Essays



Cutting through the noise: Assessing tools that employ artificial intelligence

Leticia Antunes Nogueira 匝

University Library, Norwegian University of Science & Technology, Norway

Stine Thordarson Moltubakk

University Library, Norwegian University of Science & Technology, Norway

Andreas Fagervik

University Library, Norwegian University of Science & Technology, Norway

Inga Buset Langfeldt

University Library, Norwegian University of Science & Technology, Norway

Abstract

The popularization of artificial intelligence (AI) represents a significant business opportunity for private actors developing tools and services aimed at research and higher education. Academic libraries are often at the receiving end of sales pitches for new tools and could benefit from guidance on how to assess them. Libraries' assessment of tools is a valuable service to library stakeholders, many of whom may not have sufficient time, the necessary competencies or the inclination to explore the landscape of innovations promising to support their information needs and research endeavours. This article offers concrete guidance concerning what to consider when assessing whether to adopt, endorse and/or invest in innovative information and research tools that make use of AI. The main areas proposed for reflection concern (a) tool purpose, design and technical aspects; (b) information literacy, academic craftsmanship and integrity; (c) ethics and the political economy of AI.

Keywords

Academic libraries, AI tools, artificial intelligence, library acquisitions, library services, public procurement

Introduction

Since the days of Melvil Dewey, the role of librarians has evolved from being 'keepers of books' to becoming service providers and information facilitators (Rice-Lively and Racine, 1997). That is, librarians enable the connection between users and the information resources that they need (Burke, 2002). The ways by which librarians fulfil this role has changed alongside the introduction of new information media and changes in which technologies are available to manage and access information resources. Libraries have often been early adopters of new information technologies (DWR Boden, 1993; Westin, 2023), and embracing change has historically been considered a crucial part of staying relevant. The proliferation of digital technologies and the emergence of consumer artificial intelligence (AI) are no different (Nogueira et al., 2024). These innovations are bound to influence how academic libraries are run and which services and expertise library users demand from staff (DWR Boden, 1993). Some argue that, rather than taking charge and being ahead of times, libraries and librarians have been reactive when it comes to AI (Cox et al., 2018). While the issue of





Corresponding author:

Leticia Antunes Nogueira, University Library, Norwegian University of Science & Technology, Høgskoleveien 12 (Statsarkivet), office 206 – 7030, Trondheim, Norway. Email: leticia.a.nogueira@ntnu.no

resources and work capacity is a clear barrier to bold leadership, the matter of staff competencies is also crucial (Gasparini and Kautonen, 2022).

One concrete area in which the need for competency has recently emerged concerns the assessment of innovative tools marketed to research and higher education stakeholders, which increasingly apply some sort of AI.¹ AI, it should be noted, is not a single technology, but rather a collection of multiple automation technologies that simulate human behaviour and cognition, and that can be applied in various ways across different sectors (MA Boden, 2018; Cox and Mazumdar, 2024; Nogueira et al., 2024). These technologies rely on,

complex algorithms and significant computational power to process and analyse large amounts of data used in a variety of problem-solving and decisionmaking tasks. Fundamental characteristics of AI are autonomy and adaptability, which means that computer programs can perform tasks with little human guidance and have the ability to improve their own performance via feedback loops comparable to human learning. (Nogueira et al., 2024: 8)

For example, machine learning, neural networks, deep learning and natural language processing (NLP) are all different concepts related to the broader umbrella of AI^2 (Cox and Mazumdar, 2024; Samoili et al., 2020). Librarians, who in earlier times had the role of making sure the right book was put in the hands of the right user, now have also taken the role of supporting stakeholders as they navigate the increasing complexity of the information structures. The increasing variety and volume of information resources available have meant that librarians have also taken on the role of trying and disseminating new tools for discovering, accessing, evaluating and using information.

As a result, academic libraries find themselves in a privileged position to carry out structured assessments of innovative tools, and thereby perform a valuable service to library stakeholders. Whereas some faculty members show interest in trying out new tools, few have time to prioritize this exploration amidst the many other administrative tasks that pile up on their teaching and research time (James, 2011). Moreover, simply trying out new tools does not make for a thorough assessment. The latter requires multifaceted competencies that information professionals have long cultivated and exercised. Finally, not all stakeholders in higher education and research are inclined to explore and assess new tools, but most could benefit from a trustworthy, structured assessment of new tools that is informed by core ethical values of access to information, responsibility, integrity, independence, neutrality and transparency.

In addition to the demand from users for competencies relevant to new AI tools, libraries often find themselves at the receiving end of sales pitches for tools that promise to optimize, if not revolutionize, how researchers work and how students learn. While libraries – as well as public organizations more broadly – are well versed in negotiating acquisitions and subscription contracts with publishers and other suppliers, the scope of innovative tools brings new challenges that can transform librarianship and research administration. Libraries can and do rely upon their existing competencies with acquisitions; however, new aspects might need to be considered when assessing whether to endorse an innovative tool and/or purchase a license or subscription. Furthermore, new information technologies (including AI) represent such a significant business opportunity for private actors, and the enthusiasm is such, that this business and the companies active in this space have gained their own catchy designations: edutech and eduventures (Encoura, 2022).

Innovative tools can be grouped in two broad categories: multi-purpose tools (Microsoft Copilot and ChatGPT are examples) and task-specific tools (e.g. tools for automatic de-duplication of papers in literature reviews, AI-powered tools for bibliometric anaor tools for discovery of information lysis resources). Although both are relevant for librarians, the first type of tool has a broader range of applications and thus is more conducive to collaboration with others in the university ecosystem that have complementary expertise (e.g. from information technology (IT), legal and/or purchasing departments); the latter type of tool is more specialized and therefore contingent on librarians' own evaluations. Besides, not all eduventures target institutional actors as their main customers, and some target users directly through business models based upon individual subscriptions or credit purchases. Thus, when a student or a member of the university staff asks for advice about task-specific tools that are not a part of the preapproved institutional arsenal, libraries must be prepared to assist.

With this background, this article stands on the premise that it is useful for academic libraries to have concrete guidance concerning what to consider when assessing innovative AI tools, and it seeks to support them in making these assessments. Most existing resources in this area explore the potential uses of AI in library operations (Kautonen and Gasparini, 2024), but do not offer concrete guidance for assessing AI tools. Among the few that do, despite the valuable insights they offer, these are either quite broad (Hennig, 2023) or too specific on single-purpose tools (Upshall, 2022). Even though all information professionals could benefit from gaining deeper insight into how to assess AI tools, in this article, we address more specifically those who are tasked with this kind of assessment and with providing leadership with advice for decision-making. With the popularity of large language models (LLMs), many academic libraries have instituted dedicated working groups whose mandates involve keeping up with AI development and how it can be expected to affect library operations and services. It is primarily these groups that this article addresses.

Our advice is built upon our experiences as researchers, instructors and librarians in research and higher education, and our engagement with assessing new tools and formulating institutional guidelines for generative artificial intelligence (GenAI). Because consumer AI is at present the object of much hype and exaggerated claims (both cheerful and gloomy), we think it is useful to be transparent that we adopt an open but cautious standpoint on the matter. That is, while we welcome technological change and embrace innovations that we consider add value to library and academic work, we are also attentive to undesirable side effects, mindful of ethical implications and concerned about the fact that the short-term benefits are easier to estimate than long-term shortcomings associated with technological change. As a result, we consider that the missions and values of academic libraries imply that it is important to adopt sober scepticism and the precautionary principle of responsible innovations (European Commission, 2000; Stilgoe et al., 2013). This means resisting pressure to embrace tools if the arguments for doing so rely on little more than the fact that they are technically feasible, or the claim that everyone else is using them and we do not want to be left behind. It is with this attitude that this article offers guidance for assessing AI tools.

The article proceeds as follows: the following section contextualizes the issue with brief considerations about public procurement (assuming the viewpoint of academic libraries in public universities), its role in promoting or hindering innovation and how the particularities of AI tools might challenge these processes as they are intended. With this scaffolding in place, the article proceeds to discuss and justify what critical aspects to consider when assessing an AI tool. We organize our discussion on the following themes: (a) tool purpose, design and technical aspects; (b) information literacy, academic craftsmanship and integrity; (c) ethics and the political economy of AI. Next, these discussions become more tangible by in the form of concrete guiding questions for reflection that readers can refer to when assessing AI tools. We emphasize, however, that answering these questions will not automatically lead to a clearcut decision on whether or not to adopt a tool. Nor is answering all the questions a necessary condition for using new tools. Rather, concrete decision-making will depend on the context, on different user groups and on resources available to support implementation. A given tool could be considered unsuitable in a particular context, but appropriate elsewhere with necessary preparations or guardrails. Our ambition is to lay the groundwork for a principled reflection on innovative tools and services for research in higher education in the age of the automation of knowledge work.

Public procurement, innovation and digital technologies

The public sector – including universities – has a responsibility to manage resources effectively and in alignment with their public service missions. To ensure that the acquisition of goods and services is transparent and based upon broad and fair competition, procurement regulations seek to institutionalize the process, and thereby prevent protectionism, undue influence and other maladministration of public resources.

Procurement processes can take time, and statutory routines tend to be rigid. Therefore, it is not uncommon for public organizations to establish a range of binding framework agreements that covers a range of ad hoc purchasing needs.³ Moreover, it is also a common practice to establish an expenditure ceiling, and so long as an order stays below this threshold, no structured procurement process is required. For instances in which procurement is necessary, it is important to prepare clear and precise requirements for the wanted product or services. This includes specifying things such as functionality, costs, contract terms and competition form that determines how the procurement will be carried out and which procedural rules apply. Additional concerns to be included in tender specifications refer to climate and the environment. What follows is an open market competition, in which the supplier that best answers the tender is awarded the contract. The process of determining the best supplier is based on various case-specific factors. Before a public organization can add new suppliers to its list of agreements, it may need to set additional requirements. For example, if the services involve collecting and storing personal data from people in Europe, a data processor agreement is

required, in accordance with the General Data Protection Regulation (GDPR, 2016).

Public procurement is essentially a process driven by the internal needs of public organizations. However, in the context of digital tools, particularly those involving AI, demand arises not out of the identification of prior needs, but from the contemplation of what opportunities innovations afford. This changes the dynamics between buyers and sellers.

Because public actors can have a propelling effect over the demand for certain products or services (Edler and Georghiou, 2007), obtaining a contract with a public organization can matter a great deal for private businesses. In many cases, public actors are the most important if not the only customer of a supplier (known in economics as an oligopsony). In the context of innovations, a public actor could be the first customer of a supplier. As well as ensuring a stable and foreseeable revenue stream, a public contract might guarantee visibility and confer trustworthiness. Hence, public procurement might strengthen path-dependencies, particularly in contexts that are more risk-averse, where institutional stability is important or where ill-considered or reckless experimentation might have devastating consequences.

Nonetheless, public procurement has also been pointed out as important for innovation policy and the advancement of a range of societal interests (Edquist and Zabala-Iturriagagoitia, 2012; Uyarra et al., 2020). By creating and/or directing demand for things that market mechanisms alone are not able to foster (Nogueira et al., 2023), public procurement can help steer the market towards (or away from) a certain direction and enable innovation and structural change (Uyarra et al., 2020). That is because the strategic priorities reflected in the procurement strategy have a signalling effect, which matters in addition to the potential economic weight of public contracts. In other words, simply by endorsing a certain solution or technology, public organizations can influence their developmental paths and that of their competitors. As a result, 'public sector procurement faces competing priorities, such as cost-efficiency, legal conformity, the advancement of environmental protection and the promotion of innovation' (Patrucco et al., 2017: 269).

The matter of public procurement as being more than the simple acquisition of goods and services, but as a strategic instrument that affects the dynamics of the marketplace, has been much debated in the context of sustainability transitions (Krieger and Zipperer, 2021; Rainville, 2022). With increasing digitalization and the emergence of consumer AI, the ways by which public organizations position themselves towards innovative tools and the companies that develop these is crucial to determine how the adoption and diffusion of these tools will take place. While this represents an opportunity, we cannot take for granted that solutions presented by suppliers to the needs of public organizations are directly comparable, nor that it is adequate to establish price-based competition as the key criterion. Solutions such as AI tools for use in research and higher education are not commodity products. Thus, competition for university contracts or endorsements in this context involves increased complexity. Academic libraries can contribute to assessing what value these tools offer and, in this manner, untangle complexity.

What to consider when assessing innovative AI tools and services

Right from the start, we hope to make clear that assessing new AI tools does not require technical expertise in software development and coding. It does, however, require a certain degree of technological literacy. For this reason, some information professionals might find that assessing AI tools stretches their existing expertise, and that it is in itself a competency-building activity. It can also be worthwhile to consider performing these assessments in teams rather than individually. On the theme of libraries and LLMs, Lorcan Dempsey has made the point that,

as the library positions itself as a source of advice and expertise about use of AI tools in information and data related tasks, it is important to develop understanding of how LLMs and associated tools work, what tools are available, what they are good for and what their limits are. (Dempsey, 2023)

The sentiment can be extended to technologies other than LLMs or GenAI.

Despite the overarching rhetoric of our time, machines do not in fact think. It is unnecessary to equate what machines do with thinking to appreciate how impressive their capabilities can be. Treating digital tools (with or without AI) in a matter-of-fact way, rather than feeding extravagant expectations of human obsolescence, is useful for making us aware of surprises concerning which kinds of tasks these tools perform well and which they do not (Dempsey, 2023). Telling these apart requires close examination.

It is crucial for library and information professionals to be knowledgeable on how AI tools work, in order to be able to consider purported functionalities, underlying algorithms and matters pertaining to data – both the data that users provide as they use these tools and the data used for development in the first place. In this section, we present what we consider to be important considerations for assessing AI tools.

Considering tool purpose, design and technical aspects

Librarians and other information professionals might be reluctant to try and understand technical matters of AI tools, with the justification that such matters entail sophisticated IT competencies that they lack. Nonetheless, people who do not understand how their tools work are not qualified to judge the extent to which these tools are fit for purpose, how they perform and what limits apply to their use. For example, the fact that ChatGPT is able to write code in Python came as a surprise not only to users but also to its developers, since the original purpose of LLMs was focused on producing natural language, not computer code. This unexpected application of the tool was a welcome addition to its capabilities, and it makes sense given LLMs' ability to identify patterns in strings of words. However, other unintended uses that have emerged, such as generating recipes for new dishes or providing medical advice, are not fit for purpose, although they are technically feasible. This example highlights the importance of having a general understanding of the purpose and functioning of AI tools when assessing their potential and limitations.

Librarians do not need to be able to program an AI model (unless they are inclined to doing so), but they do need to grasp how they work if they are to adopt or make recommendations in favour or against these tools. Moreover, while librarians may not have specialized competencies to assess technical choices pertaining to the design of a tool, they do have specialized competencies relevant to the context in which many of these tools are to be employed. This is valuable and cannot be substituted.

The first thing to consider is what task the tool proposes to do. Considering what ambitions originally informed the design of tools is of crucial importance. This matters not only for assessing where and how they can be employed, but also how they perform in light of the proposed use and what their limitations may be. This is especially the case when users find different uses than originally envisioned. It is increasingly the case that tools are stretched outside their original specifications and employed for applications not originally foreseen. A quintessential example of this concern are LLMs, which were designed with capabilities that allow the automation of certain communication tasks (i.e. as stochastic parrots, according to Bender et al., 2021), but that can easily be misused if treated as information sources (Nogueira and Rein, 2024; Nogueira et al., 2024). Hence, what became known as chatbot 'hallucination' is not exactly a side effect, but a consequence of them not having been designed to provide users with reliable information. Chatbots, such as ChatGPT, are neither oracles nor truth machines, and uses that differ from the capabilities afforded by their design need to take this into consideration. This should not be taken to mean that tools must only be used as originally intended. Rather, the point is that information professionals (as well as users more generally) must be aware of the purpose, design and modus operandi of a tool for making informed judgments on how they can be creative in using them differently than intended.

Another matter pertaining to the technology and design concerns an understanding of whether and how AI is employed by the tool in question. Although LLMs are a novelty to the wide public, people have been using tools that contain some kind of AI for several years without special remark. A few examples are auto-complete functions in search engines, recommendations of related content in a journal's web page and pattern matching in plagiarism detection tools. To be clear, concrete tools in these applications can have different degrees of technical sophistication; some may employ nothing more than well-established algorithmic decision trees, while others may employ complex prediction models. In some cases, the presence of AI may not be obvious, but other cases might instead feature exaggerated claims about the extent to which the technology du jour is employed - be it GenAI, smart devices, blockchain verification or something else entirely. It is therefore crucial to consider how sophisticated a tool is, and to what extent marketing claims are merely signalling virtues developers think users have come to value.

In addition, some vendors may be secretive about details regarding how their tools work or deflect this type of inquiry altogether. Reasons for this vary. It could be that their technology is not (or cannot be) protected by formal intellectual property rights, or it could be because transparency would shed light on objectionable practices like ghost work (when a task is made to appear automated but it is carried out by a human, such as content moderation or labelling data) (Gray and Suri, 2019) or models trained on copyrighted material without explicit permission or acknowledgement (Vincent, 2022). It could also be that particular properties of the technology (such as deep neural networks or other types of sub-symbolic AI) do not lend themselves to insight into the

reasons or steps behind a given output. In any case, it matters whether information is unavailable because of lack of disclosure or for legitimate reasons, and this consideration should inform the assessment of a tool.

Matters pertinent to privacy and information security are also crucial when it comes to consumer AI tools. Even though this kind of assessment requires competencies that are perhaps more at home in IT departments than libraries, a general understanding is not only useful but ever more critical for librarians. This involves understanding what potential economic interests apply to data that AI tools might have access to, and what this means for both users and vendors. Consider, for example, the implications of sharing data relevant to a piece of research or a dataset that could lead to a potentially lucrative patent. Non-economic interests are also crucial, particularly when they entail a duty to protect sensitive, or otherwise personal or confidential information, as is often the case in research and higher education. In summary, a libraries' general understanding of information security entails grasping how vendors manage privacy and data security, as well as the relevant interests that apply to the data in question.

Finally, as the library interacts with vendors and their tools, it is crucial that they be able to tell what various claims and functionalities actually mean, what consequences they pose for users and for information environments and what the concerns relevant to privacy and information security are.

Considering information literacy, academic craftsmanship and integrity

Since the early days of bibliographical instruction tied to specific media such as physical library catalogues and encyclopaedias, librarians in higher education have been central actors in disseminating information literacy. Amidst a changing landscape of resource types and means of accessing them, librarians' focus in this area has shifted towards supporting users' development of robust critical thinking skills applied to information, particularly regarding identifying and evaluating sources as well as using and referencing information appropriately (Grafstein, 2002).

According to the Chartered Institute of Library and Information Professionals (CILIP), information literacy concerns 'the ability to think critically and make balanced judgements about any information we find and use' (CILIP, 2018: 3). Digital technologies – especially AI – bring a new dimension to information literacy. That is, being information literate is fundamental for engaging effectively with new tools as well as sources of information and disinformation. At the same time, the variety and amount of information sources and tools for managing them are likely to affect users' information skills. We must also remember that far beyond text, there are other means to communicate information. From infographics to podcasts, the medium of information is ever more diverse, and AI tools can support different types and channels of knowledge dissemination. While this is a good thing, it also demands sharper abilities for navigating systems of information, whose integrity may have been compromised.

AI tools can both support and challenge academic work. On the one hand, AI-powered tools enable new kinds of data analysis and improved ways to discover and manage information. This can, in principle, support new discoveries and faster and more efficient workflows, which could increase the pace of scientific advancement. On the other hand, the side effects of this new pace and the costs of increased presumed efficiency are unknown. The risk is that, in academia - a context already plagued by heightened competition, a strict regime of performance metrics and a pervasive structure of incentives that foster varying degrees of misconduct (Edwards and Roy, 2017) - AI can fuel harmful practices even faster. The risk of a reproducibility crisis is also a challenge and may lead to a loss of confidence in science (Ball, 2023).

For librarians assessing new consumer AI tools, this means taking into consideration that many tools might have a certain declared use, but also allow for dubious applications that do not conform to academic values. Chatbots are the most obvious example, but not the only one. Take the case of tools that employ recommendation algorithms upon a database to suggest content. They may propose new ways of searching and discovering new sources, but they may also disincentivize in-depth reading. A tool that reliably suggests sources based upon inputted text might save the reader time searching but might also offer the possibility of cutting corners and citing the sources that are suggested without having read them properly.

While this might seem to some like a small price to pay for convenience, earlier instances of bad citation practices illustrate the dangers of spreading misinformation. In the case of how uncritical citation of a source that did not even constitute empirical research, Leung et al. show that,

[A] five-sentence letter published in the Journal [i.e. *The New England Journal of Medicine*] in 1980 was heavily and uncritically cited as evidence that addiction was rare with long-term opioid therapy. We believe that this citation pattern contributed to the North American opioid

crisis by helping to shape a narrative that allayed prescribers' concerns about the risk of addiction associated with long-term opioid therapy. (Leung et al., 2017: 2194).

Although it is not the role of the library to control how exactly users will apply the tools made available to them, it is the library's role to advise on such matters so that principles of (research) ethics can be upheld. Overall, the more convenience a tool offers, the more important it is to be mindful of users' preexisting level of competency and academic proficiency, as well as the possible consequences for users' opportunities to practice and strengthen information literacy and research skills.

In this way, the relationship between AI and information literacy is a two-way street. Users need information literacy to engage with the outputs of AI tools, be they synthetic text,⁴ source recommendations or other outputs; but users' ability to engage effectively with information is also influenced in turn by information environments in which AI is prevalent. This is not merely about being able to trust that the outputs yielded by these tools can be taken to be valid and truthful. It also relates to being confident about how these tools will affect academic craftsmanship.

Although AI use in academic work of the kind we discuss here is too recent to have allowed for the study of its effects on students' learning or on researchers' competency, we know from earlier instances of automation in the context of intellectual and creative professions that computers affect how people engage with their work (Carr, 2015). Going from actively performing a task to overseeing a computer perform a task (or quality-checking the computer's output) can lead to automation bias (i.e. undue and excessive trust that software's outputs are objective and flawless) and complacency (i.e. disengagement from work due to this belief), ultimately resulting in a deterioration of skills (Carr, 2015). Nonetheless, some research has started to emerge on this issue, such as Abbas et al. (2024), who suggest that students' use of these tools tends to increase when they feel pressured for time, and who also found a connection between students' use of ChatGPT and increased procrastination and memory loss. Coping with pressures related to time and performance are commonplace among both students and researchers. The pitch of many AI tools claiming to offer efficiency gains must be weighed against a reflection on what might be lost by adopting them. Should efficiency come at the cost of students' learning, the well-being of students and researchers and trust in the quality of research, then such a trade-off is unjustifiable. It would also do us well to remember that much of creativity and innovation comes about through friction and serendipity; favouring tools that eliminate these elements in favour of a seamless user experience may therefore be counterproductive.

There are many important critical questions to be asked regarding what is gained and what is lost when it comes to academic craftsmanship when you choose to employ AI tools. AI tools may make some tasks easier to do, especially if they are tedious. The risk is gradually losing the ability of carrying out a task independently or never developing that skill at all. With this in mind, we recommend that AI should ideally be used to enhance human capabilities and open new avenues for knowledge development. This could include analysing patterns in large datasets or powering search databases for non-textual sources. AI can also effectively automate tasks that add little value compared to the effort they entail, such as consolidating data from various systems, preparing datasets for analysis or categorizing (meta)data. Finally, AI is also useful in areas where individual users already possess substantial expertise, thereby mitigating the risk of deskilling. It is crucial to carefully consider tasks that, if delegated to an AI system, might hinder users from developing their own skills in those areas.

These instances require a case-by-case assessment. For example, AI enables users who lack coding skills to interact with sophisticated systems that previously necessitated such abilities. This lowers the threshold of necessary competencies and democratizes access to resources that were once limited to those with coding expertise. But we should think twice before deciding there is no need to learn coding any longer or allowing our coding competencies to deteriorate leaving us alienated and vulnerable. People's individual consideration of whether AI affects their decision to invest time and resources learning code should depend on how critical such expertise is to that person, or how much value it adds to their skillset. The same goes for writing and a range of other complex and nuanced competencies involved in academic craftsmanship.

Considering matters of information literacy and academic craftsmanship does not necessarily mean that libraries will be the final arbiter on whether or not users adopt these tools. Many developers adopt business models that target end users directly, meaning that it becomes important to focus on (a) building a baseline of competency for users to be able to assess tools themselves, and (b) strengthening the confidence of users in the library's advice and the reasoning underlying that advice. Information professionals can support students and scholars as they grapple with these questions themselves, which is particularly valuable when they are under pressure and may make choices that, although helpful in solving a short-term problem, come at the cost of compromising their best interests in the long term.

Information literacy and academic craftsmanship are not binary competencies that one either has or does not have. They are the result of experience and are things that need to be cared for continuously. This is because changes in the environment necessitate adaptation, and the information landscape of our times is changing at a fast pace. Traditional rules of thumb for source evaluation, although still valuable and worth teaching students about, are insufficient to deal with the 'current strategies of persuasion and deception' (Gigerenzer, 2021: 244) introduced by surveillance capitalism and strengthened by a widespread climate of distrust in cornerstone institutions. As a result, individuals must cultivate a learning attitude and critical thinking skills to navigate the evolving information environment. Libraries can assist with that and have been doing so for a long time.

Considering ethics and the political economy of AI

Ethics in the context of AI is a hotly debated topic. Some of this discussion entails existential risks to humanity, artificial general intelligence (i.e. the postulation that AI will at some point surpass human cognitive capabilities) and the so-called singularity (i.e. a hypothesized situation in which the development of AI has advanced beyond human control and which entail unpredictable and irreversible consequences) (for an overview, see MA Boden, 2018); yet, we contend that a more pragmatic approach to more immediate ethical challenges is advisable. Certainly, events with low probability but catastrophic impact ought to be well assessed, but it would be unfortunate to use these scenarios to distract from pressing ethical matters, and even more unfortunate to use potential risks as an argument to create entry barriers that reinforce the current power of incumbent technology companies (Hanna and Bender, 2023; Stilgoe, 2024).

Hence, when it comes to libraries' awareness of ethical challenges, we highlight the following matters as deserving special attention: (a) the consequences of uncritical AI use for research ethics and academic integrity; (b) the ramifications of opacity and low explainability in AI algorithms for different use areas; (c) the impact of AI tools on users' competency, autonomy and worldview; (d) the trade-offs involved between preventing hate speech and promoting censorship; (e) potential bias, discrimination in both inputs and outputs of AI tools; (f) unacknowledged human labour in AI tools that make them seem more sophisticated than they really are (e.g. ghost work); (g) the environmental costs of training and using GenAI; and (h) disputes on what constitutes fair use of data protected by intellectual property rights for training AI. The upcoming section presents concrete questions for reflection on these themes; note that not all questions are concentrated under the umbrella of ethical considerations, since ethical matters permeate all other aspects of developing and using AI tools.

Whereas matters of ethics in general and in research ethics in particular are deeply embedded in libraries' values and practices, the ways in which ethical considerations permeate the politics and economics of information and knowledge production are not always clearly within librarians' traditional skillset. Nonetheless, libraries' long-standing support to values such as intellectual freedom, transparency and advocacy in favour of equitable access to information through open access, open source and open licenses (American Library Association (ALA), 2017; International Federation of Library Associations and Institutions (IFLA), 2012) indicate that librarians are more embedded in the development of the political economy of modern knowledge societies than they often recognize.

It is not news that access to information and the means of organizing it have been the object of commercial interests for a long time. Academic publishing is widely recognized for being outrageously lucrative (Buranyi, 2017), and libraries engage continuously in negotiations that try to balance providing users with access to resources they deem important with the need to manage a limited budget. If data are the new oil (The Economist, 2017), the economic significance of information cannot be taken for granted. Moreover, because eduventures often adopt datadriven business models that might conflict with either the public service mission of academic libraries in public universities or with commercial interests of private, information-intensive organizations such as research institutes, the economic implications of data and information pose intricate challenges for libraries committed to managing knowledge as a public good. Libraries must navigate the tension between commercial interests of publishers/vendors and their commitment to knowledge dissemination, equitable access and societal benefit.

The proliferation of AI tools intensifies this challenge, as knowledge organization and access to information resources become increasingly mediated by AI tools developed and marketed by technology companies. In the words of Nicholas Carr, 'If we don't understand the commercial, political, intellectual, and ethical motivations of the people writing our software, or the limitations inherent in automated data processing, we open ourselves to manipulation' (2015: 208). Hence, a general appreciation of the power relationships and competitive dynamics between different technology developers and vendors in the AI landscape, as well as the business models they adopt (i.e. the political economy of AI), becomes worthy of consideration for information professionals.

In practice, this means seeking information about vendors and developers of technologies and appreciating the way they relate to one another. This can be seen not as an assessment for each individual tool, but a general understanding of the economic dynamics surrounding technology companies, from big ones such as Meta (parent company of Facebook) or Alphabet (parent company of Google) to small suppliers that fall within the umbrella of edutechs. Much has been said about how innovative technologies follow a probable path of hype cycles, and how most of GenAI applications are rapidly moving from the stage known as 'peak of inflated expectations' toward the 'trough of disillusionment' (Chandrasekaran and Ramos, 2023). For information professionals, the point is to be clear on the fact that, despite the rhetoric of disruption that accompanies both enthusiasts and sceptics, not all AI technologies have been creating value at present. Yet, much of the hype surrounding commercial AI is due to commercial actors finding themselves in an arms race for market dominance. In this process, these companies not only seek recognition as pioneers, but also to make a profitable business out of AI, thereby recouping their investments in developing the technology with interest.

It is also worth remembering that the concept of openness, that is highly valued by the library community, can be hijacked by technology companies that do not truly license their products according to opensource principles, despite using the term (Gibney, 2024). The availability of AI models at no cost for use as components of tools is not a sufficient criterion for labelling them as open source, especially when the use of these models is prohibited 'in several areas that might be highly beneficial to society' (Maffulli, 2023). Just like sustainability ideals are subverted in greenwashing, ideals of openness can also be subverted in 'openwashing' (Doctorow, 2023). Information professionals must be attuned to this.

Much of marketing materials on AI tools propose a range of applications for this technology. Nonetheless, it is not sufficient to state, for instance, that AI can revolutionize literature search, if the details of how this is supposed to happen are left obscure. AI is not a single technology, and these materials often fail to explain how AI is applied and why it adds value. In meeting with these proclamations, many information professionals may feel that their lack of understanding about how AI is supposed to advance these proposed ambitions stems from their own lack of competency in AI. In fact, it is more likely that the reason for their lack of understanding can be attributed to vague explanations from developers and vendors. By appreciating both the ethics and the political economy of AI (and of knowledge economies more broadly), information professionals can be empowered to ask vendors for the information they need, rather than assuming they should know it in the first place. Thus, it is up to potential users of these tools and customers of these vendors to assess whether they create value, and if this value is commensurate with the tool's price and other intangible costs that can be more difficult to measure.

Next, we present guiding questions for reflection in the assessment of tools that employ AI.

Guiding questions for reflection when assessing innovative AI tools and services

Information professionals interested in exploring AI tools experience a range of drivers for assessing new tools and services. At one extreme, they are motivated by concrete needs and requirements they must meet (need-oriented), and at the other end of this continuum, they are presented with new tools whose usefulness and adequacy they must assess (opportunity-oriented). In the marketplace fight for the preference of librarians, researchers and students, most (if not all) developers offer some free demonstrations of their tools. This may be available in a range of ways, including freemium models (in which basic features are made available free of charge, and premium features require a subscription), as a time-restricted free access (such as the week or month being free) or demonstration workshops. We recommend that librarians take advantage of these opportunities for trying tools by coming prepared to these free trials or demos.

Table 1 offers foundational guiding questions for assessing AI tools, and Table 2 presents an additional set of questions that require more in-depth engagement with the tool and more advanced competencies to assess. These questions emerge out of our practical experiences and professional insights gained through assessing and working with AI tools in academic settings, as well as the challenges and opportunities we have encountered in our roles as researchers, librarians and instructors. A preliminary version of these questions was designed by the first author, and quality-

Table 1. Foundational guiding questions for reflection when assessing innovative AI tools.

Overarching question	In-depth questions (consider as applicable)			
 What is the declared benefit of the tool? What problems does it solve? 	 1.1. What kind of benefit does the tool offer? (i) Speeding up existing tasks, making them more convenient or reducing costs. (ii) Improving existing tasks, achieving higher quality or minimizing mistakes. (iii) Complementing current workflows by performing tasks differently. (iv) Expanding users' capabilities, broadening the range of tasks or analyses available. (v) Eliminating the need to perform certain tasks. 1.2. How does the tool change (or fit with) the current standard workflow? 1.3. How does the tool balance personalization with generalizable or replicable results? 1.4. How do we know that the tool performs well? What indicators (or qualitative signs) can be used to measure performance? Is it possible to measure the accuracy/error rate (or rate of false positive/negative where applicable)? Is it possible to measure how often atypical, anomalous or mistaken results appear? 1.5. Does the tool enhance human capabilities or substitute the need for human input or judgment? 1.6. What training or support do users need (from vendors, from the library, etc.) before and during use? 			
	 1.7. Does the tool integrate with other software (from the same vendor and others)? In what way? 1.8. If the service were to be terminated, what are the consequences for transferring data, analysis, etc., to other suppliers (compliance with compatibility standards)? 			
2. About the range of applications: Is the tool single-purpose or muti-purpose by design?	 2.1. Is there a primary or most indicated use? In what ways other than the primary indication can we imagine users might employ the tool? Is there clear information about unsupported use cases? 2.2. Are all use areas in accordance with principles of (research) ethics? 2.3. Does the tool perform equally well in all areas? 2.4. What are the limits of use? 			
3. How does the tool work?	 3.1. Do developers claim to use Al? What kind of Al is used? Examples can be optical character recognition, natural language processing, speech recognition, read-aloud, sentiment analysis. To what end is Al used? Examples can be to categorize data, to recognize patterns, to recommend other matches or sources, to detect anomalies, to predict behaviour or events, to generate content. 3.2. How transparent and detailed are vendors about how the tool works, its limitations and vulnerabilities? Do vendors balance correct terminology with clarity, or is technical jargon used in a way that discourages further questioning? Is information communicated primarily in earnest or in language that seeks to safeguard against legal disputes or liability? Is information about the modus operandi of the tool easy to access? Are the limitations or vulnerabilities inherent to the tool discussed in a transparent way? 			
For tools combining GenAl capabilities with factual data/information sources or other services:	 3.3. Which foundational model is used (OpenAl's, Google's, Meta's, other)? 3.4. Which data or information sources support the service (e.g. Open Alex, Semantic Scholar, proprietary scholarly databases)? 			
What ethical considerations are involved in using this tool?	 4.1. Does the tool provide recommendations or prescribe actions? To what extent do users need to rely on trust versus knowledge and understanding? 4.2. How much does the tool oversimply (or rely on oversimplified) ambiguities or uncertainties of the world? 4.3. Does the tool reinforce or rely on the assumption that the future will be like the past (induction)? 4.4. Have developers established guardrails (for instance, to prevent the generation of bigoted or otherwise dangerous outputs)? To what extent do these guardrails steer meanings or implement censorship? 			

checked independently by the other authors in a practical exercise assessing an AI tool. Based on their feedback concerning the practical issues in using the framework, a revised version was formulated. In parallel, the authors sought feedback from an experienced professional in software development, whose feedback was also incorporated.

Combined, Table 1 and Table 2 address the concerns discussed in the previous section. It is important to be mindful that not all questions will apply to all sorts of tools, and that the absence of information (or difficulty in accessing it) can in itself be information to consider.

Some of the questions can be assessed simply by using the tool, but others - especially in Table 2 will require seeking other sources of information beyond that which is provided by vendors. As it can be expected, it is unlikely that answers will be easily accessible in the form that they are asked, and that reflecting upon these issues will require that users dig into not only the marketing information available for a tool, but also into the fine print terms and conditions. Furthermore, some questions will trigger the need for searching other sources than the vendor themselves. All this effort can admittedly be daunting and intimidating, but the effort pays off through increased competency in navigating a rapidly changing information environment and empowers academic libraries to keep up with technological change.

We remark that the questions in Table 1 should not be taken up as a hindrance to lower stakes experimentation. A less structured attitude is suitable to the goal of becoming familiar with the kinds of tools being launched and with what they offer, and also as an early encounter with new tools to be selected for further assessment. The structured approach we suggest is useful for those who either benefit from the support it can offer and for those who need to carry out more thorough evaluations. That is, a structured approach helps in ensuring that the tools information professionals choose to work with and recommend meet a variety of specific needs. It also supports libraries in becoming prepared to offer robust advice on the potentials and limitations of tools, as well as on what conditions should be met for tools to be used in accordance with principles of research ethics and academic integrity. Finally, a structured assessment also facilitates the comparison of functionalities across competitors, which is particularly valuable as they seek to differentiate themselves and position their tools as unique.

One limitation to the guiding questions in Tables 1 and 2 is that they are primarily informed by the authors' practical experiences and professional insights and would benefit from empirical validation. As such, the framework may require adaptation to specific institutional contexts, as not all questions will apply universally across different tools or settings. Additionally, the rapid pace of technological development in AI may mean that new considerations emerge over time, which could warrant further refinement of the questions presented.

Conclusion

AI has the potential to simplify many tasks; it does not, however, make the world any less complex (Nogueira, 2023; Nogueira et al., 2024). On the contrary, a seemingly endless stream of advanced AI tools has emerged and is being marketed, sometimes quite aggressively, which raises the competency bar required of academic libraries. In this article, we have discussed what are important things for librarians and other information professionals to consider when assessing innovative AI tools and services. This assessment is crucial not only as a service offered to library users, but also a competency-building activity that sets libraries up for keeping their continued value and relevance in a technologically changing world.

Public organizations, including universities and academic libraries, have a significant responsibility when it comes to the adoption of innovative tools. The way they position themselves towards these tools can influence both their development and diffusion. However, public procurement regulations may not be sufficiently well equipped to handle matters other than those driven by internal needs, as is the case of new opportunities afforded by technological development and innovation. Amidst a landscape characterized by hype cycles (Chandrasekaran and Ramos, 2023), where AI can be used as a buzzword to advance sales, libraries are particularly well-suited to take on the role of assessing innovative tools. Libraries are trusted as non-commercial institutions, and librarians are regarded as impartial information professionals. It is crucial to maintain the legitimacy and independence of libraries, as well as to uphold the trust placed in them by the public.

We must also keep in mind that libraries are not immune from user demands and from pressures that emerge from the so-called fear of missing out. Despite the collaborative environment of research and higher education, universities do compete for students, for talented personnel, for prestige and for funding. In this competitive landscape, the ability to signal one's innovativeness can be crucial, and adopting AI tools is instrumental in this regard. This kind of

Table 2.	More advanced	guiding que	stions for r	reflection when	assessing innovativ	e Al tools.

Overarching question	In-depth questions (consider as applicable)
5. How has the tool been developed?	 5.1. What information is available about the algorithms underlying the tool? 5.2. What information is available about the data used for developing/training the tool? 5.3. Is there any information available about how the tool prioritizes between eventual trade-offs either in development or under use?
	Examples are trade-offs between accuracy and interpretability, bias and variance, data privacy and model performance, model complexity and computational resources, model complexity generalizability, etc.
	5.4. What is the carbon footprint of developing (and using) the tool?What other sustainability aspects matter for the development and/or use of the tool?5.5. Is there any legal or moral dispute surrounding the development of this tool (e.g.
	Copyright breach, use of ghost work)? Did the development of the tool entail questionable labour practices (such as ghost work)? How can this affect future work produced by using the tool?
6. Who has developed the tool and what business model is adopted?	6.1. Did the vendor develop the tool themselves or were there third parties involved?Who has been involved and how?
	6.2. Is the vendor a startup company or an incumbent actor? If the vendor is a startup, how has it been funded (venture capital, private/public research grants, private capital investment, debt, etc.)?
	If the vendor is an incumbent: Has the vendor develop the tool themselves or were there third parties involved? Who has been involved and how?6.3. What is the reputation of the vendor/developer in relation to:
	Functionality and support (including for other tools/services)? Business and competitive practices?
	Ethics, lobbying activity and other societal interests?
	6.4. Do vendors focus on institutional customers, individual customers or both?6.5. Who are the vendors' main competitors? What solutions do they offer and how do they differ?
	Are there open-source alternatives?
	Do licenses that brand themselves as open source actually comply with open-source principles?
	6.6. What business model supports the tool?
	(through tokens or credits), supported by advertising, supported by mining user data. 6.7. To what extent do prices adequately reflect actual technological sophistication and
	proposed benefits in contrast to the hype surrounding the technology?
7. What are relevant	7.1. Is the tool a cloud-based application or does it involve the download of a software?
considerations to privacy and information security?	What are the implications for functionality and information security?
	7.2. How does the vendor manage privacy and data security?
	Are there previous occurrences of security breaches, data leaks, data tempering? What
	happened and how has each of these events been handled/mitigated?
	What are the known technical vulnerabilities to the software?Do third parties have access
	to the data? Is this legitimate and necessary for the tool to work?
	Examples are publicly available data, private user data, personal data from others than the user, sensitive data, confidential data.
	7.4. Can users opt out from sharing data?
	What consequences are there for the performance of the tool if data are not shared?
	7.5. What are the relevant interests (economic or otherwise) pertinent to data?
	From the viewpoint of users and the data they access from users.
	Do vendors use user data for improving the tool, such as training models? In what ways?

competitive pressure plays a role in determining which new technologies and practices are adopted. This, combined with the fear of missing out, may lead to hasty decisions and rapid implementation of new tools without thorough evaluation or adequate preparation for supporting subsequent needs of users. Institutions may rush to keep up with their peers and to satisfy students and staff who question them about why seemingly 'everyone else' has access to a given tool but them. This kind of competitive pressure plays a role, and it is important for institutions to be deliberate about their approach to this issue and avoid feeding hype cycles.

Our recommendations take as a point of departure the perspective of academic libraries associated with public universities. As a result, it is important to verify how they apply in other contexts. Nonetheless, we consider that other organizations - including public and private research organizations - that have intangible assets to protect and to whom information competencies are critical can also benefit from this article. Professionals other than librarians can also benefit from reflecting upon these questions as they consider what types of AI tools could benefit them. Robust information competencies are crucial in modern knowledge societies, regardless of the most recent technological frenzy (Nogueira, 2023). They endow people to make informed decisions about their information needs, as well as about how innovative tools can be employed thoughtfully and intentionally to their benefit. Librarians' long-standing work facilitating access to information and supporting critical thinking and information literacy remains key for demystifying technology and cutting through the noise, as consumer AI becomes more pervasive.

Acknowledgements

The authors would like to thank colleagues at NTNU Library for interesting conversations about how AI is challenging academic libraries, particularly in the context of Project Laibro. We are also grateful to Heine Skipenes, Bodil Åberg Mokkelbos and Silje Reiten Blichfeldt for the opportunity of gaining insights into how they worked in assessing the possible adoption of Copilot for Microsoft 365 at NTNU, and to Odracir Antunes Júnior and Alexandra Angeletaki for their feedback. We also acknowledge Jan Ove Rein for pointing out the connection between uncritical citations and the opioid crisis.

Declaration of conflicting interests

The authors declared no potential conflicts of interest.

Disclosure

This article reflects the authors' independent work. AI assistance has been employed (through Microsoft Copilot

365, as approved and provided by the Norwegian University of Science and Technology) exclusively to the following ends: (a) to translate selected sentences written preliminarily as a draft in languages other than English, (b) to reformulate selected English sentences that included jargon or that earlier feedback pointed to as lacking clarity, (c) as a thesaurus. An example of prompt is 'please suggest five alternative formulations to the sentence that begins after the colon, respecting the original meaning and style: ...'. In all cases, both inputs and outputs from the tool have been meticulously refined by the authors before being included in the manuscript. The authors assert their moral rights as the creators of this piece, as well as their responsibility for it.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

ORCID iD

Leticia Antunes Nogueira (b) https://orcid.org/0000-0002-8842-3790

Notes

- In addition to newcomers, suppliers of established tools that earlier used other vocabulary to describe their services (e.g., data analytics) have now rebranded themselves and state that they too apply artificial intelligence. This posed additional challenges for professionals without technical background who need to engage with these tools in an informed manner.
- Artificial intelligence, particularly of the generative kind (GenAI), is at the moment the subject of much hype, since it generates new content, but other kinds of AI are useful for performing routine tasks related to data and information processing, where the processes are repetitive and consistent, and the need for judgment is minimal (Cox and Mazumdar, 2024; Nogueira et al., 2024).
- 3. These agreements make simpler the process of ordering frequent items or ad hoc services; yet these agreements are also put in place through open market competition, and contracts are awarded for a predetermined duration before a new tender is issued for the same service.
- 4. Synthetic text refers to text that is yielded as outputs by chatbots or other tools, and which has not been authored by humans (Bender et al., 2021).

References

- Abbas M, Jam FA and Khan TI (2024) Is it harmful or helpful? Examining the causes and consequences of generative AI usage among university students. *International Journal of Educational Technology in Higher Education* 21: Article 10.
- American Library Association (ALA) (2017) ALA code of ethics. Available at: www.ala.org/tools/ethics (accessed 1 November 2024).
- Ball P (2023) Is AI leading to a reproducibility crisis in science? *Nature* 624(7990): 22–25.

- Bender EM, Gebru T, McMillan-Major A, et al. (2021) On the dangers of stochastic parrots: Can language models be too big? In: Proceedings of the 2021 ACM conference on fairness, accountability, and transparency, virtual event, 3–10 March 2021, pp. 610–623. New York: Association for Computing Machinery. DOI: 10.1145/3442188.3445922.
- Boden DWR (1993) A history of the utilization of technology in academic libraries. Available at: https://files.eric.ed.gov/ fulltext/ED373806.pdf (accessed 20 June 2024).
- Boden MA (2018) Artificial Intelligence: A Very Short Introduction. Oxford, UK: Oxford University Press.
- Buranyi S (2017) Is the staggeringly profitable business of scientific publishing bad for science? *The Guardian*, 27 June. Available at: www.theguardian.com/science/2017/jun/27/ profitable-business-scientific-publishing-bad-for-science (accessed 1 November 2024).
- Burke L (2002) The future role of librarians in the virtual library environment. *The Australian Library Journal* 51(1): 31–45.
- Carr N (2015) The Glass Cage: How Our Computers Are Changing Us. Reprint edn. New York: W. W. Norton & Company.
- Chandrasekaran A and Ramos L (2023) Hype cycle for generative AI, 2023. Available at: https://www.gartner.com/en/ documents/4726631 (accessed 1 November 2024).
- Chartered Institute of Library and Information Professionals (CILIP) (2018) CILIP definition of information literacy 2018. Available at: https://www.cilip.org.uk/news/421972/ what-is-information-literacy.htmI (accessed 13 December 2024).
- Cox AM and Mazumdar S (2024) Defining artificial intelligence for librarians. *Journal of Librarianship and Information Science* 56(2): 330–340.
- Cox AM, Pinfield S and Rutter S (2018) The intelligent library: Thought leaders' views on the likely impact of artificial intelligence on academic libraries. *Library Hi Tech* 37(3): 418–435.
- Dempsey L (2023) Generative AI and libraries: Seven contexts. Available at: www.lorcandempsey.net/generative-ai-andlibraries-7-contexts/ (accessed 1 November 2024).
- Doctorow C (2023) 'Open' 'AI' isn't: Not 'open', nor 'artificial', and certainly not 'intelligent'. Available at: https://pluralistic. net/2023/08/18/openwashing (accessed 1 November 2024).
- Edler J and Georghiou L (2007) Public procurement and innovation— Resurrecting the demand side. *Research Policy* 36(7): 949–963.
- Edquist C and Zabala-Iturriagagoitia JM (2012) Public procurement for innovation as mission-oriented innovation policy. *Resarch Policy* 41(10): 1757–1769.
- Edwards MA and Roy S (2017) Academic research in the 21st century: Maintaining scientific integrity in a climate of perverse incentives and hypercompetition. *Environmental Engineering Science* 34(1): 51–61.
- Encoura (2022) Higher education technology landscape: A categorization of technology providers. Available at: www. encoura.org/eduventures-research/evolve-your-online-strategy/ tech-landscape/2022/ (accessed 1 November 2024).
- European Commission (2000) Communication from the Commission on the precautionary principle. Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex% 3A52000DC0001 (accessed 1 November 2024).
- Gasparini A and Kautonen H (2022) Understanding artificial intelligence in research libraries – Extensive literature review. LIBER Quarterly: The Journal of the Association of European Research Libraries 32(1): 1–36.
- GDPR (2016) General data protection regulation, Pub. L. No. (EU) 2016/679, 88. Available at: https://eur-lex.europa.eu/

legalcontent/ EN/TXT/PDF/?uri=CELEX:32016R0679 (accessed 15 July 2024).

- Gibney E (2024) Not all 'open source' AI models are actually open: Here's a ranking. *Nature*. DOI: 10.1038/d41586-024-02012-5.
- Gigerenzer G (2021) How to Stay Smart in a Smart World: Why Human Intelligence Still Beats Algorithms. Dublin: Penguin Books.
- Grafstein A (2002) A discipline-based approach to information literacy. *The Journal of Academic Librarianship* 28(4): 197–204.
- Gray ML and Suri S (2019) *Ghost Work: How to Stop Silicon* Valley from Building a New Global Underclass. Boston, USA: Houghton Mifflin Harcourt.
- Hanna A and Bender EM (2023) AI causes real harm. Let's focus on that over the end-of-humanity hype. *Scientific American*, 12 August. Available at: https://www.scientificamerican.com/ article/we-need-to-focus-on-ais-real-harms-not-imaginary-existential-risks/
- Hennig N (2023) Evaluating generative AI tools for purchase: A checklist. Available at: https://nicolehennig.com/evaluatinggenerative-ai-tools-for-purchase-a-checklist/ (accessed 1 November 2024).
- International Federation of Library Associations and Institutions (IFLA) (2012) IFLA code of ethics for librarians and other information workers (long version). Available at: https://repository. ifla.org/handle/123456789/1850 (accessed 1 November 2024).

James A (2011) Too many tasks. Nature 475(7355): Article 257.

- Kautonen H and Gasparini AA (2024) B-Wheel Building AI competences in academic libraries. *The Journal of Academic Librarianship* 50(4): Article 102886.
- Krieger B and Zipperer V (2022) Does green public procurement trigger environmental innovations? *Research Policy* 51(6): Article 104516.
- Leung PTM, Macdonald EM, Stanbrook MB, et al. (2017) A 1980 letter on the risk of opioid addiction. *The New England Journal of Medicine* 376(22): 2194–2195.
- Maffulli S (2023) Meta's LLaMa 2 license is not open source. Available at: https://opensource.org/blog/metas-llama-2license-is-not-open-source (accessed 1 November 2024).
- Nogueira LA (2023) Trivsel i ikke-informasjonens tid [Thriving in times of non-information]. Available at: www.bokogbibliotek. no/aktuelt/trivsel-i-ikke-informasjonens-tid/197985 (accessed 1 November 2024).
- Nogueira LA and Hegle J, Kalkvik S. (2024) Artificial intelligence, real library - Final report for Project Laibro. Available at: https:// hdl.handle.net/11250/3165570 (accessed 21 november 2024).
- Nogueira LA, Lindeløv B and Olsen J (2023) From waste to market: Exploring markets, institutions, and innovation ecosystems for waste valorization. *Business Strategy and the Environment* 32(4): 2261–2274.
- Nogueira LA and Rein JO (2024) Chatbots and large language models: The case for not citing chatbots as information sources. In: The scholarly kitchen. Available at: https://scholarlykitchen. sspnet.org/2024/06/20/the-case-for-not-citing-chatbots-asinformation-sources-part-ii/ (accessed 1 November 2024).
- Patrucco AS, Luzzini D, Ronchi S, et al. (2017) Designing a public procurement strategy: Lessons from local governments. *Public Money Manage* 37(4): 269–276.
- Rainville DA (2021) Stimulating a more circular economy through public procurement: Roles and dynamics of intermediation. *Research Policy* 50(4): Article 104193.
- Rice-Lively ML and Racine JD (1997) The role of academic librarians in the era of information technology. *The Journal of Academic Librarianship* 23(1): 31–41.

- Samoili S, Lopez Cobo M, Gome Gutierrez E, et al. (2020) AI WATCH: Defining artificial intelligence: Towards an operational definition and taxonomy of artificial intelligence. Available at: https://publications.jrc.ec.europa.eu/repository/ handle/JRC118163 (accessed 13 December 2024).
- Stilgoe J (2024) Technological risks are not the end of the world. Science 384(6693): Article eadp1175.
- Stilgoe J, Owen R and Macnaghten P (2013) Developing a framework for responsible innovation. *Research Policy* 42(9): 1568–1580.
- The Economist (2017) The world's most valuable resource is no longer oil, but data. 6 May. Available at: www.economist. com/leaders/2017/05/06/the-worlds-most-valuable-resourceis-no-longer-oil-but-data (accessed 1 November 2024).
- Upshall M (2022) An AI toolkit for libraries. *Insights* 35: Article 18.
- Uyarra E, Zabala-Iturriagagoitia JM, Flanagan K, et al. (2020) Public procurement, innovation and industrial policy: Rationales, roles, capabilities and implementation. *Research Policy* 49(1): Article 103844.
- Vincent J (2022) The scary truth about AI copyright is nobody knows what will happen next. Available at: www.theverge. com/23444685/generative-ai-copyright-infringement-legalfair-use-training-data (accessed 1 November 2024).
- Westin M (2023) The 1970s librarians who revolutionised the challenge of search. Available at: https://aeon.co/essays/the-1970s-librarians-who-revolutionised-the-challenge-of-search (accessed 1 November 2024).

Author biographies

Leticia Antunes Nogueira is the Head of Section for Resources and Digital Services at the NTNU Library (Norwegian University of Science and Technology). She earned a PhD in Innovation Economics from Aalborg University, Denmark, in 2018. Nogueira is enthusiastic about exploring the interdependencies between sociotechnical change, the economy and society. In recent years, her work has focused on the implications of AI technologies for information literacy and academic integrity.

Stine Thordarson Moltubakk is a Research Librarian at the NTNU Library (Norwegian University of Science and Technology). They hold a Master's degree in Science and Technology Studies (STS). Moltubakk is passionate about the ethical and environmental implications of emerging technologies and is dedicated to advancing research within these fields.

Andreas Fagervik is a Research Librarian at the NTNU Library (Norwegian University of Science and Technology). He holds a Master's degree in Educational Science and works with teaching and guidance related to literature search, reference management and source criticism. Fagervik is particularly interested in the intersection between pedagogy and learning technologies.

Inga Buset Langfeldt is a Head Librarian at the NTNU Library (Norwegian University of Science and Technology). She holds a Bachelor's degree in Library and Information Science and is currently writing her master's thesis. Langfeldt is passionate about information and media literacy and the implications of current and emerging technologies on these competencies.