

Putting the AI into social science

How artificial intelligence tools are changing and challenging research in the social sciences

Johannes Breuer

1. Introduction

The recent rapid announcements, developments and releases in the realm of artificial intelligence (AI), especially within the domain of large language models (LLMs), have not only received a lot of public attention, but also sparked a surge of discussion, research and other activities among the scientific community, including the social sciences. Similar to other digital technologies, such as the internet (cp. Breuer 2022), AI has multiple relationships with science. It is a) an outcome or product of scientific research, b) an object of study across many different disciplines, and c) a powerful tool that affects how research is done. This chapter focuses on the third function and discusses how AI tools have been changing how social science research is conducted and what the future may hold in this regard. The discussion within this chapter will address both the potentials as well as the challenges and risks associated with the use of AI (and tools based thereon) in the social sciences.

Notably, AI can have – and already, in many cases, has – an impact on all elements of social science research. There are different ways in which the (typical) research process in the social sciences (and similar disciplines) can be structured. Common phases can, e.g., be structured as follows: 1) idea generation (e.g., formulation of research questions or hypotheses), 2) discovery (e.g., searching for and exploring existing literature, data, analysis methods, etc.), 3) study design and planning (e.g., deciding what methodology and sample to use), 4) data collection (e.g., via surveys, interviews, web scraping), 5) data processing (e.g., cleaning the data, getting it ready for analysis), 6) data anal-

ysis, 7) interpreting results, 8) reporting, publishing, and sharing.¹ Of course, in practice, these phases are often overlapping or not clearly distinguishable and do not necessarily occur in this order (and there may also be recursions). For example, in the case of an exploratory study, researchers might discover something in the data analysis phase that leads them to collect additional data, come up with new research questions, or reconsider their analysis methods. While AI can affect all of these phases, the degree to which this is the case and the ways in which this influence manifests itself differ between the individual steps. After clarifying a few important preliminaries that need to be kept in mind when dealing with the use AI in the social sciences at this time, this chapter will discuss how AI and AI-based tools have been or can be used in the various phases of social science research and the promises and potentials as well as the pitfalls and perils associated with these practices.

2. Preliminaries

Before discussing the practices, potentials, promises, pitfalls and perils of the use of AI in the social sciences, it is necessary to lay out a couple of important preliminaries. The first one relates to the terminology used in this chapter. Similar to the term big data, artificial intelligence has different definitions and, hence, can be a somewhat fuzzy concept. Oftentimes, AI is used interchangeably with machine learning (ML), or at least the distinction becomes blurry. However, as Kühl et al. 2022 point out: “ML’ and ‘AI’ are not terms that should be used interchangeably (...) ML is an important driver of AI, and the majority of modern AI cases will utilize ML. However, (...) there can be cases of AI without ML (e.g., based on rules or formulas)” (2241). Another important distinction for this chapter as well as the collected volume which it is part of, is the one between symbolic and subsymbolic AI. According to Ilkou and Koutraki (2020), the key differences between these two types of AI are the following: “(1) symbolic approaches produce logical conclusions, whereas sub-symbolic approaches provide associative results. (2) The human intervention is common in the symbolic methods, while the sub-symbolic learn and adapt to the given data. (3) The symbolic methods perform best when dealing with relatively small and precise data, while the sub-symbolic ones are able to handle large and

1 Notably, these phases as outlined here are quite generic. Most of them are, hence, also valid for other empirical disciplines (e.g., from the medical and natural sciences).

noisy datasets” (1). The currently predominant type of AI – and also the focus of this book – is subsymbolic AI.² According to Ilkou and Koutraki (2020), “sub-symbolic AI includes statistical learning methods, such as Bayesian learning, deep learning, backpropagation, and genetic algorithms” (2). As these methods, especially also deep learning, are often discussed as belonging to the area of machine learning, it becomes apparent how terminological ambiguities between AI and ML may arise in an applied context. Taking this into account, the chapter will not discuss to what degree different techniques and tools are best described as AI or ML or to what degree AI applications can be classified as symbolic or subsymbolic. What is more important for the present chapter is that, outside of computer science or other fields involved in the development of LLMs and other types of AI, the use of or interaction with AI occurs via tools. While – under the hood – these tools often make use or offer access to methods that can be seen as belonging to the area of ML, the tools are often labelled or described as AI-based. Although this may not always be (fully) appropriate and often done (primarily) for marketing reasons, for the purpose of this chapter, if they are labelled/presented as AI tools, they will also be discussed as such here.

The application area of AI tools that this chapter focuses on are the social sciences. Core disciplines in this field include sociology, political science, or communication science.³ As stated in the introduction, however, many of the prototypical phases in social-scientific research are also common in other fields. Likewise, many of the methods and tools discussed in this chapter are also used there. Regardless of the definition of the category of social sciences, the focus of this chapter is on empirical research. More specifically, while many of the methods and tools covered in the following can also be used for quali-

2 Ilkou and Koutraki (2020), however, note that in-between methods that combine symbolic and subsymbolic AI have become more common. Among other things, the rise of the concept of explainable AI has contributed to the resurgence of symbolic methods, which were the dominant approach until the 1980s.

3 There are, of course, also other disciplines that can be classified as social sciences as well as different ways of classifying disciplines. Besides, there are some disciplines for which there are different views on whether they can be seen as belonging to the social sciences, such as psychology or economics. As much of the tasks and topics covered in this chapter should also be relevant beyond the social sciences, these differences in the definition of social sciences and the classification of disciplines should not matter.

tative research, an emphasis will be on quantitative empirical research in the social sciences.⁴

A final important thing to consider for this chapter is that the field of artificial intelligence is currently developing rapidly following the release of powerful large language models (LLMs) and their quickly increasing use for all sorts of applications, including scientific research. Especially since the release of *ChatGPT* by *OpenAI* in November 2022, the development of AI applications has gained a lot of momentum. While the development and release of LLMs and tools based thereon had already been quite fast-paced before, this has been massively sped-up in the first half of 2023, with new models and tools being released daily. Accordingly, it is almost impossible to keep up with all developments. Although the timeframes of academic research (especially if it is empirical) are not fully compatible with the speed of current technology developments, the academic community has been trying to keep up by conducting timely studies and publishing them in the form of preprints. Notably, these publications are not peer-reviewed. Still, given their timeliness and relevance, such preprints will be considered in this chapter. Against this background, it should be noted that the methods and tools, as well as the scientific publications investigating their use are likely to become updated and amended or outdated, invalidated, or even replaced in the near future. Consequently, this chapter can only provide a snapshot from the time of writing (April to June 2023), and the practices of using AI and AI-based tools and methods as well as the associated potentials, promises, pitfalls, and perils can be expected to change substantially over the course of the upcoming months and years. Another thing to note is that this chapter is certainly not the only and also not the first discussion of how AI is changing scientific research. Besides the project “How is Artificial Intelligence Changing Science? Research in the Era of Learning Algorithms”, (<https://howisaichangingscience.eu/>) from which the book, that this chapter is part of, originated, there are at least two other noteworthy recent publications in this context. The first one is the preprint “Friend or Foe? Exploring the Implications of Large Language Models on the Science System” by Fecher et al. (2023), in which the authors present the results of “a Delphi

4 This is partly due to the background of the author but also because the use of ML and AI for data collection, processing, and analysis is more common in the quantitative paradigm. In fact, the use of ML and AI methods is one of the defining criteria of the rapidly growing field of computational social science (cp. Hox 2017).

study involving 72 experts specialising in research and AI”, in which the author of the present chapter also participated. Based on the expert opinions, the manuscript discusses the applications and (transformative) potential, as well as limitations, risks and ethical and legal implications of the use of LLMs in science. The second relevant recent publication is a preprint by Ziems et al. (2023) entitled “Can Large Language Models Transform Computational Social Science?”, that presents the results of evaluations of different LLMs for various typical tasks in computational social science (CSS). The present chapter is essentially situated between these two publications. While, similar to Fecher et al. (2023), it also addresses the applications of AI (tools), taking into account their potential as well as limitations and associated risks, it focuses on the social sciences, specifically on empirical research in this field which follows a specific process from idea generation to publication. Hence, compared to the work by Ziems et al. (2023), the perspective of this chapter is broader, considering not only typical CSS applications, such as automated text classification or other annotation and explanation tasks, but also addressing usage in the context of traditional data collection methods, such as surveys or experiments as well as more general practices, e.g., in phases of discovery and data analysis.

3. Practices

Scientific research has always been based on the use of tools. These tools can either be specifically designed for scientific purposes, such as a microscope or telescope, or designed for other or more general purposes and used by scientists for their research, such as tweezers or a shovel. This is the same for AI(-based) tools. Another important distinction from a practical perspective is whether tools are commercial or free and maybe even open source.⁵

Importantly, tools are not neutral. They shape the research process, define possibilities and boundaries. The concept of Maslow’s hammer describes this in a pointed manner: “If the only tool you have is a hammer, it is tempting to treat everything as if it were a nail.” (Maslow 1966: x). Of course, scientists typically do not just use a single tool, but a combination of different tools (for different purposes). Especially in the digital realm, these combinations are often referred to as tool stacks. Ideally, the tools within individual tool stacks are

5 Of course, these characteristics can change over time. E.g., tools that are initially free to use may eventually require a paid subscription.

used for one or multiple specific task(s) with little or no redundancies, compatible, and complement each other. A couple of years ago, the project *Innovations in Scholarly Communication* situated at the University of Utrecht distinguished between traditional, modern, innovative, and experimental tools (cp. Bosman/Kramer 2015).⁶ Following this distinction, most of the AI tools mentioned in the following can be classified as innovative or experimental. In their analysis, Bosman and Kramer (2015) diagnose “an avalanche of tools” and describe choosing appropriate tools and keeping up with the development of (new) tools as a challenge for researchers. This issue is even more pronounced in the current explosion of the development of AI and its applications, with new tools or versions thereof being released almost daily.

Generally, tools and tool stacks enable scientists to conduct research in the first place or at least facilitate the process and make it more efficient. Besides these potentials, however, tools and tool stacks also bring their own challenges and limitations. While the use of tool stacks widens the possibilities and space for research, they also have or create specific boundaries. In addition, the reliance on tool stacks creates dependencies. Scientists depend on them for conducting their research and tools may also depend on each other to work properly within a given tool stack. These dependencies can break if the functionalities or the availability of tools change. This illustrates that the impact of AI on research in the social sciences is not limited to the quantitative dimension. While it does, e.g., facilitate the handling of large(r) amounts of data, by altering the range of possibilities, it also affects the qualitative aspects of social-scientific research.

As noted before, these changes in the quantitative and qualitative properties of social science research run through all phases of the research process. However, the number and type of AI tools that are used and the impact they have had on research practices so far, differs between each individual phase. Given the rapid development of AI and AI-based tools, the purpose of this section is not to provide a complete list of all tools that have been or can be used for the different steps in the social science research process. Instead, the aim of this section is to provide a couple of examples of how AI has been used in the social sciences to demonstrate its qualitative impact on the tasks typically

6 Regarding the different phases of the research process for which the tools can be used, the categories by Bosman and Kramer (2015) are similar to the ones suggested in the present chapter: discovery, analysis, writing, publication, outreach, and assessment.

undertaken within the different phases.⁷ In the following, these impacts will be discussed for each of the eight (proto)typical phases listed above. Importantly, many of the available AI-based methods and tools cannot be exclusively mapped to one phase. While some methods and tools have been designed for very specific tasks, others have a broad(er) range of possible applications in the social sciences.

3.1 Discovery & idea generation

While they may be separated for analytical purposes, in practice, the phases of idea generation and discovery are usually intertwined. Formulating meaningful research questions and hypotheses requires a certain familiarity with existing literature, methods and data. This knowledge is necessary at the latest for the specification of the research questions and/or hypotheses. A large number of AI-assisted tools that have come into existence over the last few years target the discovery phase. Examples include *Semantic Scholar* (<https://www.semanticscholar.org/>), *scite* (<https://scite.ai/>), *ResearchRabbit* (<https://www.researchrabbit.ai/>), *Consensus* (<https://consensus.app/>), or *elicit* (<https://elicit.org/>).⁸ The focus of all these tools lies on discovering and exploring relevant literature. All of them allow to assess (and visualize) relationships between publications (via citations or similarity) and some offer additional functionalities. For example, *scite* can provide information on how often a paper has been supported, contradicted, or just mentioned in a citing publication, *Consensus* delivers additional information about journals and publications as well as relevant quota-

7 There are many lists and discussions of AI tools as well as short recommendations and tutorials on using AI-based tools for scientific research tasks available online. Large parts of this discourse have been happening on Twitter (although there, e.g., also are websites and YouTube videos that cover these topics). Two accounts on Twitter that have produced a large amount of content on this subject are Mushtaq Bilal (<https://twitter.com/MushtaqBilalPhD>) and Ilya Shabanov (<https://twitter.com/Artifexx>). There are also thousands of accounts that specialize in covering news on AI developments, tools, and research in general, many of which have only been created or shifted their topical focus and started to receive increased attention (and a quickly growing follower base) fairly recently.

8 Of course, there are many other services and apps for the discovery and exploration of scientific publications, such as *Google Scholar* (<https://scholar.google.com/>), *Researcher* (<https://www.researcher-app.com/>), *Inciteful* (<https://inciteful.xyz/>), *Litmaps* (<https://www.litmaps.com/>), or *Connected Papers* (<https://www.connectedpapers.com/>). However, those do not explicitly state or advertise that they employ AI-based methods.

tions from the latter, and *elicit* also offers versatile discovery functionalities via the creation of bespoke tasks. To identify relevant literature from a large corpus for a systematic literature review with the help of AI/ML methods, researchers can also use the free-and-open-source (FOSS) tool *ASReview* (<https://asreview.nl/>).

Once the relevant publications have been identified, the next task is for the researcher to read them and extract relevant information for their own research. There also are AI-based tools that can assist with that. Besides the functionalities by *Consensus* and *scite* described above, there are services like the article summarizer by *scholarcy* (<https://article-summarizer.scholarcy.com/>), *Explainpaper* (<https://www.explainpaper.com/>), or *ChatPDF* (<https://www.chatpdf.com>) that can aid with extracting information from scientific publications. While all of the other tools and services listed before were specifically created for research purposes, this is not the case for *ChatPDF*. As the name indicates, *ChatPDF* is based on *ChatGPT* by *OpenAI*, and the latter has also become a popular multi-purpose tool for research(ers) in the social sciences. Among other things, researchers have also suggested using *ChatGPT* for the idea generation phase (Dowling/Lucey 2023).

3.2 Study design & data collection

There also are several AI-based tools and methods that can be used for the study design data collection phases. Two of the most widely used data collection methods in the quantitative social sciences and related fields are surveys and experiments (which can also be combined in the form of survey experiments; cf. Mutz 2011). Surveys contain a number of questions or items that are designed to assess certain attributes, attitudes, or behaviors. Researchers often use existing items and scales that, ideally, have been validated before. However, these may not always be available, or existing scales may have to be modified. (Re-)Formulating, and refining survey items is one of the many possible uses of LLMs like *ChatGPT* or *GPT-4* by *OpenAI* and interfaces to those, such as *Microsoft Bing Chat*, in social science research. Through proper prompts, researchers could, e.g., ask LLM-based chatbots to come up with suggestions for novel questionnaire items tapping into specific concepts or optimize the wording of existing questions/question drafts. What is helpful in this regard as well as for all other research-focused uses of general-purpose chatbots like *ChatGPT* is the use of so-called priming, which describes the process of interacting with the LLM to provide some context and ensure that it understands the tasks

before prompting it to get the targeted output, such as (reformulated) question items.

Another common task in survey-based research is the translation of existing items into other languages. This can also be done or supported through LLMs or with the help of AI-based translation tools, such as *DeepL* (<https://www.deepl.com/translator>) or *Microsoft Bing Translator* (<https://www.bing.com/translator>). Research from the area of psychometrics and survey methodology has already investigated the potentials and limitations of such uses (see, for example, Behr 2023 or Kunst/Bierwaczzonek 2023).

A recent methodological innovation in the area of survey research is the use of chatbots for so-called conversational surveys, “where a chatbot asks open-ended questions, interprets a user’s free-text responses, and probes answers whenever needed” (Xiao et al. 2020: 1). The use of chatbots for such conversational surveys has the potential to increase participant engagement as well as response quality (cp. Xiao et al. 2020). Of course, while the method of conversational surveys falls into the category of quantitative social science research, chatbots could also be used for qualitative research, e.g., in interview studies.

Experimental research in the social sciences typically makes use of different kinds of stimulus materials serving as experimental treatments. These can be textual (e.g., in so-called vignettes), visual, or a combination thereof.⁹ LLMs can also be used to create textual stimuli for experimental research in the social sciences. Likewise, text-to-image tools, such as *Midjourney* (<https://www.midjourney.com>), *Stable Diffusion* (<https://stablediffusionweb.com/>), *Microsoft Bing Image Creator* (<https://www.bing.com/create>), or *Lexica Aperture* (<https://lexica.art/aperture>) can be used to create visual stimulus material for experimental studies.

Another area within the study planning and data collection phases where AI tools are helpful for social science research is simulation. Work by Argyle et al. (2023) suggests that LLMs “can be studied as effective proxies for specific human sub-populations in social science research” (2) and allow the simulation of responses to closed survey items (scales) as well as open-ended questions (free-form text responses). A similar approach was followed in a recent study by Chu et al. (2023) in which the authors trained a language model on media diets and

9 Many experimental studies also use audio or video stimuli. However, as AI tools for creating those based on text input are not yet so far developed, those will be covered in the following sections on potentials and promises and pitfalls and perils related to the use of AI in the social sciences.

found that it can be used to predict public opinion. A simpler and more playful but still interesting application of simulating responses is the website *GPTrolley* (<https://www.gptrolley.com/>) which uses *ChatGPT* to respond to user-generated versions of the ethical dilemma of the trolley problem that is often used in social-scientific, especially psychological, research.¹⁰

3.3 Data processing & analysis

For the processing and analysis of data in the social sciences, the writing of code has become increasingly common. While the use of commercial statistical software, such as *SPSS* or *Stata*, is still widespread, programming languages like *R* or *Python* are being used by a steadily increasing number of social scientists. In addition, even when commercial solutions are used, the exclusive reliance on graphical user interfaces (GUIs) and ‘point-and-click’ pipelines has become much rarer, which contributes to increasing reproducibility and transparency according to the principles of open science. To facilitate the writing, testing, and optimization of code for different programming languages, there are several dedicated AI-based tools available that are also of interest for social scientists, such as *GitHub Copilot* (<https://github.com/features/copilot>) or *replit Ghostwriter* (<https://replit.com/site/ghostwriter>). Notably, general-purpose LLM tools, such as *ChatGPT* can also be used for generating computer code via natural-language prompts. In addition, researchers can make use of these models to adapt or optimize existing code or to translate between programming languages.

With the rise of computational social science, it has become increasingly common for social scientists to work with large amounts of text data. Most of the methods used for processing and analyzing such data belong to the category of natural language processing (NLP) or ML, and the boundary to AI can become blurred here (e.g., if deep learning is used). Typical tasks in the processing and analysis of (large) textual data are annotation and classification. Recent research has demonstrated that LLMs like *ChatGPT* can, e.g., be used for identifying hate speech (cp. Huang/Kwak/An 2023), detecting psychological constructs, such as sentiment, emotions, and offensiveness in multilingual text corpora (Rathje et al. 2023), and may even outperform human crowd-

¹⁰ Notably, the reasoning abilities of LLM have also inspired research on questions like whose opinions LLMs reflect (Santurkar et al. 2023) or how to assess psychological profiles of LLMs (Pellert et al., 2022).

workers (cp. Gilardi/Alizadeh/Kubli 2023). However, another study indicates that “ChatGPT’s classification output can fall short of scientific thresholds for reliability” (Reiss 2023: 1). Likewise, Pangakis, Wolken and Fasching (2023) note that “Automated Annotation with Generative AI Requires Validation”. Besides analyzing text from online sources, LLMs, such as BERT, have also been used for classifying open-ended survey responses (Gweon/Schonlau 2023).

Some research in the social sciences makes use of audio data (e.g., from interviews). For the automatic transcription of audio files, a powerful speech-to-text model is *Whisper* by *OpenAI* (see <https://openai.com/research/whisper>), for which implementations exist for the programming languages like Python (<https://github.com/openai/whisper>) and R (<https://github.com/bnosac/audio.whisper>), which are popular in the social sciences. Once the audio data has been transformed to text, the methods and tools described previously for textual data can be applied.

3.4 Writing & dissemination

For writing tasks, researchers in the social sciences and other disciplines can make use of the options described for the formulation and translation of survey items in section 3.2 as well as other general AI-assisted writing support tools, such as *Microsoft Editor* (<https://s.unhb.de/mseditor>), *Grammarly* (<https://www.grammarly.com/>), or ones specifically designed for academic writing, such as *jenni* (<https://jenni.ai/>) or *Paperpal* (<https://paperpal.com/>). These tools can be used for all sorts of writing tasks, including generating text, editing, summarizing, paraphrasing, and translation.

AI tools can also be useful when it comes to sharing research data. As data in the social sciences is usually personal and can also be sensitive, different approaches have been developed in order to create a balance between openness on the one side and data privacy on the other. One solution is the creation of synthetic data that has comparable properties with the original data. So far, the creation of synthetic data sets (e.g., using the *synthpop* package for R; Nowok/Raab/Dibben 2016) has largely been limited to numeric data. Approaches as the ones described in the paper by Argyle et al. (2023), however, also allows for the creation of synthetic text responses to open-ended questions using LLMs.

4. Potentials & promises

As the examples in the previous sections illustrated, AI generally has the potential to facilitate and improve research in the social sciences and make the lives of researchers easier. It can increase the efficiency of research and, thus, also lead to an increase in output (publications, data, code and software, as well as other resources).¹¹ Especially AI-based tools for writing (both text and code) and no-code data collection and analysis solutions can also be beneficial for inclusivity, e.g., with regard to non-native English speakers or researchers with limited or no programming skills.

The use of AI can also reduce costs and the risk of human errors, e.g., for annotation and classification tasks (cp. Gilardi/Alizadeh/Kubli 2023). AI tools can further add to the reliability and validity of research results in the social sciences if it is used to enhance methods like multiverse analysis in which the robustness of results is assessed by systematically varying sets of processing and analysis parameters (for an example, see Pipal/Song/Boomgaarden 2022).

There are a few new developments and application areas that can be expected to become (more) interesting for the social sciences in the near future. One key area is the use of AI for images, audio and video. While text is still the much more dominant type of data in the social sciences, there is an increasing body of (computational) research that makes use of (large amounts of) image (cf. Webb Williams/Casas/Wilkerson 2020) and also video data (see Dietrich 2020 or Jürgens/Meltzer/Scharkow 2022 for exemplary applications). Besides the use of AI for detection/recognition and classification tasks for text, images, audio, and video, another relevant task for social science research is the generation of these types of content, e.g., as stimuli for experimental studies. While, as stated before, powerful models and tools already exist for generating text and images, options for generating audio (text-to-speech) or video (text-to-video) are not yet as widely available, although this can be expected to change in the near future.

11 As researchers are often already struggling to follow, filter, and digest the huge amounts of information on findings, methods, tools, etc. this increase may be seen as a mixed blessing. In a somewhat circular fashion, the increase in output may require researchers to also rely more on AI-based tools for making sense of the increased output by filtering and summarizing relevant content.

5. Pitfalls & perils

As with all innovations and transformations, in science and beyond, the use of LLMs for research in the social sciences does not only create new possibilities but also brings along challenges that need to be taken into account. Mirroring the potentials and promises, there are numerous pitfalls and perils associated with the (increasing) use of AI. In practice, this means that there are different practical, legal and ethical questions that social scientists need to be aware of and be able to address.

Key legal questions relate to privacy, copyright, together with terms of service (ToS) and other contractual agreements. Especially when using online services or application programming interfaces (APIs), it is often not fully clear where and how user inputs are stored and, depending on the type of input and the storage and processing pipeline, this may not be compatible with data protection regulations, such as the General Data Protection Regulation (GDPR) in Europe. On the other hand, platform or API ToS may also restrict the usage of outputs. Both of these issues can be(come) particularly problematic when working with research data which contains responses from study participants. Another legal domain where uncertainty exists, is that of copyright. While its application for academic research is typically treated differently than the one of commercial use, this issue becomes particularly salient when it comes to sharing research materials (e.g., experimental stimuli) in the spirit of open science. A related question is that of recognition of contributions and authorship when AI tools, such as *ChatGPT* have been used to generate text for publications.¹²

A general risk associated with the use of AI (tools) is the reliance on commercial companies and products, such as the services and APIs offered by *OpenAI*. The services, their ToS, or the underlying business model and pricing may change. The recent history of CSS research using social media data can serve as a good example for the risk of relying on APIs offered by private companies (cp. Bruns 2019; Freelon 2018). What also comes with the reliance on commercial services is the problem of intransparency, as transparency is usually not that compatible with competition and for-profit orientation. For that reason, the use and support of free and open-source (FOSS) projects in the area of LLMs, such as *Open Assistant* by *LAION* (<https://github.com/LAION-AI/Open-Assista>)

12 Via their blog, the American Psychological Association (APA), whose publication guidelines are widely used in the social sciences, has already put forth suggestions on how to cite ChatGPT (cp. McAdoo 2023).

nt), *HuggingChat* by *Hugging Face* (<https://huggingface.co/chat/>) or *GPT4All* by *Nomic AI* (<https://gpt4all.io/index.html>) becomes particularly important from the perspective of academic research(ers) in the social sciences as well as other disciplines. Regardless of the underlying governance or business model, however, a general issue leading to a lack of transparency is the black-box character of most subsymbolic AI models (Sudmann 2019; 2020). Together with the fact that they rely on stochastic processes, this can be detrimental to the aim of ensuring that social science research is reproducible and replicable.

A challenge that has been widely discussed is the introduction and proliferation of biases in LLMs and other AI applications. Although AI tools can be employed to counter human errors and biases, e.g., in the processing and analysis of data, they can create new and less directly transparent forms of bias, often introduced through training data (cp. Ferrara 2023). The (over-)reliance on AI tools might lead to 'bias cascades', as research has shown that biased AI systems can produce or increase bias in human decisions (cp. Glickman/Sharot 2022). There are, however, strategies for identifying and mitigating biases in AI, and the biases can also be made use of productively for social science research as the research by Argyle et al. (2022) and their concept of 'algorithmic fidelity' for simulating responses from specific subpopulations shows.

Another important topic is the question of trust. Different LLMs have been repeatedly shown to make up things (a process often referred to as hallucinating) and, thus, producing misinformation. Combined with the transparency issue(s) discussed above, this can lead to AI-assisted research potentially becoming less instead of more trustworthy. Related to this, there is concern in the academic community that the use of AI tools can lead to a reduced quality of peer review as well as increase in fake or junk papers, academic spam and scams, and predatory journals and conferences. This can also be seen as the flipside of an AI-fueled increase in efficiency and research output.

Finally, there are the broader societal implications of using AI tools which researchers also need to take into account, such as the risk of creating or supporting (quasi-)monopolies or oligopolies, (indirectly) supporting exploitative working conditions, e.g., for the creation of training data (cp. Perrigo 2023) and the energy consumption and environmental effects of training and maintaining LLMs and other forms of AI.

6. Conclusion

The use of AI tools and methods has already begun to transform research practices in the social sciences and will continue to do so. These changes affect all phases of the typical research process, however, not all phases are affected to the same degree. As the examples in this chapter have shown, there are a lot of AI tools that can be used in the discovery phase and quite a few that are useful for data collection, processing, and analysis. The formulation of meaningful research questions and hypotheses and the interpretation of results, by contrast, are tasks that AI tools are less suited for and require human expertise.

There is an internet idiom that goes “go away or I will replace you with a simple shell script” (see <https://s.unhb.de/shellreplace>). These days, the shell script might be replaced with an LLM. It is, however, highly unlikely that an LLM (or another form of AI) can replace human social scientists anytime soon. For now, the AI of our times seems to agree. When I asked *ChatGPT*, “Is it possible that there will be AI social scientists in the future?” it replied that “it is unlikely that AI systems will be able to completely replace human social scientists. Social science research involves a wide range of qualitative research methods, such as participant observation, interviews and case studies, that require human interpretation and understanding of social context, historical factors, and the nuances of human interactions.” (OpenAI 2023). Maybe it just wants to lull us into a false sense of security, but I agree with the assessment by *ChatGPT* as well as the conclusion drawn by Ziems et al. (2023) that “LLMs can significantly reduce costs and increase efficiency of social science analysis in partnership with humans” (1), with the emphasis being on the phrase “in partnership” here.

Nevertheless, besides making use of its potential, social scientists also need to be aware of and able to deal with the risks and challenges associated with the use of AI for their research. While they may not be replaced by AI, they certainly need to adapt to using it in a productive and ethical way, e.g., by developing new skills, such as AI literacy, or knowing how to write and optimize prompts to achieve desired results.¹³ If this is achieved, AI can support social scientists and AI tools can serve as valuable additions to established methods which can, ultimately, contribute to improving the quality of social science research.

13 With the explosion of LLMs, prompt engineering has become a relevant topic, and many resources have been created with the aim of teaching users how to write optimal prompts (<https://learnprompting.org/>) or to provide examples of useful prompts (e.g., <https://flowgpt.com/>).

List of references

- Argyle, Lisa P./Busby, Ethan C./Fulda, Nancy/Gubler, Joshua R./Rytting, Christopher/Wingate, David (2023): "Out of One, Many: Using Language Models to Simulate Human Samples." In: *Political Analysis* 31/3, pp. 337–351.
- Behr, Dorothee (2023): What to Consider and Look out for in Questionnaire Translation, GESIS Survey Guidelines, Mannheim: GESIS – Leibniz Institute for the Social Sciences (https://doi.org/10.15465/GESIS-SG_EN_043).
- Bosman, Jeroen/Kramer, Bianca (2015): "101 Innovations in Scholarly Communication: How Researchers Are Getting to Grip with the Myriad New Tools.", November 11, 2015 (<https://blogs.lse.ac.uk/impactofsocialsciences/2015/11/11/101-innovations-in-scholarly-communication/>).
- Breuer, Johannes (2022): "Www – Die Funktionen des Internets für die Wissenschaft." In: Sandra Hofhues/Konstanze Schütze (eds.), *Doing Research – Wissenschaftspraktiken zwischen Positionierung und Suchanfragen*, Bielefeld: transcript, pp. 417–423.
- Bruns, Axel (2019): "After the 'APocalypse': Social Media Platforms and Their Fight Against Critical Scholarly Research." In: *Information, Communication & Society* 22/11, pp. 1544–1566.
- Chu, Eric/Andreas, Jacob/Ansolabehere, Stephen/Roy, Deb (2023): *Language Models Trained on Media Diets Can Predict Public Opinion*, arXiv Preprint (<https://doi.org/10.48550/arXiv.2303.16779>).
- Coscia, Michele/Rossi, Luca (2022): "How Minimizing Conflicts Could Lead to Polarization on Social Media: An Agent-Based Model Investigation." In: *PLoS ONE* 17/1, e0263184.
- Dietrich, Bryce J. (2021): "Using Motion Detection to Measure Social Polarization in the U.S. House of Representatives." In: *Political Analysis* 29/2, pp. 250–259.
- Dowling, Michael/Lucey, Brian (2023): "ChatGPT for (Finance) Research: The Bananarama Conjecture." In: *Finance Research Letters* 53, 103662.
- Fecher, Benedikt/Hebing, Marcel/Laufer, Melissa/Pohle, Jörg/Sofsky, Fabian (2023): *Friend or Foe? Exploring the Implications of Large Language Models on the Science System*, Zenodo Data set (<https://doi.org/10.5281/ZENODO.8009429>).
- Ferrara, Emilio (2023): *Should ChatGPT Be Biased? Challenges and Risks of Bias in Large Language Models*, arXiv Preprint (<https://doi.org/10.48550/arXiv.2304.03738>).

- Freelon, Deen (2018): "Computational Research in the Post-API Age." In: *Political Communication* 35/4, pp. 665–668.
- Gilardi, Fabrizio/Alizadeh, Meysam/Kubli, Maël (2023): ChatGPT Outperforms Crowd-Workers for Text-Annotation Tasks, *arXiv Preprint* (<https://doi.org/10.48550/arXiv.2303.15056>).
- Glickman, Moshe/Sharot, Tali (2022): Biased AI Systems Produce Biased Humans, *Open Science Framework Preprint* (<https://doi.org/10.31219/osf.io/c4e7r>).
- Gweon, Hyukjun/Schonlau, Matthias (2023): Automated Classification for Open-Ended Questions with BERT, *arXiv Preprint* (<https://doi.org/10.48550/arXiv.2209.06178>).
- Hox, Joop J. (2017): "Computational Social Science Methodology, Anyone?" In: *Methodology* 13 (Supplement 1), pp. 3–12.
- Huang, Fan/Kwak, Haewoon/An, Jisun (2023): "Is ChatGPT Better than Human Annotators? Potential and Limitations of ChatGPT in Explaining Implicit Hate Speech." In: *WWW '23 Companion: Companion Proceedings of the ACM Web Conference 2023*, Austin, TX, USA, pp. 294–297.
- Ilkous, Eleni/Koutraki, Maria (2020): "Symbolic Vs Sub-Symbolic AI Methods: Friends or Enemies?" In: *Proceedings of the CIKM 2020 Workshops*, Galway, Ireland (<https://ceur-ws.org/Vol-2699/papero6.pdf>).
- Jürgens, Pascal/Meltzer, Christine E./Scharnow, Michael (2022): "Age and Gender Representation on German TV: A Longitudinal Computational Analysis." In: *Computational Communication Research* 4/1, pp. 173–207.
- Kühl, Niklas/Schemmer, Max/Goutier, Marc/Satzger, Gerhard (2022): "Artificial Intelligence and Machine Learning." In: *Electronic Markets* 32/4, pp. 2235–2244.
- Kunst, Jonas R./Bierwiazzonek, Kinga (2023): Utilizing AI Questionnaire Translations in Cross-Cultural and Intercultural Research: Insights and Recommendations, *PsyArXiv Preprint* (<https://doi.org/10.31234/osf.io/sxcyk>).
- Maslow, Abraham Harold (1966): *The psychology of science: a reconnaissance*, New York: Harper & Row.
- McAdoo, Timothy (2023): "How to cite ChatGPT.", April 7, 2023, (<https://apastyle.apa.org/blog/how-to-cite-chatgpt>).
- Mutz, Diana C. (2011): *Population-Based Survey Experiments*, Princeton: Princeton University Press.

- Nowok, Beata/Raab, Gillian M./Dibben, Chris (2016): "Synthpop: Bespoke Creation of Synthetic Data in R." In: *Journal of Statistical Software* 74/11 (<https://doi.org/10.18637/jss.v074.i11>).
- OpenAI (2023): ChatGPT, version from April 12, 2023 (<https://chat.openai.com/chat>).
- Pangakis, Nicholas/Wolken, Samuel/Fasching, Neil (2023): Automated Annotation with Generative AI Requires Validation, *arXiv Preprint* (<https://doi.org/10.48550/arXiv.2306.00176>).
- Pellert, Max/Lechner, Clemens M./Wagner, Claudia/Rammstedt, Beatrice/Strohmaier, Markus (2022): AI Psychometrics: Using Psychometric Inventories to Obtain Psychological Profiles of Large Language Models, *PsyArXiv Preprint* (<https://doi.org/10.31234/osf.io/jv5dt>).
- Perrigo, Billy (2023) "OpenAI Used Kenyan Workers on Less Than \$2 Per Hour to Make ChatGPT Less Toxic.", January 18, 2023 (<https://time.com/6247678/openai-chatgpt-kenya-workers/>).
- Pipal, Christian/Song, Hyunjin/Boomgaarden, Hajo G. (2022): "If You Have Choices, Why Not Choose (and Share) All of Them? A Multiverse Approach to Understanding News Engagement on Social Media." In: *Digital Journalism* 11/2, pp. 155–175.
- Rathje, Steve/Mirea, Dan-Mircea/Sucholutsky, Ilia/Marjeh, Raja/Robertson, Claire/Bavel, Jay Joseph van (2023): GPT Is an Effective Tool for Multilingual Psychological Text Analysis, *PsyArXiv Preprint* (<https://doi.org/10.31234/osf.io/sekf5>).
- Reiss, Michael (2023): Testing the Reliability of ChatGPT for Text Annotation and Classification: A Cautionary Remark, *Open Science Framework Preprint* (<https://doi.org/10.31219/osf.io/rvy5p>).
- Santurkar, Shibani/Durmus, Esin/Ladhak, Faisal/Lee, Cinoo/Liang, Percy/Hashimoto, Tatsunori (2023): Whose Opinions Do Language Models Reflect?, *arXiv Preprint* (<https://doi.org/10.48550/arXiv.2303.17548>).
- Sudmann, Andreas (2019): "On the Media-Political Dimension of Artificial Intelligence. Deep Learning as a Black Box and OpenAI." In: Andreas Sudmann (ed.), *The Democratization of Artificial Intelligence. Net Politics in the Era of Learning Algorithms*, Bielefeld: transcript, pp. 223–244.
- Sudmann, Andreas (2020): "Künstliche neuronale Netzwerke als Black Box: Verfahren der Explainable AI. Medienwissenschaftliche Perspektiven." In: Peter Klimczak/Christer Petersen/Samuel Schilling (eds.), *Maschinen der Kommunikation*, Wiesbaden: Springer Fachmedien, pp. 189–199.

- Webb Williams, Nora/Casas, Andreu/Wilkerson, John D. (2020): *Images as Data for Social Science Research: An Introduction to Convolutional Neural Nets for Image Classification*, Cambridge: Cambridge University Press.
- Xiao, Ziang/Zhou, Michelle X./Liao, Q. Vera/Mark, Gloria/Chi, Changyan/Chen, Wenxi/Yang, Huahai (2020): "Tell Me About Yourself: Using an AI-Powered Chatbot to Conduct Conversational Surveys with Open-Ended Questions." In: *ACM Transactions on Computer-Human Interaction* 27/3, pp. 1–37.
- Ziems, Caleb/Held, William/Shaiikh, Omar/Zhang, Zhehao/Yang, Diyi/Chen, Jiaao (2023): *Can Large Language Models Transform Computational Social Science?*, arXiv Preprint (<https://doi.org/10.48550/arXiv.2305.03514>).

