&T) Auto-Transcription Policy for Meetings & Events

Additional Resources

- "The bargain we are being asked to ratify" – AI as technological bribery
- AI in Academic Libraries, Part One: Concerns and Commodification
- AI in Academic Libraries, Part Two
- A survey of AI tools in library tech
- Dispelling Myths of AI and Efficiency
- ITHAKA S&R Generative AI Product Tracker
- Indiana University Center for Women in Technology's Ethical AI Podcast
- Library IT vs the AI bots
- Testing AI Academic Search Engines part 1
- Bellingcat: Don't Get Scammed! Tips For Spotting AI-Generated Fake Products Online
- Protect Stanford Digital Assets from Online Video Scrapers
- Why say no to generative AI infographic
- Workforce Decoded: AI, Skills and the Future of Hiring
- Zooming Out: WebinarTV's Rampant Scraping of Online Meetings

## Credits

Many wonderful LFP members helped with this guide–from researching to writing to careful reviewing to copyediting. Some of them include: María Mejía, Kelly McElroy, Mark Swartz, Edward Junhao Lim, Jajwalya Karajgika, Marisol Moreno Ortiz, Jennifer Keyser, Reanna Esmail, Tess Wilson, Alison Macrina, and Kimberly Springer.

Thank you for organizing with LFP!