



Can AI help save the world?

Artificial intelligence can be a powerful force for good, enabling numerous technological breakthroughs – but we must ensure the benefits are shared equitably

DOROTHY CHOU
Head of Public Affairs, DeepMind



How do we solve the incredibly complex problems facing our world? There's a huge amount of data out there that can help. But the scarcest resource we have is our intelligence – our ability to convert that information and data into actionable knowledge.

DeepMind's goal is to solve intelligence to advance science and benefit humanity. What we're trying to do is increase the amount of intelligence available, and the capabilities we have, in order to make huge breakthroughs in science.

For example, we created a system called AlphaFold, which predicts the complex 3D shapes of proteins. It used to take scientists years and years to figure out the structure of a single protein. Now we've developed an AI system that can predict the structure of every known protein in the world.

That speeds up science in all sorts of ways. It can help us better understand diseases, potentially leading to new or better treatments. Another interesting application is waste-reducing enzymes. Landfills are full of plastics that are really difficult to break down; so if we can develop new enzymes to do that, the environmental impact would be significant.

We've also worked with Google to automate their data centre cooling, reducing their costs by about 30%. If you think about how that could potentially be scaled across all the data centres in the world – which account for up to 1.5% of global electricity use – that could lead to significant cost and energy savings.

And we're using AI to look at various images like mammograms and eye scans, to help doctors detect abnormalities faster, and over time, to work towards better treatments for patients.

So there are lots of different applications of artificial intelligence that we believe could benefit humanity.

Building more inclusive systems

The other big challenge is figuring out how to ensure that people are benefiting equitably from this technology.

If AI can be built in a way that is inclusive, and takes all populations around the world into account, that would be a really important step forward.

So for example, a lot of the most popular developments in AI right now are around language models. What if we could use these models to catalogue and preserve languages that are increasingly dying out? That would be a significant preservation of culture.

Equally, if you think about how these models are applied, lots of people still find that when they try and speak to their phones or use voice activation, it won't recognise their voice. We don't all use the same language; we don't all speak the same way. So companies that are developing language models need to make an extra effort to be inclusive in the way they develop them.

The issue is that technology will typically reflect who's building the system and who's labelling the data. For example, I was recently playing around with an index of AI-generated imagery. When you type in Scarlet Johansson, you get a pretty accurate version of Scarlet Johansson. But when you type in Serena Williams, you get these fairly brutalist depictions of a black woman. And when you type in Gemma Chan, you just get Asian fetish images.

It's incredibly discouraging to see that that's the way that these systems are

“What we're trying to do is increase the amount of intelligence we have, and the capabilities we have, in order to make huge breakthroughs in science”



being built. But it's a reflection of who's building them at this point. That will have to change over time.

What future do we want?

It's really important to think about – and have a dialogue about – the role AI can play in society.


What are the things that we think can happen in computational time rather than human time (with human oversight)? And what kind of recourse do we have if and when something goes wrong? We need to have real conversations about the ways that AI can help benefit society, and the risks that we need to be able to manage and mitigate.

Ultimately, we need to be thinking through how technology is fundamentally shifting our lives.

There needs to be a renegotiation of the social contract: what do we want this technology to do? How should it be integrated into our daily lives? How will this change social norms? And how can we use policy to ensure more equitable outcomes?

All of this requires building collective imagination around what we want society to look like. After all, that is surely the whole point of democratic institutions and debate: to come to a collective view of how we want the future to look. Then we can think about the role that technology plays in helping us get there.

“There needs to be a renegotiation of the social contract: what do we want this technology to do? How should it be integrated into our daily lives?”



DeepMind's AlphaFold predicts the 3D shape of proteins



However, that's a really hard conversation to have. And right now, politics has become very regressive. We don't have many politicians or leaders giving us a vision of what the future should look like or facilitating a conversation about it. Without that, it's really hard for tech companies to know what role they should play.

A huge leap forward

Humanity's problems have become way too complex for people to solve alone. Most are not just related to one single field: solving them requires us to range across multiple fields. We need more intelligence to knit all this together, understand the patterns and convert all this information to actual knowledge.

In our own policy work, we have two priorities. The first is closing the knowledge gap: to help people understand the applications of this technology in its best form; and then to facilitate collective imagination about the future.

(For example, we recently partnered with the Central Saint Martin's design school to think about what a future robot should look like. Perhaps unsurprisingly, eight of the 10 designs they came up with had to do with the environment and sustainability.)

And the second priority is closing the participation gap: bringing people to the table who traditionally may not have had a seat, so we can talk about how AI impacts their communities.

In the next five years, we'll see the capabilities of this technology take huge leaps forward. It's incredibly important that as we build these capabilities, they are not just optimised for the same people who have always had these privileges.

I'm optimistic that we can get there; but we may need to change the way we think about markets, about policy, about social norms. And that can't be done by any one institution alone: done by any one institution alone. It will require partnerships, and making space for public discourse and dialogue across multiple sectors and fields.

“It's incredibly important that as we build these capabilities, they are not just optimised for the same people who have always had these privileges”

Dorothy Chou leads Public Affairs at DeepMind, an artificial intelligence company. She has spent her career building social justice, ethics, & accountability structures at technology companies. She was previously with Uber, where she was responsible for policy development on issues related to consumer protection, safety, & self-driving cars. Before that she worked in various roles across communications and public policy at Dropbox and Google.

Outside of work, she is working toward a Master's in Bioethics at the University of Oxford and is an angel investor with Atomico, a leading European venture capital firm.



20/30 Visions is a series of interviews with global thought-leaders, exploring how we build a more sustainable and inclusive world in the next decade

BRIDGES
Fund Management

