

# CURRENT HISTORY

January 2025

*“To give people a say over issues with such profound impact on their lives, AI governance ought to be democratized on all levels, including the global level.”*

## The Democratic Challenges in Global Governance of AI

EVA ERMAN AND MARKUS FURENDAL

At the heart of the current artificial intelligence boom is the steadily repeating mantra that we live in extraordinary times. Depending on who is asked, we seem to be just a few years away from unleashing AI technologies that will boost overall productivity, solve medical enigmas, turn politics on its head, or dispose of humankind.

This kind of sentiment was expressed by AI’s poster boy Sam Altman, chief executive of OpenAI, when he argued with characteristic gravity in a July 2024 *Washington Post* opinion piece that we currently “face a strategic choice about what kind of world we are going to live in”:

Will it be one in which the United States and allied nations advance a global AI that spreads the technology’s benefits and opens access to it, or an authoritarian one, in which nations or movements that don’t share our values use AI to cement and expand their power? There is no third option—and it’s time to decide which path to take.

Keeping regimes like those of Russia and China at bay, Altman claimed, requires the United States to invest heavily in both digital infrastructure and human capital, as well as to help set up global AI institutions akin to the International Atomic Energy Agency (IAEA) or the Internet Corporation for Assigned Names and Numbers. This would benefit the US economy, while creating “a world shaped by a democratic vision for AI.” The right strategy would not only ensure that “democratic

AI” defeats “authoritarian AI,” but also help create a more democratic world.

Although he has his own corporate interests at stake in the matter, Altman is right to conceptualize the governance of AI technology as a multi-level issue. And there are two key reasons not to underestimate the importance of global regulatory initiatives.

The first is that the AI industry is a truly global phenomenon, in the sense that it is driven by large multinational companies like Microsoft, Alphabet, and Meta. They recruit talent from all over the world, train AI models on data collected from the Internet, and release their products in markets across jurisdictions. Local and national regulatory initiatives—like the recently debated California Senate Bill 1047 and US President Joe Biden’s October 2023 executive order on AI—may end up being toothless against these giants. AI companies could simply decide to withhold their products from the more strictly regulated markets, and even move their headquarters if need be. This might, in turn, trigger a race to the bottom in which legislation is successively weakened in each jurisdiction until an equilibrium is reached where AI is regulated less extensively than most people would want. Consider also that AI breakthroughs are happening in the open-source community, which is even more amorphous and difficult to assign to any particular jurisdiction.

A second factor that makes AI governance a truly global issue is that the technology’s disruptive effects are not confined by geographical boundaries. Just like air pollution and the release of greenhouse gases, AI models create substantial value for some, but serious and tangible problems for many others. Economists would thus describe

---

EVA ERMAN is a professor of political science at Stockholm University. MARKUS FURENDAL is a researcher at the Department of Political Science at Stockholm University and at the Institute for Futures Studies.

AI technology as having significant externalities—costs and benefits that affect people outside the industry. International relations experts would add that well-designed global institutions can play key roles in resolving the problem. So Altman is correct in suggesting that global AI governance institutions could act in everyone’s interest by preventing a regulatory race to the bottom and promoting a fair distribution of AI’s positive and negative effects, just as the IAEA was created to promote the safe, secure, and peaceful use of nuclear technology.

At the same time, Altman’s analysis is overly simplistic. He exploits the rhetorical power of terms like “democratic” and “authoritarian” AI without explaining what they mean. He assumes that we can identify each of these types by asking if it was developed in a democratic state or an authoritarian one, implicitly suggesting that “Western” means democratic, and democratic means good AI. Yet there are plenty of examples where such an equivalence breaks down. China is often criticized for conducting AI-powered automated surveillance of its minorities, but it arguably has the capacity to build such systems only because American chipmakers like Nvidia have willingly done business with Chinese firms.

Nor is privacy-violating surveillance restricted to authoritarian states. The American company ClearviewAI has scraped personal pictures from the Internet to build powerful facial recognition software, marketed its services to Western law enforcement agencies, and tried to shut down investigative reporting on its practices. It was recently fined by the Netherlands’ data protection authority under the European Union’s General Data Protection Regulation for including images of Dutch people in its software training data without their consent. The company claims that the fine is unenforceable since it does not offer its services to customers in the Netherlands or the EU.

Being clear and precise about the concept of democratic AI is crucial. Aside from being used for rhetorical effect by influential people like Altman, the idea features in increasingly common calls for “democratization of AI.” On closer inspection, such calls appear to combine several distinct claims.

One is the assertion that it is desirable to increase diversity among AI developers, assuming that AI development will be more in line with what

people want if there is greater similarity between users and creators. Another is that AI technology should be available for more people to use, with “democratic” as a placeholder for “inclusiveness” and “equal access.” Altman’s infrastructural proposals belong to a third variant, which has less to do with democracy as an ideal than with how to leverage AI as a potent technology in the global struggle for power between democracies and non-democracies.

More attention should be paid to what is arguably the most important sense in which we can speak of democratic AI, namely, that the development and deployment of AI technology ultimately should be democratically controlled. In other words, we need a closer focus on AI *governance*. Even if only a fraction of the predictions about how society will be changed by AI technology come to pass, it is clear that people will be affected both as private individuals and as citizens. Technological development, after all, is not an independent and unstoppable force; it can be partly guided on the basis of certain values and

toward particular goals. Democratic AI, in this view, means that those who are significantly affected by the advent of AI technology—which is all of us—should have a say in how it is governed.

---

*Most AI development is currently  
beyond democratic control.*

---

## WHO CONTROLS AI TODAY?

Although the era of AI has just begun, it is obvious that the technology already affects our daily lives. Just how this happens depends, of course, on who you are. Early adopters are probably reaping the benefits of incorporating AI before most others, but you do not have to use AI yourself to be affected by it.

If you have regular interactions with bureaucracies, you are most likely already subjected to automated decision-making, with far-reaching consequences for your chances in life. If you work in a profession that was previously thought to be difficult to automate, recent advances in generative AI might raise concerns about your job safety. If you are a student or a teacher, you may be confused about how to use AI as a tool; its arrival has raised questions about what the point of education is. Even skeptics who may be hesitant to try AI applications are not exempt: photos of their faces have likely been included without their knowledge in the masses of data used to train

those applications. And whether we like it or not, we find ourselves in a society where the culture, norms, and social expectations are being transformed by the new technology, much like what happened with the arrival of the smartphone or social media in the past couple of decades.

Despite recent efforts by governments and organizations to regulate the AI industry, it is fair to say that most AI development is currently beyond democratic control. At the very least, some people—like venture capitalists, Silicon Valley CEOs, and AI engineers—exert much more influence over AI development than others, leaving the rest of us to react and adapt to their disruptive technology.

We are not claiming that democratic AI governance requires each design decision to be made by committee. This not only would be infeasible, but also would halt what is often highly desirable technological progress. Moreover, democratically controlled AI would not necessarily mean that *you* get to decide how AI influences the world. It is inherent to democracy that some people end up in the minority and see their preferences disregarded in favor of the majority's. Democratic governance is a way to ensure that this happens in a legitimate way. There is an important difference between not having your views taken into consideration because you lack a seat at the table, and having an opportunity to state your views but ending up in the minority.

Spelling out what democratic AI governance means requires us to ask a set of complex questions, which Altman's dichotomy cannot capture. These include *what* should be democratically controlled, *who* should have a say, and *how* this should happen. Our answers to these questions depend, in turn, on *why* we value the democratic ideal and what reasons we have for extending it to the AI domain.

## WHAT SHOULD BE DEMOCRATIZED?

What needs to be democratized in order to help ensure that those affected by AI have influence in decision-making about it? Although AI is often described as such a fast-moving technology that it cannot be regulated, it is more accurate to say that there is, by now, an emerging "regime complex" governing its development. So far, there is no specially designed international institution of the kind Altman envisages, but existing organizations have developed standards, guidelines, and principles to which AI companies can more or less voluntarily commit. In what is commonly called a turn from soft law to hard law, there are now also

legislative efforts like the 2024 California bill and the EU's 2024 Artificial Intelligence Act, which include formal constraints, monitoring, and economic sanctions for AI developers who do not abide by the rules.

As was the case earlier in other emerging policy areas, like Internet governance, the global AI regime complex is characterized by a lack of central institutions and hierarchies. Different actors have developed partly overlapping or even conflicting legislation. Effective governance requires a legal and institutional framework in which more specific aspects of AI deployment and development may be democratically controlled.

Before these recent developments, much of the discourse on governance was centered on "AI ethics," often developed by well-funded think tanks, tech companies, and international institutions. At first glance, these ethical frameworks seem to promote the ideal of democracy, since they promise the kinds of outcomes that would be expected from democratic governance.

UNESCO's 2021 Recommendation on the Ethics of Artificial Intelligence and the OECD's AI Principles, first adopted in 2019, both claim that governance can effectively achieve social justice by requiring that AI systems follow ethical principles and human rights standards. This has also led to a focus on accountability, stressing the importance of establishing mechanisms to secure public access to key documents governing AI, which in turn could help prevent corruption among decision-makers.

Similarly, the popular notion of "alignment" is often used to describe AI systems that perform in accordance with what their creators prefer, but it also captures the idea that the outcomes of democratic decisions should align as far as possible with people's interests or with what they think is important. Some technology optimists go so far as to suggest that democratic decision-making could achieve this through AI applications that can track citizens' preferences and help experts make informed judgments.

But critics suspect that business-initiated corporate social responsibility efforts could amount to merely performative "ethics washing." A number of objections have been raised against these self-regulatory frameworks, such as a lack of enforceability due to their voluntary nature, their tendency to empower private actors who prioritize profit-driven goals over ethical concerns, and their lack of mechanisms for redress in cases of harm or breaches of standards.

Setting these concerns aside for now, democratic theory leads us to believe that there is an additional pitfall with these efforts: they depend on the assumption that what we could call *output* aspects of democracy can replace or fully compensate for the lack of important *input* aspects. Even if business-led efforts end up perfectly tracking citizens' interests or aligning with their preferences—securing output aspects like accountability—they will always fall short with regard to the input aspect of granting citizens a say in decision-making, and thus cannot lead to a real democratization of AI governance. This raises critical concerns about the inclusivity and legitimacy of the emerging regime complex.

Holding decision-makers accountable for their actions is a virtue in any form of governance. But to meet the threshold of democratic accountability, these decision-makers must somehow have been authorized to make decisions by those who are expected to abide by them. AI governance will be more democratic when people have agency in the form of the ability to approve or authorize decision-makers and political bodies to design and implement regulation. Governance efforts initiated by non-state actors and international organizations may be laudable, but there is an additional layer of democratic accountability in authorized entities—like governments that are empowered through elections and international organizations that are empowered by states—that is essential for ensuring democratic AI governance.

This is intimately related to another important issue relating to the “what” aspect of democratizing AI governance. In the debate over the lack of democratic control in AI governance, there has been a tendency to focus on particular decisions in specific policy domains, such as privacy protection and bias mitigation, and especially AI safety, given the fears that a sufficiently powerful model could pose a threat to humans. Initiatives like the EU AI Act and the OECD AI Principles are typically structured around such policy issues. Various stakeholders, such as states, tech companies, academics, and civil society organizations, are invited to provide input and feedback on regulatory or ethical guidelines.

Though these efforts are valuable for making decision-making more inclusive, the drawback is

that they tend to fall short in addressing the more foundational democratic problem of who gets to decide what to put on the AI regulation agenda in the first place. In his 1961 book *Who Governs?* the political scientist Robert Dahl argued that agenda-setting—the process by which certain issues are prioritized, framed, and given attention in policy discussions—is an overlooked but essential process to consider in assessing the democratic character of a society. Without democratic influence in agenda-setting—influence over which questions are to be decided—any downstream democratization in decision-making remains fundamentally limited and potentially skewed.

To a large degree, agenda-setting shapes what kind of society people want to live in, as well as our common national, supranational, and international institutions. In the context of AI governance, agenda-setting determines which aspects of AI development and deployment are considered problematic and worthy of regulation and oversight. It sets the direction for dealing with AI's societal impact.

---



---

*Giving the private sector a role in AI governance is problematic.*

---



---

## WHO SHOULD HAVE A SAY?

This naturally leads us to the question of *who* should be present in a satisfyingly democratic form of AI governance, and on what grounds. There is a long-standing concern that global governance in general suffers from a democratic deficit. In response, scholars have suggested that apart from states and international organizations, non-state actors such as nongovernmental organizations (NGOs), advocacy groups, and social movements could play a central role. These actors represent citizens' interests and seek to ensure that they are considered in the decision processes of international organizations and institutions. They also function as watchdogs, holding those who wield power accountable. In recent decades, international organizations have opened up and expanded their interaction with civil society groups in several policy areas, such as global environmental and health governance.

With regard to democratizing the global governance of AI, however, we need to consider the fact that many of the most influential non-state actors in this policy domain are not civil society watchdogs, but rather the very multinational AI companies that are subject to this governance. When the

US Department of Homeland Security announced a new AI safety and security board in April 2024, 14 of the 22 members were CEOs of large tech companies. Similarly, the US State Department partnered with Amazon, Anthropic, Google, IBM, Meta, Microsoft, Nvidia, and OpenAI to launch the Partnership for Global Inclusivity on AI in September 2024, with the stated aims of promoting sustainable development and improved quality of life in developing countries, and using AI tools to advance democracy.

There might be good reasons to include AI developers in governance discussions, such as their deep technical expertise, their innovation capacities, and the fact that they are directly responsible for building and implementing AI systems. But giving the private sector a role in AI governance is problematic from a democratic point of view. It grants a few large tech companies disproportionate influence in decision-making, which they tend to use to set agendas and promote governance frameworks skewed in favor of their corporate interests. Even if these non-state actors were to promote the output aspects of democracy, they cannot promote any input aspects, since none of them has received a democratic mandate, through processes of authorization, to make the decisions on behalf of those significantly affected by AI.

The importance of agenda-setting and the role of non-state actors are illustrated by two central concerns in the scholarly debate around AI. One is the possibility that AI models may achieve capacities that allow them to threaten human life and property, ultimately posing an existential risk. The other has to do with near-term risks of AI implementation, such as data privacy breaches and algorithmic bias, as well as broader structural effects, like labor displacement and intensification of socioeconomic inequalities.

There is no principled reason why a sensible discussion could not address both concerns. But scholars from each camp have accused the other of distracting the public from what really matters. And influence in this academic debate arguably translates into agenda-setting power.

Consider the California bill, SB 1047. The fact that it focused exclusively on long-term AI safety issues concerning existential risk could be taken to indicate that one of these camps managed to exercise greater agenda-setting power and influence legislators' understanding of the risks surrounding AI technology. The bill would have required

developers of the largest class of AI models to adopt security measures—not to prevent the risk of bias or economic disruption, but to mitigate the risk that the models themselves might engage in conduct leading to substantial direct harm to humans or the economy, including cyber, nuclear, or chemical attacks. This is likely to strike many as an odd priority for lawmakers, given that these risks are hypothetical, whereas many other harms from AI are already visible throughout society.

In announcing his eventual decision to veto the bill in September 2024, Governor Gavin Newsom did not reject the notion that AI must be regulated to prevent catastrophic risks. Instead, he objected to the bill's focus on large models, arguing that smaller models could pose the same risks. He also echoed the AI industry's mantra, repeated in its lobbying efforts against the bill, that excessive regulation might stifle innovation.

We should note that initiatives like SB 1047, even if they were to be implemented, do not necessarily mark a step toward democratized AI governance. They can also be examples of what happens when a highly invested interest group manages to shape the governance agenda in a way that reflects its particular understanding of what is at stake. This is not to deny that AI safety is important, or that the electorate might start to care more about the issue if it became better informed. The point is that when we ask “who governs” AI and what it means for AI governance to be democratic, we should not simply take decision-making on a set of fixed issues as a given—we should look at who is influencing the public's understanding of the values at stake and the agenda to be decided.

## MAKING GOVERNANCE MORE DEMOCRATIC

In light of our analysis of the “what” and “who” of AI governance, how could we go about democratizing it? In short, we believe the shift from a soft-law approach to hard regulation is welcome from a democratic point of view.

Granted, there is clear value in ethics guidance documents, strategies, and policies authored by intergovernmental organizations, multinational tech companies, and international NGOs. Studies of these documents reveal that they stress key values and principles deeply connected to democracy. But the participatory and authorizing aspects of democracy are missing. Without formal and inclusive processes of decision-making—in which those significantly affected by AI technologies have

a say in the most fundamental matters, such as the overall direction of AI development and its intended role in society—self-regulatory frameworks will not contribute to the democratization of AI governance.

Needless to say, the development of hard law in global AI governance faces many challenges. It is difficult for inherently slow legislative processes to keep up with rapid technological advances. Whereas AI is a global phenomenon, hard law has so far been enacted at national and regional levels. But despite all their possible flaws, international regulations like the EU AI Act have key democratic features. In democratic societies, lawmaking involves elected representatives and public consultations. If they can be protected from the influence of industry interests, and bolstered by civil society engagement, they can play essential roles in establishing a robust overall institutional structure for democratizing AI governance.

It is up to ongoing and future empirical research to determine whether and to what extent voluntary commitments to soft-law regulation prevent the development of stronger mechanisms, or whether they could be advanced in tandem. Hard law can promulgate uniform standards that apply across jurisdictions, reducing the risks of disintegrated regulatory frameworks. International agreements, and the kinds of international institutions

suggested by Altman, could help harmonize rules on AI use concerning cross-national issues like data privacy, surveillance, and algorithmic bias. Moreover, hard law such as international treaties and legislation imposes binding obligations on all parties involved and provides enforcement mechanisms.

Let us end by stressing again why this matters. If what the AI developers are telling us about the technology they are developing is true, it promises to be a powerful new tool that will impact the way we work, live, and interact. It may also reinforce or upend economic cooperation and power relations. Yet AI development is being spearheaded by a tiny minority of the world's population, and shaped much more by their conceptions of what seems to work as a product or profitable business model than by the political preferences of the majority.

Most of us are being invited to engage with AI technology as consumers, but lack influence over AI as citizens. To give people a say over issues with such profound impact on their lives, AI governance ought to be democratized on all levels, including the global level. To achieve the democratization of AI governance, we need to constantly remind ourselves that we are all in the same boat. Even though AI developers are rowing, democratic AI governance can ensure that we all get to decide what direction to steer in. ■