



IFLA

Entry point to

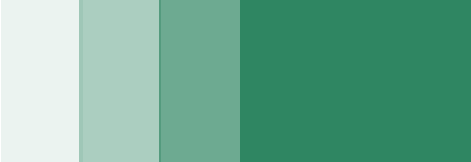
Libraries and AI

Internet Manifesto Annex I


May 2025



www.IFLA.org



AI is a promising set of technologies with the potential to promote some of libraries' principal values around knowledge access and creation. However, it is also a deeply controversial technology, and some current developments are actively creating harm. It is essential that any benefits of AI are equitably distributed, responsibly achieved and in the widest sense sustainable.



Library values and ethics are highly relevant to the responsible use of AI in society. They seek to promote equal access to knowledge for all and the right to free expression. They argue for openness and accountability and yet also respect individual privacy, organisational confidentiality, and intellectual property rights.

In some ways the term AI is very broad and often it may be more useful to talk about specific technologies such as machine learning or large language models. However the wider public debate and public policy is centred on the more general term, and so it is important that libraries also position themselves on it. We are therefore defining AI here in a broad way to encompass some well established and familiar technologies, not solely generative AI.

The premise of this document is that as a profession we have a duty to ensure the potential benefits of AI for knowledge access and creation are achieved in a responsible way and that harmful aspects are challenged. The document provides librarians with a list of key considerations in evaluating and discussing the benefits and risks of AI. It is intended to be a concise, helpful tool which informs professional debate around AI, primarily in the form of a set of questions rather than answers, which ultimately will also need to reflect the breadth and rapid evolution of the technologies.

While this document is intended to be a tool to support librarians in the assessment of ethical AI use; we do not recommend using it as a one step solution to defining its possible uses and applications definitively. It is not designed to serve as a standalone decision-making instrument. Because of that, we strongly encourage individuals to engage in broader discussions with colleagues, peer-organisations and other libraries to ensure context-sensitive approaches to AI use.

Libraries' response to AI will recognise the importance of IFLA's other policies^[1], other library sector statements on AI^[2] and the wider international debate about AI ethics and regulation^[3]. Given the speed of change in this domain we expect to have to update it frequently.

These guidelines are part of a series of thematic publications to be released in 2025-2026 and that follow the publication of IFLA's revised Internet Manifesto in 2024.

The library's role

Libraries will work with other stakeholders to define responsible AI use.

They can exercise influence at six different levels:

1. Developing or licensing AI for library services
2. Inputting to the development and training of (public interest) AI
3. Advising library users on the choice and usage of safe and responsible AI services
4. Highlighting data, algorithmic and AI literacy as a dimension of information and wider digital literacy among users and the public
5. Advising on the use of AI within the wider organisation within which the library sits
6. Advocating for responsible uses of AI in society and appropriate regulation

The possible benefits of AI

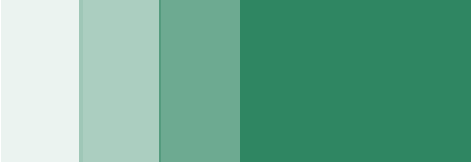
In 2025 the term AI is closely identified with generative AI, but for the purposes of these guidelines we are using it as an umbrella term for multiple digital technologies, some of which have been used in libraries for some time, and which when developed responsibly have potential benefits to access to knowledge, through:

- Digitisation, such as turning text, handwritten text, oral, multimedia and visual media into machine readable forms
- Description of content at scale, including metadata creation
- Recommendation, adaptivity, personalisation and filtering
- Summarisation and synthesis
- Data analysis
- Translation
- Interactivity

In addition, AI applications have potential to improve library service delivery and make some library operations simpler and more reliable through automation of routine tasks.

AI also has potential to offer benefits to particular populations of users, for example through enhancing accessibility such as for disabled users.

Librarians have a responsibility to maximise the benefits of AI for access to information, in an equitable, responsible and sustainable way - as well as to challenge uses that are fundamentally unethical.



Possible library related applications include:

- Implementing a library book or article recommendation system
- Supporting users' responsible use of generative AI tools in the search process, as well as other processes across the research across the lifecycle
- Providing access to AI-augmented services
- Using AI to enhance access to library collections
- Developing a library or organisational level chatbot to answer reference questions
- Developing a Retrieval Augmented Generation (RAG) application
- Providing data to train models
- Using generative AI in the accomplishment of professional tasks such as summarizing texts or drafting documents and training plans

The possible risks of AI

Achieving these benefits equitably, responsibly and sustainably can only be realised if the recurrent issues with AI technologies are addressed. These include:

1. **Over-statement of proven benefits and hype.** There remain limited use cases in the library field with demonstrated cost benefit. Some claims are simply hype. There can be simpler, better and fairer solutions to problems than AI that might be overlooked because of the hype. Harm is being created by some uses.
2. **Inaccurate and misleading information.** AI works on calculations of likelihood. Often this results in errors such as in recommending or filtering. Generative AI has well known problems in terms of out of date, plausible but incorrect information, lack of citation of sources and the invention of sources. In contexts where accuracy is critical this is a key concern.
3. **Bias.** As well as bias arising from inappropriate algorithms, AI outputs can reflect uncontrolled bias in training data. They can reproduce damaging stereotypes, which in turn can alienate some communities from using AI.
4. **Threats to cultural and linguistic diversity, and respect for diverse knowledge systems.** AI developed in non-inclusive ways increases the risk of reproducing colonial power structures and undercutting library efforts to decolonise collections and practices.
5. **The abuse of AI for disinformation, misinformation and censorship.** In the wrong hands AI can have damaging impacts on information cultures.
6. **Inequitable access.** Those with limited digital access, skills or confidence may be disadvantaged unless consideration is given to their needs in the design of systems. Given the cost of AI systems they are often only accessible to privileged states, institutions and users.
7. **Lack of explainability and transparency.** AI is inherently hard to fully understand and is often operating in the background unnoticed. Many AI companies lack transparency about critical issues such as what training data was used. It can be

unclear who is responsible if AI makes errors. When and how AI is being used should always be transparent to the user.

8. **Threats to privacy and security.** There is a pattern in current development of AI being trained partly through extracting data from users without informed consent or opt outs. The collection, use and sale of data may compromise user security.
9. **Threats to copyright and author rights.** There is a pattern in current development of AI being trained on copyright material. There is not always legitimate access to the original data. It is also against the law in some jurisdictions. There may also be hidden costs of how AI is trained, such as the costs to libraries of data being scraped from library repositories.
10. **Lack of consultation of stakeholders.** There is a pattern of AI being developed without full involvement and consultation with the widest range of stakeholders, including those who may be represented through AI.
11. **Threats to human agency.** AI can make life easier, but can also lead to a loss of skills and confidence, creating technology dependence. AI is often being used in invisible or opaque ways, which can result in the reduction of human choice.
12. **Threats of job displacement or exploitative employment relations.** AI is already being used by some employers to displace creative expertise. AI development often relies on low paid, precarious workers to perform difficult, even traumatic aspects of data labelling and filtering.
13. **Significant environmental impacts.** Digital technologies have environmental impacts through the materials and power used to manufacture and transport devices and the power demands to train and use AI. AI, especially, generative AI makes significant demands on power and fresh water to cool datacentres.
14. **Undue power of unregulated, monopolistic Tech companies and the emergence of a damaging international AI arms race.** Rather than being developed for the benefit of all humankind in sustainable ways, many aspects of the development of AI rest in unaccountable hands. AI should be developed locally, in culturally sensitive ways. Open, non proprietary approaches are to be preferred.

14 Questions librarians should ask about AI

In seeking to ensure the benefits of AI while combatting the risks, librarians may ask the following questions about a specific AI service that might be used by a library (depending on their context and the service in question).

The questions are meant to be used as a self-assessment tool to evaluate the possible risks and benefits of AI use within the library. The image below the questions can be used for a group game or dynamic to address these questions in a more interactive way.

1. What is the nature of the benefits concerning knowledge access and creation? Is the effort/cost of delivering them justified in relation to library priorities and alternative solutions?
2. What is the level of accuracy, timeliness and impartiality of information output by AI and how is the level of accuracy explained at the point of use?

3. What safeguards are in place to minimise bias and the reproduction of damaging stereotypes and assumptions?
4. How are diverse knowledge systems and cultural and linguistic diversity protected/promoted through the AI?
5. Is the availability of AI to some bad actors promoting misinformation or censorship?
6. Are the benefits available equitably, including aspects of digital inequality and accessibility?
7. Is the application transparent in a socially meaningful way? Is it defined who owns responsibility for errors made by AI and their impacts?
8. How are users' data managed to ensure their privacy and security?
9. Has the AI service been developed in a way that has not made use of illegitimately accessed works?
10. How have local and societal stakeholders been (or will they be) involved in the development and implementation of the AI, ensuring the right of communities to data sovereignty (i.e. to maintain ownership, control, and governance over data about themselves and their cultural heritage)?
11. How is it ensured that users retain agency in their interactions with AI? What is the impact on skills and confidence of users?
12. What are the impacts on human employment, including potential job displacement or the use of precarious labour?
13. What is the environmental impact of the whole lifecycle of the AI service, including power demands and water use?
14. If the developer is a third party, what is their business model and their approach to responsible operations?

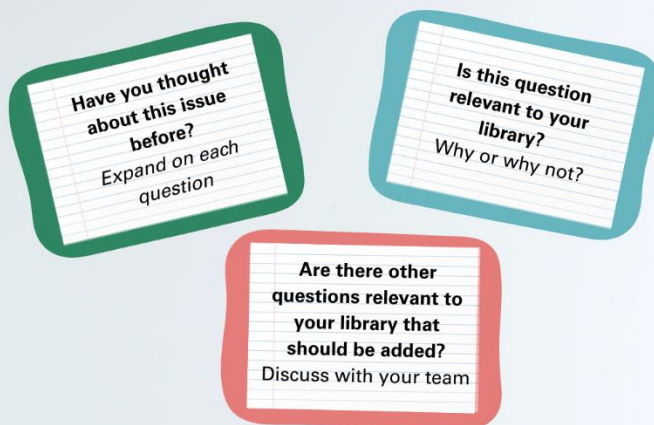
Let's Talk About AI in Libraries!

Step 1:

Check out the 14 questions in page 5: "Questions Librarians Should Ask About AI"

Step 2:


Gather your team and use these 3 cards to kick off a discussion:



✓ **Score 1 point** for each question you've already addressed with a clear plan.

🎯 Your Score:

- **11–14 points:** Great! You're AI-ready.
- **10 or less:** Revisit and plan before moving forward.



Scenarios

The following scenarios have been created to be used or adapted for discussion in the context of professional development. They extend an earlier set of scenarios which remain relevant^[4,5]. We recommend using them not only for individual reflection but also as catalysts for collective discussion and learning.

We encourage librarians to use them for interactive activities within their libraries, for instance while organising a roundtable discussion where participants read through each scenario and respond to the questions from different perspectives.

Our tips for facilitation include: **1)** Ensuring that everyone understands the nature of the questions asked, **2)** Including a diverse group of participants across roles, departments and even disciplines, **3)** Assigning roles or viewpoints to help explore the scenario from diverse angles, **4)** Documenting key reflections and areas of uncertainty, **5)** Comparing your experiences with other libraries who have made use of the guidelines.

Scenario 1. Supporting library users to adopt generative AI responsibly

A study of library user behaviour suggests that some significant shifts are happening linked to generative AI. Many users are turning to free generative AI chatbots as a first source of information. A common comment is: “it’s easy and quick. It’s making me more efficient”. Some users appear to find them particularly valuable, for example, neurodiverse users. Users are aware about some issues around information accuracy and technology dependence; but do not seem very aware or concerned about wider social issues.

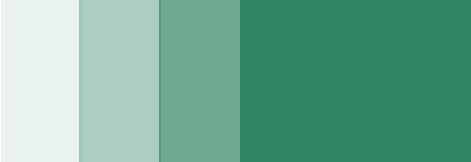
Some users are also paying for access to better models for greater functionality. Some are asking to use specialist AI for tasks such as literature reviewing and data analysis.

There is also some evidence that the use of library databases is falling, with users saying that they just don’t match up to the ease of use of free chatbots. At the same time, you cannot afford to subscribe to all the new features appearing in library databases.


- What ethical and value issues do you identify?
- What would you do?
- How realistic is the scenario? Are there aspects of the current situation that shift the issues?

Notes

1. What is the nature of the benefits to access to knowledge? Is the effort/cost of delivering them justified in relation to library priorities and alternative solutions? Users feel they are benefiting from enhanced access to knowledge, but it appears that the library role is partly displaced. The library may be able to reposition itself.
2. What is the level of accuracy, timeliness and impartiality of information output by AI and how is the level of accuracy explained at the point of use? The evidence suggests users are only partly aware of some of the information quality issues with generative AI, especially when one considers the widening range of apps in use.
3. What safeguards are in place to minimise bias and the reproduction of damaging stereotypes and assumptions? Issues of bias may be less apparent to users than problems with accuracy. The library will need to develop educational materials that help users identify these issues and improve prompts to avoid them.
4. How are diverse knowledge systems and cultural and linguistic diversity promoted? Staff will develop educational materials on recognizing AI biases related to cultural representation and improving prompts to avoid them.
5. Is the availability of AI to some bad actors promoting misinformation or censorship?
6. Are the benefits available equitably, including aspects of digital inequality and accessibility? Recommending specific services would tend to reduce inequality.
7. Is the application transparent in a socially meaningful way? Is it defined who owns responsibility for errors made by AI and their impacts?
8. How are users' data managed to ensure their privacy and security? While users have been made aware of data protection issues, they may not have considered all the implications in the context of generative AI.
9. How are the creators and owners of content used to train AI consulted, recognised and rewarded? Users may need educating on this issue, particularly as it might relate to uploading their own or copyright content such as for summarisation or translation.
10. How have or will local and societal stakeholders been involved in the development and implementation of the AI, ensuring the right of communities to data sovereignty, that is to maintain ownership, control, and governance over data about themselves and their cultural heritage? Users may need educating in this issue.
11. How is it ensured that users retain agency in their interactions with AI? What is the impact on skills and confidence of users? Users perceive AI to make them simply more "efficient" but what are the wider impacts on their skills, creativity and agency?
12. What are the impacts on human employment, including potential job displacement or the use of precarious labour? Users may be somewhat aware of this issue.
13. What is the environmental impact of the whole lifecycle of the AI service, including power demands and water use? There is growing social awareness of this issue. The library will want to ensure its commitment to sustainability is consistent in its use of AI.

- 
14. If the developer is a third party, what is their business model and their approach to responsible operations? Users may well need educating in this issue in relation to services particularly those from big tech companies.

Scenario 2. Library staff use of generative AI



As a manager you are increasingly noticing that some of your staff are using generative AI to help them perform tasks such as writing emails, creating marketing material, summarising and translating documents and even as a first stop to answering reference queries. In the process they are gaining useful skills.

Other staff refuse to use generative AI for anything. They cite issues such as accuracy and “ethical issues”, including environmental impact.

You are planning a workshop to bring staff practices into alignment and perhaps to generate a policy for AI use.

- What ethical and value issues do you identify?
- What would you do?
- How realistic is the scenario, are there aspects of the current situation that shift the issues?

Notes

1. What are the nature of the benefits to access to knowledge? Is the effort/cost of delivering them justified in relation to library priorities and alternative solutions? It appears that the library is benefiting from staff use of AI both in terms of internal efficiencies and potentially end user access to information, at no apparent cost. But consideration should be given to the less obvious impacts as revealed by answers to other questions.
2. What is the level of accuracy, timeliness and impartiality of information output by AI and how is the level of accuracy explained at the point of use? Staff should be careful about checking the quality of information generated by AI. It may be useful to create generic guidelines and help with prompts to maximise effective use. Perhaps there are some uses that should be considered out of scope.
3. What safeguards are in place to minimise bias and the reproduction of damaging stereotypes and assumptions? Staff may need training in detecting such biases.
4. How are diverse knowledge systems and cultural and linguistic diversity promoted? Staff may need training in how to evaluate outputs for cultural biases and design prompts to counteract them.
5. Is the availability of AI to some bad actors promoting misinformation or censorship?
6. Are the benefits available equitably, including aspects of digital inequality and accessibility? Training may be needed to ensure all staff have the skills to make use of generative AI effectively.

7. Is the application transparent in a socially meaningful way? Is it defined who owns responsibility for errors made by AI and their impacts? As the library may be responsible in the eyes of its users for errors arising from its use of generative AI, the full implications should be considered.
8. How are users' data managed to ensure their privacy and security? This could be a concern if staff are inputting data about the library or even about users.
9. How are the creators and owners of content used to train AI consulted, recognised and rewarded? There may be alternatives to the familiar generative AI services that have more responsible approaches. There are copyright issues potentially around the use of generative AI produced images such as for marketing.
10. How have or will local and societal stakeholders been involved in the development and implementation of the AI, ensuring the right of communities to data sovereignty, that is to maintain ownership, control, and governance over data about themselves and their cultural heritage? There may be alternatives to the familiar generative AI services that have more responsible approaches to the involvement of stakeholders in AI development and implementation.
11. How is it ensured that users retain agency in their interactions with AI? What is the impact on skills and confidence of users? There may be an issue with the degradation of staff skills through the use of AI. Particular care should be taken to avoid the loss of decision-making to AI systems.
12. What are the impacts on human employment, including potential job displacement or the use of precarious labour? There may be alternatives to the familiar generative AI services that have better models of labour practices.
13. What is the environmental impact of the whole lifecycle of the AI service, including power demands and water use? There may be alternatives to the familiar generative AI services that have a better environmental record.
14. If the developer is a third party, what is their business model and their approach to responsible operations? There may be alternatives to the familiar generative AI services that have more responsible approaches to the social impact of AI.

Scenario 3. Using AI to support metadata creation for library material

A library is contemplating a pilot project with a third party AI provider to explore how AI can assist in the creation of metadata. The immediate driver is to deal with the huge backlog of books and other resources for cataloguing. However, there is thought to be potential to extend this to large-scale, digital content in special collections and research data. Discussions with other institutions indicate that only part of the metadata creation process can be easily automated, and that it will continue to require staff to check quality.

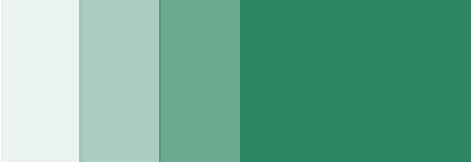
Some staff view the whole proposed project with suspicion: they say it will mean a significant loss of metadata quality and suggest it is setting the library on a pathway towards wholesale automation.

- What ethics and value issues do you identify?
- What would you do?

- How realistic is the scenario, are there aspects of the current situation that shift the issues?

Notes

1. What are the nature of the benefits to access to knowledge? Is the effort/cost of delivering them justified in relation to library priorities and alternative solutions? The benefit to access to knowledge needs clarification. There are conventional approaches to solving the backlog, such as hiring a new member of staff. There may be simpler technologies that can achieve benefits with fewer implications.
2. What is the level of accuracy, timeliness and impartiality of information output by AI and how is the level of accuracy explained at the point of use? The impact on accuracy of metadata needs to be evaluated.
3. What safeguards are in place to minimise bias and the reproduction of damaging stereotypes and assumptions? There seems to be significant risk that AI developed in another context will not adequately represent diverse perspectives.
4. How are diverse knowledge systems and cultural and linguistic diversity promoted? Workflows will seek to include community verification of AI descriptions. Library staff will need to critically evaluate AI outputs for cultural biases and Western-centric perspectives.
5. Is the availability of AI to some bad actors promoting misinformation or censorship? Is AI promoting misinformation or censorship?
6. Are the benefits available equitably, including aspects of digital inequality and accessibility?
7. Is the application transparent in a socially meaningful way? Is it defined who owns responsibility for errors made by AI and their impacts? The library needs to take reasonable responsibility for errors and be transparent about how AI is being used.
8. How are users' data managed to ensure their privacy and security?
9. How are the creators and owners of content used to train AI consulted, recognised and rewarded?
10. How have or will local and societal stakeholders been involved in the development and implementation of the AI, ensuring the right of communities to data sovereignty, that is to maintain ownership, control, and governance over data about themselves and their cultural heritage? It is essential that a wide range of user groups are consulted within the project.
11. How is it ensured that users retain agency in their interactions with AI? What is the impact on skills and confidence of users?
12. What are the impacts on human employment, including potential job displacement or the use of precarious labour? Existing staff roles may be simplified but this could imply a loss of expertise and status. Consultation about these issues should be central to the project.
13. What is the environmental impact of the whole lifecycle of the AI service, including power demands and water use? Assurances should be sought from the AI provider about aspects of sustainability.

- 
14. If the developer is a third party, what is their business model and their approach to responsible operations? The library will have a view on the ethical compass of their potential partner.

An additional resource to support discussion would be Lee's checklist^[6].



Turning insight into action

While these guidelines are just a starting point, the questions and scenarios included in it are meant to spark critical thinking and collaborative decision-making about the role of AI in library settings. We offer a set of recommendations on how to use the scenarios, but we are also aware of the value that lies in the way each library chooses to use it, whether through structured discussions, meetings, or informal chats, policy reviews or community engagement.

Because of this, we would like to know how these guidelines were applied in your context.

Did they help guide a particular decision or conversation?, What was your chosen format? and What feedback emerged from the discussions that took place?

Your experiences will be vital for understanding the real world impact of this resource and to improve future versions. We invite you to reach out to share your reflections, challenges and ideas. Your comments will help us shape a more responsive, inclusive and useful tool for the upcoming versions.

References

- [1] FAIFE (2020) IFLA Statement on Libraries and Artificial Intelligence, <https://repository.ifla.org/items/8c05d706-498b-42c2-a93a-3d47f69f7646>
- IFLA (2012) IFLA Code of Ethics for Librarians and Other Information Workers (Long Version), <https://repository.ifla.org/items/6c3d64ed-4eb8-45e4-bb60-9e178aa94cbf>
- IFLA (2017) IFLA Statement on Digital Literacy <https://repository.ifla.org/items/f47dc600-7bbb-4ae4-b59e-646cc154ba47>
- IFLA (2024) IFLA Internet Manifesto (2024 Update) <https://repository.ifla.org/items/69397ea2-f4f3-49bc-a5ef-342260f7a54c>
- IFLA (2025) IFLA Statement on Core Principles on Libraries and Copyright <https://repository.ifla.org/handle/20.500.14598/3928>
- [2] Association of Research Libraries (2024) Research Libraries Guiding Principles for Artificial Intelligence, <https://www.arl.org/resources/research-libraries-guiding-principles-for-artificial-intelligence/>
- [3] Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature machine intelligence*, 1(9), 389-399.
- Corrêa, N. K., Galvão, C., Santos, J. W., Del Pino, C., Pinto, E. P., Barbosa, C., ... & de Oliveira, N. (2023). Worldwide AI ethics: A review of 200 guidelines and recommendations for AI governance. *Patterns*, 4(10).
- [4] Cox, A. (2021). Ethics Scenarios Of Artificial Intelligence For Information and Knowledge Management And Library Professionals AFTER FEEDBACK. The University of Sheffield. Dataset. <https://doi.org/10.15131/shef.data.17081138.v1>
- [5] Cox, A. (2022). The ethics of AI for information professionals: Eight scenarios. *Journal of the Australian Library and Information Association*, 71(3), 201-214.
- [6] Lee, B. C. G. (2025). The “Collections as ML Data” checklist for machine learning and cultural heritage. *Journal of the Association for Information Science and Technology*, 76(2), 375-396.