

How is generative AI changing social science?

*Drawing on discussions from a recent workshop hosted by the Department of Methodology at LSE, **Thomas Robinson** highlights three ways in which AI is reshaping research across the social sciences and how it raises questions around the reflexivity of social science outputs in a world where AI is increasingly mainstream.*

Readers can watch highlights from the workshop on [LSE Methodology's website](#).

From discovering novel antibiotics, to generating news media and streamlining business operations, the adoption of generative AI models is rapidly permeating society. It is hard to ignore the hype and simultaneous concern about how these tools are developing, and their potential to upend established practices.

Generative AI poses challenges for social scientists. Mirroring their applications in wider society, these tools have the capacity to scale current research practises, as well as offering new ways of studying social phenomena. At the same time, the very individuals, groups, and institutions we study are themselves being affected by AI — in the workplace, in politics, and in the ways we socialise and communicate with each other.

these tools have the capacity to scale current research practises, as well as offering new ways of studying social phenomena

To address these questions the Department of Methodology at LSE recently organised a workshop bringing together leading researchers at the cutting edge of applying and studying generative AI in the social sciences.

Three key themes emerged from these discussions: when used carefully, AI can help perform research at a scale and fidelity greater than we can do otherwise; generative AI has the capability to affect social beliefs and behaviours; but, there are significant limitations to both the analytic and public acceptance of material generated by AI.

Scale

[Friedrich Geiecke](#), [Blake Miller](#), and [Melissa Sands](#), each highlighted ways in which AI can improve the scale at which we can conduct analysis. Friedrich Geiecke demonstrated how the interactive nature of generative AI (the model's ability to facilitate conversations and probe the responses of individuals) can enable researchers to reach many more subjects than would be otherwise possible, without using quantitative surveys.

Another key advantage of recent generative models is their ability to process more than just text (often referred to as "multi-modal" models). For example, Blake Miller's research uses this multi-modality to digitize text from scans of original documents. He uses this technique to extract text from historic papal bulls (official documents issued by a pope). In turn, this digitisation means that many documents can be rapidly converted into an analysable format, enabling a better understanding of how religious institutions implemented social policies across the Modern era.

Another key advantage of recent generative models is their ability to process more than just text (often referred to as "multi-modal" models)

Similarly Melissa Sands demonstrated how AI can infer subjective information from images. Her work compares the ability of models like ChatGPT to give subjective judgements of the safety of an area, using open-source, 360-degree photographs of streets in Detroit. They find promising correlations between the subjective scores given by real human responses and those scores generated by AI. Like Friedrich Geiecke's work on conversations, the fidelity of these models offer a promising way of doing something that typically could only be done on a small-scale, given the need for qualitative judgements.

Persuasion

The workshop also highlighted the capacity of generative AI to persuade individuals politically. [Christopher Summerfield](#) showed, in collaborative work with Google Deepmind, how generative [AI can be more effective at finding consensus between groups of individuals](#) than those individuals themselves. Crucially, the AI here acts as an arbiter: collating and balancing the individual opinions of citizens to find common ground.

Relatedly, [Lisa Argyle](#) provided experimental evidence that AI models can be persuasive at shifting voters towards certain viewpoints. These papers collectively show the promise, but also serious ethical issues, of using generative AI: how do we safeguard our social interactions in contexts we know automated AI tools can influence opinions and behaviour?

Limitations

Despite these exciting advances and use cases, the workshop also highlighted limitations to the use of generative AI. PhD student [Elif Akata](#) provided evidence that while AI models may appear proficient in reasoning tasks, under the surface they do not (yet) seem to be performing the sorts of processes we think are characteristic of human cognition. In his keynote, [Arthur Spirling](#) also pointed to our use of highly-structured data, which can limit the benefits of using complex, AI models to learn about the social world. Beyond research practise, too, Spirling presented evidence that the general public remain sceptical of AI-generated content, even if they find it plausible.

Common across the day was an implicit acceptance that researchers and the general public remain central to the effective use of generative AI

Common across the day was an implicit acceptance that researchers and the general public remain central to the effective use of generative AI. From preventing 'hallucinations' to providing the data to fine-tune models, the inputs and outputs of these models make the most sense when there are humans in-the-loop. While AI will let social scientists perform tasks at scales and precisions that were until recently exceptionally costly (if not impossible), they do not remove the need for the foundational social science skills.

Of course, the development of generative AI is happening at a rapid pace. Staying abreast of the capacities, and limitations, of these models requires a combination of computational, scientific, and social skill. For that reason, the use of generative AI is inherently interdisciplinary. This workshop helped demonstrate precisely how collaboration across disciplines can help make sense of these developments, and the fundamental challenges they pose to our understanding of the world.

The Generative AI in Social Science Research workshop was held on 7 June 2024 at the LSE, co-organised by Dr Thomas Robinson and Dr Daniel de Kadt, and was kindly supported by the Department of Methodology and LSE's Research Impact Support Fund.

The content generated on this blog is for information purposes only. This Article gives the views and opinions of the authors and does not reflect the views and opinions of the Impact of Social Science blog (the blog), nor of the London School of Economics and Political Science. Please review our [comments policy](#) if you have any concerns on posting a comment below.

Image Credit: Emily Rand & LOTI, Better Images of AI, [AI City](#), ([CC-BY 4.0](#))
