Building an AI Governance Framework – Companion Document

Establishing your library's AI ethics and values

Developing a custom AI ethics statement is a critical first step because it serves as your library's "north star" for all subsequent technology decisions. Artificial intelligence itself is neither inherently ethical nor unethical; rather, it should be intentionally developed and used in an ethical manner. Therefore, your AI principles should not exist in a vacuum; they should align seamlessly with your library's existing institutional values, codes of conduct, and commitments to intellectual freedom.

Establishing a framework for responsible AI requires organizations to commit to several interrelated core values, starting with accountability to ensure that humans remain answerable for the outcomes and impacts of AI systems. Transparency is equally essential, meaning that AI operations and decision-making processes should be understandable and explainable to those who use or are affected by them. Furthermore, fairness must be prioritized so that AI models treat all individuals equitably and do not perpetuate or amplify harmful societal biases. Organizations must also guarantee reliability and safety, rigorously monitoring and testing systems to ensure they operate consistently as intended without posing any physical or psychological harm. Protecting user data is a non-negotiable requirement, making privacy and security a foundational cornerstone of any ethical AI deployment to defend against malicious threats and unauthorized access. Finally, a commitment to inclusiveness ensures that AI solutions are purposefully designed to empower and engage a diverse range of people, accommodating all backgrounds and abilities.

A cornerstone of responsible AI use is the "human-in-the-loop" requirement, which is not merely a technical necessity but an ethical imperative.

AI should be viewed as "augmented intelligence" that supplements and maximizes human capabilities rather than replacing human workers and their judgment. Ethically deploying AI means encouraging use cases that actively reduce barriers for your community rather than reinforcing existing inequities.

Conducting an AI readiness assessment

AI introduces significant organizational risks, including security vulnerabilities like shadow AI and unauthorized data harvesting, as well as trust-related issues such as confident hallucinations, deepfakes, copyright infringement, and the amplification of harmful biases and disinformation.

Before you can successfully scale AI or implement new initiatives in your library, you need to take inventory of where your institution is right now. A readiness assessment establishes your baseline, allowing you to accurately gauge existing risks, understand your infrastructure gaps, and measure the true benefits of any future AI adoption. To effectively integrate AI into your organization, conducting employee surveys or baseline discussions is a critical first step in the governance process. By directly asking staff how they have been using AI and how they are feeling about the technology, leadership can uncover unofficial or "shadow" usage and gauge the overall comfort level within the organization. This baseline assessment ensures that your resulting AI governance frameworks and training programs are grounded in reality and tailored to address the specific needs, practices, and concerns of your employees.

The first and often most surprising part of a readiness assessment is identifying the current uses of AI already happening within your library. You should create an AI tool inventory. It is highly likely that your staff or patrons are already using AI in "invisible" ways—what we call "shadow AI". For example, a staff member might be using a public, unvetted generative AI tool to draft marketing copy, or a vendor may have silently rolled out an AI-powered search feature in your discovery layer. Without an inventory, you cannot manage the risks associated with data privacy or copyright.

A comprehensive assessment also evaluates your infrastructure readiness. For AI tools to be effective and safe, they need a solid foundation. This means looking closely at your data quality and availability, your security and access controls, and your cloud capacity or APIs. If your library's data is siloed, messy, or lacks clear governance, AI tools will struggle to provide accurate or secure outputs.

Readiness is not just about technology; it is about your people. A key component of this assessment is evaluating the AI literacy and culture of your workforce. You need to understand your staff's current comfort levels, their anxieties about AI, and their competency in evaluating AI outputs. AI is a people project first, and assessing readiness means understanding where your staff needs structured support, training, and permission to experiment safely.

Identifying use cases and tools

Transitioning from assessing readiness to taking action requires a structured approach to soliciting and managing AI use cases. When identifying where to deploy AI, you should balance "Efficiency AI," which optimizes current processes, with "Opportunity AI," which helps reimagine services and create new value.

Good use cases are often highly visible, involve complex or changing tasks, and provide 24/7 utility. Furthermore, the best targets are areas where humans are in the loop but may lack specific experience or knowledge, or where a process needs high customization to the individual user. When selecting use cases, it is crucial to have a tolerance for some errors

and ensure that the outcomes are highly measurable. Consider both "quick win" use cases and "moonshots" that stretch your capabilities with potentially big rewards.

To manage these ideas effectively, you need to define screening criteria to evaluate proposals based on expected impact, student and employee benefit, feasibility, data readiness, privacy and ethics risks, regulatory exposure, and cost. Once a use case is prioritized, the next step is determining the right tool. Your governance framework should include a formal review and approval process for all AI tools considered for institution-wide adoption. Each proposed tool should be evaluated for its educational and operational value, data privacy and security standards, accessibility and inclusivity, and legal or contractual risks.

You will also face "build versus buy" decisions; these require carefully considering the total cost of ownership, data control, intellectual property concerns, system integrations with existing platforms, and strategies to avoid vendor lock-in. Ultimately, this process should result in an approved AI tools catalog that provides clear usage guidance for faculty, staff, and students.

When matching use cases to tools, you should scale your governance to the risk level. A helpful framework is the "human-in-the-loop ladder". Low-risk use cases, such as using AI for drafting assistance with emails or agendas, simply require the human to edit and own the final product. Medium-risk tasks, like summarization and synthesis of policies, require cross-checking outputs against source documents. Higher-risk use cases that involve decision-shaping outputs—such as student discipline or financial aid guidance—demand formal oversight, documentation, and explicit guardrails against automation bias. The goal is to automate the routine administrative "sludge" while protecting the "craft zones" where human empathy, mentoring, and complex stakeholder negotiations are fundamentally required.

Developing AI/data policies & compliance

You should establish the rules of the road. Developing a robust AI and data policy framework ensures that your library complies with regulatory requirements while supporting thoughtful, human-centered adoption of AI technologies. A common misconception is that libraries must write entirely new policies from scratch; in reality, many of your existing non-AI policies—such as those governing acceptable use of IT resources and data privacy—already apply to AI and simply need to be clarified or updated.

However, there are specific, core policies your institution should establish and maintain. These include an overarching AI Acceptable Use Policy for staff and patrons, an AI Development and Procurement Policy for evaluating vendors, guidelines for shadow AI detection and mitigation, and an actively managed AI Tool Whitelist and Blacklist. These policies should be reflect your institution's existing mission, ethics, and values. When drafting these documents, it is crucial to use plain language and make the policies

accessible and inclusive. An effective policy is not just a list of restrictions; it should clearly articulate what is forbidden, what is allowed, what is required, what is encouraged, and what is celebrated. This level of clarity enables innovation because your staff will have explicit boundaries and guidelines to safely operate within.

Data governance is the foundation upon which safe AI use is built. Your framework should establish strict data governance principles that define exactly which categories of institutional data may and may not be used with AI tools. You should map your library's data classifications to permitted AI uses, creating clear examples of "allowed," "restricted," and "prohibited" data. For example, you should clarify the restrictions on entering sensitive student, employee, research, or patron data into third-party or generative AI systems, as this poses significant privacy and model-training risks. Furthermore, your data policies should address the complex legal risks of training models on copyrighted data, clarify who owns the intellectual property of generative AI outputs, and determine how your library will handle "Right to be Forgotten" requests when patron data has already been baked into a trained model's weights.

AI policies won't help if staff are not aware of them. Ensure that your policies are published where they can be easily found. You should also periodically conduct training or review sessions to ensure staff fully understand them and have opportunities to ask questions. Awareness also means knowing who is accountable for AI. If the use of AI leads to harm, it should be clearly-defined who is accountable for those actions so that person can take action to remediate it.

For an AI governance framework to successfully transition from written policy to everyday practice, organizational leaders must actively and visibly model the behaviors they expect from their staff. When library directors and department heads openly engage with AI tools while transparently adhering to established ethical guidelines, data privacy rules, and "human-in-the-loop" verifications, it signals that responsible AI use is a core institutional priority rather than just an administrative hurdle. By sharing their own learning processes, discussing how they critically evaluate AI-generated outputs for bias or inaccuracy, and consistently prioritizing compliance, leaders demystify the technology and cultivate a culture of trust where staff feel empowered to innovate safely within the established guardrails.

Building a culture of AI proficiency

Integrating AI into your library is fundamentally a culture problem first and a technology problem second. True AI adoption isn't just about rolling out new software tools; it requires a profound shift in behavior, work habits, and expectations. Therefore, defining what "AI literacy" means in your specific library context is essential.

Traditional leadership competencies map directly to effective AI usage. For instance, just as you wouldn't expect perfection from a new hire after giving them a vague assignment,

effective communication with AI requires "prompt engineering"—providing clear, concise instructions rich in context to get the right results. Delegation becomes the strategic assignment of tasks, identifying which duties are best handed off to an AI agent and which require human nuance. Coaching translates to iterative refinement, where you provide the AI with feedback on its initial drafts to guide it toward a better final product. Meanwhile, your ability to decompose problems allows you to sequence massive projects into smaller, step-by-step prompts that the system can successfully process. Similarly, time management and broader problem-solving skills allow you to recognize when AI creates true efficiency and to design workflows that overcome operational bottlenecks.

However, building this level of proficiency should start at the top. Before deploying AI tools to the wider workforce, senior leadership should understand and actively use the technology themselves. They should also identify other leaders in the organization who can serve as "AI champions" to help promote and support AI adoption.

When library directors and managers visibly use AI tools in their daily workflows, it grants staff "psychological permission" to explore the technology without fear of being penalized for mistakes or appearing replaceable. Importantly, this modeling should not just highlight your successes. It should also involve the public exercise of critical judgment, where you actively share instances where AI failed, hallucinated, or produced biased results. By doing this, you reinforce the "human-in-the-loop" requirement and prove to your staff that their professional expertise is irreplaceable.

As leaders, you should also be highly sensitive to the emotional load that AI carries. Staff often experience "AI anxiety"—driven by the fear of obsolescence—or "change fatigue" from the constant barrage of new tools. You can mitigate these concerns by framing AI not as a replacement, but as "augmented intelligence" that automates the administrative sludge so they can focus on the "craft zones" where human empathy and mentoring are required.

Finally, you cannot mandate innovation; it must be cultivated. Shift away from top-down mandates and adopt a coaching approach that empowers your employees to identify their own productivity bottlenecks and apply AI to solve them. You should provide structured support and dedicated time for your staff to learn. Exploring these tools cannot become an unfunded mandate piled on top of their existing deliverables. Encourage the creation of communities of practice where staff can share successful "prompt libraries," discuss their successes, and safely troubleshoot challenges together.

Communication, promotion, and transparency

Governance is only effective if your community understands it. The final pillar of your framework should focus on outward-facing communication, marketing, and the promotion of your AI initiatives. You need to strategically communicate your library's official AI narrative, including its goals, values, and ethical boundaries. Ensure you coordinate

closely with your marketing and legal teams for message consistency, as this narrative establishes your library as a responsible steward of AI.

True transparency requires that your policies are accessible, inclusive, and plainly written. You should be open about new and emerging AI use cases to generate conversation and provide transparent documentation of your system processes. Furthermore, it is critical to establish feedback mechanisms that allow your stakeholders to report concerns or inconsistencies found in AI-generated outcomes.

A practical first step is developing a central web presence or "AI hub" for your institution. This hub should act as a single source of truth that hosts an overview of your AI efforts, an inventory of available AI tools, recent news, and your official governance statements and guiding principles.

Finally, you should foster engagement through proactive listening and promotion. Plan and conduct structured "sharing and listening" visits across different library units to collect examples of AI aspirations, pilots, and concerns. Identify recurring themes from these visits to inform your ongoing messaging and policy promotion.

You can also organize AI-focused events—such as showcases, demo days, speaker panels, or community conversations on AI's impact—that are explicitly geared toward awareness and engagement rather than formal training. By capturing and sharing the content from these events, you reinforce a culture of responsible AI and demonstrate that your governance framework is collaborative and responsive to your community's needs.