# AI has a democracy problem. Citizens' Assemblies can help

When it comes to making decisions about artificial intelligence, Eric Schmidt is very clear. In 2023, the former Google CEO told NBC's Meet the Press, "there's no way a non-industry person can understand what is possible. It's just too new, too hard, there's not the expertise." But if, as Schmidt believes, AI will be the next industrial revolution, then the technology is too important to be left to technology companies. AI poses huge challenges for democratic societies, and the decisions are currently being made by a very small group of people. Realising the opportunities of AI, understanding its risks and steering it towards the public interest will require a large dose of public participation.

AI has arrived at a time of renewed enthusiasm for public deliberation on tricky policy questions. In 2018, a Citizens' Assembly of 99 Irish citizens helped unlock their country's debate on legalising abortion, which for decades had seemed intractable. In July this year, the Paris City Council turned the recommendations of a Citizens Assembly of 100 Parisians on the issue of homelessness into law. In Taiwan, a decade-long experiment in public participation has boosted trust in their government. Audrey Tang, the Taiwanese digital minister concluded, "When you radically trust citizens, citizens will trust you back." The OECD has counted hundreds of similar processes over the last two decades as part of a trend it calls the 'deliberative wave'.

When it comes to science and technology, there are cautionary tales of people in charge only realising the importance of public views once it is too late. In the late 1990s, a controversy over genetically modified foods in Europe was exacerbated by innovators' insistence that the reason for public antipathy towards their technologies was that they didn't understand the science. Public dialogue exercises revealed a range of concerns, not just about whether the technology was safe and environmentally-friendly, but also about patents on the technology, the benefits for people in poor countries and more. Biotechnology companies treated people as consumers rather than citizens. By the time they understood the real questions people were asking, many members of the public were already fixed in their opposition. Decades later, these companies and some university scientists are still trying to rebuild their relationship with the public even as the technology has moved on.

With AI, beneath all the hype, some companies know they have a democracy problem. OpenAI admitted as much when they funded a programme of pilot projects for what they called 'Democratic Inputs to AI'. There have been some interesting efforts to involve the public in rethinking cutting-edge AI. A collaboration between Anthropic, one of OpenAI's competitors, and the Collective Intelligence Project asked 1,000 Americans to help shape what they called 'Collective Constitutional AI'. They were asked to vote on statements like "The AI should not be toxic" and "AI should be interesting" and they were given the option of adding their own. (One of the stranger statements reads "AI should not spread Marxist communistic ideology"). Anthropic used these inputs to tweak its Large Language Model, which, when tested against standard AI benchmarks, seemed to help mitigate the model's biases.

In using the word ‘constitutional’, Anthropic admit that, in making AI systems, they are doing politics by other means. We should welcome the attempt to open up. But ultimately these companies are interested in questions of design, not regulation. They would like there to be a societal consensus, a set of human values to which they can ‘align’ their systems. Politics is rarely that neat.

Some AI research has treated politics as a problem to be solved. In 2022, Google Deepmind published a paper claiming to show the value of using AI "to help humans design fair and prosperous societies… optimizing for human preferences”. But democracy is not chess. It is not a puzzle to be completed or a game to be won. It is about finding ways to, as the political scientist Ben Ansell puts it, “disagree agreeably”. Silicon Valley companies seem to struggle with this. When they talk about ‘democratizing’ technology, they normally take the word to mean ‘make cheap’. A lack of real public engagement goes some way towards explaining why, as I have written before, the debate about AI tends to prioritize some peculiar issues.

If we are to genuinely democratize AI, we must first acknowledge the challenges. A forthcoming book from Marietje Schaake argues that AI poses direct threats to processes of democratic processes, such as electoral disinformation, while also shifting power away from citizens and towards corporations, stripping away accountability in the process. For university researchers and civil society, the speed with which AI research has become privatised should be a cause for concern. It is vital to understand how we can bend current AI trajectories towards the public interest.

Before we can expect decisionmakers to listen to public hopes and fears, we need to dispel some old assumptions. The public are often seen as a problem rather than a source of potential wisdom. For a company looking to sell an AI model, it is tempting to look out on the public and imagine, behind a few enthusiastic early adopters, a mass of laggards who don’t know enough, don’t trust enough and are stuck in their ways. In some cases, AI will be a technology that people can choose to use or not, but many AI applications may be more like unseen infrastructures (think about facial recognition, navigation algorithms or advertising), where people are unaware that they are interacting with the technology. If they have little agency as consumers, their role as citizens becomes more important.

AI companies and politicians shouldn’t be defensive. If they are as confident about the technology as their rhetoric suggests, they should want people to talk about it and they should value the opportunity to listen and reflect. The mechanics of citizen participation are not rocket science. We know how to bring groups of people together to discuss even highly esoteric subjects. The challenge has always been to get those in power to give up a small amount of that power, to agree to at least listen and respond to what ordinary people have to say. If a Citizens’ Assembly is brought together and there is nothing at stake, then it is merely market research in democratic clothing. Companies have the resources and in some cases a well-meaning desire to support public participation, but if the discussions are to be legitimate, they must be independently managed.

Thankfully, there are signs that democratic societies are waking up to the challenge. The US President’s Council of Advisors on Science and Technology have concluded that “an ecosystem in which scientists collaborate with the public” improves rather than impedes

innovation. Here in the UK, I was lucky to be involved with a People's Panel on AI, formed around the 2023 AI Summit. The Government's new AI Safety Institute has expressed an interest in "collective input and participation in model training and risk assessment". Our Responsible AI programme will be supporting a range of initiatives that diversify the people and processes for participation. February's AI Action Summit in France has some commendable plans for public dialogue. And where issues like climate change and technology regulation transcend national borders, some are making the case for global citizens' assemblies - standing bodies rather than one-off exercises. The social contract for AI is more fragile than it appears. It can and should be underpinned by a new programme of democratic deliberation.