

Artificial Intelligence

By David Bann¹ and Liam Wright¹

This Methods Futures Briefing focuses on Artificial Intelligence (AI), a rapidly developing technology that will purportedly revolutionize society and science. We give an overview of recent advances, outline some potential future scenarios, and discuss the opportunities and challenges for social science that AI presents—focusing particularly on generative AI (GenAI). This briefing was written jointly by the authors; readers can compare with briefings entirely written by GenAI systems [here](#); and listen to an AI generated podcast based on this briefing [here](#).

What is AI in the current context?

The recent explosion in AI capability has centred around the development of large language models (LLMs). LLMs are a type of AI that generate outputs (usually text) based upon prompts and are typically accessed via chatbot interfaces (e.g., [ChatGPT](#), [DeepSeek](#)). Initially limited to text, [leading AI models](#) are now multimodal and work with text, image, and voice prompts. They are also increasingly integrated into software, connected to the internet, capable of writing and executing code and of analysing large documents. This means they can be used to search the internet, perform *and* interpret data analyses, summarise and query new information, and be used to write and edit documents.

LLMs are trained on the vast database on content containing a large chunk of the internet, including academic publications, and currently in the order of trillions of words. Leading LLMs pass versions of the ‘Turing test’: producing text responses in a way that are indistinguishable from humans. Major LLMs now meet or exceed average human performance in a range of specialised cognitive tasks (Perrault and Clark, 2024), including those related to language, logical problem solving, maths and code generation; the very skills which social researchers require. Recent models have also improved at reasoning (Guo et al., 2025) and there are AI systems which have been optimised for specific purposes, e.g., for drug development, solving mathematical problems, and so on.

So far, LLMs capabilities have scaled with the size of training runs, and [enormous sums](#) are spent on their development. However, even if no more powerful models were trained, we may expect further improvements as models are ‘[unhobbled](#)’ through the discovery of algorithmic and other improvements lowering costs and releasing latent capabilities.

Possibilities for social science

AI to help undertake research

AI tools can be incorporated—with varying degrees of success—into all aspects of the social scientist’s workflow. Errors notwithstanding, there is already great **potential to increase productivity** in research (see Resources section). AI can speed up learning new methods (Yan et al., 2024), writing tasks (Noy and Zhang, 2023), and coding tasks (Cui et al., 2024). Meanwhile, tools for literature reviews are developing, yet ‘hallucinations’ (erroneous references) remain a concern. AI can also aid in more quotidian tasks, such as grant applications and funder reports. The tantalising prospect is a future where researchers—who, some estimates suggest, spend [40% of their time](#) on admin at present—spend more time focusing on research and undertake more creative, ambitious projects. Additionally, social scientists are increasingly expected to have broader societal impact from their research—a drive which is not without criticism (Bann et al., 2024). AI can help here too: for example, by creating blogs, podcasts, or data visualisers with simple plain English

¹ University College London. Contact author: david.bann@ucl.ac.uk

prompts. Again, this may enable greater focus on research.

Current tools are akin to capable colleagues who often need careful direction and may, like all human colleagues, make mistakes. Increasing progress seeks to increase their utility and reduce error rates further. Unlike human colleagues, their response is effectively immediate, available 24/7, and at negligible cost. Ambitious projects which in the past required lengthy grant applications to fund the labour of junior research assistants may be looked back upon as archaic; the **promise of AI is more efficient and cheaper research**, enabling ambitious research to be undertaken by junior scholars or those in resource-poor settings. In-person access to experts or senior mentors is another resource which is inequitably distributed in academia—AI can be used as a colleague to seek advice on issues ranging from statistics or career advice, and to stimulate the [creative process in research](#).

Humans are currently required to be ‘in-the-loop’ to check AI outputs and guide their implementation. The degree of oversight required varies across tasks, though investing time improving prompts (*prompt engineering*; (Schmidt et al., 2024; Sahoo et al., 2024) yields better results. Eventually, particularly if AI advances continue, large fractions of the scientific workflow—idea generation, literature review, analysis, write-up and review—could be semi- or even fully-automated. Recent work uses multiple AI ‘agents’ to form **AI research assistants** (Schmidgall et al., 2025) or [co-scientists](#). **Fully automated systems** have been recently used in computer sciences and seemingly yield moderate-quality full publications (Lu et al., 2024). **Artificial General Intelligence** (AGI) would likely increase total research productivity, but could reduce demand for (human) researchers, and substantially change the job market (see below).

AI as a tool for social scientific inquiry

AI models can also be used to analyse the increasing vastness of data we obtain in the social sciences (‘omics, wearable technology, etc.). Some of the major scientific findings in other fields have arisen from AI-driven computational approaches to **analyse very large datasets** (e.g., the solving of protein folding) (Jumper et al., 2021). LLMs are beginning to be adopted as a statistical method for analysing social scientific datasets: for instance, the Life2Vec team (Savcisen et al., 2024) trained an LLM on life course sequences in administrative data and achieved accurate predictions.

As algorithms that turn text into numerical representations, LLMs can also be used to bridge the gap between qualitative and quantitative analyses, converting text into variables – a recent analysis (Wolfram et al., 2024) of ~10k essays in a large birth cohort used GPT-4 to represent the content of these essays, data that was then used in regression models. Similar methods for extracting meaning from text have also been used to create tools that identify semantically similar items across surveys (e.g., where [different languages were used](#)), holding promise for cross-context or harmonisation work.

AI for research infrastructure and data collection

AI can also increase the efficiency of tasks involved in designing and delivering the underlying data for social science research. For example, **testing and designing** questionnaires, and participant-facing materials for surveys. Prosaically, this may be by searching existing literature to identify widely adopted and psychometrically valid survey instruments used in a domain. More inventively, AIs can be used to generate practice participants prior to a studies’ roll-out; a recent study used AI generated characters that behaved similarly to humans in experiments (Mei et al., 2024).

With online experimental platforms, AGI may enable entire experimental protocols to be designed and delivered with simple plain English instructions. Further, **AIs could be used as (low cost) interviewers** themselves, asking relevant follow-up questions based on prior responses (including, in longitudinal studies, from prior survey sweeps) and a dynamically updating model of uncertainty about the interviewee. Just as AI is being used to help [design efficient computer chips](#), so AI can be used to help us design more efficient and robust systems for undertaking quantitative social science—including infrastructure such as Trusted Research Environments.

Concerns for social science

Career incentives can reward publishing an abundance of papers. For science, the challenges of AI include a potential further proliferation in scientific papers—arriving into a system which is struggling with growth of paper submissions (Hanson et al., 2024) far outstripping growth of human scientists. An **increase in low quality papers** may also occur—as occurred previously with the growth of poor quality meta-analyses (Ioannidis, 2016) and Mendelian Randomisation papers (Tobias et al., 2024). These risk damaging the scientific record and

trust in science. Nevertheless, profound errors abound in the literature already, including irreproducible (Open Science Collaboration, 2015) or irrepliable (Trisovic et al., 2022) results, errors not caught by reviewers (Schroter et al., 2008), p-hacking and publication bias (Bruns and Ioannidis, 2016). Mitigations include correctly aligning incentives towards higher quality, rather than volume, of scientific outputs across the world. Familiarity with Open Science practices will continue to aid future methodologists to undertake robust, replicable research.

The nature of quantitative research continues to change—from laborious manual calculations of underlying algebra to computer-calculated estimates in extremely complex analyses obtained almost instantly. Where researchers may have previously spent time performing calculations manually, they now spend it editing and deploying computer code; are any of these the optimum use of researcher time, or should they instead spend more of their time thinking and interacting with data at a higher level? For example, asking a LLM in plain English, 'How does X relate to Y in this dataset?'

Where computers exceed human level competence, we typically invest our time elsewhere—for instance, the calculator preserves mental effort adults would otherwise spend on manual arithmetic. A deeper understanding of (1) AI performance and (2) which forms of training are truly needed for longer-term learning will inform how we modify training in future.

Responsible, efficient and reflective uses of AI tools are new parts of social science training we should consider adding to curricula.

For social scientists, it is challenging to plan when developments are rapid and the scope of change is uncertain. The window in which any skill is useful (including prompt engineering) may now be much shorter and, by increasing overall productivity, there may be less need for as many researchers. This is especially true if one major source of funding, student tuition fees, declines as fewer people attend university, instead opting, far more cheaply, to use AIs as personal instructors.

Training LLMs uses a great deal of energy and energy consumption for simple requests, such as internet searches, can require 10-30 times more energy than a conventional search engine. More complex queries require more energy. When making methodological and ethical choices researchers will need to understand and

consider the environmental and sustainability dimensions of GenAI.

Even where a current AI tool is deeply flawed, dismissing its future utility appears unwise. Some tasks are more imminently amenable to AI automation than others, particularly repetitive tasks or those with large quantities of training data. Being nimble and keeping abreast of developments may help researchers keep track of these changes and inform where we allocate our development. Social scientists are needed to both understand and deploy AI to better achieve our goals.

Future

The implications for social science and methodologists depends on AI's future progress: whether development plateaus or human-like ability is achieved on all cognitive tasks (Artificial General Intelligence, AGI). Well informed researchers differ on their forecast for when (or if) AGI will be achieved; many predict over 50% chance that it occurs within the next 5 years. At some point, developments could generate an 'intelligence explosion', with AIs used to spur further improvements, leading to an exponential increase in intelligence far beyond human capability (Artificial Super Intelligence, ASI).

Selected resources and additional reading

Current resources:

Major LLMs, which can be used to generate and edit text or analytical code using plain language prompts:

- Open source: [DeepSeek](#), [Qwen](#), [Llama](#), [NVLM series](#), [Gemma](#), [Mistral](#).
 - See the [hugging face](#) resource.
 - To use locally, see [Ollama](#).
- Closed source: [Chat GPT](#), [Gemini](#), [Claude](#), [Grok](#).
- Podcasts creation: [notebookLLM](#)
- Code generation: [Co-Pilot](#) (note integration within R Studio and VS Code).
- Practical guides to using such tools: [Ethan Mollick](#), [Prompt Engineering Guide](#), [OpenAI Guide](#) and [Claude Guide](#), [Google Cloud](#)

Future AI development resources:

- Ethan Mollick ([One Useful Thing, Substack](#); [Co-Intelligence, Book](#))
- Zvi Mowshowitz ([Don't Worry About the Vase, Substack](#))
- Dan Shipper ([AI & I, Podcast](#))
- Dwarkesh Patel ([Dwarkesh Podcast / The Lunar Society, Podcast](#))
- Leopold Aschenbrenner ([Situational Awareness, Monograph](#))
- Alexander Kruel ([Axis of Ordinary, Substack](#))

References

- Bann D, Courtin E, Davies NM, et al. (2024) Dialling back 'impact' claims: researchers should not be compelled to make policy claims based on single studies. *International Journal of Epidemiology* 53(1): dyad181.
- Bruns SB and Ioannidis JP (2016) P-curve and p-hacking in observational research. *PloS One* 11(2): e0149144.
- Cui KZ, Demirer M, Jaffe S, et al. (2024) The Productivity Effects of Generative AI: Evidence from a Field Experiment with GitHub Copilot.
- Guo D, Yang D, Zhang H, et al. (2025) Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Hanson MA, Barreiro PG, Crosetto P, et al. (2024) The strain on scientific publishing. *Quantitative Science Studies*. 1-21.
- Ioannidis JP (2016) The mass production of redundant, misleading, and conflicted systematic reviews and meta-analyses. *The Milbank Quarterly* 94(3): 485-514.
- Jumper J, Evans R, Pritzel A, et al. (2021) Highly accurate protein structure prediction with AlphaFold. *Nature* 596(7873): 583-589.
- Lu C, Lu C, Lange RT, et al. (2024) The ai scientist: Towards fully automated open-ended scientific discovery. *arXiv preprint arXiv:2408.06292*.
- Mei Q, Xie Y, Yuan W, et al. (2024) A Turing test of whether AI chatbots are behaviorally similar to humans. *Proceedings of the National Academy of Sciences* 121(9): e2313925121.
- Noy S and Zhang W (2023) Experimental evidence on the productivity effects of generative artificial intelligence. *Science* 381(6654): 187-192.
- Open Science Collaboration (2015) Estimating the reproducibility of psychological science. *Science* 349(6251): aac4716.
- Perrault R and Clark J (2024) Artificial Intelligence Index Report 2024.
- Sahoo P, Singh AK, Saha S, et al. (2024) A systematic survey of prompt engineering in large language models: Techniques and applications. *arXiv preprint arXiv:2402.07927*.
- Savcisen G, Eliassi-Rad T, Hansen LK, et al. (2024) Using sequences of life-events to predict human lives. *Nature Computational Science* 4(1): 43-56.
- Schmidgall S, Su Y, Wang Z, et al. (2025) Agent laboratory: Using llm agents as research assistants. *arXiv preprint arXiv:2501.04227*.
- Schmidt DC, Spencer-Smith J, Fu Q, et al. (2024) Towards a catalog of prompt patterns to enhance the discipline of prompt engineering. *ACM SIGAda Ada Letters* 43(2): 43-51.
- Schroter S, Black N, Evans S, et al. (2008) What errors do peer reviewers detect, and does training improve their ability to detect them? *Journal of the Royal Society of Medicine* 101(10): 507-514.
- Tobias JH, Samaras K, Ivell R, et al. (2024) The proliferation of derivative and redundant studies in endocrinology due to the application of Mendelian Randomisation and other methods to open databases. *Frontiers in Endocrinology* 15: 1400583.
- Trisovic A, Lau MK, Pasquier T, et al. (2022) A large-scale study on research code quality and execution. *Scientific Data* 9(1): 60.
- Wolfram T, Tropf FC and Rahal C (2024) Large Language Models Predict Cognition and Education Close to or Better than Genomics or Expert Assessment. *OSF Preprints*.
- Yan L, Greiff S, Teuber Z, et al. (2024) Promises and challenges of generative artificial intelligence for human learning. *Nature Human Behaviour* 8(10): 1839-1850.

Acknowledgements

The authors would like to thank Robert Meckin, Mark Elliot, and members of the Centre for Longitudinal Studies paper review group for their comments on an earlier version.

National Centre for Research Methods
Social Sciences
University of Southampton
Southampton, SO17 1BJ
United Kingdom.

Web	http://www.ncrm.ac.uk
Email	info@ncrm.ac.uk
Tel	+44 23 8059 4539
X	@NCRMUK