# What's in the Chatterbox?

## Large Language Models, Why They Matter, and What We Should Do About Them

**IMPLICATIONS FOR THE SCIENTIFIC LANDSCAPE**


Johanna Okerlund
Evan Klasky
Aditya Middha
Sujin Kim
Hannah Rosenfeld
Molly Kleinman
Shobita Parthasarathy



GERALD R. FORD SCHOOL OF PUBLIC POLICY
SCIENCE, TECHNOLOGY, AND PUBLIC POLICY
UNIVERSITY OF MICHIGAN


**TECHNOLOGY ASSESSMENT PROJECT** REPORT

# Contents

About the Authors 3

About the Science, Technology, and Public Policy Program 5

Executive Summary 6
- Understanding the LLM Landscape 6
- How Do Large Language Models Work? 7
- The Implications of LLM Development 8
- The Implications of LLM Adoption 10
- Case Study: Transforming Scientific Research 13

Transforming the Scientific Landscape 15
- LLMs will transform scientific practices 16
- Scientific Evaluation by LLMs will Create Crises of Credibility 20
- LLMs will Reinforce Public Myths about Science 21
- LLMs will Hurt Open Access Movements 22
- LLMs Will Reinforce Anglo-American Scientific Dominance 23

Policy Recommendations 26

Recommendations for the Scientific Community 30

# About the Authors

**Johanna Okerlund** is a Human-Computer Interaction researcher with a background in Computer Science and additional training in Science Technology Studies and Public Policy. She has a PhD in Computing and Information Systems from the University of North Carolina at Charlotte, where she studied makerspaces relative to their promise of democratization. As a postdoc at U-M working with the Science, Technology, and Public Policy program and the Computer Science department, Johanna has been developing ways to bring ethics and justice into CS courses and contribute to ongoing research about the societal implications of emerging technology. She plans to continue approaching technology from a critical interdisciplinary perspective.

**Evan Klasky** is completing their Master's degree in Environmental Justice from the University of Michigan's School for Environment and Sustainability, along with a graduate certificate in Science, Technology, and Public Policy from the Ford School for Public Policy, in May 2022. Their research has focused on the biopolitics of agricultural technology. They hold a BA in Political Science from Haverford College, where he researched regime transformation in Venezuela. In the fall of 2022, he plans to enter a doctoral program in Geography.

**Aditya Middha** is an undergraduate student in Computer Science at the University of Michigan College of Engineering, with a minor in Public Policy, graduating in May 2022. Previously, he contributed to research on an ethical computer science curriculum, as well as risk-limiting election audits. On campus, he helped co-found D2 Map, a mobile platform to expand the reach of local community organizers, and serves as a weekly volunteer for the Downtown Boxing Gym. After graduation, Aditya will be a Product Manager for Microsoft and has plans to enter into the educational technology space in the near future.

**Sujin Kim** is completing her BA in Political Science from the University of Michigan, where she will graduate with honors and distinction in May 2022. She is interested in the politics of the congressional legislative process, and American political institutions more broadly. She has worked on research projects spanning a range of topics, including public health harm reduction legislation, cybersecurity policy, and congressional oversight capacity. Following graduation she will be pursuing a PhD in American Politics, and hopes to apply her experience with Congress and the legislative process in practice on the Hill.

## Hannah Rosenfeld

earned a Master of Public Policy degree from the University of Michigan, where she also received graduate certificates in Science, Technology, and Public Policy and Diversity, Equity, and Inclusion. Hannah was an author of the first Technology Assessment Project report *Cameras in the Classroom: Facial Recognition Technology in Schools* (2020) and conducted research on COVID-19 testing and medical technology innovation. She worked in the tech industry for over seven years developing consumer products and medical diagnostic tools before moving into technology regulation and led the New York City chapter of the LGBTQ+ non-profit Out in Tech before becoming the Head of Diversity, Inclusion, and Belongingness for the international organization. In April 2022, she will continue developing policy for emerging technology at the Food and Drug Administration, focusing on digital health.

## Molly Kleinman

serves as the Managing Director of the Science, Technology, and Public Policy program at the University of Michigan. In this role, Molly oversees the day-to-day management and provides strategic direction for STPP. Molly brings over 15 years of experience across several areas of higher education, with much of her work centering on educational technology, access to information, and intellectual property. Molly received her Ph.D. in Higher Education Policy from the University of Michigan Center for the Study of Higher and Postsecondary Education, her M.S. in Information from the University of Michigan School of Information, and her B.A. in English and Gender Studies from Bryn Mawr College.

## Shobita Parthasarathy

is Professor of Public Policy and Women's Studies, and Director of the Science, Technology, and Public Policy Program, at the University of Michigan. She conducts research on the political economy of innovation with a focus on equity, as well as the politics of evidence and expertise in policymaking, in comparative and international perspective. Her research topics include genetics and biotechnology, intellectual property, inclusive innovation, and machine learning. Professor Parthasarathy is the author of multiple scholarly articles and two books: *Building Genetic Medicine: Breast Cancer, Technology, and the Comparative Politics of Health Care* (MIT Press, 2007) and *Patent Politics: Life Forms, Markets, and the Public Interest in the United States and Europe* (University of Chicago Press, 2017). She writes frequently for public audiences and co-hosts *The Received Wisdom* podcast, on the relationships between science, technology, policy, and society. She regularly advises policymakers in the United States and around the world, and is a non-resident fellow of the Center for Democracy and Technology.



PDF: Click to return to top

# About the Science, Technology, and Public Policy Program

The University of Michigan's **Science, Technology, and Public Policy (STPP) program** is a unique research, education, and policy engagement center concerned with cutting-edge questions that arise at the intersection of science, technology, policy, and society. It is dedicated to a rigorous interdisciplinary approach, and working with policymakers, engineers, scientists, and civil society to produce more equitable and just science, technology, and related policies. Housed in the Ford School of Public Policy, STPP has a vibrant graduate certificate program, postdoctoral fellowship program, public and policy engagement activities, and a lecture series that brings to campus experts in science and technology policy from around the world. Our affiliated faculty do research and influence policy on a variety of topics, from national security to energy.

# Executive Summary

Large language models (LLMs)—machine learning algorithms that can recognize, summarize, translate, predict, and generate human languages on the basis of very large text-based datasets—are likely to provide the most convincing computer-generated imitation of human language yet. Because language generated by LLMs will be more sophisticated and human-like than their predecessors, and because they perform better on tasks for which they have not been explicitly trained, we expect that they will be widely used. Policymakers might use them to assess public sentiment about pending legislation, patients could summarize and evaluate the state of biomedical knowledge to empower their interactions with healthcare professionals, and scientists could translate research findings across languages. In sum, LLMs have the potential to transform how and with whom we communicate.

However, LLMs have already generated serious concerns. Because they are trained on text from old books and webpages, LLMs reproduce historical biases and hateful speech towards marginalized communities. They also require enormous amounts of energy and computing power, and thus are likely to accelerate climate change and other forms of environmental degradation. In this report, we analyze the implications of LLM development and adoption using what we call the analogical case study (ACS) method. This method examines the history of similar past technologies–in terms of form, function, and impacts–to anticipate the implications of emerging technologies.

This report first summarizes the LLM landscape and the technology's basic features. We then outline the implications identified through our ACS approach. We conclude that LLMs will produce enormous social change including: 1) exacerbating environmental injustice; 2) accelerating our thirst for data; 3) becoming quickly integrated into existing infrastructure; 4) reinforcing inequality; 5) reorganizing labor and expertise, and 6) increasing social fragmentation. LLMs will transform a range of sectors, but the final section of the report focuses on how these changes could unfold in one specific area: scientific research. Finally, using these insights we provide informed guidance on how to develop, manage, and govern LLMs.

## Understanding the LLM Landscape

Because LLMs require enormous resources in terms of finances, infrastructure, personnel, and computational power, only a handful of large tech companies can afford to develop them. Google, Microsoft, Infosys, and Facebook are behind the prominent LLM developments in the United States. While a few organizations (such as EleutherAI and the Beijing Academy of Artificial Intelligence) are developing more transparent and open approaches to LLMs, they are supported by the same venture capital firms and tech companies shaping the industry overall. Meanwhile, although there are many academic researchers in this area, they tend to depend on the private sector for LLM access and therefore work in partnership with them. Government funding agencies, including the National Science Foundation, support these collaborations. This tightness in the LLM development landscape means that even seemingly alternative or democratic approaches to LLM development are likely to reinforce the priorities and biases of large companies.

## How Do Large Language Models Work?

LLMs are much larger than their predecessors, both in terms of the massive amounts of data developers use to train them, and the millions of complex word patterns and associations the models contain. LLMs also more closely embody the promise of "artificial intelligence" than previous natural language processing (NLP) efforts because they can complete many types of tasks without being specifically trained for each, which makes any single LLM widely applicable.

Developing an LLM involves three steps, each of which can dramatically change how the model "understands" language, and therefore how it will function when it is used. First, developers assemble an enormous dataset, or "corpus", of text-based documents, often taking advantage of collections of digitized books and user-generated content on the internet. Second, the model learns about word relationships from this data. Large models are able to retain complex patterns, such as how sentences, paragraphs, and documents are structured. Finally, developers assess and manually fine-tune the model to address undesirable language patterns it may have learned from the data.

After the model is trained, a human can use it by feeding it a sentence or paragraph, to which the model will respond with a sentence or paragraph that it determines is appropriate to follow. Developers are under no obligation to disclose the accuracy of their models, or the results of any tests they perform, and there is no universal standard for assessing LLM quality. This makes it difficult for third parties, including consumers, to evaluate performance. But publicly available assessments of GPT-3, one of the largest language models to date, suggest two areas for concern. First, people are not able to distinguish LLM-generated text from human-generated text, which means that this technology could be used to distribute disinformation without a trace. Second, as suggested earlier, LLMs demonstrate gender, racial, and religious bias.

We add two more concerns, related to the emerging political economy of LLMs. As noted above, there are only a handful of developers working on these technologies, which means that they are unlikely to reflect much diversity in need or consideration. Developers may simply not know, for



PDF: Click to return to top

*We add two more concerns, related to the emerging political economy of LLMs. Because there are only a few developers working on these technologies, they are unlikely to reflect much diversity in need or consideration. And, because the vast majority of models are in English, they are unlikely to achieve their translation goals. Taking these dimensions together, they could exacerbate global inequalities.*

example, the limitations in their models and corpora and thus, how they should be adjusted. Additionally, the vast majority of models are based on English, and to a lesser extent Chinese, texts. This means that LLMs are unlikely to achieve their translation goals (even to and from English and Chinese), and will be less useful for those who are not English or Chinese dominant. Taking these dimensions together, they could exacerbate global inequalities.

We have divided the findings of our ACS analysis into two categories. The first focuses on the implications of LLM design and development, examining the social and material requirements to make the technology work. The second identifies how LLM applications and outputs might transform the world.

## The Implications of LLM Development

### Exacerbating Environmental Injustice

LLMs rely on physical data centers to process the corpora and train the models. These data centers rely on massive amounts of natural resources including 360,000 gallons of water a day and immense electricity, infrastructure, and rare earth material usage. As LLMs become widespread, there will be a growing need for these centers. We expect that their construction will disproportionately harm already marginalized populations. Most directly, data centers will be built in inexpensive areas, displacing low-income residents, as US highways did in the 1960s when planners displaced over 30,000 Black and immigrant families per year. In the process of accommodating LLMs, tech companies will turn a blind eye to similar community disruption. Meanwhile, those that continue to live near data centers will



PDF: Click to return to top

be forced to deal with an increased strain on scarce resources and its subsequent effects. Already, residents near Google and Microsoft data centers on the West Coast have expressed concerns about the companies' overconsumption of water and contribution to toxic air pollution. Unfortunately, it is unlikely that these concerns will influence siting decisions; like oil and gas pipelines, we expect that data centers will be legally classified as "critical infrastructure". Attempted protests will be treated as criminal offenses.

## Accelerating the Thirst For Data

As we note above, LLMs are based on datasets made up of internet and book archives. The authors of these texts have not provided consent for their data to be used in this way; tech developers use web crawling technologies judiciously to stay on the right side of copyright laws. But because they collect enormous amounts of data, LLMs will likely be able to triangulate bits of disconnected information about individuals including mental health status or political opinions to develop a full, personalized picture of actual people, their families, or communities. We expect that this will trigger distrust of LLMs and other digital technologies. In response, users will use evasive and anonymizing behavior when operating online which will create real problems for institutions that regularly collect such information. In a world with LLMs, the customary method for ethical data collection–individual informed consent–no longer makes sense.

We are also concerned that LLM developers will turn to unethical methods of data collection in order to diversify the corpora. As noted above, researchers have already demonstrated how LLMs reflect historical biases about race, gender, religion, and sexuality. The best way to address these biases is to ensure that the corpora include more texts authored by people from marginalized communities. However, this poses serious risks of unethical data extraction such as when Google attempted to improve the accuracy of its facial recognition technology by, in part, taking pictures of homeless people without complete informed consent.

At the same time, LLMs will enhance feelings of privacy and security for some users. Disabled people and the elderly, who often depend on human assistants to fulfill basic needs, will now be able to rely on help from LLM-based apps.

## Normalizing LLMs

We expect that in order to ensure that LLMs become central to our daily lives, developers will emphasize their humanitarian and even empowering features. At present, most people know nothing about the technology, except for tech news watchers aware that Google fired two employees due to their concerns about equity and energy implications. In this environment, developers will emphasize the technology's modularity: that it can be tuned to serve specific purposes. This emphasis on flexibility will be reminiscent of the early days of the auto industry, when car manufacturers promoted broad social acceptance of the automobile by encouraging skeptical farmers to use the technology as a malleable power source. We also expect developers to quickly integrate the technology into crucial and stable social systems, such as law enforcement.

Finally, developers will emphasize the accuracy of LLMs and attempt to minimize any errors and deflect blame for them. This was already clear in the Google episode, when the company asked their employees to remove their names as co-authors from a research paper critical of LLMs. But this is a common approach, especially at early stages of a technology's deployment. One particularly high-profile example is the Boeing 737 MAX plane. After Boeing quietly installed the Maneuvering Characteristics Augmentation System (MCAS) system onto its planes and an Indonesian airliner crashed, the company insisted that the pilots were at fault. Only after a second plane crash in Ethiopia did corrective action take place. LLM development could follow a similar path, deflecting blame away from the technology until problems become too big to ignore or until affected parties learn about one another and build a coalition in response.

# The Implications of LLM Adoption

## Reinforcing Inequality

Trained on texts that have marginalized the experiences and knowledge of certain groups, and produced by a small set of technology companies, LLMs are likely to systematically misconstrue, minimize, and misrepresent the voices of historically excluded people while amplifying the perspectives of the already powerful. But fixing these problems isn't just a matter of including more, better data. LLMs are built and maintained by humans who bring values and biases to their work, and who operate within institutions, in social and political contexts. This will shape the LLM issues that developers perceive, and how they choose to fix them.

*Trained on texts that have marginalized the experiences and knowledge of certain groups, and produced by a small set of technology companies, LLMs are likely to systematically misconstrue, minimize, and misrepresent the voices of historically excluded people while amplifying the perspectives of the already powerful.*

Our analysis shows that LLMs are likely to reinforce inequalities in a few ways. In addition to producing biased text, they will reinforce the inequitable distribution of resources by continuing to favor those who are privileged through its design. For example, racial bias is already embedded in medical devices such as the spirometer, which is used to measure lung function. The technology considers race in its assessment of "normal" lung function, falsely assuming that Black people naturally have lower lung function than their white counterparts. This makes it more difficult for Black people to access treatment. Similarly, imagine an LLM

*Credit: Ian Warburton / © BBC / Better Images of AI (CC-BY 4.0)*

app designed to summarize insights from previous scientific publications and generate health care recommendations accordingly. If previous publications rely on racist assumptions, or simply ignore the needs of particular groups, the LLM's advice is likely to be inaccurate too. We expect similar scenarios in other domains including criminal justice, housing, and education where biases and discrimination enshrined in historical texts are likely to generate advice that perpetuates inequities in resource allocation. Unfortunately, because the models are opaque and appear objective, it will be difficult to identify and address such problems. As a result, individuals will bear the brunt of them alone.

Meanwhile, LLMs will reinforce the dominance of Anglo-American and Chinese language and culture at the expense of others. We are particularly concerned that the corpora are composed primarily of English or Chinese language texts. While some developers have argued that LLMs could help preserve languages that are disappearing, LLMs are likely to function best in their dominant training language. Eventually this will reinforce the dominance of standard American English in ways that will expedite the extinction of lesser-known languages or dialects, and contribute to the cultural erasure of marginalized people. Furthermore, because they are based on historical texts LLMs are likely to preserve limited, historically suspended understandings especially of the non-American or Chinese cultures represented in its corpora.

## Remaking Labor and Expertise

Most people studying the impact of automation on labor warn of job losses, particularly for those in lower skilled occupations. In the case of LLMs, we expect job losses to be more prevalent in professions tightly coupled with previous technologies;

LLMs will completely eliminate certain kinds of tech-based work such as content moderation of social media while creating new kinds of tech-based work. But our analysis suggests that LLMs are also likely to transform labor. In particular, we expect that with widespread adoption LLMs will perform mundane tasks while shifting humans to more difficult or damaging tasks. This will even happen in high-skilled professions. Consider genetic counselors, who began helping people assess their and their families' genetic risks in the early 20th century. With the recent rise of genetic testing, consumers are increasingly learning about their risks through private companies such as 23andMe. But genetic counselors are still working; they just handle the more complex, urgent, and stressful cases.

Professions that heavily use writing (e.g., law, academia, journalism) will have to develop new standards and mechanisms for evaluating authorship and authenticity. For example, the invention of the typewriter led to the creation of the "document examiner" position to determine the provenance of typed text; we could imagine a similar job for LLM-based text. Finally, we expect widespread use of LLMs to trigger labor resistance. There is a long legacy of technology-driven labor unrest including the Luddites of the 19th century. More recently, the United Food and Commercial Workers International Union's developed public campaigns against Amazon's cashierless grocery store model. LLMs will incite similar resistance from workers and consumers based on fear of job loss, violations of social norms, and reduced income taxes.

## Accelerating Social Fragmentation

While LLMs may be used primarily in the workplace, we also expect a variety of public-facing apps, including those that summarize medical information and help citizens generate legal documents. Such apps are likely to empower some communities in important ways, even allowing them to mount successful activism against scientific, medical, and policy establishments. But, because LLM design is likely to distort or devalue the needs of marginalized communities we worry that LLMs might actually alienate them further from social institutions. We also expect social fragmentation to arise elsewhere, as LLMs will allow individuals to generate information that aligns with their interests and values and erode shared realities further.

Finally, as LLMs get better at writing text that is indistinguishable from something a human could have written, they will not only challenge the cultural position of authors but also trust in their authorship. For example, many schools and universities today use plagiarism detection technologies to prevent student cheating. However, this has triggered a technological arms race. A variety of services have emerged to help students cheat while evading detection by Turnitin, from websites full of how-to advice to paid essay writing services. LLMs will trigger a similar dynamic. The more writers of all kinds use LLMs for assistance, the more efforts to authenticate whether they "really" wrote their article or book, and the more writers will find new ways to take advantage of LLM capabilities without detection. In the long run, this will create cultures of suspicion on a massive scale.

## Case Study: Transforming Scientific Research

Overall, this report focuses broadly on the social and equity impacts of LLMs, and we have suggested that the technology will affect a range of professions. In the final substantive section of the report, we provide an example of how LLMs will affect just one: scientific research. First, because academic publishers, such as Elsevier and Pearson, own most research publications, we expect that they will construct their own LLMs and use them to increase their monopoly power. While LLMs could be extremely valuable tools for disseminating knowledge, publishers' LLMs will concentrate knowledge further and most people will be unable to afford subscriptions. While researchers may try to construct alternative LLMs that provide accessible and egalitarian access to scholarly research, these will be extremely difficult to build without targeted assistance from both the scientific community and government funders.

In addition to shaping access to knowledge, we expect that LLMs will transform scientific knowledge itself. Technologies, from the microscope to the superconducting supercollider, have long shaped the substance of research, and LLMs will be no exception. We expect that fields that analyze text, including the digital humanities, to be the most affected. Researchers will need to develop standard protocols on how to scrutinize insights generated by LLMs and how to cite LLM output so that others can replicate the results. LLMs are likely to have profound impacts on the nature of scientific inquiry as well, by encouraging recent trends that focus on finding patterns in big data rather than establishing causal relationships.

LLMs are also likely to transform scientific evaluation systems. Editors currently struggle to find peer reviewers, and LLMs could help. However, LLMs are likely to be rigid and systematically biased. Institutional review boards, which evaluate the ethics of scientific research, have been repeatedly criticized for reducing ethical assessments to legal hurdles, and we expect a similar outcome if LLMs are used for peer review. For example, LLMs will probably not be able to identify truly novel work, a task that is already quite difficult for human beings. Given these likely outcomes, we suspect that scientists will come to distrust LLMs.

Finally, we expect that LLMs will help some researchers improve their English or Chinese writing skills and increase their publications in top journals. The technology will likely be particularly useful for scholars from British Commonwealth countries whose language may differ only slightly from standard English. However, we expect translation in and out of other languages to be poor and researchers unfortunately may not always be aware of such limitations at the outset. Meanwhile, the more common LLMs become as a scientific tool, the more they will reinforce English as the lingua franca of science. This will likely also mean that the values and concerns of the English-speaking world–particularly the United States and Britain–will dominate global scientific priorities. And yet, these political implications may remain hidden because LLMs will be promoted as a technology that will be able to truly globalize science.

# Transforming the Scientific Landscape

**KEY POINTS**

- LLMs will transform both the kind of research scientists do, and how they do it.
- Academic publishers are likely to develop LLMs to maintain their monopoly power over most scientific literature.
- Using LLMs to conduct scientific evaluation will generate controversy among scientists.
- LLMs will reinforce Anglo-American dominance in science.

Throughout this report, we have anticipated the social, political, and equity implications if LLMs are adopted across a range of sectors. In this chapter, we examine how LLMs might transform one sector in particular: science. In this analysis, it is crucial to remember that the major LLMs currently under construction are based on corpora composed primarily of open access texts available online. But, most recent research publications–particularly scientific journal articles–are owned by academic publishing companies such as Elsevier and JSTOR. Therefore, they are not part of these corpora. We expect that these publishers might develop their own LLMs that leverage their proprietary text databases, particularly at a moment when universities are frustrated by their high fees (Resnick & Belluz, 2019). These proprietary LLMs are likely to be of greatest interest to the scientific community because they will be the most up-to-date, in contrast to publicly available LLMs that may contain slightly older scientific knowledge. As they become more important to academic researchers, universities may be forced to maintain their subscriptions. Less likely is that academic publishers will sell their texts to the large companies for inclusion in their corpora, because it would make their texts essentially available to everyone.

In this new environment, LLMs will transform scientific practices, including authorship and citations. They may also transform peer review systems, which have increasingly come under scrutiny. LLMs will also reinforce Anglo-American dominance in science. While they may help some scientists from low and middle income countries participate more actively in the international scientific community and engage in cross-national collaboration, the English and Chinese language dominance of the corpora will limit efforts to "decolonize" science. Finally, LLMs will limit the power of the open access movement, as academic publishers are likely to have more resources than governments, non-profit organizations, and individuals to generate LLMs.

## LLMs will transform scientific practices

### Remaking scientific authorship and methods

Given their capacity to process and summarize huge amounts of text, we expect LLMs to have a profound impact on authorship and scientific methods as well as evaluation. As we describe in more detail below, researchers in non-English speaking countries are likely to use LLMs to more accurately translate texts or check their grammar or spelling. This might make it easier for them to publish in top journals, which are invariably published in English. Even English-dominant researchers might use LLMs to generate more generic parts of scientific texts, including materials and methods, and parts of introductions and conclusions. As we discuss in Section 5, we expect that these uses will trigger questions about rightful authorship.

We also expect LLMs to profoundly shape scientific practice. The development of particle accelerators in the 1930s allowed physicists to investigate the structure of the atomic nucleus, and more recently to investigate subatomic particles (Ishkhanov, 2012). The polymerase chain reaction technique, which makes millions of copies of small pieces of DNA, transformed genetics and biotechnology research and enabled mapping and sequencing the human

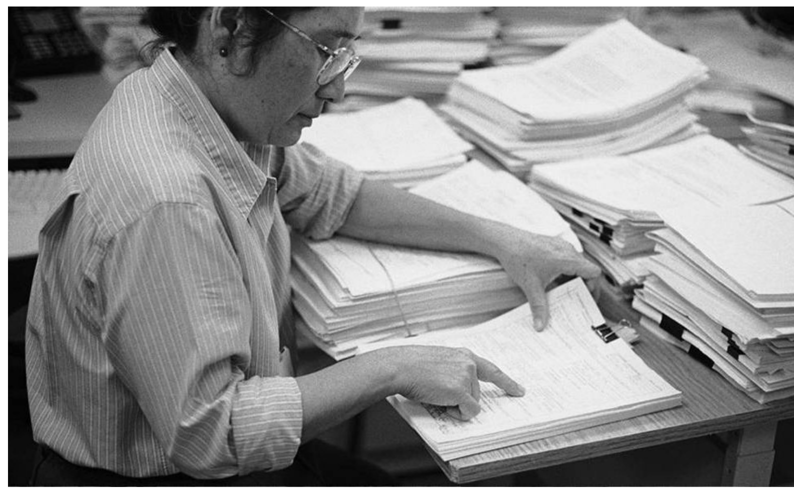

*Credit: National Institutes of Health (CC0)*

genome, the study of ancient DNA, and gene manipulation including CRISPR gene editing (Rabinow, 2011). And the internet has already had profound impacts on research. It has made it easier for scholars to read research across fields, and thus promote interdisciplinary thinking (Herring, 2002). It has also helped researchers contact a wider array of potential subjects, whether for clinical trials or for surveys and interviews. Social scientists, for example, use email, social media, and even the "crowdworking" platform Mechanical Turk (MTurk) owned by Amazon to publicize their studies and recruit subjects. MTurk allows researchers to access a fairly representative population for a small fee (less than half of minimum wage) (Fort et al., 2011).

LLMs will similarly enable new forms of research, perhaps most notably in the humanities. Historians and scholars of English literature will be able to quickly generate summary information about historical texts or genres in the major corpora or new texts they wish to consider. However, scholars may be reticent to use these sources for two reasons. First, scholars accustomed to using archives and carefully documenting the provenance of texts are likely to be wary of LLMs as data sources at least initially, because of the lack of transparency about the texts contained in the corpora and the inability to cite them specifically. Scholars and academic publications will likely have to develop conventions about whether and how LLMs are used and documented. Wikipedia, for example, has become an important source introducing scholars to a particular topic, but is generally not acceptable as a reference in serious scholarly work (Chen, 2009). Second, because corpora predominantly include dominant and privileged voices, they may be of less utility in fields that are increasingly trying to capture the perspectives and experiences of those who have been historically marginalized.

LLMs will also continue to transform the nature of scientific inquiry. In recent years, there has been an explosion in enormous datasets and the computing power needed to process them. As a result, scientists can now use algorithms to identify correlations in huge datasets rather than starting with hypotheses (Huang, 2018; Kitchin, 2014). However, these correlations tell them neither about causality nor how such relationships emerge. In addition, just because a correlation appears in the data doesn't mean it is real or meaningful (Zhang, 2018). Researchers could also use LLMs as a new tool for data analysis, using them to extract insights from or summarize large amounts of text. Qualitative researchers are often constrained by the laborious manual processes of thematic coding, for example, but LLMs would allow them to analyze greater quantities of data or draw insights from data sources such as social media posts that were previously too large to consider as research sources. Psychologists and political scientists could use data from the corpora to assess public attitudes and concerns. Given academic pressures to publish ("or perish"), we expect the proliferation of articles identifying data correlations. However, without changing statistical methods, this could also increase the production of spurious data that cannot be reproduced.

## Scientific Credit Systems will Change

Scientists identify the lineage of their interests, theories, and methods through explicit citations to earlier work. This is an important method of providing credit. It has also become crucial to measuring scholarly impact. Scientists use "citation counts" to decide whether a publication is worth reading, or citing in their own publications.

Hiring, tenure, and promotion committees use these indicators to judge a scientist's impact. Meanwhile, journals have developed "impact factors" based on the average number of times their articles are cited; these impact factors in turn affect scientists' decisions where to publish and university decisions on how to evaluate employees and applicants. However, citation practices are also highly political; white men tend to be the most cited across fields (Caplar et al., 2017; Dworkin et al., 2020).

We expect LLMs to reduce citations overall, and ultimately reinforce existing biases in research fields. While LLMs currently do not have the technical capability to identify which text from the corpus informed the generated text, if a future LLM is able to provide citations along with the text summaries, we expect it to privilege highly cited articles which are not likely to represent the field's diversity or its most novel findings. But in the more likely scenario, scientists might query an LLM about the prevailing knowledge related to a particular phenomenon and simply treat the output as general knowledge that doesn't need to be cited. Consider the recent controversy over sharing data about COVID-19 genomic variants. Western scientists advocated putting this information into an open database that could be used across the world, to facilitate quicker understanding of disease progression and development of prophylactics, diagnostics, and treatments (Van Noorden, 2021). However, scientists from Southern countries protested, arguing that the open approach would rob them of the opportunity to receive credit for their hard work identifying variants such as Omicron (Maxmen, 2021). They worried further that scientists from wealthy nations would publish papers based on—but not citing—their results, because they had the resources to do further analysis, write up their findings, and submit them for publication. More generally, they were frustrated that as soon as they had begun to build expertise and resources to participate in the transnational world of science, Western leaders seemed to be changing the game. Similarly, marginalized scientists might worry that LLMs will make it more difficult for them to receive credit and for their ideas to become recognized as part of a mainstream corpus of knowledge.

*Marginalized scientists might worry that LLMs will make it more difficult for them to receive credit and for their ideas to become recognized as part of a mainstream corpus of knowledge.*

## Transforming Peer Review

We also expect research funding agencies, scientific publishers and editors, and even patent systems to consider incorporating LLMs into their review processes. These institutions depend on technical experts to assess the novelty of a study or invention, the appropriateness of the methods, and the plausibility of findings. Invariably, these experts also advise researchers how to consider and address counterfactuals, strengthen their claims or findings, or simply improve their writing. But peer reviewers are unpaid, and as academic pressures increase it is difficult to find good peer reviewers; editors say that they spend an enormous amount of time searching, and even then the reviewers may be uninformed, provide insufficient evaluation, or take too long and delay publication (Benos et al., 2007; Severin & Chataway, 2021). LLMs could solve many of these problems. Developers could create algorithms based on the backlists of all scholarly publications, or smaller ones targeted to a particular field or a particular journal, in order to identify high-quality publications and even advise authors how to improve their publications or fit better with the journal's standards. In fact, researchers have already begun to develop algorithms that claim to predict the grantability of patent applications, and even which patents are likely to be the most consequential (Candia & Uzzi, 2021). The next step would be to use LLMs to determine patentability, a particularly attractive option as patent offices struggle to hire and retain their personnel.

In the short term, editors might use LLMs as a half-measure, to help identify peer reviewers. They might ask the LLM: "who is an expert in X topic?" Editors have long used email and the internet in this way, which has allowed them to diversify their pool of reviewers. However, because LLM corpora are composed of historical texts, this use might actually eliminate the gains in reviewer and field diversity made in recent years. Unless the LLM is used very carefully, and with additional checks, this use could also affect a field's trajectory. An LLM might define reviewer expertise in terms of the number of citations in a particular journal (or set of journals), which may not represent a field's cutting edge.

If humans begin to use LLMs to conduct peer review itself, this could become a bigger problem. LLMs are likely to produce conservative peer reviews. We expect editors to use LLMs to scaffold parts of the peer review process—that is, to train the technology to look for particular elements in a paper, such as particular methods—to ensure quality reviews. However, this scaffolding could produce inflexible standards and slower recognition of truly novel results. It could also transform scientific practices. Consider the history of the IRB, in which narrow definitions of risk, benefit, and generalizable research have become hurdles for researchers (White, 2007). Or, educators in K-12 schools, who have increasingly had to twist their instructional strategies to accommodate standardized testing (Shelton & Brooks, 2019). Overall, LLMs might be good at evaluating papers in a field where the conventions, materials, and methods are well-established. However, it is hard to imagine how a corpus based on historical texts could adequately evaluate new and evolving science (Kuhn, 1962); we already know that this is a challenge for human reviewers (Pontis et al., 2017). As a result, widespread use of LLMs for primary peer review could limit creativity. It could also perpetuate biases against certain types of investigation, such as on structural racism or systemic inequality (Hoppe et al., 2019).

## Scientific Evaluation by LLMs will Create Crises of Credibility

LLM-based scientific evaluation systems could also erode trust both within and beyond science. Today, peer review is the predominant form of scientific evaluation. Experts in a subfield review grant applications and scientific publications, and validate the ideas or findings as credible and worthy of funding or further circulation through scholarly journals or academic presses (Latour, 1987). Media outlets and governments often expect research to be peer reviewed before reporting on it or using it as the basis for policymaking. But this approach to evaluating scientific results is not natural or self-evident; it is the product of social negotiations and settlement. And it could certainly be otherwise. In the 17th century, wealthy gentlemen were assumed to be trustworthy–and producing credible scientific findings–because they were free from economic pressures (Shapin, 1995). They maintained their credibility by employing probabilistic discourse and minimizing precision, so as to avoid direct conflict with their peers. Scientists also trusted others' findings because they could witness the experiments themselves (Shapin & Schaffer, 1985). As the scientific enterprise grew, witnessing became "virtual", through standardization of methods, research publications, and peer review (Baldwin, 2018). These changes, however, came from within the scientific community, invariably when they concluded that they needed to establish credibility among new audiences.

*Credit: Philadelphia College of Pharmacy and Science (CC BY 4.0)*

In fact, professional communities respond quite poorly to externally imposed evaluation systems, and these external impositions tend to be less successful when the community is powerful. For example, in 1836 the US Congress passed a law requiring the Patent Office to employ examiners with science and engineering backgrounds, to replace the clerks who had previously handled patent applications. It was concerned that the bureaucracy was issuing too many patents based on old, unoriginal, and non-workable ideas, and believed that highly trained technical experts would solve the problem (Swanson, 2009). However, when these new examiners applied scientific standards for novelty and nonobviousness, they found that very few applications should be granted. Patent agents and lawyers, who were accustomed to a bureaucracy that had only legal criteria for granting patents, protested vigorously and threatened that if no patents were granted, the fledgling US economy would fail. They were ultimately successful; Patent Office administrators negotiated with the new examiners to lower their standards. Physicians launched similar protests when the United States began to consider a national health care system in the mid-20th century, because they worried that it would lead to new forms of oversight and evaluation (Starr, 1982).

Especially because many scientists have already begun to criticize the business models of academic publishing—and ultimately distrust their intentions—we expect that if these companies build LLMs to replace peer review it will create a similar crisis among scientists. Scientists will not trust the technology to replace their judgment, and will likely point out the types of limitations that we have outlined above. We also expect publics to question scientific results that LLMs have evaluated, particularly in the early days of the technology or in response to the publication of particularly controversial ideas. And if communities don't trust evaluation systems then they will challenge the institutions promoting them. Prescription drug recalls have engendered not only mistrust in the US Food and Drug Administration, but hesitancy towards vaccines (Goldenberg, 2021). Similarly, distrust in the US Centers for Drug Control and Prevention has exacerbated resistance to mask wearing and other protection measures during the COVID-19 pandemic.

## LLMs will Reinforce Public Myths about Science

As we have discussed in earlier sections of this report, we expect LLMs will increase the trend towards open and free information facilitated by the internet. Patients will be able to query disease symptoms and receive summaries of related medical articles. Curious individuals can generate lay summaries about the most technical topics, from astrophysics to artificial intelligence. In many respects, this will, as developers argue, democratize access to knowledge.

But as the technology presents complex scientific findings in comprehensible language, we expect that it will flatten important nuance, caveats, error rates, and uncertainties. This, we fear, will reinforce the illusion that scientific findings are objective, stanceless, value-free, and are generated with a view from nowhere. Ultimately, this could exacerbate public skepticism of science. We have seen this with previous efforts to popularize science. Scientific journalism, for example, tends to minimize what scholars call the "translational gap": the amount of additional research needed before scientific findings can lead to better medical practice (Summers-Trio et al., 2019). Instead, they tend to overestimate the importance of early stage studies. For example, many early biomedical studies are performed on mice. This can provide general indicators about the safety or effectiveness of a particular treatment, or shape of a particular phenomenon, but


PDF: Click to return to top

mice are quite different physiologically than humans. However, media articles still report these results with breathless excitement, creating false expectations about the imminence of treatments and the power of science (Chakradhar, 2019). Similarly, museums and other exhibitions such as World's Fairs tend to produce idealized images of cultures and countries, reinforcing distorted public understandings with real geopolitical consequences (Swift, 2019). We expect LLMs to reinforce a similarly idealized image of science, which will leave publics bewildered and frustrated when they confront its realities. Ultimately, this could exacerbate problems of public trust and alienation particularly among publics already questioning scientific findings (Funk, Kennedy, & Tyson, 2020; Funk, Kennedy, & Johnson, 2020).

## LLMs will Hurt Open Access Movements

Finally, we expect LLMs to become another tool for academic publishing giants to maintain their control over scientific knowledge. In recent years, researchers have become increasingly concerned about how journal subscription costs hurt access to knowledge. This, they argue, limits who can participate in scientific knowledge production and ultimately, the quality of science itself. In response, universities are canceling huge journal subscriptions (Resnick & Belluz, 2019). Researchers are sharing preprints on their own websites, or on portals such as Sci-Hub and ArXiV.org (Nicholas et al., 2019). They are publishing in "open access" journals. Journals may implement new forms of monetization by charging LLM developers who use their university subscriptions to incorporate journal articles into training corpora. But we believe that LLMs will increase the attractiveness of Elsevier and other academic publishers themselves. Given their financial resources and monopolies over huge volumes of scientific texts, publishers could create their own LLMs for researchers and bundle them in their services to academic institutions. They might even require universities to purchase all of their journals in order to access their LLM. Indeed, companies frequently leverage emerging technologies to maintain or enhance their monopoly power. Monsanto spliced "terminator gene" technology into its genetically modified crops in order to prevent them from replicating (Masood, 1998). This meant that farmers could not replant their seeds after the growing season, which they had done for hundreds of years. Similarly, academic publisher JSTOR, in conjunction with MIT, used its internet surveillance capabilities to track down and stop excessive downloads of journal articles it owned. An MIT student activist Aaron Swartz downloaded these articles in order to promote their open access; he was later criminally charged for this act and died by suicide (Schwartz, 2013).

Given the vitality of the open access movement, we expect scientists to resist by creating grassroots LLMs. They might build on the work of non-profit initiatives such as Eleuther AI and rely on pro bono expertise and donated pre-prints and other text to develop apps. Scientists made similar attempts to gather data about disease-causing mutations in genes linked to breast and ovarian cancer (known as the BRCA genes), to compete with biotechnology company Myriad Genetics' virtual monopoly on BRCA gene testing in the United States (Conley et al., 2014). Myriad used its testing monopoly to build a proprietary database of information about the genomic variants discovered, their association with disease, as well as individual and family health histories. Even though it lost its US testing monopoly in 2013 after patients, physicians, and scientists contested its patents (Parthasarathy, 2017), Myriad maintained its intellectual property through this database; patients and physicians preferred to use Myriad's testing service rather than others because the database could provide better interpretations

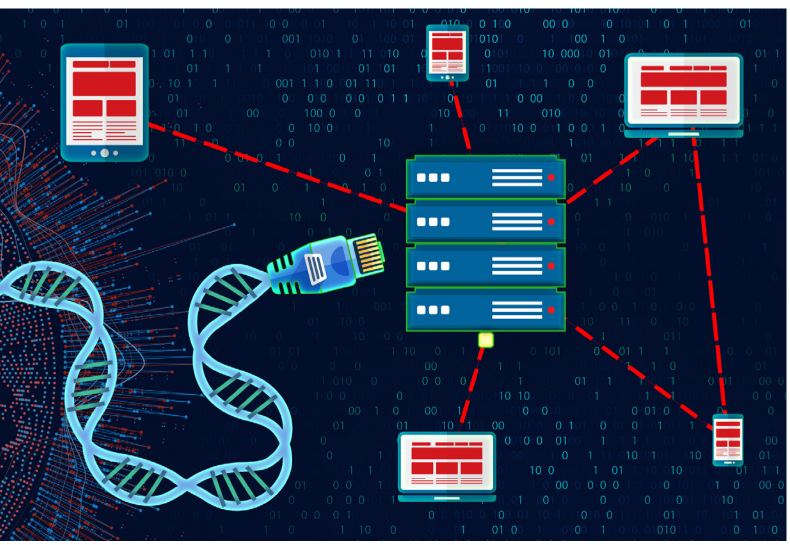

Credit: Ernesto Del Aguila III, NHGRI

about the implications of the genetic variants for disease. In order to build their alternative, scientists had to rely on word of mouth, and voluntary submissions of test results and other information from patients and physicians. This made it virtually impossible to build a database as powerful or useful as Myriad's, which in turn made it difficult to challenge the company's monopoly. We expect scientists developing grassroots LLMs to confront similar challenges, even if they have access to adequate technical expertise and financial resources.

## LLMs Will Reinforce Anglo-American Scientific Dominance

Like the telephone and the internet, LLMs may facilitate global scientific communication and even cooperation. However, given the technology's capacity to summarize and translate text, some may assume that it could facilitate real international inclusion and even the "decolonization" of science. Consider how the internet has changed science. Internet search engines, scientific databases, and social media have helped scientists learn about and build upon one another's work, regardless of where they are in the world. Email has facilitated communication, allowing researchers to contact one another and even collaborate despite living in different time zones or on distant continents. Indeed, there is evidence that international scientific collaboration has increased significantly in recent years, allowing scientists to share project costs, gain access to expansive or unique physical resources, share more data, and enhance creativity (Matthews et al., 2020). And yet, technology-mediated communication also increases misunderstandings. Whereas previous collaborations may have required scientists to visit laboratories for extended periods of time to learn methods, now such collaborations can occur without any in-person contact. This makes it much more difficult to transfer tacit knowledge–intangible scientific practices–which is essential for proper collaboration (Collins, 1992). However, scientists may not be aware that this knowledge is lost.

In the abstract, LLMs could allow scientists across countries to read texts in their native languages, facilitating communication. In practice, however, the picture already looks more complicated. As we have noted repeatedly throughout this report, LLM corpora–particularly those being built by the major companies–are primarily in English, and to a lesser extent, Chinese. This is crucial when considering the impacts of LLMs for international scientific cooperation; it means that the technology's translation capabilities are likely to be poor, particularly for the languages where there are fewer digitized texts. While scientists in non-English speaking countries may initially use them for translation purposes, the outputs will likely be filled with errors and this practice will stop. However, we do expect scientists to use LLMs to improve their English writing, to facilitate journal publication. While scientists in former British or US colonies could also use them to gain easier access to knowledge, they may still not have access to the proprietary LLMs sold by academic publishing companies. Thus, while LLMs may help some scientists in low and middle income countries, the prevailing political economy of science is likely to prevent true mutual learning and engagement.

*While LLMs may help some scientists in low and middle income countries, the prevailing political economy of science is likely to prevent true mutual learning and engagement.*

Instead, we expect LLMs to reinforce Anglo-American dominance in science while also helping Chinese scientists. In fact, it may also promote international collaboration between the two. Our research suggests that most efforts to promote mutual understanding across nations cannot escape geopolitical power struggles. Consider the World's Fairs, international platforms to showcase national scientific and technological achievements and facilitate cultural exchange, which began in the late 18th century. Cities hosting these yearly events brought global attention to their activities, and the sites also usually featured themed pavilions from a variety of countries that allowed them to showcase themselves and perhaps even develop grounds for collaboration (Molella & Knowles, 2019). However, countries used these as opportunities to advance their priorities. In 1993, South Korea's fifth largest city Daejeon hosted a Specialized Expo which produced international investment, and brought attention to another region beyond the large and prosperous city of Seoul (Knowles, 2019 p. 207). Similarly, while both the United States and Soviet Union focused on similar themes of technological progress and cultural diversity in the 1958 World's Fair, the United States took a less serious approach in order to downplay the perception of its strength and power during the Cold War (Swift, 2019 p. 38). Similarly, *Nature* has always characterized itself as a premier scientific journal that explicitly serves an international community despite its British base. However, in its early decades it saw the world through a British lens (Baldwin, 2015). Contributors adopted a voyeuristic approach to foreign science, and often used it as a foil to comment on national affairs.

The more common LLMs become as a scientific tool, the more they will reinforce English as the lingua franca of science. This will likely also mean that the values and concerns of the English-speaking world–particularly the United States and Britain–will dominate global scientific priorities. Furthermore, knowledge produced in English may be viewed as more generalizable than knowledge produced in other languages. And yet, these political implications may remain hidden because LLMs will be promoted as a technology that will be able to truly globalize science.

In this section, we have explored the range of implications that LLMs will have on scientific knowledge and practice. We expect LLMs to transform scientific priorities and practices, and systems of authorship, credit, and evaluation. This may produce crises of credibility, not only within science and beyond. It will also strengthen the power of scientific publishers, despite growing frustration about their knowledge monopolies. Finally, while we are hopeful that LLMs could facilitate international cooperation and inclusion, we fear that this will not materialize unless the corpora become much more diverse.

# Policy Recommendations

LLMs have great potential to benefit society. However, the priorities of the current development landscape make it difficult for the technology to achieve this goal. Below, we articulate how both LLMs (the models themselves, corpora, and output) and LLM-based apps must be regulated in order to maximize the public good. We also recommend greater scrutiny of LLMs' impacts on labor and the environment. Finally, we recommend that the National Science Foundation (and similar science funding agencies around the globe) invest more heavily in research related to LLMs and their impacts, to balance attention in an area currently dominated by the private sector.

1

RECOMMENDATION 1

The US government must regulate LLMs, for example through the Federal Trade Commission. This should include:

a. Clear definition of what constitutes an LLM.

b. Evaluation and approval of LLMs based on: 1) process of corpus development and ongoing procedures for maintenance and quality assurance; 2) diversity of the corpus; 3) LLM performance including accuracy particularly in terms of output related to marginalized communities; 4) transparency of the corpora and algorithms; and 5) data security.

c. Evaluation of efforts to diversify corpora. Government should monitor data extraction practices to ensure that efforts to diversify the corpora are ethical.

d. A complaint system that allows users to document their negative experiences with an LLM. These complaints should be publicly available. Developers must articulate in writing how they have addressed all complaints.

e. Ongoing oversight and monitoring of LLMs. Developers must make the corpora available to regulators for periodic testing. This should include both basic accessibility and comprehensibility to someone with a basic understanding of data and computer science.

f. Requirement to label all LLM output as such and include information about the developer.

## POLICY RECOMMENDATIONS (CONTINUED)

2

### RECOMMENDATION 2

The US government must regulate all apps that use LLMs, for example through the Federal Trade Commission, according to their use. The more consequential the LLM output, the greater the regulatory scrutiny (e.g., LLM-based apps related to criminal justice and patient care receive more extensive evaluation). Evaluation should consider:

a. Whether app developers are using the right LLM for their needs.

b. Likelihood that the app will generate false or dangerous results.

c. Potential benefits for the user.

d. Social, equity, and psychological implications, including potential harms to end users.

3

### RECOMMENDATION 3

Either a national or international standard setting organization (e.g., National Institute for Standards and Technology, International Standards Organization) must publish yearly evaluations of LLMs. They should assess: 1) diversity of the corpora; 2) performance; 3) transparency; 4) accuracy; 5) data security; and 6) bias towards marginalized communities.

4

### RECOMMENDATION 4

The US government must enact comprehensive data privacy and security laws.

5

### RECOMMENDATION 5

Under no circumstances should LLM-based apps deployed by the government (e.g., chatbots that provide information about social services, pre-trial risk assessment apps in criminal justice proceedings) harvest personally identifiable information.

## POLICY RECOMMENDATIONS (CONTINUED)

6

### RECOMMENDATION 6

The agencies that regulate LLMs and LLM-based apps, those that incorporate LLMs into its services, and all standard-setting bodies (e.g., the National Institute for Standards and Technology) must employ full-time advisors in the social and equity dimensions of technology. This "Chief Human Rights in Tech" Officer would advise procurement and technology evaluation decisions, monitor the technology once it is used and flag problems, and address disparate impacts.

7

### RECOMMENDATION 7

Both national and international intellectual property authorities (e.g., the US Copyright Office, the World Intellectual Property Organization) must develop clear rules about the copyright status of LLM-generated inventions and artistic works.

8

### RECOMMENDATION 8

All environmental assessments of new data centers must evaluate the impacts on local utility prices, local marginalized communities, human rights in minerals mining, and climate change.

9

### RECOMMENDATION 9

The US government must work with other governments around the world (perhaps under the auspices of the United Nations) to develop global labor standards for tech work (including minerals mining).

10

### RECOMMENDATION 10

The government must evaluate the health, safety, and psychological risks that LLMs and other forms of artificial intelligence create for workers, e.g., reorienting them towards more complex and often unsafe tasks. The Occupational Safety and Health Administration can perform this role, but it will require new regulations for workplace safety and an expansion of its purview to include psychological risks.

## POLICY RECOMMENDATIONS (CONTINUED)

11

### RECOMMENDATION 11

The US government must develop a robust response to the job consolidation that LLMs, and automation more generally, are likely to create. At a targeted level this should include job retraining programs and at a broad level, a guaranteed basic income and universal health care.

12

### RECOMMENDATION 12

The National Science Foundation must substantially increase its funding for LLM development. This funding should prioritize:

a. Developing alternative corpora and models, especially those driven by the needs of low-income and marginalized communities (and in partnership with them).

b. Meetings that establish standards for making corpora representative and for incorporating the knowledge of citizens (particularly low-income and marginalized communities)

c. Supporting updates and maintenance of existing corpora and models (in contrast to just making more new models).

d. Support research into building new types of models that are more easily updated and maintained.

e. Research into evaluation of fit between model and use.

f. Research on the equity, social, and environmental impacts of LLMs.

# Recommendations for the Scientific Community

We urge all professions to develop rules and guidelines to accommodate the rise of LLMs. Because we focused our attention on how LLMs might affect science (Section 7), we offer recommendations specific to this community. We hope this will guide researchers, journal editors, scientific publishers, and universities, as they contend with this emerging technology.

## Development of LLMs by the scientific community

- If scientific publishers develop LLMs, they should:
    - Provide users with information about how output is generated (i.e., the composition of the corpora and the logic of the algorithm).
    - Ensure that the LLM is accessible to and accurate for non-English speakers.
- The National Science Foundation should support the development of an LLM that includes publicly available journal articles and all results generated from their funding. It should deliberately include texts across all fields. To ensure that it captures the nuances of a variety of fields, experts from multiple disciplines–from the natural sciences to the humanities–should test it before deployment.
- All authors should be permitted to opt-out of their texts' inclusion in LLM corpora.

## RECOMMENDATIONS FOR THE SCIENTIFIC COMMUNITY (CONTINUED)

### LLM use for evaluation

- If scientific journals and academic publishers use LLMs to evaluate the quality of manuscripts, they must be transparent about this use. This includes clear explanations on the publisher's website so that prospective authors can be fully informed about LLM use before submission.
- Scientific journals and academic publishers should not rely completely on LLMs for "peer review". LLMs are likely to produce conservative evaluations–and therefore be more critical of novel findings and ideas–because they are based on historical texts.

### Research using LLMs

- Scientific journals and academic publishers must develop rules for how they–and peer reviewers–will evaluate research conducted using LLMs.
- All publications that rely on LLMs for text analysis should provide detail about the corpora and algorithms on which the results are based.

### Scientific communication using LLMs

- Scientific communicators should help publics understand how to use LLMs to interpret science. This includes evaluating which LLMs are the most appropriate for their needs, and how to understand the credibility of LLM output.
- Scientific communicators and publics should test LLMs before deployment to ensure that outputs related to scientific topics are accurate, credible, and comprehensible.

**VIEW THE FULL REPORT**

# myumi.ch/LLMReport

If you would like additional information about this report, the Technology Assessment Project, or University of Michigan's Science, Technology, and Public Policy Program, you can contact us at **stpp@umich.edu** or **stpp.fordschool.umich.edu**.

GERALD R. FORD SCHOOL OF PUBLIC POLICY
SCIENCE, TECHNOLOGY, AND PUBLIC POLICY
UNIVERSITY OF MICHIGAN

**Technology Assessment Project**
**Science, Technology, and Public Policy Program**
Gerald R. Ford School of Public Policy
University of Michigan
735 S. State Street
Ann Arbor, MI 48109

(734) 764-0453
stpp.fordschool.umich.edu
stpp@umich.edu



Brain/circuit icon: Freepik