

Think

new things

Make

new connections

Conference Summary

The impact of AI on work and education

22-23 February 2024

DITCHILLY

Conference Summary

EXECUTIVE SUMMARY

It must be likely that our economies will, before too long, be fundamentally changed by AI technologies. The sheer amounts of money and competitive ambition over the production of microchips alone will have profound economic and geopolitical effects.

Public discourse about how our societies will engage with AI is not sufficiently advanced, neither is the research and analysis in the organisations and institutions tasked with developing public and education policy. This conference agreed on the urgent need to improve understanding and awareness of the uptake of AI and the impact it is already having on economies, labour markets, in education and in areas such as defence and security. How will our societies collectively prepare for and shape the change coming so that potential gains in human capabilities, prosperity and life satisfaction can be made widely available and the negative effects actively mitigated?

Context and why this was important

Since the launch of Open AI's ChatGPT in late 2022, there has been an acceleration in generative AI technology, huge general interest and, inevitably, lots of hype. Last year saw further issues with AI chatboxes as several major tech companies brought out their own models with promises of further and more advanced releases in the pipeline. Concerns about the risks of this technology were also being voiced, especially at the international AI Safety Summit convened by the UK government in October 2023, which led to the setting up of the UK's AI Safety Institute.

Initially, the impact of generative AI for work was thought to be on white-collar jobs as a result of the text and image-generating capabilities, but as advances in AI are made in other areas such as robotics, autonomous vehicles, multimedia output and AI-optimised hardware, it is clear that AI will be integrated across many areas of the economy in forms of automation and as a result of an intensive drive for productivity gains. These forces, unchecked, will inevitably change many kinds of jobs (and activities). The convergence of technologies with uses of big data, cloud computing, IoT, blockchain and others are expected to transform the processes of knowledge and skills acquisition, as well as the ways we work.

There is a pathway for many of these current capabilities to grow at an ever-increasing rate, thereby becoming significantly more powerful over the next couple of years. Massive financial investment and extreme competitive commercial pressure reinforces this trajectory. As the conference observed, we are on an exponential curve and moving fast. Will we stay on, or will we fall off? There are many disruptive forces in the world that could push this trajectory off course or slow its pace, not least the sheer cost of further innovation — ChatGPT-4 cost \$100m, while ChatGPT-5 is said to cost \$2.5bn and ambitions for the global semi-conductor industry are reported to cost around \$7tn. A lack of sufficient global compute power and active political choices to regulate could slow the pace, as might a geopolitical or climate-induced shock. Preparation is advisable.

AI is already beginning to change job roles and skills. But information about how jobs are changing is not being systematically collected, or at least public awareness of current and likely change seems low. Productivity gains, for example, made by companies because of savings on graduate recruitment may seem attractive in the short term, but could turn into shortages in skilled human capital over the longer term. Private sector investment in workforce training is already low; such decisions seem short-sighted and could have wider social effects.

What might the integration of AI-based technologies in work mean for education? Is there a modern-day equivalent of the introduction of compulsory and universal education made in earlier centuries in response to a rapid need for greater skills (in the United States) and to tackle the exploitation of child labour (in the UK)? What should be put in place now to assure the collection of data and evidence with the goal of contributing to the raising of collective literacy on AI and to inform innovation and policy change in education in ways that serve individual and national interests?

AI is already here, and its impacts are being felt in the economy, but the radical change is yet to come. What are the practical steps that can be taken now to prepare? Several actions were outlined:

- 1) Firstly, to gather information and evidence of the ways in which responses to AI are being made within bureaucracies, businesses, schools, universities, research settings, as well as in defence, military, and other settings. There are differences in current responses and learning to be had on immediate outcomes of the choices already made.
- 2) Secondly, to incentivise businesses to share their data on the impact of AI integration and automation to allow wider learning and policy development. AI is already being used within companies and is having direct effects on labour markets.
- 3) Thirdly, to raise public awareness about the impact of AI applications in areas such as virtual assistants, chatboxes and self-driving cars, but also in medical research, forms of data collection and over ethical issues such as privacy. AI is changing the lives of children now and they should be better informed about what AI-curated social media content is and the business models that deliver it.

This conference was framed to bring the issues of work and education together. Education relates to, but is not determined by, the demands of the labour market and the economy. The labour market responds to people's skills and aptitudes as shaped by education. Both are being changed by powerful advances in AI.

People

This conference brought together education leaders, teachers, university professors, serving politicians and civil servants with journalists, social scientists and technology and labour market analysts from the UK, the United States, Canada and Singapore.

Analysis

FULL REPORT

Despite awareness of the hype over AI and the huge uncertainties about rates of adoption and the ways that AI would be operationalised, there was a sense from this conference that the impact of AI would be profound and foundational. However, that's not to say it would be automatic, straightforward or necessarily desirable. We were reminded of the risks of overestimating short-term consequences while underestimating long-term consequences. Full adoption of technologies (such as electricity) can take 30-40 years. (Although LLMs are different in that they can be used straight out of the box by opening a browser window.) Fully enabling AI to reach its maximum potential will likely require 'a re-wiring' of our economy that might take many years. An AI revolution is taking place within the shell of an incomplete or now creaking industrial revolution (balanced on fraying Victorian infrastructure) and, in the case of the UK, in the context of a stagnant economy with low productivity and talent shortages. The change brought about by the integration of AI will likely be patchy and uneven and moving at varying speeds. AI will also be operationalised at a time when our demographics are changing, with ageing populations and low fertility rates bringing structural

challenges for labour markets. Is this also a moment to be ambitious about the potential for change and our ability to shape the progress for individual capabilities, greater prosperity, better health, wealth and quality of life?

Accepting that change is coming, this conference discussion turned on a balance between risk and opportunity. Can potential benefits be shared and structural reinforcement of inequalities avoided? Should we call 'fire' over AI to focus attention on all the short-term risks: job displacement; pressures on job security; skills mismatches; increased social security costs; loss of privacy and intrusion of unethical practices, as well as increases in inequality and regional disparities because of an uneven distribution of the economic benefits. There are also ethical concerns, such as the role of AI in decision-making, for example in job recruitment, or the tensions that might arise if older workers are reluctant to change working practices? What happens to social cohesion if there is over supply within labour markets with downwards pressures on pay?

Alongside a review of potential risks, there was also optimism that over the longer term the integration of AI will bring huge societal opportunities to create many new highly skilled roles (previous waves of technology adoption have led to demand for new skills). People aided by AI would become effective in ways not seen before. For example, would the already highly capable reach new feats of success with super CEOs voraciously devouring multiple senior roles?

Would personalised co-piloting or the augmenting of human abilities reduce or increase the demand for human-delivered work? Some participants had a firm belief that people want to interact with other humans in relationships supporting care, social activities and the nurturing elements of education, and that the demand for human work will remain. The assumption being that parents won't want to send their children to school in the metaverse and that people prefer care homes and cafes staffed by humans rather than robots. Assumptions that may easily erode over time and even, in fact, by next week. That said, if human work is in dealing with unsolved problems (and there are plenty of those in the world), then demand is secure.

If the short-term outlook seems rocky while the longer-term outlook might seem more optimistic, what are the mechanisms to recognise, understand and manage the immediate and intermediate risks? And, if the nature of the skills and aptitudes likely needed for future work are not yet known, how best should our societies prepare? How best can education systems optimise for adaptability in the face of uncertainty? A belief in a foundational role for education seemed a safe bet. Good teaching, discovery, curiosity, joy and the basics of strong independent thinking were considered to be core to modern education, but the need for innovation is nevertheless unavoidable. It would be complacent to assume that the value of a university education as the preeminent signal of high quality and skill will remain unchallenged or that what universities teach and the way they do it won't need to change. Other routes and other credentials are surely coming.

AI is already here, and its impacts are being felt in the economy, but the radical change is yet to come. What are the practical steps that can be taken now to prepare? Several actions were outlined:

- 4) Firstly, to gather information and evidence of the ways in which responses to AI are being made within bureaucracies, businesses, schools, universities, research settings, as well as in defence, military, and other settings. There are differences in current responses and learning to be had on immediate outcomes of the choices already made.
- 5) Secondly, to incentivise businesses to share their data on the impact of AI integration and automation to allow wider learning and policy development. AI is already being used within companies and is having direct effects on labour markets.
- 6) Thirdly, to raise public awareness about the impact of AI applications in areas such as virtual assistants, chatboxes and self-driving cars, but also in medical research, forms of data collection and over ethical issues such as privacy. AI is changing the lives of children now and

they should be better informed about what AI-curated social media content is and the business models that deliver it.

This conference discussion accepted the seemingly unstoppable forces that will drive the operationalisation of AI, albeit in ways that are patchy and uneven. There was debate over the rate and nature of its adoption. Overall, there was more techno-optimism than techno-dystopianism. Even so, the conference ended with important and challenging questions for our societies. Alongside the calls for bridging mechanisms to help manage short-term labour market transitions, there are more profound implications for public policy and for the broad set of educational organisations.

Adoption of AI will not manifest as a single big moment, its impact will be silent in part and more disruptive in other parts but, for the moment, will take place within our existing economic norms, policies and frameworks for tax, business regulation and education systems, so how is the legitimacy of radical action (if that is what is needed) to be created? How will explicit efforts to protect truth, trust and epistemological integrity be properly asserted? AI is not the only disruptive force in play; its integration intersects with the green transition and with uncertain geopolitical dynamics.

In addition to the UK's AI Safety Institute, what other new institutions will be needed? What kinds of skills and training will future civil servants need to increase state capacity in uses of AI for public services? The conference heard about early efforts to create new resources for schools in the form of safe LLMs or an education cloud. What other measures can be taken to protect fundamental liberal democratic content underpinning education? How can the chronic and long-term underinvestment in training by UK business be reversed? Could tax measures provide the same benefits to the self-employed as they do to corporates for investing in training, and can full expensing be offered for human capital investment not just capital investment? Would forms of 'data tax' or sovereign wealth help pay for life-long learning? Will labour protection and other regulation work?

Conference participants split into three Working Groups to consider AI in the world of work, AI capabilities within education, and education for self-fulfilment, society and democracy.

AI in the world of work

There are sections of the economy that are already being affected by the integration of AI; these include financial services, insurance, software engineering, law and medicine. It was considered likely that there would be changes in the ways people deliver particular tasks, and there may also be higher expectations for what people will be able to achieve through their work. The convergence of new technologies with availability of big data, cloud computing, IoT, blockchain and so on was thought likely to have a transformational effect on the workplace. It is already clear that labour productivity gain is being used to ease labour shortages. The conference heard that some employers are considering dramatic reductions in their graduate recruitment, while for more established sectors it was thought that professions with strong professional associations and unions, such as in medicine and law, may be able to protect their working practices. Nevertheless, large areas of work are exposed to AI (60% will have incorporated elements of AI by 2025), especially as compared to previous waves of technology adoption. This will likely affect both high and lower skill sectors. Jobs will be lost, and jobs will be changed. Automation and augmentation will be major factors driving changes in jobs, but also the reinvention and replacement of jobs. These changes will disrupt all sectors, white collar and blue. The next couple of years is likely to be a period of experimentation. Most knowledge workers will likely make use of various AI companions with access to LLMs as a matter of routine.

Demographic change also means that retraining is a priority as a shortage of younger workers increases reliance on older ones. Where is the widespread institutional support for retraining and upskilling older workers? Or will well adapted businesses that prioritise innovation be able to offer shorter workdays and a lowering of the retirement age?

Automation and augmentation will likely amplify human outputs. In other words, humans could achieve much more in shorter time. Will this save time for other activities and support choices to combine raising a family with maintaining a career, or will it drive ever more rapacious expectations and exploitative work practices? Reports from major international economic research institutions highlight a further push towards the gig economy, with the loss of traditional jobs, high unemployment and possibly increases in poverty. What might happen in a potential jobless future? Could there be 'national emergencies' the result of major job layoffs? Would the state become an employer of last resort or does the state adopt AI, become efficient and redistribute in different ways? Will social security effectively become a form of universal basic income?

In the past, low skilled workers tended to work longer hours than higher skilled workers. This has changed with higher skilled workers apparently 'choosing' to work longer hours. These decisions are wrapped up in questions of competition, self-fulfilment and identity.

AI is also likely to be incorporated as a management and surveillance tool and to collect more data on workers. At what point will workers' rights be infringed? The fastest pace of change is expected in countries with less regulated labour markets. And what might happen if automated roles manage non-automated jobs, i.e. human workers, and what will the relationship between managed and manager look like? An automated manager will be able to manage unlimited reports, but accountability might be problematic! (The decades-long experience of UK sub-postmasters increases queasiness over trust in some kinds of tech.)

These predicted scenarios are, however, very different from current realities. Present challenges include shortages of workers in care and health services. It is not clear how AI might provide quick fixes here. Although, over time a nurse paired with AI could have much greater access to medical expertise and an ability to combine greater knowledge with nursing skills.

The record of industrial transitions is not always a happy one. The story within coal mining communities was not one of a successful move from an old to a new industry. Many in these communities have endured long-term structural unemployment felt across generations. What can the state do to mitigate devastating social consequences? Massive retraining, strengthened labour protection, UBI? The role and responsibilities of the private sector also need much more scrutiny. Private sector investment in training has been allowed to fall off. Where does the future balance lie between publicly funded education and privately funded retraining? What are the incentives that would support business investment in human capital and workforce retention?

The question of whether and how AI becomes integrated is dependent in large part on decisions made by the tech companies. The major seven are American. U.S. geopolitical interests are best served by rapidly developing AI within relatively relaxed labour laws. Without a billion-dollar tech company, the UK lacks influence and is a taker of decisions made elsewhere that serve other interests. Will attempts at regulation in the EU affect development in the United States or limit potential within the EU to exert control? The existing major U.S. companies currently have the power to dictate norms and to some extent the distribution of benefits and costs. Although others argued that most of the benefits of technological change are in the end passed on to consumers and that insurgent consumer uses ultimately challenge the power of corporations. How should we think about incumbents vs. insurgents, if incumbents are the major tech companies and the developed world, and the insurgents are new start-ups and the Global South? How might opportunities for open-source AI systems inform these dynamics?

AI capabilities within education

Schools will inevitably use AI, teach about how AI works and develop the teaching of AI skills. Opportunities and risks are everywhere. The use of accredited AI tools can support lesson planning, grading, assessment, and save teacher time. (Will this contribute to teacher retention, higher pay and more job satisfaction?) Risks include increased threats to privacy from insights extracted from

new kinds of data and facilities such as handwriting analysis. Other downsides build on the existing division. Private schools with more resource could quickly outrun the state sector in offering learning platforms, adaptive and personalised learning. State schools without such resources will make use of LLMs. (At which point, the question of their underlying veracity and the world view buried within these models becomes important. The issue of protected education IP is under active consideration.)

Further out, by 2030 other AI-driven capabilities will become available. Use of facial recognition to gain access to buildings, uses of AI to conduct interviews, personality assessments and mental health risk detections alongside much greater surveillance. Already teachers can see who in their class has accessed lesson reading materials. Uses of education data were said to be underdeveloped, especially compared to health data. The data collected by Ed Tech companies, by institutions and by government is not currently being shared. Procurement frameworks were thought to offer a tool to drive data collection to inform policy and should be combined with a national AI plan and data protection.

Whole new forms of experiential learning could become available. VR allows for scenarios and simulations and close to real world learning with new forms of self-assessment of skills related to empathy, negotiation, collaboration and adaptivity. Games can be seen as teaching machines providing multiple streams of information and real-time feedback. Gaming and education have not yet been effectively combined for curriculum-based learning. In general, Ed Tech has a poor record, but it is surely coming, especially if innovations in pedagogic practice are delivered in ways that work with technology.

Overall, the opportunity is to raise educational attainment across the board, but the gap between the lowest and highest achievers also has the potential to widen. At the edges, new models are emerging that could challenge the ancient model of fixed time, variable learning. Synthesis ‘for kids who think for themselves’ is Elon Musk’s sponsored school, offering subscription for access to virtual super tutors. Personalised learning takes away the need for fixed time for lessons. Modern competency-based models allow for fixed learning, but in variable time. Already sustained engagement with LLMs can deliver dramatic results and effective personalised learning. How long will it be before children challenge their parents over the need to go to school? What role then do schools and universities have as centres of foundational learning, microcosms of democratic values and practice in a world where polarisation is intensifying and free speech is under pressure? Schools are an engine of social integration and social mobility, as well as places for discovery and exploration. They are places in which to think critically and creatively and to learn to use technologies to do all these things.

Post-secondary, the relationship between universities and further education colleges has yet to be organised in effective and complementary ways (something that was called for at Ditchley’s Modern Education conference held in 2019). Neglected and underfunded, but in many ways more innovative, FE colleges are well placed to respond in the changing economy and demand for new skills. Collaborative forums with industry already take place with the design of new employer-recognised credentials. Universities can provide educational foundations, whilst FE can provide short-term courses to orientate people for more specific skills.

Universities provide the highest quality foundational learning to prepare people for a lifetime in work and beyond. They remain institutions that provide a signal to employers for recruitment of people with motivation, aptitude and intelligence, and continue to prepare people with foundational skills and for future uncertainty. But they cannot continue without innovation and investment. Complacency is a risk. It was suggested that, over time, AI/machine learning will get better at extracting the tacit knowledge and skills from universities. There must be a renewed commitment to foundational education in a context of uncertainty. Universities will also invest in AI for roles and work related to research, science labs and in extending applications of data science, as well as for

learning resources across the board.

Education for self-fulfilment, society and democracy

The fundamental values of education were reiterated in this discussion. As well as providing a basis in factual knowledge, scientific method and analytical skills, education was said to be about building character, inculcating civic values, and supporting wider critical appraisal of both threats to and improvements of democratic systems. These values need to be protected as AI tools are introduced. Students must be able to trust their institutions. There is already suspicion over potential uses and misuses of new tools. Teaching about AI systems and uses of data and privacy protection will be essential in an education system that could ramp up AI-driven data collection.

Uses of AI could also enhance the basic tenets of education. Augmentation of learning and knowledge processes via AI companions might allow more time for deep thinking, analysis, personal growth, and human interaction. VR might support new forms of scenario-based learning. What kind of new brain capabilities should education be developing? These might relate to assimilating information, interaction with AI agents, ever more advanced mathematical, scientific and philosophical conceptual frameworks, and new methods for knowledge generation. Fundamentally, we will need people who can challenge outputs, apply brakes, and think against the overwhelming bias and self-directed pathways advanced AI systems might take. How can universities as independent entities help society to guard against the potential of AI-driven existential risk? How will universities manage the risks that come if digital intelligence begins to replace biological intelligence?

Whilst the fundamental values of education might be held dear by educators with safe institutional-based careers, an appreciation of what 'self-fulfilment' might mean for students facing precarious economic futures may be different. Within the context of economic insecurity and reduced prospects for home ownership, what kinds of trade-offs, for example, between job satisfaction or home ownership might people make?

The idea of adult or life-long learning was recognised as a 'permanent national necessity' in the UK more than 100 years ago in a 1919 report for the then-government. As an ageing population intersects with the arrival of AI, will the concept of life-long learning finally be properly and substantially renewed? Will student loans be replaced with lifelong learning accounts linked to digital records of employment, voluntary experience, caring roles, credentials, and personalised learning plans?

Finally, inequality both within developed economies and between western and poorer countries is compounded by technology, access to it and uses of it. Will AI help to reduce inequality? A significant proportion of the global population, likely several billion people, do not currently have access to the internet. Can AI be a mechanism that allows parts of the world to catch up? Will it, much like the mobile phone, allow some poorer economies and countries to leapfrog in some areas? There will be divergences between countries that have access to compute power and those that don't. The use of facial recognition in refugee camps that don't have regular provision of electricity is particularly stark, albeit the justification is to support the safety, distribution of food and overall management.

The assessment made in this conference of the impact of AI on work and education was unavoidably shaped by our existing economies, corporate structures, educational and government systems and by our understanding of broader geopolitics. It is difficult to think beyond the world we think we see and inhabit. Much of the discussion was about how much this existing framework could be changed to better manage AI impacts; how far our systems will themselves be changed by AI or whether our systems will adapt and fundamentally deliver outcomes that majorities in our societies can support.

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

Participants

Professor Jeremias Adams-Prassl
Professor of Law and Associate Dean (Research), Law Faculty, University of Oxford.

Baroness Diana Barran MBE
Minister for the School System and Student Finance, Department for Education.

Ms Tamzin Booth
Partner in the Technology, Media and Telecoms (TMT) sector team, Brunswick Group.

Ms Sinead Bovell
Strategic foresight advisor and founder, WAYE.

Mr Frank Bowley
Head, Unit for Future Skills, Department for Education.

Dr Alex Curtis
Head of School, Choate Rosemary Hall.

Mr Dan Fitzpatrick MA (Dunelm), PGCE (IOE Lond), PGDIP
Director, The AI Educator and Thirdbox Ltd.

Dr Mary Flanagan
Sherman Fairchild Distinguished Professor in Digital Humanities and Professor, Department of Film and Media Studies, Dartmouth College.

Dr Michael Fung
Executive Director, Institute for the Future of Education, Tecnologico de Monterrey.

Miss Lizzie Guinness
Head of AI and Digital Strategy in post-16 education, Department for Education.

Jan Hall OBE
Founding Partner, No 4.

Ms Claudia Harris OBE
CEO, Makers.

Ms Brittan Heller
Lecturer, Stanford University; Senior Non-Residential Fellow, Digital Forensics Research Lab, Atlantic Council.

Mr Josh Hillman
Director of Education, Nuffield Foundation.

Ms Rachel Ingram
Founding Partner, Cadmium Partners.

The Rt Hon The Lord Johnson of Marylebone
Executive Chairman, FutureLearn. An Honorary Governor and former Governor, The Ditchley Foundation.

Mr Paul Kett
Senior Adviser and Global Director Education and Skills, PwC.

Miss Resham Kotecha
Global Head of Policy, ODI.

Dr James Kuht MBE
Founder, Inversity.

Mr Stephen Lecce
Minister of Education, Province of Ontario.

Chauncy Lennon PhD
Vice president for learning and work and senior strategy advisor, Lumina Foundation.

Mrs Brittany Masalosalo
Senior Vice President, Chief Public Policy Officer, HP Inc.

Ms Sam Miller
Co-founder, Google DeepMind Institute.

Mr Rajay Naik
Chief Executive Officer, Skilled Education; Chairman, UK Commission on Lifelong Learning.
A Governor and a member of the Programme Committee, The Ditchley Foundation.

Mr Ashley Ramrachia
Founder and CEO, Academy.

Ms Anne Rimmer
Deputy Director, HE Student Finance Policy, Department for Education; Co-Chair, Minister's Advisory Council on Higher Education and Skills, Government of Alberta, Canada.

Mr Nestor J. Rivera
Senior Vice President and Deputy General Counsel, Trust & Privacy, HP Inc.

Mr Matthew Sanders
Education Director, Meta.

Mr Imran Shafi OBE

Director for AI Policy, Department for Science, Innovation and Technology..

Mr Roger Taylor

Advisor to the Responsible AI team, as part of Accenture's Luminary program.

Mr John Thornhill

Innovation Editor, Financial Times.

Professor Stephen J. Toope OC, LL.D., FRSC

President and CEO, Canadian Institute for Advanced Research (CIFAR).

The Rt Hon. the Lord Willetts FRS

President, Resolution Foundation.