EXPLORING THE EFFECTS OF GENERATIVE AI ON ACADEMIC LIBRARIES

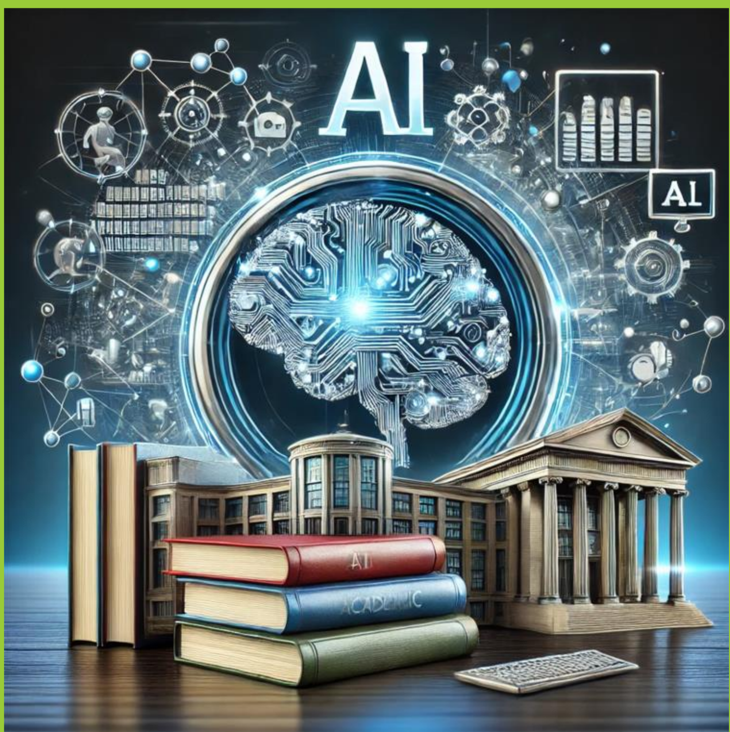

White paper; Pleiade Management & Consultancy

Maurits van der Graaf; 26 August 2024

CONTENT

Exploring the effects of Generative AI on Academic libraries

Maurits van der Graaf

Pleiade Management & Consultancy

https://pleiade.nl/

Email: m.vdgraaf@pleiade.nl

2024

Image on frontpage created by ChatGPT with the prompt 'exploring the effects of generative AI on academic libraries'

PREFACE

WHAT WILL BE THE EFFECTS OF GENERATIVE ARTIFICIAL INTELLIGENCE (AI) ON ACADEMIC LIBRARIES?

This is a key question for everyone involved in academic libraries. It might be impossible to give a complete answer to this question at this moment in the development of AI.

Although various types of AI have been around for quite some time, the AI developments have accelerated significantly since the launch of publicly available chatbots based on generative AI since November 2022. At the time of writing this White Paper, nearly two years have passed and at least some contours of the answer to the above-mentioned question are becoming clearer. In this White Paper, we attempt to sketch some of these contours. Of course, it needs emphasising that this sketch certainly is incomplete and may be incorrect when speculating on future developments. Nevertheless, we believe that it is worthwhile to attempt to understand where we are at this point in time: it will help monitoring the developments in AI and its applications for academic libraries and hopefully it also might assist academic libraries to make informed choices about what to prioritize.

1 WHAT DO YOU NEED TO KNOW ABOUT AI AS A LIBRARIAN?

In order to drive a car, you do not need to understand exactly what happens under the hood. This is often said, but in practice, it is not entirely true: every driver knows quite a few things about a car that are necessary to drive and maintain it. This chapter aims to describe this type of essential knowledge about generative AI for librarians.

1.1 TYPES OF AI

There are several forms of AI. A common form is *narrow AI*, which refers to AI applications that carry out a single task, such as image recognition or speech recognition. Another common form is *generative AI*: these are AI systems that can perform various tasks by generating text, images, or other media in response to commands (known as 'prompts') from users. This enables a conversation between the user and the AI system, which of course is highly appealing.

Generative AI is based on so-called *foundation models.* These are models trained by *machine learning* using large data sets. These *foundation models* have learnt the patterns and structure of their input training data. This is then used as a 'foundation' for generative AI systems to generate new data with similar characteristics, such as text, imagery, video, audio, and computer code.

The best-known form of a *foundation model* is the so-called *Large Language Model* (LLM), which generates language. Meanwhile, there are also systems that generate images or computer code, and these have been merged in ChatGPT, for example. However, the developments in AI do not stop. There are also *small language models* (trained on a smaller, more focused dataset and requiring less computer power), which are used in so-called *edge AI*, where AI processing takes place on devices with limited computing power, such as sensors or mobile phones.

These LLM models are data hungry: the models become more accurate with more data. This 'the more, the better' phenomenon means that the producers of AI chatbots are looking for more data. For example, OpenAI has converted the audio of YouTube videos into transcripts to be used for its AI model[1]. Most of the data in the AI models so far have been harvested from the internet, but owners of websites (such as publishers) increasingly are prohibiting this or selling it for considerable amounts[2,3,4]. Some Big Tech companies have attempted to solve this by creating so-called synthetic data, but this has led to the collapse of AI models[5]. This means that the search for more human-created data will continue. This might also have consequences for library collections (see chapter 7.1).

Finally, it is important to realise that AI is a system technology. System technologies are ubiquitous, constantly improving technologically and leading to complementary innovations. Other system technologies are the internal combustion engine and electricity. Because it is a system technology, AI will have a great impact on many aspects of our society.

1.2 THE DOWNSIDES OF AI AND RESPONSIBLE AI

DOWNSIDES OF AI

The downsides of generative AI can be grouped under four themes, summarised as the four B's: bias, black box, Big Tech, and bad actors:

- **Bias:** The *foundation models* of generative AI are built on large amounts of structured and unstructured data from various sources. However, these datasets have significant problems: they are often not representative of the diversity within our society or of how we would like to see society. Specific groups in society are underrepresented in these data, leading to a distorted picture. The datasets reflect the past and contain outdated stereotypes that are now considered unacceptable. Moreover, the quality of the data often leaves much to be desired: it contains erroneous and incorrect data. Finally, the way the data are collected sometimes violates copyright and privacy laws[6].
- **Black box:** Generative AI applications work with billions of parameters[6]. As a result, their responses are often not traceable and thus these systems effectively function as a black box. This raises questions about reliability and the reproducibility of these models. Are the algorithms based on scientific knowledge or on pseudoscience? Is the distinction between correlation and causation properly made? The inner workings of these AI systems (for example, ChatGPT-3 has 175 billion parameters[7]) are incomprehensible, even to their developers. This Black Box phenomenon is even more troubling because AI systems sometimes can 'hallucinate', making up facts and references but presenting them very convincingly. Therefore, the question of explainability of the results of AI applications is of utmost relevance (see also 2.2).
- **Big Tech:** Generative AI requires large datasets and enormous computing power. As a result, only a handful of Big Tech companies, all based in the United States[8], are capable of developing and managing this technology. For what purposes are these applications optimised? The answer is that these are not necessarily socially desirable purposes. For example, social media have AI-driven perverse mechanisms designed to capture attention, leading to addictive behaviour in users[9]. So, there is a concentration of market power that makes people and organisations dependent on these Big Tech companies and erodes their digital autonomy.
- **Bad actors**: Bad actors can misuse AI applications to mass-produce disinformation and deepfakes that are barely distinguishable from the real thing. There is a strong case for a watermark or another recognition mechanisms, but there is no universal solution to this yet in sight. Additionally, the use of AI for cybercrime, such as voice cloning for fraudulent purposes, is a growing problem.

RESPONSIBLE AI

The downsides of AI highlight the need for *responsible* AI. Responsible AI focuses on the development and deployment of AI applications aligned with human values. It aims to harness the potential of AI while minimising risks, such as excluding specific groups in our society. The key characteristics of responsible AI are:

- Transparency and reliability
- No *bias* that puts people at a disadvantage
- Safeguarding of privacy
- Accessibility and inclusiveness.

Responsible AI is the course chosen by the European Union, which distinguishes Europe from the USA (with a more capitalist approach) and from China (with a focus on state control)[10].

Responsible AI includes both the development of AI applications and their practical implementation and the impact/effects on users or society. In a recent scientific article[11], the authors provide an overview of all the elements of responsible AI (see Table 1) and also elaborate on corresponding scientific research questions.

Dimension	Description
Fairness	AI systems should enable inclusion and diversity and not lead to discriminatory outcomes.
Transparency	AI systems should be open and transparent regarding processes and outcomes and facilitate traceability
Accountability	AI systems should be developed considering the responsibility and accountability of their outcomes with ethics and principles.
Robustness and safety	AI systems should be developed with a preventative approach to risks and in a manner that they behave as intended while minimising unintentional and unexpected harm.
Data governance	AI systems should ensure that adequate data governance covers the quality and integrity of the data used throughout the entire lifecycle.
Laws and regulations	AI systems should adhere to the respective laws and regulations that dictate their functioning.
Human Oversight	AI systems should generate tangible benefits for people and always stay under human control.
Societal and environmental well-being	AI systems should promote sustainability

Table 1 Dimensions of responsible AI [11]

1.3 SOCIETAL ASPECTS OF GENERATIVE AI

- **Investments in AI**: Hundreds of billions of dollars are currently being invested in AI, with the USA and China taking the lion's share[12]. While a lot of this money probably will turn out to be wasted, you can expect many innovations to come out of it as well.
- **Effects on labour market**: A report by McKinsey[8] looked at the potential to automate 2,100 work activities of about 850 occupations by generative AI. The report concluded that 60%-70% of these work activities could be automated by AI. However, this transition will not happen overnight: they estimate that around half of these activities will indeed be automated around 2045. Generative AI applications will mainly affect higher-skilled knowledge workers - unlike other automation waves that have hit mainly executive, lower-skilled workers. To illustrate this, knowledge workers currently spend an average of 1 day per working week searching and gathering information. Tasks that generative AI applications easily can take over.

- **Energy consumption:** Generative AI consumes lots of energy. For instance, a ChatGPT-powered search uses nearly 10 times as much electricity as a Google search. Big Tech companies are now constructing many new data centers, and their energy consumption is putting a strain on the electricity grid. Apparently, Big Tech is even looking into energy projects such as nuclear fusion to meet the future energy demands of AI[13].

1.4 COPYRIGHT

At the AI4Libraries conference in 2023, D. Hansen gave a very interesting talk about copyright and generative AI[14]. He made the following statements:

1. The purposes of copyright clauses in American law are: (1) to promote science and useful arts; (2) the protection of copyright lasts for the life of the author plus 70 years; (3) the protection is about the creativity in the text, not the ideas expressed in the text; (4) the main exclusive rights granted by copyright are to reproduce the work and to prepare derivative works.
2. Regarding generative AI and copyright, there are three main questions**:**
    (1) Is it permissible to use copyrighted work as training data? At the moment, this question is currently being legally disputed in many lawsuits in the USA. However, for libraries the exception for teaching, scholarship and research in the copyright law is important.
    (2) Are AI outputs infringing on the copyright? This is referred to as the ‘Snoopy problem’: if you ask an AI system to draw pictures of Snoopy (a copyrighted character) it does recreate it and thus the copyrighted work.
    (3) Are generative AI outputs protected by copyright? No, to be an author and claim copyright, you must be human. This was the outcome of the lawsuit regarding the selfie taken by a monkey (see picture)[1].

- What is the significance of this for libraries?
    - Text- and datamining (TDM) for academic research has a strong basis in the law because of the earlier mentioned exception for scholarship, teaching, and research.
    - Metadata generated by AI cannot be copyrighted, at least in the USA.
    - Libraries could help its users by reviewing the terms of use of AI tools used by researchers and students in order to guarantee that the uploaded information cannot be reused, and the privacy of the user is protected.
    - In addition, it is noteworthy that for sophisticated, elaborate prompts, it is probably possible to claim copyright.

[1] This is not necessarily true for other jurisdictions, such as in the UK[15]

1.5 REGULATION IN EUROPE: THE AI ACT

The European Commission's AI Act is expected to be fully implemented around 2026[16] [17]. The AI Act addresses risks from AI and distinguishes four domains:

- **Prohibited applications**: Certain AI applications, such as social scoring, are explicitly banned.
- **High-risk applications**: A group of AI applications classified as high risk must comply with specific regulations before and during use. These include transparency, human oversight, data accuracy, and monitoring risks to security and fundamental rights. The list of high-risk applications covers eight domains[18].
- **Virtually unregulated applications**: AI applications that interact with individuals must make themselves known as AI systems.
- **Foundation models**: This has become a separate category with specific obligations for foundation model providers, including specific transparency and data governance requirements, such as to verify the suitability of data sources.

Although the AI Act marks a new phase for AI applications at least in Europe, its success will depend on how national governments and other parties work towards clear and adequate standards, a consistent interpretation of the regulations, and an effective structure for supervision and cooperation in the field of AI[17].

2 GENERATIVE AI IN LIBRARY APPLICATIONS

In this chapter, we look at possible application types with AI and discuss one type – Retrieval Augmented Generative AI [RAG]) – as a setup for applications that seems most relevant for use in academic libraries. This against the background that the transparency and reliability of the answers of commercial AI chatbots are problematic, which is especially unacceptable in academic libraries. Therefore, the causes of this are briefly discussed as well.

2.1 APPLICATION TYPES WITH GENERATIVE AI

How can one apply the generative AI models in the practice of academic libraries? In the literature, three options have been mentioned so far:

- **Direct prompts via the interface or via the API of the chatbot**: A good example is the Leessimpel app: this Dutch application uses the API of ChatGPT. The user takes a picture of a letter from, for example the tax office, this is converted into text and sent, along with an elaborate prompt, to ChatGPT in order to generate a clear text in simple wording and a clear description what the recipient of the letter should do. The app is very successful and a great help for people with limited language skills. By the way, refining a prompt to ensure that the AI chatbot does what you want it to do, is called 'prompt engineering.'
- **Refining and retraining an off-the-shelf LLM:** This requires a lot of resources in terms of computer power and human resources for training and evaluating the model's performance. At this stage, this appears to be a bridge too far for most library applications.
- **Retrieval Augmented Generative AI (RAG):**

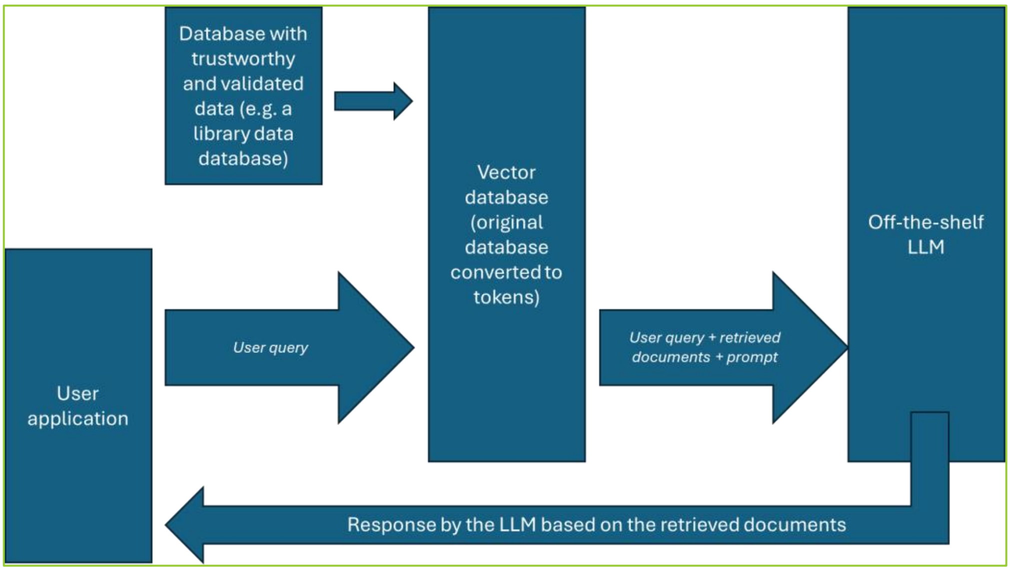

Figure 1 Simplified sketch of Retrieval Augmented Generative AI

Retrieval augmented generative AI is a combination of using LLMs as offered by the Big Tech companies and using trustworthy and validated information. In simple terms, it works as follows (see also figure 1). The user's question is fed into a vector database (which contains the trusted data converted into tokens) and retrieves all documents related to this specific question. In the next step, the user's question, together with the retrieved documents, is fed into the large LLM with a prompt instructing the LLM to use the attached documents in answering the question. A RAG-type application has been developed for example at the library of the Northwestern University[19] and is also in use at Scopus[20].

2.2 TRANSPARANCY AND RELIABILITY

The tendency of AI-chatbots based on LLM models to sometimes hallucinate casts as a shadow over Generative AI and is particularly unacceptable for use in academic libraries. Hallucinations are defined as the content generated by LLMs that seems plausible but is factually unsupported. Huang et al.[21] mention three main sources: (1) erroneous data in the model; (2) mistakes in the training of the model; (3) errors in the so-called decoding algorithm that creates the output of the AI-model.

To comprehend the causes of these hallucinations, one has to realize that LLM models do not understand a certain topic in the way humans do. The answers of a LLM model are generated by calculating stochastically which word should follow the previous one. Although the results are often amazingly accurate, they can sometimes be wrong or even nonsensical.

Regarding mitigating hallucinations, efforts are being made on all three domains (erroneous data, training mistakes and decoding algorithms).

With regard to the usage of correct data, the retrieval augmented generative AI (RAG) is a promising approach and appears to be relevant for use in academic libraries. However, hallucinations do keep occurring, especially if the retrieval process has come up with irrelevant documents[21].

Another approach to using reliable data is suggested by Breeding[22], who proposes to create a LLM solely based on scholarly information ("without the crazy stuff on the internet") that might generate objective, factual, and reliable outcomes. However, apart from the fact that there are also erroneous data in the scholarly literature, it looks like some hallucinations are inherent to LLM models. For example, Huang et al [21] mention 'long tail knowledge' as one cause of hallucinations: when on a certain topic there is little information, the statistical processes behind the LLM models have difficulties handling such paucity of data resulting in wrong outcomes. Nevertheless, the idea of an LLM solely based on scholarly information for use in science is very appealing and deserves more attention and research.

3 AI-EFFECTS ON SCIENCE

This chapter is primarily based on the 'Report to the President' on AI in science, which provides a comprehensive vision on the effects of AI on science and scientific methods[23]. This report describes three effects of AI: an acceleration of science, a transformation of science and changes in scientific infrastructure and collaboration.

3.1 ACCELERATION OF SCIENCE

AI helps prioritise the most likely scientific routes to pursue. One example of this is that AI models can sift through large amounts of data in order to evaluate hundreds of thousands of candidates for medicine or materials for engineering applications to identify the most promising solutions to a given problem. It also allows scientists to focus on research by leaving secondary tasks to AI. Although the output from existing LLM models is currently not sufficiently reliable, the authors expect more professional-quality AI writing assistants in the near future that will eventually be adopted by the scientific community. Also writing of computer code, performing literature reviews, and acquiring expertise in adjacent scientific fields might be for the larger part carried out by AI tools. The authors also predict automated laboratory procedures by AI-powered robots.

3.2 TRANSFORMING SCIENCE

The authors also foresee a transformation of science and scientific methods by AI. In this context, they mention 'digital twins': models based on AI that simulate highly complex systems such as the climate, quantum chemistry of materials, or cellular structures. These digital twins are already in use in the industry, whereby for example all processes of a factory are simulated in a model[24]. This makes it possible to intervene in the digital clone to see what the effects of this intervention are, before intervening in the real system. In the long term, such digital twins might be developed for economic or societal systems with the aim to testing the effects of potential interventions. Other aspects of the transformation of science are that multimodal foundation models will bring together multiple forms of data in order to simulate complex systems using the knowledge and insight of different scientific disciplines. Finally, the ability of AI to use large, heterogeneous, and noisy real-world data creates new opportunities for scientists for new types of research.

3.3 INFRASTRUCTURE AND COLLABORATION

The authors foresee that the self-correction mechanisms of science will also be applied to the use of AI tools in research. This will eventually lead to a culture of responsible and ethical AI practice and models to test and evaluate these practices. A key part of this will be sharing of AI models and data, thereby reducing the total cost of using AI. If such a shared AI-infrastructure is established, AI-assisted research will also become accessible for smaller institutions with fewer resources and even for members of the general public. The authors also foresee that AI will facilitate collaboration between researchers from different scientific backgrounds, helped by AI to overcome differences in language and levels of expertise.

What is the relevance of all this to academic libraries? We can distinguish between a few direct consequences and a few more speculative ones:

More direct consequences:

- Because of the data hunger of AI models, research datasets will become even more important and thereby increasing the relevance of Research Data Management services provided by academic libraries.
- As multidisciplinary research and computational research will gain prominence, the library collection and other library services have to reflect that.
- AI-assisted writing of scientific publications will affect plagiarism detection by libraries and may make the detection of fraudulent papers more challenging.

More speculative consequences stemming from the acceleration of science and AI-assisted writing are:

- More publications: This acceleration would likely result in an increase in the volume of publications. As an analogy: there are now already more images produced by AI than by traditional photography[25].
- Speed of publishing: A shorter time between submission and publication will be required, as well as faster availability of published works in library collections.
- The nature of scientific publications might change: Why publish text when data and models are key? The atomisation of the scholarly record might progress further with the publication of different kinds of data used and produced by the AI, software code, images, videos. Text publications might also change their form. For example, traditional literature overviews in article introductions might lose their significance if written by AI tools.

4 AI-EFFECTS ON USERS INTERACTING WITH THE DIGITAL LIBRARY

The behaviour of library users will change under the influence of their usage of AI tools. This will have important implications for the services of academic libraries. We foresee two effects of AI on the behaviour of library users: one effect that is happening now and one effect that is probably still several years away: (1) digital interaction becomes a dialogue; (2) personal assistants.

4.1 DIGITAL INTERACTION BECOMES A DIALOGUE

Chatbots such as ChatGPT, combined with speech and image recognition, are revolutionizing interaction with the Internet. This can be summed up as "away from the keyboard" and "away from English.” It is, and increasingly will become possible to interact with the Internet in a dialogue in your own language - written or spoken. This means that users get an answer to their questions instead of lists of information sources that they have to read to find the answer.

4.2 PERSONAL ASSISTANTS AS THE NEXT STEP

Our research into the possibilities of AI for people with visual impairments found that a range of AI-assisted wearable devices are being developed that, for instance, translate images on-the-fly to spoken text[26]. As a surprising by-catch, we found that these technical developments are also having an effect on the mainstream market: various forms of "wearables" have already been developed that can compete with the smartphone. We mention two examples: (1) Meta Smart Glasses, which are glasses with a camera and a sound system combined with an AI system. The user simply says, ‘Hey Meta’ and then asks what he or she sees. The AI model then answers, and follow-up questions are possible[27]; (2) Rabbit R1 is a device that works with voice commands and carries these commands out[28]. It can, for example, select a song on Spotify or order a cab. This Rabbit device can be seen as an early form of an AI-based personal assistant, which executes commands without the user having to open specific applications. According to Bill Gates [29], in the next five years these personal assistants will totally change the way we use computers: they will open applications themselves to execute our commands and will do so with insight into our personal preferences. This will likely have a great impact how users interact with library services. If you think this is far-fetched, remember that that Bill Gates has been right before!

5 AI-EFFECTS ON DISCOVERY

In this chapter, we delve deeper into the implications of changes in interactions with the digital world for the discovery of scientific information.

5.1 PARADIGM SHIFT IN THE DISCOVERY PROCESS: TOWARDS CONVERSATIONAL DISCOVERY

As previously mentioned, digital interaction is increasingly becoming a dialogue with AI chatbots. This transformation is so significant that it constitutes a paradigm shift in the discovery process. Internet search machines such as Google so far gave a list of information sources retrieved on the basis of the user query with search terms that sometimes are combined with Boolean logic. It was left to the user to make a selection among the retrieved information sources and read and digest these to find the answer to the question. Now, AI chatbots such as ChatGPT or Google Gemini provide immediate answers to the user's questions, formulated in natural language. And the user can follow-up with additional questions to which the AI chatbot will respond, memorizing its previous answers. This process is also called 'conversational discovery.'

For the scientific discovery process, it is of course crucial that the answers provided by AI chatbots are based on referenced scientific publications. Already the first steps in this direction are being made. Scopus – using a variant of the retrieval-augmented AI-model as described in paragraph 1.2 – provides already answers based on the most relevant articles with snippets from these articles on which the answer is based on[20 30]. Additionally, Ex Libris has announced the launch of a similar AI-feature for their Primo discovery tools[31], and undoubtedly, other scientific search engines are working on this as well.

Marshall Breeding suggests that this could potentially be taken a step further: an AI-tool that extracts and cites relevant portions of text when researchers are writing a scientific paper[22]. If such a tool indeed would be developed and would become widely used, this might in the longer term also negatively influence the significance of citation rates and citation ranking.

5.2 CROSS-LANGUAGE SEARCHING

Another aspect of this paradigm shift in discovery is that, due to the language capabilities of AI, users can search in their own language and receive answers in that language. In other words, AI supports cross-language searching. In addition, people outside academia that are interested in scholarly information for professional or other purposes prefer to read the results of their queries in layman language[32]. Several initiatives to generate layman abstracts are underway[33].

5.3 CHANGING RELATIONSHIP BETWEEN SITES WITH CONTENT AND AI-CHATBOTS/SEARCH ENGINES THAT WILL AFFECT THE USAGE OF LIBRARY COLLECTIONS

Another aspect of this paradigm shift is the changing relationship between content sites and search engines. So far, this generally was beneficial for both parties. Making your content site available for the Google index generated traffic to your site. However, when Google Gemini directly provide an answer to a user (using the content of your site), there is no longer a need for that user to visit your site[22].

If this type of discovery is integrated to scholarly search engines, we anticipate that researchers will in many cases not access the publications on which the answer is based because the provided answer is satisfactory. We think that this will especially be the case when the user query was about consolidated knowledge such as presented in books and review articles. One can assume that very recent articles still need to be read.

Once AI-assisted search engines will have become the standard for scholarly discovery, the usage of library collections, repositories of the libraries and sites of scholarly publishers by human users will likely decrease. This shift might lead to the development of new types of usage metrics: whereas COUNTER statistics currently measure views and downloads per article by end-users, they might need to include usage figures from AI-assisted search engines that are the result of the queries by human users.

6 METADATA AND DIGITAL HERITAGE COLLECTIONS

There are already quite a number of experiments regarding the creation and/or enrichment of metadata using generative AI, especially related to digital heritage collections. In this chapter, we discuss the results of these experiments in the context of 'collections as data' and linked data.

6.1 LIBRARY CATALOGUING

In an article with the title 'From ChatGPT to CatGPT', Bruzstowicz describes experiments using AI for copy cataloguing[34]. ChatGPT has apparently been trained on data from catalogues such as WorldCat, the Library of Congress, the National Library of Medicine, the British Library, COPAC, Europeana, and the HathiTrust Digital Library. This shows in the results: the accuracy of ChatGPT-generated MARC records was comparable to manually created records. The author concludes that this could significantly reduce the resources required for copy cataloguing.

However, a potential significant problem with this way of automatically creating MARC records lies in the various copyright rules. The chatbot might regenerate an exact copy of a MARC record protected by copyright, leading to copyright infringement. Additionally, as we have seen in paragraph 1.4, MARC records created by AI cannot be copyrighted. If this way of copy cataloguing would become widespread, this could negatively impact the business model of WorldCat.

6.2 DIGITAL HERITAGE COLLECTIONS

GENERATING METADATA

Librarians at the UCLA library are working on a pipeline to automatically extract metadata from their digitized library collection with the aim to including it in the library catalogue[35]. Their collection consists of scanned newspapers, photographs, and manuscripts. The procedure followed involved three phases: a tool to perform OCR on the scanned materials, a tool to identify named entities and AI chatbots to create the metadata records (Dublin core and MODS). The AI chatbots were able to produce fairly accurate descriptive metadata (title, description), but were less accurate with metadata that required any sort of external information. Their conclusion is that these tools make it possible to create metadata of passable quality, making the heritage materials discoverable.

The National Library in the Netherlands[36] and the Royal Library of Belgium[37] are experimenting with applications from Microsoft (Power Automate and AI Builder) for cataloguing on the basis of scans of the title page, colophon and back cover, of printed books. The idea is to use the AI-generated MARC21 records as the basis for cataloguers, saving them time. The National Library in the Netherlands aims to use these techniques for cataloguing their retro collection (books from the past that are not yet in their collection).

There is already a large corpus of recent literature on AI and heritage collections[38]. Discussing this in detail is beyond the scope of this paper. However, the main take ways messages are: (1) metadata creation for multimodal heritage collections is supported by generative AI tools; (2) AI tools can also enrich metadata records of heritage collections by [39]:

- Transcribing text from images and audio
- Creating structured data from unstructured text
- Translating content into other languages and for different audiences
- Detecting objects in images; generate keywords, labels, descriptions
- Detecting and linking entities such as people, places, dates, concepts
- Clustering similar images, and texts
- Improving search by expanding input keywords and images.

CHATBOTS ANSWERING QUESTIONS ABOUT HERITAGE COLLECTIONS

Quinn[19] describes a project using the Retrieval Augmented Generative AI approach (see paragraph 2.2). The idea behind this is to ground LLM interactions in data controlled by the library. To this end, they combined several systems:

- Vector database with the data from the library
- LLM (such as ChatGPT)
- An application to process the user requests.

This works as follows: the user's question is used to retrieve a list of relevant documents from the vector database. These documents, along with the user's question, are fed into the LLM with the instruction to find the answer within the documents from the library. Their digital collection consists of heritage materials. The application returns the answer to the specific question generated by ChatGPT, along with the relevant documents.

A similar project describes an audio guide for a museum that answers questions using the curated information from the museum[40].

HERITAGE COLLECTIONS AS DATASETS

These developments around AI, combined with the increase in computational research in the Humanities and other disciplines, make it increasingly attractive for libraries to make their heritage collections available as data. An interesting example is the digitization program of the National Library of Scotland. They originally digitised collections in order to make them worldwide accessible by essentially replicating their physical readability in digital form. In recent years, they have also begun presenting their collections as datasets. This shift has led to a remarkable uptick in the usage of these datasets, with academics using them for research and teachers incorporating them into their coursework. Now, the library prioritises the digitisation of collections based of their potential as a dataset[41].

6.3 LINKED DATA, BIBFRAME AND AI

Within the academic library world, a large effort is underway to transition traditional metadata formats to linked data formats. Regarding bibliographic metadata, this transition is focused on gradually replacing the MARC21 format with BIBFRAME. Providers of library management systems increasingly support these efforts. We predict that innovators such as the Library of Congress, the National Libraries of Sweden and Finland, and a few other academic libraries such as the Stanford University library will complete their conversion to BIBRAME in the coming years, with the rest of the academic libraries following suit[42]. There is however some debate about whether this bibliographic transition is feasible[43] or might be rendered superfluous by AI[22]. The main arguments behind the transition to BIBFRAME are twofold:

A. Bibliographic metadata in a linked data form connects the library holdings with other data on the internet (web of data).
B. The creation of knowledge graphs (networks of connected linked data) enables a more precise discovery. For example, a knowledge graph makes it possible to search for all books by a certain author published after a certain year. Additionally, search engines can generate better recommendations to the user based on these knowledge graphs. Knowledge graphs also facilitate visualisations of the publication landscape or the connectedness between authors, among other things.

Does AI change any of these arguments? On the positive side, one could argue that AI-systems can read knowledge graphs and deliver more reliable output as a result; AI also can help to expand knowledge graphs. Knowledge graphs are also important for providing users with relevant recommendations. In other words, linked data provides an infrastructure that allows AI to work more effectively[44]. This works out positively for the arguments A and B. In addition, AI can facilitate the conversion of MARC21 records to BIBFRAME, which will make the whole transition somewhat easier.

On the flipside, a strong point of AI models is that they can combine various types of data, both structured and unstructured. Therefore, with AI, there may be less need to connect library metadata with other data on the internet (argument A). As shown in chapter 5, we do expect the discovery process to change rather drastically, which might make the more sophisticated navigation options supported by BIBFRAME – envisioned for the traditional search interfaces – somewhat less relevant (argument B).

On balance, we believe that the transition to BIBFRAME is unstoppable and will enrich the web of data with reliable library data, which in turn will make the outcomes of LLM models more reliable and trustworthy.

7 FOUR LIBRARY ISSUES TO BE FURTHER EXPLORED

At this moment, the main implications of generative AI for academic libraries are seen in the domains of discovery and metadata. However, the effects of AI are manifold. In this chapter, we touch on four other issues that will be affected by AI and will likely receive more attention from the library community in the near future.

7.1 ACCESS TO LIBRARY COLLECTIONS FOR COMMERCIAL AI MODEL

Academic libraries, through LIBER, have campaigned for a Text and Data Mining (TDM) exception in European copyright law so that scholars would be able to carry out these TDM activities in licensed collections[45]. But what about access to library collections for the collection of data by commercial LLM models? As we have seen, publishers are increasingly closing their content to AI models or charging fees for it[2,3]. What about access to collections owned by the libraries themselves? There is a statement by ICOLC[46] and guidance by JISC[47] that focus on content licensed from publishers. The ICOLC statement clearly states that AI should be permitted for non-commercial research, teaching, learning, and equitable access to information. Another point made is that AI clauses should not restrict user action in ways that are fundamentally unenforceable. This raises the question of whether a library can make its collections open for AI models only if these models are used for non-commercial research and education. The National Library of the Netherlands has answered this question in the affirmative: they have restricted access to their heritage collections for crawlers of commercial AI providers by changing the conditions of use and by implementing some technical measures. The main reason for this is that they believe that the copyright owners are not respected by the AI chatbots. The National Library of the Netherlands explicitly states that their collections remain open for research purposes[48].

7.2 INFORMATION LITERACY

Obviously, the courses given by academic libraries must include AI literacy[49], which encompasses the responsible and ethical usage of AI tools[49–51]. Many libraries already are taking this up[52]. However, with the anticipated paradigm shift in discovery (see 5.1) and the transformation of science by AI (see chapter 3), we expect that the information literacy courses will need a significant overhaul in the coming years in order to reflect these changes.

7.3 OPEN SCIENCE

There is surprisingly little literature on the effects of AI on Open Science. We found only a report by The Royal Society that addresses this issue in some depth[53]. They highlight the dominance of a few Big Tech companies in AI, including dominance in computer infrastructure, ownership of massive datasets needed for training AI models, and top AI talent. This centralisation of AI-based science development can limit a wider participation in steering the AI research agenda and restricting the number of decision makers involved. This dominance by the industry is also seen in AI research publications[54]. This does not bode well for Open Science: commercial incentives could restrict open science practices, and non-industry scientists might not be able to contribute equitably to AI research. Additionally, one of the goals of Open Science is to increase the reproducibility of research results. For sure, Open Science practices in research using AI models are the solution[55], but the proprietary nature of some AI models might hinder these practices.

7.4 LIBRARY ORGANISATIONS

As we have discussed above, specific tasks of academic libraries will change because of the effects of AI. However, other tasks that are similar to those in other organisations will also be affected. In the back office, one can foresee further automatization of workflows aided by the integration of AI tools in office automation software packages and AI tools for the management of smart buildings, potentially including robotics. In the front office, AI-based chatbots will support or replace front desk services, as they are expected to do in other organisations as well.

8 SUMMARY

In this paper we explored the potential effects of generative AI on academic libraries using recent literature and conference contributions. Table 2 on the next page provides a comprehensive overview of the effects of AI and the possible implications for the library services as discussed in this paper.

Our exploration started with a description of what a librarian should know about generative AI (chapters 1 and 2). The essence is that AI models do not understand a topic as humans do, but weave words together based on complicated statistical calculations. To do this, AI models use vast datasets collected from a range of sources including some with bad data. Therefore, their answers are not transparent and sometimes wrong ('hallucinations'). However, academic libraries can make the answers more transparent and reliable by forcing the commercial AI chatbots to use validated data for their answers with Retrieval Augmented Generative AI applications (RAG).

Next, we assessed the effects of generative AI on the environment of academic libraries: its impact on science and on the behaviour of library users (chapter 3 and 4). AI will transform science with the usage of AI models for complex systems, bringing together multiple disciplines. Also, the pace of scientific research is expected to quicken. As a result, we anticipate more scientific publications and a further atomization of the scientific record. The behaviour of library users will in our view change dramatically as the interaction with internet becomes a dialogue, while AI-powered personal assistants are predicted to interact with software applications as proxies for the end-user.

This changed interaction with the digital world is causing a paradigm shift in discovery (chapter 5). The transition to conversational discovery with AI chatbots answering user queries will lead to (much) less usage of library collections and potentially resulting in a new type of usage figures: usage by AI chatbots on behalf of the end-user.

In the next step of this explorative tour, we discussed the potential uses AI in metadata production and heritage collections (chapter 6). Here, there are promising opportunities to automate (parts of) the metadata production process, thereby making digital heritage collections more discoverable. Another application is the use of RAG-type applications to open up heritage collections for questions-and-answers interactions.

Finally, we touched upon four other issues relevant to academic libraries that have gotten so far little attention in the literature but surely will play a role in the coming years (chapter 7):

- The contentious question of allowing access by commercial AI chatbots to library collections
- The need to redevelop library courses on information literacy.
- The potential threats to Open Science practices in AI-driven science due to the dominance of a few Big Tech companies.
- The effects of AI tools for businesses on the library organisations.

AI effects	Implications for library services
AI effects on science and scientific methods	
Data hunger of scientific AI models	RDM services become more critical
AI transforming science: more multidisciplinary and computational research	Library collection and services need to reflect these changes
AI-assisted writing of science publications	Detection of plagiarism and fraud more challenging
Acceleration of science	Increase in the number of scientific publications/year
	Shrinking acceptable time lags between submission, publication, and availability in the library collection
	Further atomisation of the scientific record
Usage of commercial AI chatbots in the library	
Retrieval augmented generative AI (RAG)	Key model for library applications that makes it possible to enforce usage of reliable and validated library data while leveraging commercial AI chatbots
AI effects on discovery	
Conversational discovery	Changes in internet search engines and scientific discovery tools
Translation capacities of AI	Cross-language searching facilitated
	Layman access to scientific information facilitated
Summaries of search results	Reduced usage of library collections, especially regarding consolidated knowledge (eBooks, review articles)
	Potential for a new type of user figures for collection usages that reflect AI chatbots usage on behalf of users
	Tools that extract and cite relevant portions of text for incorporation into research papers
	Changing relationship between content sites and discovery tools using AI
AI effects on metadata and digital heritage collections	
Copy cataloguing using prompts	Reduction of cataloguing capacity within libraries
	No copyright claims possible for AI-generated metadata, with implications for aggregators of library metadata like WorldCat
Metadata production for digitized heritage collections	To a large extent carried out by AI systems, resulting in more heritage collections that are discoverable
Chatbots answering questions about heritage collections	New type of library services using the RAG model
Data hunger of scientific AI models and spread of computational sciences	'Collections as data' become an increasingly important library service
Transformation Marc21 to BIBFRAME	Likely to make AI more accurate and library data more used by AI models
Other effects on the library and library services	
Access to collections for commercial AI models	Point of contention and ongoing discussion
Teaching AI literacy	Large parts of current curricula regarding information literacy need to be redeveloped
AI-driven science	The dominance by Big Tech in AI models, computer infrastructure, ownership of massive datasets, and in top talent might negatively impact Open Science practices in AI-driven science
AI tools for the business market applied in library organisations	Automatization of workflows, managing smart buildings and AI chatbots taking over front desk services

Table 2 Overview AI effects and implications for library services

SOURCES

1.Kang, C., Metz, C. & Thompson, S. A. Four Takeaways on the Race to Amass Data for A.I. *The New York Times* (2024).

2.Roose, K. The Data That Powers A.I. Is Disappearing Fast. *The New York Times* (2024).

3.Wiley and Oxford University Press confirm AI partnerships as Cambridge University Press offers 'opt-in'. *The Bookseller* https://www.thebookseller.com/news/wiley-cambridge-university-press-and-oxford-university-press-confirm-ai-partnerships.

4.Gibney, E. Has your paper been used to train an AI model? Almost certainly. *Nature* (2024) doi:10.1038/d41586-024-02599-9.

5.Peel, M. The problem of 'model collapse': how a lack of human data limits AI progress. *Financial Times* (2024).

6.Metz, C., Kang, C., Frenkel, S., Thompson, S. A. & Grant, N. How Tech Giants Cut Corners to Harvest Data for A.I. *The New York Times* (2024).

7.*Generative AI: Implications for Trust and Governance*. https://aiverifyfoundation.sg/downloads/Discussion_Paper.pdf (2023).

8.Chui, M. *et al. Economic Potential of Generative AI*. https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier (2023).

9.*Position Paper Ten Behoeve van Het Rondetafelgesprek over de Impact van Intensief Gebruik van Digitale Technologieën En Middelen van de Tweede Kamercommissie Voor Digitale Zaken*. https://www.rathenau.nl/sites/default/files/2023-04/Rondetafelgesprek_Intensief_gebruik_digitale_technologie_en_middelen_Position_paper_Rathenau_Instituut.pdf (2023).

10.Tempelman, I., van der Wijst, J., Oosterwijk, N. & Kuijt, S. *Mensgerichte Artificiële Intelligentie - Een Hernieuwende Oproep Voor Zinvolle En Verantwoorde Toepassingen*. https://nlaic.com/bouwstenen/mensgerichte-ai/nederlandse-ai-coalitie-schrijft-manifest-voor-mensgerichte-artificiele-intelligentie-2023/ (2023).

11.Mikalef, P., Conboy, K., Lundström, J. E. & Popovič, A. Thinking responsibly about responsible AI and 'the dark side' of AI. *Eur. J. Inf. Syst.* **31**, 257–268 (2022).

12.AI investment forecast to approach $200 billion globally by 2025. https://www.goldmansachs.com/insights/articles/ai-investment-forecast-to-approach-200-billion-globally-by-2025.

13.AI is exhausting the power grid. Tech firms are seeking a miracle solution. *Washington Post* https://www.washingtonpost.com/business/2024/06/21/artificial-intelligence-nuclear-fusion-climate/ (2024).

14.Hansen, D. AI and copyright and TDM. *ai4Libraries* https://www.ai4libraries.org/2023recordings (2023).

15.Kelly, J. An introduction to copyright law and practice in education, and the concerns arising in the context of GenerativeAI. *Artificial intelligence* https://nationalcentreforai.jiscinvolve.org/wp/2024/03/11/copyright-and-concerns-arising-around-generative-ai/ (2024).

16.Autoriteit Persoonsgegevens; Directie Coördinatie Algoritmes (DCA). *Rapportage Algoritmerisico's Nederland | Autoriteit Persoonsgegevens*. https://www.autoriteitpersoonsgegevens.nl/documenten/rapportage-algoritmerisicos-nederland (2023).

17.Rapportage AI- & algoritmerisico's Nederland (RAN) - najaar 2023 | Autoriteit Persoonsgegevens. https://www.autoriteitpersoonsgegevens.nl/documenten/rapportage-ai-algoritmerisicos-nederland-ran-najaar-2023.

18.AI-wet | Shaping Europe's digital future. https://digital-strategy.ec.europa.eu/nl/policies/regulatory-framework-ai (2024).

19.Quinn, B. Starting the conversation: enriching collection access and use with Generative AI at Northwestern university libraries. in (2023).

20.Carpenter, T. A. A Look Under the Hood of Scopus AI: An Interview with Maxim Khan. *The Scholarly Kitchen* https://scholarlykitchen.sspnet.org/2024/07/25/interview-with-maxim-khan-about-scopus-ai/ (2024).

21.Huang, L. *et al.* A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. Preprint at http://arxiv.org/abs/2311.05232 (2023).

22.Breeding, M. AI and Library systems. https://www.youtube.com/watch?v=ZTgFQaDzc9w (2023).

23.The President's Council of Advisors on Science and Technology. *Report to the President: Supercharging Research: Harnessing Artificial Intelligence to Meet Global Challenges*. https://www.whitehouse.gov/wp-content/uploads/2024/04/AI-Report_Upload_29APRIL2024_SEND-2.pdf (2024).

24. How accelerated computing is powering the next Industrial Revolution. http://www.ft.com/partnercontent/nvidia/how-accelerated-computing-is-powering-the-next-industrial-revolution.html.

25. AI Image Statistics: How Much Content Was Created by AI. https://journal.everypixel.com/ai-image-statistics (2023).

26. Velleman, E., Regeer, D. & van der Graaf, M. *AI En Mensen Met Visuele Beperkingen - Een Verkennend Onderzoek Met Mogelijke Handelingsperspectieven*. (2024).

27. Chen, B. X. & Isaac, M. Meta's Smart Glasses Are Becoming Artificially Intelligent. We Took Them for a Spin. *The New York Times* (2024).

28. Rabbit r1 - order now. https://www.rabbit.tech/.

29. Gates, B. AI is about to completely change how you use computers. *gatesnotes.com* https://www.gatesnotes.com/AI-agents (2024).

30. Scopus AI Release: May 2024 | Elsevier Scopus Blog. https://blog.scopus.com/posts/scopus-ai-release-may-2024.

31. Ex Libris Launches a Beta Program of AI-Powered Library Research Assistants. *Ex Libris* https://exlibrisgroup.com/announcement/ex-libris-launches-a-beta-program-of-generative-ai-powered-primo-research-assistant/.

32. van der Graaf, M. *Toegangsvormen Tot Wetenschappelijke Publicaties Voor Onderzoeksinstellingen En Het Brede Publiek*. https://ukb.nl/wp-content/uploads/2024/01/Toegangsvormen-tot-wetenschappelijke-publicaties-voor-onderzoeksinstellingen-en-het-brede-publiek.pdf.

33. STM's New Dotcoms to Watch by Martijn Roelandse. https://www.morressier.com/o/event/655b8f316515150019d5f90b/sessions/6569c788444a81001272f01d?eventId=655b8f316515150019d5f90b.

34. Brzustowicz, R. From ChatGPT to CatGPT: The Implications of Artificial Intelligence on Library Cataloging. *Inf. Technol. Libr.* **42**, (2023).

35. Kohler, A., Allen, K., Phan, L. & Tucker, Z. Leveraging AI tools for automating metadata extraction. in (2023).

36. Buesink, M. KB Onderzoekskroniek: AI-gedreven ondersteuning voor en door catalogiseerders – IP | Vakblad voor informatieprofessionals. *IP* https://informatieprofessional.nl/kb-onderzoekskroniek-ai-gedreven-ondersteuning-voor-en-door-catalogiseerders/ (2024).

37. Lowagie, H. Demonstration on the automatic creation of metadata. https://www.linkedin.com/posts/hanneslowagie_powerautomate-aibuilder-activity-7113592369010069504-XbFO?utm_source=share&utm_medium=member_desktop (2024).

38. AI4LAM/awesome-ai4lam. AI for Libraries, Archives, and Museums (2024).

39. Ridge, M. Evolutionary Innovations: Collections as Data in the AI era. (2024) doi:10.5281/zenodo.10795641.

40. Carré, M. AI in Cultural Heritage: opportunities and challenges. in (2024).

41. Byrne, I. Beyond digitization: collections as dataset at the National Library of Scotland. in (2024).

42. van der Graaf, M. *Etude sur les implications pratiques de la transition bibliographique pour les bibliothèques de l'ESR*. https://abes.fr/wp-content/uploads/2023/08/Les-implications-pratiques-de-la-Transition-bibliographique-pour-les-bibliotheques-dESR-1.pdf (2023).

43. Edmunds, J. BIBFRAME Must Die. (2023).

44. Marciano, A. Linked library ecosystem powered by AI. in (2023).

45. LIBER. A copyright exception for Text and Data Mining. (2020).

46. ICOLC Statement on AI in Licensing | ICOLC Website. https://icolc.net/statements/icolc-statement-ai-licensing.

47. Attewell, S. Guidance on resisting restrictive AI clauses in licences. *Artificial intelligence* https://nationalcentreforai.jiscinvolve.org/wp/2024/06/19/guidance-on-resisting-restrictive-ai-clauses-in-licences/ (2024).

48. Kleppe, M. Statement over commerciële generatieve AI. https://www.kb.nl/ai-statement.

49. Hoh, A. Do we need to develop AI literacy or is it part of digital literacy?! *Maastricht University Library* https://library.maastrichtuniversity.nl/news/do-we-need-to-develop-ai-literacy-or-is-it-part-of-digital-literacy/ (2024).

50. Bali, M. Where are the crescents in AI? *LSE Higher Education* https://blogs.lse.ac.uk/highereducation/2024/02/26/where-are-the-crescents-in-ai/ (2024).

51. Shaw, B. Research Guides: Artificial Intelligence (AI) and Information Literacy: Home. https://lib.guides.umd.edu/c.php?g=1340355&p=9880565.

52. Developing a library strategic response to Artificial Intelligence. *IFLA* https://www.ifla.org/g/ai/developing-a-library-strategic-response-to-artificial-intelligence/ (2024).

53. *Science in the Age of AI | Royal Society*. https://royalsociety.org/news-resources/projects/science-in-the-age-of-ai/ (2024).

54. Ahmed, N., Wahed, M. & Thompson, N. C. The growing influence of industry in AI research. *Science* **379**, 884–886 (2023).

55. Gundersen, O. E., Gil, Y. & Aha, D. W. On Reproducible AI: Towards Reproducible Research, Open Science, and Digital Scholarship in AI Publications. *AI Mag.* **39**, 56–68 (2018).

56. van der Graaf, M. *Effecten van AI Op Digitale Inclusie En Digitaal Burgerschap*. https://digitaalsamenleven.nl/app/uploads/2023/12/Trendrapport_Effecten-van-AI-op-digitale-inclusie-en-digitaal-burgerschap-1.pdf (2023).

ABOUT THE AUTHOR AND THIS WHITE PAPER

A SHORT CV OF MAURITS

Maurits van der Graaf started Pleiade Management and Consultancy in 2000 with the aim to specialize in libraries, archives, cultural institutions, and academic publishers. Over the years, he has led numerous projects in the Netherlands for university libraries (such as the University libraries of Groningen, Amsterdam, Eindhoven and Leiden) and for the national consortium UKB and the national association of universities (UNL), Maurits also worked on various projects for international organisations (LIBER, ISSN, DRIVER, IFLA, Knowledge Exchange, Publishing Research Consortium, SpringerNature and Elsevier) and for institutions in France (ABES, ADBU, Science Po), in Switzerland (RERO, Renouvaud, Swiss National Library) and Germany (ZB Med). He is author of over 75 public reports and articles in the information domain. Before Pleiade, Maurits held various management positions in publishing and in library and documentation institutes. The positions he has occupied include those of Product Manager EMBASE, Director of the Dutch Agency for Current Research Information and Deputy Director of the Netherlands Institute of Scientific Information Services.

A PERSONAL WORD ABOUT THIS WHITE PAPER

In the second half of 2023, I carried out a study for the *Dutch Alliantie Digitaal Samenleven* with the aimed at writing a Trend Report on AI and digital inclusion. Part of that study involved trying to understand what generative AI exactly was. I interviewed many experts and read numerous reports. In that period, the hype of generative AI was already in full swing, and it was difficult to keep track of all the reports that were published with AI as a topic. With our Dutch-language report[56], we added to the growing pile of these reports. But the potentials of AI had piqued my interest.

Beginning of 2024, Eric Velleman asked me to build upon this report with a study on AI and visually impairments for the Bartimeus Foundation. We held group discussions and collected publications about this issue. This has resulted in another Dutch-language report[26]. And I was even more intrigued by what AI could do.

This spring, I broke my ankle due to a misstep when walking in the dunes. Consequently, summer holidays had to be cancelled and I was left with a lot of time on my hand but confined to home. So, I decided to use this time to explore the possible effects of AI on academic libraries and the result of that exercise lies before you.