

American Library Association

Guidance on the Use of Artificial Intelligence in Libraries

Contents

Introduction	2
Library Values and AI	3
1. Public Good	3
2. Intellectual Freedom	5
3. Privacy	6
4. Sustainability	9
5. Diversity, Equity, Inclusion, and Access (DEIA)	10
6. Labor	10
General Guidance on AI Use	12
Conclusion	13
APPENDIX:	15
AI Ethics Framework	15

Introduction

The purpose of this document is to provide ALA guidance for libraries developing or revising local artificial intelligence (AI) use policies, procedures, and vendor-review practices. This guidance identifies professional values and decision points for deciding when and whether to use AI systems, AI-enabled features, and vendor integrations.

AI applications, including generative AI, agentic AI, predictive systems, recommender systems, automated decision systems, and AI-enabled vendor features, are changing how libraries operate and how patrons interact with information. For purposes of this guidance, “AI” refers to software systems that use statistical, machine-learning, or related computational methods to generate, classify, rank, recommend, summarize, predict, automate, or assist decisions. This guidance applies to stand-alone AI tools and to AI-enabled features embedded in library databases, discovery systems, catalog systems, productivity tools, school platforms, campus systems, municipal systems, consortial systems, and vendor integrations. AI may support instruction, operations, accessibility, discovery, or public education when library workers have identified a clear public-service purpose, reviewed benefits and risks, preserved human judgment, and established appropriate safeguards. AI must not replace staff judgment or accountability.

ALA recommends that all library organizations evaluate candidate AI systems, AI-enabled features, and vendor integrations against core professional values, ethics, local policy, applicable law, labor impacts, community needs, and public trust, and act on behalf of libraries to address AI defects with vendors. The review should identify the public-service purpose, likely benefits, significant risks, required safeguards, staff capacity, community effects, and appropriate AI and non-AI options. The outcome of the review may be adoption, limited use, pilot testing, delay, pause, non-adoption, removal, or discontinuation. Adoption or limited use is appropriate when the library can document a clear service benefit and adequate safeguards.

A considered decision to delay, omit, remove, or discontinue use of a product is an affirmative professional choice grounded in library values, not a failure of adoption. It is appropriate when local policy, vendor controls, staffing capacity, or available review cannot adequately mitigate risks to privacy, intellectual freedom, equity, labor, sustainability, access, accuracy, accessibility, or public trust.

AI use guidance builds on existing ALA policies rather than operating as a separate or conflicting layer. The ALA Artificial Intelligence Policy Working Group, using the ALA Core Values¹ of Public Good, Intellectual Freedom, Privacy, Sustainability, Diversity, Equity, Inclusion and Accessibility as a starting point, has developed this Guidance on the Use of Artificial Intelligence in Libraries. Where applicable, this document incorporates an ethics framework (see Appendix) based on the work of philosopher Luciano Floridi, serving as a foundation for consideration of moral questions relevant to the use and misuse of AI.

¹ <https://www.ala.org/news/2024/02/american-library-association-updates-core-values>

Library Values and AI

The following values guide how library workers evaluate, use, limit, refuse, remove, or discontinue AI systems in practice.

1. Public Good

Public good means working to improve society and protect the rights to education, literacy, and intellectual freedom. AI may help libraries improve access, instruction, operations, accessibility, and public education. It can also erode human self-determination and mental well-being, deepen social inequality, and contribute to environmental harm. AI evaluation and use in libraries should be guided by the common good, intergenerational sustainability, and meaningful human agency.

Promoting Human and Environmental Wellness

Libraries should use AI only for documented public-service purposes, including reducing access barriers for underserved populations, when the benefits and risks have been reviewed. Local policies should prohibit AI uses that harm, misinform, deceive, or deepen social inequality.

Preserving Human Decision-Making and Self-Determination

AI will complement rather than replace human intelligence, reasoning, deliberation, and critical thinking; humans remain accountable for AI-automated decisions and their consequences.

Library users and workers should have access to human assistance for library services that use or are affected by AI. Libraries should maintain sufficient human expertise to provide that assistance and should give staff authority to question or override algorithmic recommendations that appear arbitrary, inaccurate, or unreasonable.

All library workers and communities have the right to explanation: if AI was used to make a decision that affected a patron, that patron has the right to know why the automated system made that decision.

Libraries should not use AI to justify staff reductions, deskilling, or replacement of professional judgment. AI decisions should preserve human expertise, meaningful library work, and publicservice capacity.

Democratic Participation and Community Trust

Decisions about AI adoption, expansion, limitation, or discontinuation should be guided by consultation with affected library workers, appropriate community input, local governance, and documented public-service purposes. Libraries should scale participation to the nature and risk of the decision, with broader community engagement when AI systems affect public-facing services, patron data, discovery, accessibility, access to services, or community trust.

ALA recommends that library workers teach AI literacy through a public-service and harmreduction lens. AI literacy should help patrons, students, staff, and community members understand where AI appears in everyday information environments; what benefits AI may offer; what risks AI may create for privacy, accuracy, labor, equity, accessibility, mental health, intellectual freedom, and the environment; how to verify AI-generated claims; how to recognize AI-generated content; how to protect personal information; and how to make informed choices about use, limited use, or refusal. AI literacy programs should avoid shaming people who are curious about, already using, required to use, skeptical of, or trying to avoid AI.

ALA recommends that libraries take the following actions:

- Define meaningful human review for AI-assisted workflows, specifying who reviews AI output, what training and subject knowledge reviewers need, when review occurs, and what authority staff have to question, correct, reject, or escalate AI output.
- Require meaningful human oversight before using AI-generated content for public services, communications, or administrative decisions.
- Train staff to document errors and grant them the clear authority to reject or escalate AI recommendations.
- Prioritize staff well-being and service quality by using AI efficiency gains to improve working conditions and community services, rather than justifying staff reductions.
- Assess AI's impact on labor, consult affected workers, and implement safeguards against deskilling, increased surveillance, and uncompensated work.
- Use technology to enhance staffing capacity, support training, empower library workers, and improve services.

Collections

- Review collection development policies and licensing practices to address known or reasonably identifiable AI-generated and AI-assisted materials, including books, articles, images, audio, video, summaries, reviews, metadata, and vendor database content. Local policies should specify how library workers evaluate quality, disclosure, creator labor, accessibility, research integrity, community relevance, metadata, and patron transparency before acquiring, licensing, describing, or promoting these materials. Professional collection judgment should guide these decisions, not vendor convenience, automated production, or platform defaults.
- Choosing not to acquire primarily AI-generated materials is a legitimate exercise of professional collection judgment and fiscal stewardship when those materials do not meet local standards for quality, accuracy, originality, transparency, or community relevance.

Public Services

- Prioritize teaching users how to critically evaluate AI tools by identifying their strengths, weaknesses, and limitations as information sources.
- Incorporate AI literacy into the reference interview by helping patrons ask better questions, understand tool limits, compare AI output with trusted sources, and verify claims for accuracy and bias.

- Use AI to support accessibility — such as draft plain-language summaries, draft alt text, language-support aids, or translation support — when the library has identified a clear access need and reviewed privacy, accuracy, cultural, and accessibility risks. Staff or qualified reviewers with relevant language, cultural, subject, or accessibility expertise should review outputs before public use when the output affects user understanding, official communication, or access to services. Libraries should not rely on AI as a substitute for qualified translation, interpretation, accessibility review, culturally responsive communication, or professional judgment when accuracy, safety, legal obligations, or community trust are at stake.

Technical Services

- Application of AI tools to cataloging, classification, and acquisitions processes must include human-in-the-loop configuration and oversight including full control to accept, edit, or reject suggestions.

Administration

- Create meaningful opportunities for affected staff and community members to provide input during AI planning, deployment, maintenance, and reassessment.

2. Intellectual Freedom

ALA actively advocates and educates in defense of intellectual freedom — the rights of library users to read, seek information, and speak freely as guaranteed by the First Amendment. Intellectual freedom is a core value of the library profession and a basic right in a democratic society.² AI and other automated information systems can threaten intellectual freedom when they obscure how information is ranked, extract public or library-created data without clear public benefit, or optimize engagement at the expense of attention, reflection, and independent judgment.

Library workers should provide programming and reference services that help patrons build information-seeking resilience and critical evaluation skills for AI-generated content. Programs and reference services should help patrons understand how AI tools can reflect social biases, encourage overreliance, and weaken users' ability to assess information independently.

ALA recommends that libraries take the following actions:

- When an automated library system is deployed, clear, accessible documentation should explain its function, purpose, accountability, and influence on user outcomes.
- Users should be able to opt out of AI-automated products when possible.
- Libraries will require vendors to offer equivalent services that do not include AI assistance, or document the limitation and identify alternatives.

² <https://www.ala.org/advocacy/intfreedom>

Public Services

Prioritize human creativity, dignity, inclusion, and patron rights while maintaining human oversight for AI-assisted public services and communications.

Technical Services

- Review AI-enhanced discovery, ranking, recommendation, and personalization features at implementation, after major configuration or vendor changes, and on a scheduled basis for patterns that reduce visibility, distort representation, or limit access for particular topics, authors, viewpoints, languages, or communities. When a vendor, consortium, campus, municipality, or school district limits audit access, document the limitation and review available evidence, including outputs, configurations, vendor documentation, user feedback, and staff observations.
- Preserve predictable, transparent, non-personalized, and explainable access paths alongside AI-enhanced discovery features, such as known-item search, subject-based browsing, chronological lists, or other locally documented discovery options.

AI may be used to assist technical services workflows only for defined support purposes, such as:

- Draft descriptions,
- Candidate terms,
- Metadata enrichment, or
- Initial classification suggestions.

Final authority over metadata decisions should remain with staff who have relevant cataloging, language, cultural, subject, or community expertise, especially for subject headings, classification, authority work, non-English or non-Latin script materials, local history, community materials, and sensitive topics. Local procedures should define review methods, sampling practices where appropriate, escalation steps, documentation requirements, and error-correction processes based on the scale and risk of the workflow.

Administration

- Recognizing that AI development is controlled by a small number of companies whose incentives are not aligned with libraries, intellectual freedom, or the public good, include questions about intellectual freedom in procurement for AI tools, such as how ranking, recommendation, or moderation functions operate and can be controlled.
- Provide complete transparency about library use of AI if it is implemented.

3. Privacy

Library workers have a responsibility to understand how AI tools work, how they handle personal information, and what risks they may create for patrons. The ALA Code of Ethics³

³ <https://www.ala.org/tools/ethics>

affirms: "We protect each library user's right to privacy and confidentiality with respect to information sought or received and resources consulted, borrowed, acquired, or transmitted." School librarians must also consider federal K–12 privacy laws, including the Family Educational Rights and Privacy Act (FERPA) and the Children's Online Privacy Protection Act (COPPA), along with district-level data-sharing restrictions. Many school librarians work without dedicated administrative support, so AI guidance should account for those capacity limits. Libraries of all types must carefully review AI-enabled tools for potential risks to patron privacy, data security, and equitable access.

Working with Vendors on AI and Privacy

ALA recommends that library workers, administrators, governing bodies, consortia, and parent institutions apply ALA vendor privacy guidance when evaluating AI systems, AI-enabled features, and vendor integrations. Patron data, institutional data, local history, staff workflows, and usage logs should remain under library, public-sector, or community control. Libraries must not be required to trade data control for access to AI functionality.

Vendor review should identify:

- whether AI features are enabled by default;
- whether the library can disable them;
- what data the system collects;
- whether prompts, searches, account activity, reference interactions, or staff inputs are logged;
- whether data are used for model training or product improvement;
- where data are stored;
- what subprocessors are used;
- how long data are retained;
- how data can be deleted;
- What audit or exit rights the library has. AI-related privacy information should be written in clear, plain language for patrons and staff.

When third-party AI services collect or process patron or staff data, influence search results or recommendations, or track users beyond the library, libraries should disclose the practice, explain available choices, provide a meaningful non-AI or minimally automated path when the library can provide one, and document limitations when a vendor, consortium, school district, campus, or municipality controls the system.

Libraries should help patrons understand when AI tools may track reading, searching, or uploaded content beyond the library, and educate them about their legal rights related to access, correction, and deletion of personal data.

Staff Education and Responsibilities

Staff should be trained comprehensively about the use and potential impact of artificial intelligence if adoption is chosen. Whether or not a library opts in or out of AI use still means all staff members should receive professional development.

- Library workers will be trained on data ownership and the privacy consequences of AI usage. This includes training them to avoid disclosing sensitive data and ensuring that patron names, contact information, or other private details are never entered into AI systems during research or reference support.
- Libraries will advocate for clear data retention and sunseting policies while ensuring users are informed about the specific AI tool providers involved.
- Libraries will prioritize protective measures such as managing permission settings and advocating an "opt-in" model that ensures users have clearly consented to AI service usage.

ALA recommends that local policies prohibit library workers from entering the following into AI systems unless the system has been formally approved for that use

- personally identifiable information;
- sensitive personal information;
- borrowing histories;
- reading histories;
- reference interactions;
- account information;
- staff records;
- prompts containing identifiable details; or
- other non-public records.

Approval should include privacy review, security review, legal review where applicable, a documented library purpose, data minimization, retention limits, vendor training-use restrictions, staff guidance, and required public notice. When a library cannot determine how an AI system collects, stores, retains, shares, deletes, or uses data, the library should delay, limit, refuse, or discontinue that use.

- Review AI-enhanced discovery, ranking, recommendation, and personalization features at implementation, after major configuration or vendor changes, and on a scheduled basis for patterns that reduce visibility, distort representation, or limit access for particular topics, authors, viewpoints, languages, or communities.
- When a vendor, consortium, campus, municipality, or school district limits audit access, document the limitation and review available evidence, including outputs, configurations, vendor documentation, user feedback, and staff observations.
- Preserve predictable, transparent, non-personalized, and explainable access paths alongside AI-enhanced discovery features, such as known-item search, subject-based browsing, chronological lists, or other locally documented discovery options.

Administration

- Establish local policies that require regular privacy review of AI-enabled systems, including data collection, retention, vendor access, staff use, public notice, incident response, and changes introduced through system updates.

4. Sustainability

Sustainability, as defined in the ALA Core Values, means making choices that are good for the environment, make sense economically, and treat everyone equitably. Sustaining our environment means seeing nature itself as intrinsically valuable, not only for its natural resources but also for the diversity of life it produces and sustains. The decisions we make today regarding AI use have lasting consequences for future generations. All humans have the basic right to inherit an environment that nourishes and supports their health and well-being.

To reduce environmental impact, libraries should prefer AI tools, configurations, or service models with documented lower environmental impact, such as smaller or task-specific models where evidence shows lower energy use, water use, emissions, hardware demand, or e-waste. Libraries should consider open-source or community-governed tools separately when transparency, inspection, local control, or public-interest governance is the goal.

Library workers should advocate for responsible management of AI data center construction and operations that may harm their communities by straining utilities and resources, increasing electricity bills, and causing pollution that contributes to health risks and social and environmental inequality.

Libraries and Using AI Sustainably

Libraries should assess and account for the full environmental lifecycle of potentially chosen AI systems — including energy use, data center emissions, electronic waste, and resource extraction; procurement decisions should weigh environmental impact alongside efficiency gains.

Library workers should prefer AI tools, configurations, or service models with documented lower environmental impact. Open-source or community-governed tools may support transparency, inspection, and public-interest governance, but open-source status alone should not be treated as evidence of lower environmental impact.

When comparing options:

- Require vendors to disclose carbon impact estimates, data-center energy sources, and water usage as a condition of procurement evaluation. Prefer smaller, task-specific models over massive general-purpose ones.
- Avoid promoting tools with generative features that represent an undue environmental burden while remaining consistent with human needs.

When considering the full lifecycle of AI systems:

- Develop procurement checklists that include energy consumption, expected hardware lifespan, e-waste considerations, and upgrade cycles.
- Request vendors to provide recycling or takeback plans for obsolete AI hardware.

When supporting sustainable AI development:

- Advocate for open-source, community-governed AI tools when possible.
- Work with vendors to encourage models that support conservation, renewable energy management, or sustainability research.

Intergenerational Justice: Protecting Institutional and Professional "Ecological Capital"

Libraries should consider how current AI deployment decisions affect the professional knowledge, institutional capacity, and community relationships available to future librarians and patrons; short-term efficiency gains should not deplete long-term institutional viability.

Library expertise, community trust, and democratic participation are irreplaceable forms of capital that should not be eroded by automating too many services, diminishing staff skills, or creating long-term dependency on vendor-controlled AI systems.

5. Diversity, Equity, Inclusion, and Access (DEIA)

The American Library Association affirms that, as part of the library's key mission to provide access to information, the library must provide an environment where people can access artificial intelligence technologies while recognizing the need for diversity, equity, and inclusion in its implementation. As AI systems become increasingly integrated into libraries and the communities they serve, library workers have a responsibility to ensure these tools serve all community members equitably and do not perpetuate or amplify existing social inequities.

Bias and Equitable Access

Library workers will evaluate AI systems for potential bias in datasets, algorithms, and outcomes before adoption and throughout their lifecycle. This includes conducting regular audits of AI tools used in cataloging, reference services, content recommendation, and patron interactions. Staff should be trained to recognize algorithmic bias and understand its impact on library users, particularly those from historically marginalized communities.

Library workers should:

- Prioritize equitable access to AI-enabled products and services, ensuring that marginalized and underserved communities are not excluded from their benefits. This commitment requires proactive identification and removal of barriers that may prevent full participation.
- Seek solutions that accommodate diverse languages, literacy levels, and technological capabilities.
- Require AI-enabled services that the library selects, deploys, or configures — including chatbots and virtual assistants — to meet recognized accessibility standards, support users across languages and literacy levels, and provide a clear path to human assistance.
- Use or contribute to open-source AI tools and datasets that explicitly aim to reduce bias and increase representation where they meet operational needs.

6. Labor

This guidance upholds human agency over artificial intelligence and automation. Across sectors, employers have used AI and automation to justify restructuring, job reductions, deskilling, and

workload shifts. Libraries should evaluate AI adoption with attention to labor impacts, worker expertise, staff autonomy, and the public-service value of human judgment.

Protecting Human Expertise, Opportunity, and Worker Well-Being

Libraries should not replace reference, readers' advisory, instructional services, or community support roles with AI recommender systems or AI chatbots.

Tasks and decisions that require empathy, judgment, and subject-specific knowledge should not be replaced with AI automation.

Libraries should recognize that library work is tied to professional identity and self-esteem; displacement carries psychological costs beyond economics and should weigh in decisions about automation.

Including library workers in AI deployment decisions

Library workers should have meaningful input into decisions about which AI systems affect their work. Workers understand their own job demands and should shape how, or whether, automation occurs in their roles.

- Use a staff advisory group, existing staff governance process, or other locally appropriate mechanism to review proposed AI tools.
- Identify which work should remain human-led, which work may be assisted by AI, and what harms to staff autonomy or service quality could result.

Human Decision-Making and Deliberation

- Consider how AI automation shifts the decision-making process. When a workflow is automated, it shifts most decision-making to the end of the task. This can be extremely detrimental, especially in areas, such as cataloging, where careful decisions must be made about what subject terms should be used.

Patron Data as Labor

Libraries should recognize that patron data and digital activity used to train or improve AI systems raise labor and equity concerns. Patron contributions — including prompts, ideas, questions, searches, and analysis — are often uncompensated, uncredited, and extracted as free labor.

Evaluate AI Vendors for Labor Ethics

Before purchasing, renewing, or deploying AI-enabled tools, libraries should:

- Request documentation about data-labeling and content-moderation labor conditions, including outsourced annotation work.
- Request documentation about worker protections, compensation, and mental health support for data annotators and content moderators.
- Request supply chain information, including training data origins and moderation workflows.
- Avoid purchasing, renewing, or deploying AI tools when vendors cannot provide sufficient documentation of ethical labor practices.

- Document limitations and raise unresolved concerns through procurement, governance, consortial, or vendor-review channels when the library does not control the purchase, renewal, or platform selection.

Protect staff from Displacement and Devaluation

When AI systems materially affect library employment, job tasks, staffing models, or service workflows, ALA recommends that libraries:

- Assess labor impacts before adopting, renewing, expanding, or reconfiguring AI systems.
- Consult affected library workers, staff governance bodies, and bargaining units where applicable before changing workflows, duties, or service models.
- Identify which work should remain human-led, which work may be assisted by AI, and what harms to staff autonomy, service quality, privacy, or user trust may result.
- Use verified efficiency gains to improve working conditions, staffing capacity, training, compensation where applicable, or community services rather than to justify staff reductions, deskilling, increased surveillance, or uncompensated review work.
- Provide paid time and institutional support for workers to learn how AI functions, what its strengths and weaknesses are, how AI ethics apply to library work, and when AI may or may not be appropriate.
- Preserve core library expertise — including cataloging, subject knowledge, access services, instructional design, reference, readers' advisory, and community support — even when AI can assist parts of those tasks.
- Establish retraining, transition support, or other locally appropriate supports before using AI to substantially change or eliminate job tasks.

Current AI decisions should preserve professional knowledge, career pathways, and meaningful library work for future librarians and library workers.

General Guidance on AI Use

This recap summarizes general practices that follow the values above. It is not a substitute for local policy, legal review, or vendor review. Libraries should adapt these practices to local governance, staffing, community needs, and the level of risk involved.

When libraries evaluate or use AI-related products and systems:

- Use AI only where it serves clear user-centered purposes (improving access, instruction, workflows, or services), not simply because a tool is available; consider environmental and cultural costs as part of decision-making.
- Provide instruction and reference support that meets patrons where they are in their AI knowledge. Use a harm-reduction lens to help patrons recognize risks — including misinformation, privacy concerns, and mental health concerns — and make safer, more informed decisions about AI use.

- Provide programming and resources to help increase AI literacy for patrons, guided by an AI Literacy Framework.⁴
- Clearly disclose when and how AI is used in public-facing services, instruction, and content development. Make disclosure easy to find and consistent across channels, based on a locally defined threshold such as use beyond basic grammar or spelling support. Obtain informed consent before implementing patron-facing AI-enabled tools. Explain what data are collected, how data are used, and how the tool may affect patron privacy. Provide examples of appropriate use and clear counterexamples where AI is not a good fit, such as legal or medical advice.
- Treat AI tools — especially cloud-based and “free” services — as potential data processors. Avoid entering personally identifiable information, sensitive topics, or nonpublic records into tools whose data practices are unclear. Limit staff use to institution-approved tools that have been vetted for privacy and security, including account creation and data handling.
- Recognize that many AI systems are trained on corpora that underrepresent non-English languages, dialects, and marginalized communities, which can reproduce bias and erasure; acknowledge these limitations and collaborate with multilingual community members to test outputs and provide education materials in multiple languages and formats.⁵
- Treat AI outputs as drafts requiring human review, not authoritative answers, especially for complex reference, local context, or community-sensitive issues. Explain that AI may support routine tasks in some contexts, but human expertise remains necessary for nuance, originality, contextual judgment, accountability, and trust.
- Oppose AI tools that flag, suppress, deprioritize, label, or recommend removal of materials in ways that undermine intellectual freedom, narrow viewpoint diversity, or disproportionately affect marginalized voices.
- Provide a clear path to human assistance and, where available, a non-AI or minimally automated option for public-facing services.

Identify a point of contact for questions or concerns about AI use, privacy, accessibility, or human assistance.

Conclusion

The ALA Guidance on Artificial Intelligence Use in Libraries helps address the needs of libraries of all types as they develop AI policies in a substantive, functional manner grounded in ALA's Core Values. This guidance is meant to evolve as the landscape of artificial intelligence and the needs of libraries, library workers and patrons change over time. ALA commits to supporting

⁴ https://www.choice360.org/wp-content/uploads/2024/08/TaylorFrancis_whitepaper_08.28.24_final.pdf

⁵ <https://news.stanford.edu/stories/2025/05/digital-divide-ai-llms-exclusion-non-english-speakers-research>

libraries not only through this guidance but through direct advocacy with AI vendors and developers, convening coalitions, and intervening in industry conversations on behalf of the profession. Individual libraries cannot generate the systemic pressure needed to hold vendors accountable alone. This guidance gives library workers a values-based starting point for creating, revising, and explaining local AI policies as technology, vendor practices, community needs, and library responsibilities change.

APPENDIX:

AI Ethics Framework⁶

Library workers should consider three principal areas of concern when evaluating AI use:

- Ethical issues raised by the design choices made by AI creators (e.g., privacy issues arising from AI's dependence on personal data for training; intellectual property issues from data harvested at scale)
- Ethical issues raised by how humans choose to use AI (e.g., defining standards for responsible use versus overuse, or humans intentionally using AI to harm one another)
- Ethical issues raised by the social and environmental impacts of AI (e.g., using automation to replace human labor; the energy costs of AI and its contributions to the climate crisis)

Keeping these areas of concern in mind, we recommend the following ethical principles that library professionals should consider when evaluating AI systems and use. Each principle is directly reflected in the Library Values and Use Cases that follow.

1. The Common Good of Humanity and the Planet

AI is best used by libraries and patrons to achieve outcomes that benefit the well-being of humanity and the planet. Library workers will ensure that the benefits of AI enable human flourishing — including mental and social wellness — and that those benefits are shared equitably with all members of the community to further accessibility.

To apply this principle, ALA recommends that libraries:

- Use AI to reduce access barriers for underserved populations.
- Apply a sufficiency mindset to AI use to mitigate environmental harm.
- Consider the full environmental lifecycle — including energy, water, and e-waste — when making procurement decisions.

2. No-Harm Principle, Privacy, and Protection

AI should not be used to harm others, whether intentionally through misuse or indirectly through the creation or proliferation of misinformation. Libraries will not adopt AI tools that fail to protect individuals' privacy, personal data, and security. Library workers will make concerted efforts to protect marginalized communities from algorithmic oppression and surveillance, and to lessen the impacts of algorithmic bias.

To apply this principle, ALA recommends that libraries:

- Rigorously evaluate AI tools for privacy and security risks before adoption.
- Ensure patron data is not used to train AI models without consent; training staff to recognize and respond to algorithmic bias.
- Disclose when third-party AI services collect or process patron data.

⁶ <https://link.springer.com/article/10.1007/s11023-018-9482-5>

3. Human Dignity

AI use should not undermine people’s personal, creative, or professional aspirations. Human dignity includes the right to social connection, creative expression, meaningful work, and opportunities for professional growth. Libraries should evaluate and use AI in ways that support human autonomy, decision-making, and self-realization rather than impede those goals.

To apply this principle, ALA recommends that libraries:

- Ensure AI complements rather than replaces human intelligence and judgment.
- Guarantee that all library workers and patrons have the right to human assistance for any library service.
- Maintain sufficient human expertise to provide that option.

4. Responsibility and Explainability

AI use should not replace human decision-making, judgment, and accountability. Library workers will require AI vendors to provide algorithmic transparency and explainability to support any automated decisions or erroneous outputs that may impact their communities. Patrons have the right to know if and why an automated system made a decision on their behalf that significantly affected them.

To apply this principle, ALA recommends that libraries:

- Require clear, accessible documentation for any deployed automated system.
- Ensure library workers can question or override algorithmic recommendations.
- Educate patrons about their rights regarding AI-influenced decisions.

5. Democracy

Communities should have meaningful input into library decisions about AI deployment and automation. Libraries should prioritize their communities' shared values to build and maintain trust. Library workers should be aware of AI best practices and make their communities aware of systems and tools that violate the principle of trust, which is the foundation of democratic freedom.

To apply this principle, ALA recommends that libraries:

- Guide AI adoption through community deliberation and democratic participation
- Build and sustain public AI literacy through programming and resources.
- Offer opt-out options for any AI-enabled service.

6. Equity, Social Justice and Working Towards a Better Future

Justice is the fair and equitable treatment of all people. It is also distributive — the concerted, equitable application of the benefits of AI to all members of society, especially those who face additional barriers. Libraries play an essential role in ensuring that AI tools and training are available to marginalized community members as part of the broader effort to bridge the digital divide. Intergenerational justice is the moral duty we have to ensure that future generations inherit a planet that sustains life and a society that fosters human dignity and the common good.

To apply this principle, ALA recommends that libraries:

- Audit AI tools for bias and discriminatory outcomes.
- Prioritize equitable access to AI-enabled services.
- Advocate for open-source and community-governed AI tools.
- Protect the institutional knowledge, community trust, and democratic participation on which future librarians and patrons will depend.