

POLICY FORUM

ARTIFICIAL INTELLIGENCE

Large AI models are cultural and social technologies

Implications draw on the history of transformative information systems from the past

By Henry Farrell¹, Alison Gopnik^{2,3,4},
Cosma Shalizi^{4,5,6}, James Evans^{4,7}

Debates about artificial intelligence (AI) tend to revolve around whether large models are intelligent, autonomous agents. Some AI researchers and commentators speculate that we are on the cusp of creating agents with artificial general intelligence (AGI), a prospect anticipated with both elation and anxiety. There have also been extensive conversations about cultural and social consequences of large models, orbiting around two foci: immediate effects of these systems as they are currently used, and hypothetical futures when these systems turn into AGI agents—perhaps even superintelligent AGI agents. But this discourse about large models as intelligent agents is fundamentally misconceived. Combining ideas from social and behavioral sciences with computer science can help us to understand AI systems more accurately. Large models should not be viewed primarily as intelligent agents but as a new kind of cultural and social technology, allowing humans to take advantage of information other humans have accumulated.

The new technology of large models combines important features of earlier technologies. Like pictures, writing, print, video, internet search, and other such technologies, large models allow people to access information that other people have created. Large models—currently language, vision, and multimodal—depend on the internet having made the products of these earlier technologies readily available in machine-readable form. But like economic

markets, state bureaucracies, and other social technologies, these systems not only make information widely available, they allow it to be reorganized, transformed, and restructured in distinctive ways. Adopting Simon's terminology, large models are a new variant of the "artificial systems of human society" that process information to enable large-scale coordination [(1), p. 33].

Our central point here is not just that these technological innovations, like all other innovations, will have cultural and social consequences. Rather we argue that large models are themselves best understood as a particular type of cultural and social technology. They are analogous to such past technologies as writing, print, markets, bureaucracies, and representative democracies. Then we can ask the separate question about what the effects of these systems will be. New technologies that are not themselves cultural or social, such as steam and electricity, can have cultural effects. Genuinely new cultural technologies—Wikipedia, for example—may have limited effects. However, many past cultural and social technologies also had profound, transformative effects on societies, for good and ill, and this is likely to be true for large models.

These effects are markedly different from the consequences of other important general technologies such as steam or electricity. They are also different from what we might expect from hypothetical AGI. Reflecting on past cultural and social technologies and their impact will help us to understand the perils and promise of AI models better than worrying about superintelligent agents.

SOCIAL AND CULTURAL INSTITUTIONS

For as long as there have been humans, we have depended on culture. Beginning with language itself, human beings have had distinctive capacities to learn from the experiences of other humans, and these capacities are arguably the secret of human evolutionary success. Major technological

changes in these capacities have led to dramatic social transformations. Spoken language was succeeded by pictures then by writing, print, film, and video. As more and more information became available across wider gulfs of space and time, new ways of accessing and organizing that information also developed, from libraries to newspapers to internet search. These developments have had profound effects on human thought and society, for better or worse. Eighteenth-century advances in print technology, for example, which allowed new ideas to quickly spread, played an important role in the Enlightenment and the French Revolution. A landmark transformation occurred around 2000 when nearly all the information from text, pictures, and moving images was converted into digital formats; it could be instantly transmitted and infinitely reproduced.

As long as there have been humans, we have also relied on social institutions to coordinate individual information-gathering and decision-making. These institutions can themselves be thought of as a kind of technology (1). In the modern era, markets, democracies, and bureaucracies have been particularly important. The economist Friedrich Hayek argued that the market's price mechanism generates dynamic summaries of enormously complex and otherwise unfathomable economic relations (2). Producers and buyers do not need to understand the complexities of production; all they need to know is the price, which compresses vast swathes of detail into a simplified but usable representation. Election mechanisms in democratic regimes focus distributed opinion toward collective legal and leadership decisions in a related way. The anthropologist Scott argued (3) that all states, democratic or otherwise, have managed complex societies by creating bureaucratic systems that categorize and systematize information. Markets, democracies, and bureaucracies have relied on mechanisms that generate lossy (incomplete, selective, and uninvertible) but useful representations well before the computer. Those representations both depend on and go beyond the knowledge and decisions of individual people. A price, an election result, or a measure such as gross domestic product (GDP) summarizes large amounts of individual knowledge, values, preferences, and actions. At the same time, these social technologies can also themselves shape individual knowledge and decision-making.

The abstract mechanisms of a market, state, or bureaucracy, like cultural media, can influence individual lives in crucial ways, sometimes for the worse. Central

¹SNF Agora Institute and School of Advanced International Studies, Johns Hopkins University, Baltimore, MD, USA.

²Department of Psychology, University of California, Berkeley, CA, USA. ³Department of Philosophy, University of California, Berkeley, CA, USA. ⁴Santa Fe Institute, Santa Fe, NM, USA.

⁵Department of Statistics and Data Science, Carnegie Mellon University, Pittsburgh, PA, USA. ⁶Department of Machine Learning, Carnegie Mellon University, Pittsburgh, PA, USA.

⁷Department of Sociology, University of Chicago, Chicago, IL, USA. Email: j.evans@uchicago.edu

banks, for example, reduced the complexities of the financial economy down to a few key variables. This provided apparent financial stability but at the cost of allowing instabilities to build up in the housing market, which central banks paid little attention to, precipitating the 2008 global financial crisis (4). Similarly, markets may not represent “externalities” such as harmful carbon emissions. Integrating such information into prices through, for example, a carbon tax can help but requires state action.

Humans rely extensively on these cultural and social technologies. These technologies are only possible, however, because humans have distinct capacities characteristic of intelligent agents. Humans, and other animals, can perceive and act on a changing external world, build new models of that world, revise those models as they accumulate more evidence, and then design new goals. Individual humans can create new beliefs and values and convey those beliefs and values to others through language or print. Cultural and social technologies transmit and organize those beliefs and values in powerful ways, but without those individual capacities, the cultural and social technologies would have no purchase. Without innovation, there would be no point to imitation (5).

Some AI systems—in robotics, for example—do attempt to instantiate similar truth-finding abilities. There is no reason, in principle, why an artificial system could not do so at some point in the future. Human brains do, after all. But at the moment, all such systems are far from these human capacities. We can debate how much to worry now about these potential future AI systems or how we might handle them if they emerge. But this is different from the question of the effects of large models at present and in the immediate future.


LARGE MODELS

Large models, unlike more agentive systems, have made notable and unexpected progress in the past few years, making them the focus of the current conversation about AI in general. This progress has led to claims that “scaling,” simply taking the current designs and increasing the amount of data and computing power they use, will lead to AGI agents in the near future. But large models are fundamentally different from intelligent agents, and scaling will not change this. For example, “hallucinations” are an endemic problem in these systems because they have no conception of truth and falsity (although there are practical steps toward mitigation). They simply sample and generate text and images.

Rather than being intelligent agents, large models combine the features of cultural and social technologies in a new way. They generate summaries of unmanageably large and complex bodies of human-generated information. But these systems do not merely summarize this information, like library catalogs, internet search, and Wikipedia. They also can reorganize and reconstruct representations or “simula-

larly, it is easy to treat markets and states as though they were agents, and agencies or companies can even have a kind of legal personhood.

But behind their agent-like interfaces and anthropomorphic pretensions, large language models (LLM) and large multimodal models are statistical models that take enormous corpora of text produced by humans, break them down into par-



tions” (1) of this information at scale and in new ways, like markets, states, and bureaucracies. Just as market prices are lossy representations of the underlying allocations and uses of resources, and government statistics and bureaucratic categories imperfectly represent the characteristics of underlying populations, so too are large models “lossy JPEGs” (6) of the data corpora on which they have been trained.

Because it is hard for humans to think clearly about large-scale cultural and social technologies, we have tended to think of them in terms of agents. Stories are a particularly powerful way to pass on information, and from fireside tales to novels to video games, they have done this by creating illustrative fictional agents, even though listeners know that those agents are not real. Chatbots are the successor to Hercules, Anansi, and Peter Rabbit. Simi-

ticular words, and estimate the probability distribution of long word sequences. This is an imperfect representation of language but contains a surprisingly large amount of information about the patterns it summarizes. It allows the LLM to predict which words come next in a sequence and so generate human-like text. Large multimodal models do the same with audio, image, and video data. Large models not only abstract a very large body of human culture, they also allow a wide variety of new operations to be carried out on it. LLMs can be prompted to carry out complex transformations of the data on which they are trained. Simple arguments can be expressed in flowery metaphors, while ornate prose can be condensed into plain language. Similar techniques enable related models to generate new pictures, songs, and video in response to prompts.

A body of cultural information that was previously too complex, large, and inchoate for large-scale operations has been rendered tractable.

In practice, the most recent versions of these systems depend not only on massive caches of text and images generated and curated by humans but also on human judgment and knowledge in other forms. In particular, the systems rely on reinforcement learning from human feedback (RLHF) or its variants: Tens of thousands of human employees provide ratings of model outputs. They also depend on prompt engineering: Humans must use both their background knowledge and ingenuity to extract useful information from the models. Even the newest “chain of thought” models regularly begin from dialogue with their human users.

The relatively simple though powerful algorithms that allow large models to extract statistical patterns from text are not really the key to the models’ success. Instead, modern AI rests atop libraries, the internet, tens of thousands of human coders, and a growing international world of active users. Someone asking a bot for help writing a cover letter for a job application is really engaging in a technically mediated relationship with thousands of earlier job applicants and millions of other letter writers and RLHF workers.

CHALLENGES AND OPPORTUNITIES

The AI debate should focus on the challenges and opportunities that these new cultural and social technologies generate. We now have a technology that does for written and pictured culture what large-scale markets do for the economy, what large-scale bureaucracy does for society, and perhaps even comparable with what print once did for language. What happens next? Like past economic, organizational, and informational “general purpose technologies,” these systems will have implications for productivity (7), complementing human work but also automating tasks that only humans could previously perform, and for distribution, affecting who gets what (8).

Yet they will also have wider and more profound cultural consequences. We do not yet know whether these consequences will be as great as those of earlier technologies such as print, markets, or bureaucracies, but thinking of them as cultural technologies increases rather than decreases their potential impact. These earlier technologies were central to the extensive social transformations of the 18th and 19th centuries, both as causes and effects. All of these technologies, like large models, supported the abstraction of information so that new kinds of opera-

tions could be carried out at scale. All provoked justified concerns about the spread of misinformation and bias, cultural homogenization or fragmentation, and shifts in the distribution of power and resources. The emergence of new communications media, including both print and television, was accompanied by reasonable worries that the new media would spread misinformation and strengthen malign cultural forces. Similarly, the categorization schemes that bureaucracies and markets deploy often embed oppressive assumptions.

At the same time, these technologies generated new possibilities for recombining information and coordinating actions among millions of people at a planetary scale. Emerging debates over the social, economic, and political consequences of LLMs continue deep-rooted historical worries and hopes about new cultural and social technologies. Orienting these debates requires both recognizing the commonalities between new arguments and old ones and carefully mapping the particulars of the new and evolving technologies.

Such mapping is among the central tasks of the social sciences, which emerged from the social, economic, and political upheavals of the Industrial Revolution and its aftermath. Social scientists’ investigation of the consequences of these past technologies can help us think about less obvious social implications of AI, both negative and positive, and to consider ways that AI systems could be redesigned to increase the positive impacts and reduce the negative. As media, markets, and bureaucratic technologies expanded in the 19th and 20th centuries, they generated economic losers and winners, displacing whole categories of workers, from clerks and typists to “human computers.” Today, there are obvious worries that large models and related technologies may displace “knowledge workers.”

There are also less obvious questions. Will large models homogenize or fragment culture and society? Thinking about this in historical context can be particularly illuminating. Current concerns resemble 19th- and 20th-century disagreements over markets and bureaucracies. Weber worried (9) about the deadening homogenizing consequences of economic and bureaucratic “rationalization,” whereas Mill (10) thought that market exchanges would expose participants to different forms of life and soften impulses to conflict (“doux commerce”).

Large models are designed to work well—to faithfully reproduce the actual probabilities of sequences of text, images, and video—on average. They therefore have an intrinsic tendency to be most accurate in situations most commonly found in their

training data and least accurate in situations that were rare in data or entirely new. This might lead large models to worsen the kind of homogenization that haunted Weber.

On the other hand, large models may allow us to design new ways to harvest the diversity of the cultural perspectives they summarize. Combining and balancing these perspectives may provide more sophisticated means of solving complex problems (11). One way to do this may be to build “society-like” ecologies in which different perspectives, encoded in different large models, debate each other and potentially cross-fertilize to create hybrid perspectives (12) or to identify gaps in the space of human expertise (13) that might usefully be bridged. Large models are surprisingly effective at abstracting subtle and nonobvious patterns in texts and images. This suggests that such technologies could be used to find patterns in text and images that crisscross the space of human knowledge and culture, including patterns invisible to any particular human. We may require new systems that diversify large model reflections and personas and produce the same distribution and diversity as do human societies.

Diversifying systems like this might be particularly important for scientific progress. Formal science itself depended on the emergence of the new cultural technologies of the 17th and 18th centuries, from coffee houses and rapid mail to journals and peer review. AI technologies have the potential to accelerate science further, but this will depend on imaginative ways of using and rethinking these technologies. By wiring together so many perspectives across text, audio, and images, large models may allow us to discover unprecedented connections between them for the benefit of science and society. These technologies have most commonly been trained to regurgitate routine information as helpful assistants. A more fundamental set of possibilities might open up if we deployed them as maps to explore formerly uncharted territory.

There are also less obvious and more interesting ways that new cultural and social technologies influence economic relationships. The development of cultural technologies leads to a fundamental economic tension between the people who produce information and the systems that distribute it. Neither group can exist without the other: A writer needs publishers as much as the publisher need writers. But their economic incentives push in opposite directions. The distributors will profit if they can access the producer’s information cheaply, whereas the producers will profit if they can get their information distributed cheaply. This ten-

sion has always been a feature of new cultural technologies. The ease and efficiency of distributing information in digital form has already made this problem especially acute, as evidenced by the crisis in everything from local newspapers to academic journals. But the very speed, efficiency, and scope of large models, processing all the available information at once, combined with the centralized ownership of those models, makes these problems loom especially large. Concentrated power may make it easier for those who own the systems to skim the benefits of efficiency at the expense of others.

There are crucial technical questions: To what extent can the systematic imperfections of large models be remedied, and when are they better or worse than the imperfections of systems based around human knowledge workers? Those should not overshadow the crucial political questions: Which actors are capable of mobilizing around their interests, and how might they shape the resulting mix of technology and organizational capacities? Very often, commentators within the technology sector reduce these questions into a simple battle between machines and humans. Either the forces of progress will prevail against retrograde Luddite tendencies, or on the other hand, human beings will successfully resist the inhuman encroachment of artificial technology. Not only does this fail to appreciate the complexities of past distributional struggles, struggles that long predate the computer, it ignores the many different possible paths that future progress might take, each with its own mix of technological possibilities and choices (8).

In the case of earlier social and cultural technologies, a range of further institutions, including normative and regulatory institutions, emerged to temper their effects. These ranged from editors, peer review, and libel laws for print, to election law, deposit insurance, and the Securities and Exchange Commission for markets, democracies, and bureaucracies. These institutions had varied effectiveness and required continual revision. These countervailing forces did not emerge on their own, however, but resulted from concerted and sustained efforts by actors both within and outside the technologies themselves.

LOOKING FORWARD

The narrative of AGI, of large models as superintelligent agents, has been promoted both within the tech community and outside it, both by AI optimist “boomers” and more concerned “doomers.” This narrative gets the nature of these models and their relation to past technological changes wrong. But more importantly, it actively

distracts from the real problems and opportunities that these technologies pose and the lessons history can teach us about how to ensure that the benefits outweigh the costs.

Of course, as we note above, there may be hypothetical future AI systems that are more like intelligent agents, and we might debate how we should deal with these hypothetical systems, but LLMs are not such systems, any more than were library card catalogs or the internet. Like catalogs and the internet, large models are part of a long history of cultural and social technologies.

The social sciences have explored this history in detail, generating a distinct understanding of past technological up-

“...large models are themselves best understood as a particular type of cultural and social technology.”

heavals. Bringing computer science and engineering into close cooperation with the social sciences will help us to understand this history and apply these lessons. Will large models lead to greater cultural homogeneity or greater fragmentation? Will they reinforce or undermine the social institutions of human discovery? As they reshape the political economy, who will win and lose? These and other urgent questions do not come into focus in debates that treat large models as analogs for human agents.

Changing the terms of debate would lead to better research. It would be far easier for social scientists and computer scientists to cooperate and combine their respective strengths if both understood that large models are no more—but also no less—than a new kind of cultural and social technology. Computer scientists could bring together their deep understanding of how these systems work with social scientists’ comprehension of how other such large-scale systems have reshaped society, politics, and the economy in previous eras, elaborating existing research agendas and discovering new ones. This would help remedy past confusions in which computer scientists have adopted overly simplified notions of complex social phenomena (14) while social scientists have failed to understand the complex functioning of these new technologies.

It would move policy discussions over AI decisively away from simplistic battles between the existential fear of a machine takeover and the promise of a near-future paradise in which everyone will have a per-

fectly reliable and competent artificial assistant. The actual policy consequences of large models will surely be different. Like markets and bureaucracies, they will make some kinds of knowledge more visible and tractable than they were in the past, encouraging policy-makers to focus on the new things that they can measure and see at the expense of those less visible and more confusing. As a result, reflecting past cases of markets and media, power and influence will shift toward those who can fully deploy these technologies and away from those who cannot. AI weakens the position of those on whom it is used and who provide its data, strengthening AI experts and policy-makers (14).

Last, thinking in this way might reshape AI practice. Engineers and computer scientists are already aware of the problem of large model bias and are thinking about their relationship to ethics and justice. They should go further. How will these systems affect who gets what?

What will their practical consequences be for societal polarization and integration? Can large models be developed to enhance human creativity rather than to dull it? Finding practical answers to such questions will require an understanding of social science as well as engineering. Shifting the debate about AI away from agents and toward cultural and social technologies is a crucial first step toward building that cross-disciplinary understanding (15). ■

REFERENCES AND NOTES

1. H. Simon, *The Sciences of the Artificial* (MIT Press, 1996).
2. F. A. von Hayek, *Am. Econ. Rev.* **35**, 519 (1945).
3. J. C. Scott, *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed* (Yale Univ. Press, 1998).
4. D. Davies, *The Unaccountability Machine* (Univ. Chicago Press, 2025).
5. E. Yiu, E. Kosoy, A. Gopnik, *Perspect. Psychol. Sci.* **19**, 874 (2024).
6. T. Chiang, *New Yorker* **9** (2023).
7. C. Goldin, L. Katz, *Q. J. Econ.* **113**, 693 (1998).
8. D. Acemoglu, S. Johnson, *Power and Progress: Our 1000 Year Struggle over Technology and Prosperity* (Hachette, 2023).
9. M. Weber, *Wissenschaft Als Beruf* (Duncker & Humblot, 1919).
10. J. S. Mill, *Principles of Political Economy* (Longmans and Green, 1920).
11. L. Hong, S. E. Page, *Proc. Natl. Acad. Sci. U.S.A.* **101**, 16385 (2004).
12. S. Lai et al., *Proc. 41st Int. Conf. Mach. Learn.* **235**, 25892 (2024).
13. J. Sourati, J. A. Evans, *Nat. Hum. Behav.* **7**, 1682 (2023).
14. S. L. Blodgett, S. Barocas, H. Daumé, H. Wallach, *arXiv:2005.14050 [cs.CL]* (2020).
15. L. Brinkmann et al., *Nat. Hum. Behav.* **7**, 1855 (2023).

ACKNOWLEDGMENTS

All authors contributed equally to this work. J.E. began a visiting researcher affiliation with Google after this manuscript was submitted.

10.1126/science.adt9819