

Think

new things

Make

new connections

Conference Summary

Transforming democracy: how can democratic states best use AI and thrive?

25-26 January 2024

DITCHILLY

Conference Summary

EXECUTIVE SUMMARY

This conference began with the clear sense that AI is an extraordinary, unprecedented technology for extending scale, pace and reach across multiple domains, potentially enabling citizens, states and private enterprises to both enrich and encroach on freedoms. One of the major challenges speakers saw facing us with the exponential growth expected in this technology over the next few years, and with the expectation that this will offer huge economic opportunities, is balancing the duelling roles of the private and public sector in building AI. Just because we can do something, doesn't necessarily mean we should. A question of managing the art of the possible.

The starting point for this discussion was the understanding that AI is unlike most technologies we have seen before in its breadth, depth and unpredictable impact. There is an inherent tension between the need for regulation (safety, data privacy etc) and the frontier mindset required for innovation. Even if the state merely wanted to be a regulator, even if it were to say there's nothing here for public sector to build, it would still face the challenge that to be a smart regulator in this tech will require at least some knowledge of it. On the other hand, a much more creative role for the state could be in imagining how AI might remake and amplify public services, i.e. working for citizens instead of on them. Examples such as healthcare and the demands of an ageing population were seen as being particularly amenable to AI solutions. This might also allow for new firms to start with an AI-native approach and take incumbents by surprise.

Participants also noted that a major risk factor for liberal democracies would be under-investing in the sector and that this would be a crucial role for the state to play going forward, both in encouraging investment and facilitating its own investment in AI. In the case of the UK, for example, it was suggested that we need a strategic investment plan that concentrates efforts in areas where we have the potential to out-perform. This would necessitate making trade-offs in deciding where funding would be allocated, but such a strategy could also provide a blueprint for other states seeking to maximise the impact of their investment in AI with limited resources. Another challenge in a world where major tech firms are paying top dollar, is how on earth the government can compete. The answer was seen to be by having a sense of mission.

The question at the core of this conference was how democratic states can best use AI and thrive and this sparked lengthy debate over whether AI has any "responsibility" for the state of democracy. In particular, the issues of misinformation and disinformation and the way they will scale over the coming months and years, was seen as one of greatest threats we face from AI in the immediate term. There was also a push-pull noted between the power bestowed by intelligence (and currently mostly concentrated in the hands of corporations), as opposed to the checks and balances on the concentration of power imposed by democracy. The challenge is that leading AI companies are turning their power into political power, so the question raised was how will the state respond to potential abuses of power.

On the other hand, some speakers noted that democracy was a culture rather than a set of institutions and the end goal was to enable the population to determine their political futures, rather than achieving certain political outcomes. There is profound unease that democracy is deeply flawed, but this cannot be blamed on AI. An optimistic take is to view the ways in which AI could help strengthen democracies, perhaps through mediating difficult conversations and political debate.

Another area of concern for the citizen is the impact AI will have on jobs. The IMF reports that 60% of jobs will likely be affected by AI, whether positively or negatively, and so the state must grapple with how it will support citizens whose jobs are impacted. Despite this concern, participants agreed that the answer was not to put the brakes on AI development, but rather to figure out how we unlock value for the public. Education was seen as having a major role to play here, both in providing new skills and development, as well as creating new opportunities. Education will also be vital in teaching people how to interact with AI models effectively and safely, and this is one way to combat or mitigate the effects of disinformation, misinformation and waning trust in the media.

Discussions, for the most part, avoided both “techno-optimism” and “techno-miserabilism” and hewed mostly towards the centre, whilst attempting to focus on concrete issues and solutions.

Context and why this was important

Developments in AI promise tremendous benefits for states that harness them effectively, and the competition for the supporting technologies, talent and infrastructure needed to realise them is intense. AI promises transformational change across society. It seems likely that AI, rather than simply accelerating economic growth or offering new tools for security and statecraft, will bring deeper structural shifts in the way that power is accrued and exercised. This offers both risks and opportunities for democratic states, which must consider how best to protect and enhance democracy, the freedom of the individual and the social contract in a rapidly evolving context. Private sector leadership of these technological developments gives the challenges a new dimension, both for the companies involved, as well as for the governments who must collaborate with them.

People

The conference brought together thought leaders in business, government, technology, think tanks and academia, including from the UK, the United States, France, Canada, Spain, South Africa and Japan. Participants included senior representatives from major tech companies and government agencies, among many others.

Analysis

FULL REPORT

The conference started by discussing what kind of AI-enabled society, democracy and state we should be aiming for, and what the major opportunities and obstacles would be in achieving this vision. Where do we want to be, in other words what are the opportunities offered by AI? And, by contrast, what are the risks? That is to say, considering what can be done versus what should be done. There is danger inherent in the power concentration of AI and the risk of abuse of that power and AI is moving so fast that, with this in mind, we will need the state to provide some kind of guidance and/or advice at some point. How do we balance regulation versus competition? The tension between governance and free markets? AI progress cannot be contained so it is vital that liberal democracies stay in front of it, whether that be in the form of investment, regulation, education or other means.

Machine learning and AI are dramatically improving our understanding in all sorts of areas. They are already able to take problems that have challenged scientists for decades and solve them in a matter of minutes, and the fields affected range from biology to meteorology and everything in between. AI can be transformational in mathematics, as well as in practical applications, such as stabilising plasma in fusion reactors, and it is also accelerating scientific discoveries by allowing hundreds of thousands of academic papers to be read and summarised in days, saving countless time for scientists. In

addition, generative models are producing multimedia content that cannot be distinguished from human-produced content.

Within two years, it was generally agreed that AI models will be smaller, cheaper and more widely available. In addition, interactions with AI will become increasingly natural to the point where it will not seem as though you are talking to a machine. It was also suggested (although not everyone agreed with this) that AI will have developed forms of reasoning within this timeframe and that there will be deeper integration with the human senses. With this exponential development in mind, participants discussed the need to determine what we want an AI-enabled future to look like and what the opportunities were ahead of us.

It was suggested that the state can shape AI, for example, by providing and encouraging investment. Encouraging sufficient capital flows into safety was seen a precondition to much of the broader development of AI. At present, it was felt that insufficient capital has been allocated to developing skills in and safety frameworks for AI. Yet the West is arguably ahead of, for example, China in this area. One risk to major democracies was thought to be states under-investing in AI, and therefore finding themselves lagging behind autocratic competitors.

With regard to further risks, there was a debate about the nature of power and whether more intelligence helps to leverage power. The challenge is that leading AI companies are turning their power into political power, so the question raised was how will the state respond to potential abuses of power. Because AI has overwhelmingly been developed in the private sector of the United States, it has primarily been governed via a corporate lens and therefore governance around it was made with shareholder risk in mind and does not take into account the needs of citizens. There is an inherent tension then between governance and free markets. What kinds of governance tools can handle the issues of societal risk?

And it is not just new forms of power, but also the knock-on effects of human alienation – will it become harder for us as humans to express our purposes as opposed to the purpose of machines? These questions are not new and yet the response on the part of the state and companies over the last decade has been slow. The deeper urge to stoke competition is holding back regulation and this is why there needs to be some coordination across democratic countries. There might also be differences and even conflict between models and views of safety arising from distinctions between and the characters of different democracies, notwithstanding the deeper differences between democracies and authoritarian states.

Another major question broached was if and how AI would change the social contract between the citizen and the state. The IMF reports that 60% of jobs will likely be affected by AI, whether positively or negatively. How then will the state support citizens whose jobs are impacted by AI and how do we navigate these next steps? Participants agreed that it is inevitable that people who work in certain fields (meteorology was given as one example) will find that their jobs change – facing both risks but also opportunities. But, despite this, participants were in agreement that this was not a reason to stop developing AI and in fact could create further and new possibilities. Instead, we must consider carefully how we navigate these next steps. What do we need to do to unlock value for the public?

It was suggested that education had a major role to play here, both in providing new skills and development, as well as opening up new opportunities. Some expressed concerns that using AI might ultimately mean that we lose the ability to think through complex layers of our problems, like a muscle that is rarely used and starts to weaken. In addition, as AI becomes more prevalent, people will need to learn how to interact with it efficiently and safely, for example understanding that whatever comes out of an LLM is not necessarily true.

One participant distinguished between AI that works “for you”, and AI that works “on you”, although others felt that these were false distinctions. “For you” AI is AI that you use with your explicit consent – you choose why, when, and how to use it. Generative AI (hypothetical AI personalised tutors and medical assistants) were suggested as examples. “On you” AI is AI that you are subjected to. A bank declining your loan application by using an algorithm, for example. We can expect to see increasing use of both types of AI in daily life in the next few years. Autocratic regimes could also utilise “on you” AI for censorship and surveillance. Our job as democracies is to emphasise cases that “amplify individual autonomy and self-determination”. Some also suggested AI that works “with you” as another alternative. In any case, AI progress cannot be contained, so it is vital that liberal democracies stay in front. This point was echoed several times by many attendees.

Regarding broader democratic ideals and touching on the debate about open-source AI, the conference asked: Should everyone have access to AI? Access to AI means potential access to the weapons of cyberattack and disinformation and there are reasons why we do not allow people to have dangerous weapons in their homes. This same consideration should apply to AI. Above all else, we need to make sure that such powerful systems do not fall into the wrong hands (the example of a possible model given here was that of nuclear regulation).

It was suggested that AI is perhaps best confined to a procedural role in the relationship between the citizen and state, rather than anything that risks infringing on individual freedoms. However, there are grey areas within this. If AI can use some personal data to, for example, reduce benefit fraud and tackle people smugglers, that might be of benefit for all citizens. But there are also major implications for social mobility – poorer people are less likely to have the fibre internet connection or high-powered laptop that would enable access to new and innovative AI-led systems enabling people to glide through daily interactions with the state and businesses. What then would this mean for global inequality, and not simply within a single country, but across the globe?

The issue of AI colonialism was also discussed in this context. AI is mostly in English at the moment and, because of this, there is a risk that either many countries will be unable to adopt AI, or that AI-producing countries will be pushing the tech influence of their culture onto others. However, some felt there was also an opportunity here to exercise soft power in and encourage collaboration with the Global South.

A lively debate on democracy established that we make the error of treating it as an object of transformation, when in reality it is a culture and not just a set of institutions. It presents a range of freedoms that are pressured in various ways. Transformation is actually a matter of thinking about what democracies can do to enable the population to determine their political futures, rather than about getting certain political outcomes. We often talk as though democracies are thriving and we need to protect them from AI, but in fact democracies appear to be falling apart and an optimistic take is to view the ways that AI could help with this, perhaps through mediating difficult conversations and political debate. There were no concrete answers offered to the problems of disinformation and misinformation, although some speakers thought the whole discussion was predicated on the (false) notion that we currently live (or used to live) in a perfect democracy where objective truths had always been shared. Others pushed back, insisting that there is a qualitative difference between a flawed society that sometimes lies and autocratic regimes that deliberately push mistruths.

AI was also seen as another tool, a more powerful technology, for amplifying and providing state services. But there are challenges in many areas of system design. Encouraging sufficient capital flows into safety is a precondition to much of the broader development of AI. This could be an area where democratic countries can develop a lead because AI regulation is not just a state-level question, but rather concerns a global technology that permeates societies and is provided mainly

by large U.S. corporations producing AI products that operate internationally. Data privacy issues and accountability are now writ large. Deployment of AI in the public sector must be done under a democratic mandate with full ability to interpret and audit decisions. The components of trusted AI are that it should be competent, transparent, responsible and accountable and it must be able to explain its actions and decision-making processes and act to increase citizen capacity. Questions of capacity included the ability of people in government – ministers and civil servants – to understand and use AI effectively as well as a critical mass of experts on hand and taskable by government as is currently being developed in the UK AI Safety Institute. One attendee also suggested that having this safety institute might function as a bit of a Trojan Horse for the UK to build an AI talent pool.

Conference participants split into three Working Groups to consider the state as protector and police officer, the state as framer of the economy and innovation, and political life and the social contract in an AI age.

The state as protector and police officer (defence, security, diplomacy, policing and immigration)

This working group discussed the role of government in AI and what the relationship is and/or should be between the government and the private sector, which has almost a monopoly on AI deployment and development. In this context, the state was seen as having a dual role as protector, i.e. in shielding people, but also as a protector of rights and in empowering people. Participants also discussed the importance of keeping pace with the rapidly evolving AI landscape and said that the state's agility in adopting and implementing AI technologies must be significantly enhanced to match the rapid advancements happening within the private sector. It was strongly felt that fostering open communication with major corporations and acknowledging their need for clear guidance and input from government stakeholders was crucial to achieving this objective.

In the UK context, solving the physical infrastructure challenge with regard to AI was seen as key. Participants felt that the UK's current lack of dedicated AI infrastructure, including centralized CPU resources, presented a significant barrier to the deployment of LLMs, particularly within sensitive sectors like healthcare, due to data location requirements. To overcome this hurdle, a holistic approach akin to the management of the national energy infrastructure was recommended. This would necessitate atomizing the various components of the "AI stack" and guaranteeing government-wide access to each element.

Another obstacle in the UK was seen as procurement timelines, which currently exceed two years in duration, and are therefore demonstrably incompatible with the rapid pace of advancement within the AI domain. The adoption of innovative models, such as the Pentagon's bi-weekly training and testing cycle, utilising unclassified data sets and incentivised participation, was deemed crucial to ensure agility and competitiveness. Attracting and retaining top talent was also noted as crucial to remaining competitive in this domain, with participants saying that current hiring practices within UK government agencies are inadequate for attracting and retaining the specialised skillsets required for successful AI implementation. A comprehensive overhaul of hiring authorities and talent acquisition strategies is essential.

The prevailing discourse surrounding AI regulation often focuses unduly on restrictive measures and control mechanisms, the group noted. While acknowledging the importance of responsible development, they felt that it was crucial to recognise the immense potential of AI to enhance citizen experience. A proactive approach, wherein the government leverages AI for public benefit, could significantly shift the conversation and foster greater public trust. In addition, building public trust in AI necessitates demonstrably safe and ethical applications. Investing in robust safety guidelines and standards, similar to established pharmaceutical regulations, is paramount. And it was suggested that prioritising

transparency throughout the development and deployment of AI systems, including clear communication of potential risks and benefits, would further contribute to public confidence.

The group aimed to focus on practical solutions and, with this in mind, asked how we can help industry and governments work better together in the context of AI. One suggestion was to identify those processes within government that are well suited (or not) to using generative AI. For example, there are 830 processes in government and the top 50 of these occupy 90 percent of all time and resources. How then, the group asked, can the government use AI to do certain specific things? It was also noted that it is easier to weigh up privacy and security risks etc if you are considering them with specific narrow examples to hand. There are case studies of applications in areas that we really care about, like emergency response, proactive health care, government websites being multilingual, policing and national security, and there is a clear public value to this.

Participants said anything that looks like manual cognitive labour — filling out or interpreting forms, reading meeting reports or diplomatic cables — is ripe for replacement with LLMs. One example given was a trial done in several NHS hospitals whereby AI was deployed to write discharge summaries (which include clinical notes, tests performed, prescriptions etc) for junior doctors. A participant explained that this is one of the most tiresome parts of a junior doctor's job and that, typically, any junior doctor will spend dozens of hours a week writing these. What was striking about the results was that it not only saved a lot of time, but the reports drafted by the LLM were consistently better than the human versions. The end-goal with such a deployment of AI would be to give the doctors both more time with patients and less gruelling work hours. On the other hand, concerns have been expressed that it is possible some learning value could be lost by automating this process.

Moving beyond paternalistic approaches to "AI literacy" education, the state must actively engage citizens in open and equal dialogue about AI's implications. This requires fostering environments where diverse perspectives are heard and valued, not simply dictating information from above. Participants also proposed exploring the feasibility and potential impact of an AI bill of rights, which would outline concrete proposals for fairness and fundamental rights protections within the AI domain, and they determined that this warranted thorough consideration. In addition, participants stressed that the ethical responsibility of democratic nations to prevent the proliferation of AI technologies that empower authoritarian regimes cannot be overstated. Close scrutiny of companies operating within sensitive markets, such as Facebook's activities in Vietnam, is essential to ensure alignment with fundamental human rights principles.

The state as the framer of the economy and innovation (economy, research and innovation, and state services)

The discussion in this group was focused on how AI changes innovation and what the state's role is in that. What are the benefits and how are they distributed? What are the current impediments to productivity? Participants also discussed the need to challenge assumptions, such as whether we are prioritising the accumulation of national GDP or improvements for an individual.

Recognising that some middle-weight powers might not have the financial muscle of other more-ambitious ones (such as the United States, China, UAE, Saudi Arabia, India, etc), the group felt that in the case of the UK we need a strategic investment plan to maximise impact. This must concentrate efforts in areas that we have the potential to out-perform, in other words only investing in things we have a chance of winning. Participants noted that this would necessitate making trade-offs and debate ensued over which of a long list of areas should get funding, for example safety, alignment, literacy, expert talent, governance, data management and infrastructure. Ultimately though, this adaptable framework, while informed by the UK-centric discussion, transcends national specifics,

and can offer guidance for any state seeking to optimise its innovation investment with limited resources.

The discussion also highlighted the need for developing comprehensive AI literacy across various groups. This is not just about the developers crafting the tools, but also includes sophisticated users wielding them in their jobs, politicians shaping informed policies and, ultimately, the entire population understanding the impact of AI-driven decisions. Only then, participants agreed, can we justify these investments, ensure an equitable allocation of resources, and reap the strategic advantages AI offers. Education needs a major overhaul, as traditional models are struggling to keep pace with rapid technological advancements and disruptions. Lifelong learning structures, championed by initiatives like Google's training programmes, require further investment and standardised accreditation to maximise their impact. Just as working with Google and Raspberry Pi helped the UK improve its IT curriculum in schools and boost computer science interest, so proactive investments in AI education, from policymakers to tech giants, are crucial to ensure widespread adoption and responsible engagement with this transformative technology.

Returning to strategy, instead of trying to compete with frontier models, the discussion suggested focusing on practical applications with achievable goals. The key, they felt, lay in empowering talent: providing a relatively small number of experts with access to open-source models and good, potentially state-owned, data. This can be achieved through moderate investments for deploying models on our own infrastructure. By fostering this environment, the UK can build competence and density in AI rapidly, even without aiming to be the world leader. In essence, it is about being a "conviction investor" in our own talented data scientists and supporting their work with open-source models, not large multinational companies. This pragmatic approach would emphasise real-world impact, rather than aiming for theoretical dominance.

The discussion also highlighted the need for software products that prioritise adoption and end-user benefit in the state-driven, AI-powered economy. While powerful new tools like AI emerge, legacy systems can hold back crucial work. We need frameworks to assess the efficacy of existing tools and then figure out how to incorporate better ones, while still dealing with the existing issues around the ability to meet deadlines and standards. Beyond model superiority, effective management and widespread diffusion of narrow AI applications will be key to geopolitical success. The focus should shift from mere cost reduction or increased efficiency to tangible improvements for end-users.

On the flipside, concerns were expressed around opening data controlled by the state to private AI developers. Security, exploitation, and short-term profit motives (both of companies and politicians) raise red flags. Valuing and governing data usage remain unsolved puzzles, prompting questions about the public perception of selling data and around the cost of doing this well. Could we use non-exclusive models of data licensing, where more than one recipient can have access and we can have multiple benefits? While promises of innovation and benefits are alluring, careful consideration is needed to avoid corporate abuse and ensure responsible, data-driven progress.

And finally, there was an exhortation to be more Victorian! Let's get back to the peak of Victorian innovation and the mindset that drove it.

The Democratic State (political life and the social contract in an AI age)

This group addressed the different ways in which the development of AI could both undermine and support democracy and, paramount among these concerns, was state abuse of these tools. As the state takes on more AI tools and these tools become realised, the potential for abuse is great. For example, the concept of 'smart cities' is present now in a way that was just conceptual before. When everything is optimised and connected our world will be easier to use and more efficient, but these tools can also easily be abused. We cannot opt out of this system, decreasing the autonomy of the individual.

Some participants pointed to an elevated sense of concern around AI and misinformation around elections, but others said the truth is that there are also benefits to the use of AI in elections. For example, generative AI would help less well funded candidates to communicate with voters. In other words, there are both beneficial and harmful models. Others added that we have been so focused on disinformation that we are not attending enough to inequality and its impact on democratic institutions and on the world at large.

Participants did, however, stress the fact that we must be vigilant about the risks of “digital authoritarianism”. The world is already at risk from increasing authoritarianism even without AI, for example President Trump’s promise of repopulating the U.S. intelligence service with those loyal to him if he returns to the presidency. This concentration of state power is unprecedented. Looking back at history, in the 1750s we had to write rules of governance for a state that became more centralised, therefore participants asked: What does this look like in a modern setting? Do we need a new form of digital institutionalisation? It was agreed that establishing norms and formal constraints is key as they will dictate what will happen in a crisis or wartime, where existing laws do not suffice. Separation of surveillance power, for example, was something worth looking at, with participants noting that we already have some examples, such a judicial overview of wire taps, but we also need new initiatives now.

Interestingly, it was pointed out that India has chosen to go with a de-centralised model in its state use of AI, as opposed to the centralised Chinese model. So, for example, if you are making a payment, the provider checks back with a central database, but the state doesn’t know what payment is being made. It is mediated by the provider and this is an intentional choice. The individual therefore has more autonomy here and data is mediated between private companies and the state. That said, participants pointed out that we should not disregard the risks of abuse in a government led by Prime Minister Modi and that this topic was worthy of more research.

In terms of the UK, participants said that making the country economically viable needed to be a priority. None of these concerns about AI will ultimately matter unless we build a competitive tech industry. It was suggested that the question of focus should be “How do we get an £100 billion dollar software company in Europe?”, and that everything else was essentially a distraction. If we do not have a well-functioning economy, they said, privacy concerns would cease to matter. And having a prosperous society is also an important part of a healthy democracy. We must get the economic engine turning, regulate minimally and take risks.

Part of this effort would be to find a balance between prosperity and regulation. This is not a hard trade-off, participants noted. We can have tech prosperity plus regulation and guidance, and minimum privacy regulation that lets us safely do these things would give our companies advantages compared to other parts of the world. We should set out to do responsible AI that turns a profit, because simply getting rid of privacy or responsibility will not result in a prosperous company.

In addition, we cannot have a vibrant democratic society without education. There are possibilities here to use AI to develop personalised learning, which could be either positive or negative, however if used to the good there is the potential to close gaps and reduce inequality. The opportunities around AI education are great, but they need to be paired with opportunities for meaningful employment. It is dangerous to educate a population but give them no opportunity to exercise that education.

In conclusion, the group asked how we set reasonable expectations for what AI can do, should be expected to do, and cannot or should not do. Participants felt that it was more likely to be a tool that could help us manage, mitigate and improve, rather than “fix” democracy for us. We need to fix the foundations of democracies ourselves and these big institutional fixes will not be achieved by AI. The

inconvenient truth is that democracy is not political technology, it requires personal relationships and we cannot escape the importance of human contact and collaboration, however AI could help frame this.

Looking ahead: Governments must take responsibility as legislators to provide the frameworks and necessary regulation and investment in regulation to protect citizens and enable AI in support of democratic states. It was also suggested there is a clear need to bring AI companies and government representatives together to figure out how to collaborate, since the area of most tension would appear to be balancing regulation with innovation, while also encouraging investment in the sector. (Perhaps, some suggested, Ditchley could act as a facilitator for this.) Another important element going forward is going to be access to cheap renewable energy, so that AI can be made more sustainable. There may also be new firms starting with an AI-native approach that could take incumbents by surprise.

Speakers also cautioned though that this debate is just the beginning and that technologies (like the metaverse) that have seeming been “on hold” as all the major players caught up with Open AI’s ChatGPT will begin to surge and bring with them their own challenges. One speaker noted, for example, that enormous amounts of biodata are captured by these systems and go into existing corporate structures and that it is not at all clear how rights to privacy apply in these instances. There is also rapid development happening in quantum computing. Elements of this will be around faster than people realise and the speaker suggested that one of the areas where this will happen is in gaming and, by extension, one of places that we are failing in the policy space is by seeing gaming as something only kids do.

In the future, some said, the way that we will interact with AI will be in the metaverse and this will raise all sorts of social issues, such as the future of crime and policing in this space. What does it look like when we inhabit a virtual space or have AI inside us? Research on how humans absorb information will inform future information creation and these conversations may be closer than we think, especially in light of the recent announcement from Neuralink, Elon Musk’s neurotechnology company, that the first human had received an implant from the brain-chip startup and was recovering well. In short, we’re in for a turbulent time ahead but, if we can work our way through it, we can get to a stronger and more stable future. So, buckle up!

This summary reflects personal impressions of the conference. No participant is in any way committed to its content or expression.

PARTICIPANTS

Dr Chloé Bakalar PhD

Chief Ethicist, Meta.

Mrs Bojana Bellamy

President, Hunton Andrews Kurth's Centre for Information Policy Leadership.

Mr Yoshua Bengio

Founder and Scientific Director, Mila; Full Professor, Department of Computer Science and Operations Research, Université de Montréal.

Ms Cassidy Bereskin

PhD researcher, Oxford Internet Institute, University of Oxford.

Dr Matthew Botvinick

Senior Director of Research, Google DeepMind.

Mr Mike Bracken CBE

Founding partner, Public Digital. Former founder and executive director, UK Government Digital Service and the UK's first Government Chief Data Officer.

Dr Cristian Canton PhD

Head of the Responsible AI organization at Meta.

Ms Sílvia Casacuberta Puig

Rhodes Scholar, studying Computer Science at the University of Oxford.

Mr Matt Clifford MBE

Prime Minister's Representative for UK Summit on AI Safety; co-founder and CEO, Entrepreneur First.

Mr Kenneth Cukier

Deputy Executive Editor and former correspondent and editor, The Economist.

Dr Catherine Cutts

Principal, KKR & Co. Inc. and former Chief Data Scientist at 10 Downing Street.

Ms Kat Duffy

Senior Fellow for Digital and Cyberspace Policy, Council on Foreign Relations.

Ms Rebecca Finlay

CEO, Partnership on AI and former Vice President, Engagement and Public Policy at CIFAR. A member of the Canadian Ditchley Foundation Advisory Committee.

Ms Kay Firth-Butterfield M.A. LL.M.

CEO, Good Tech Advisory and former Inaugural Head of AI at World Economic Forum

Mr Ben Garfinkel

Director, Centre for the Governance of AI; Research Fellow, University of Oxford.

Mr Reid Hoffman

Co-Founder and Partner, Inflection AI; Partner, Greylock; board member, Microsoft. Former co-founder, Chairman and CEO, LinkedIn.

Mr Ren Ito

Founder and Partner, Solaris Fund Management.

Mr Saul Klein OBE

Co-founder and managing partner, Phoenix Court.

Dr Pushmeet Kohli

Vice President of Research, Google DeepMind.

Mr Michael Kratsios

Managing Director, Scale AI.

Mr Chris Mairs CBE

Co-founder and CTO, Metaswitch Networks.

Mr Ken Manget ICD.D

Former Global Head of Relationship Investing, Ontario Teachers' Pension Plan Board. A Director of the Canadian Ditchley Foundation.

Mr John Marshall

Executive Director, World Ethical Data Foundation and CEO of the World Ethical Data Forum.

Professor Dame Angela McLean DBE, FRS

Government Chief Scientific Adviser and Head of the Government Science and Engineering Profession.

Ms Blaise Metreweli

Director General, Tech and Innovation, Foreign, Commonwealth and Development Office.

Mr Emran Mian CB OBE

Director General for Digital Technologies and Telecoms, Department for Science, Innovation and Technology.

Ms Sam Miller

Co-Founder and Director, Google DeepMind Institute.

Mr Sean Moriarty

CEO, Primer.

Mr Louis Mosley

Executive Vice President UK & Europe, Palantir Technologies.

Mrs Katie O'Donovan

Director of Public Policy for Google UK.

Mr Christophe Prince

Director for Data and Identity in the Home Office.

Ms Renate Samson

Interim Associate Director (Society, justice and public services), Ada Lovelace Institute.

Dr Elizabeth Seger

Researcher, Centre for the Governance of AI (GovAI).

Ms Sienna Tompkins

Analyst, Lazard Geopolitical Advisory.

Professor Alex van Someren FREng FIET

Chief Scientific Adviser, Government Office for Science.

Dr Marc Warner

Founder, Faculty.

The Hon. Robert Wills

Founder and Managing Partner, Collective Capital. A Member of the Council of Management, Finance and General Purposes Committee and a Governor of The Ditchley Foundation.