University of Southampton

WEB SCIENCE INSTITUTE

Government policy on AI:

What should the next administration focus on?

WSI Position Paper 2024-04
July 2024

Ben Hawes, Wendy Hall

About the WSI

The Web Science Institute (WSI) brings together the University of Southampton's world-class, interdisciplinary, socio-technical expertise in web science, data science and artificial intelligence (AI) to leverage the unique role of online technologies in tackling global challenges. We work to create collaborations within the University and with industry, governments and third sector organisations that bring interdisciplinary socio-technical insights and impacts to the world's most pressing problems.

https://www.southampton.ac.uk/research/institutes-centres/web-science-institute

Web Science Institute Building 32, Highfield Campus, University of Southampton, SO17 1BJ
wsi@soton.ac.uk (DOI: 10.5258/SOTON/WSI-WP011)

About the Authors

Dr Ben Hawes is a technology policy consultant and an associate director at the Connected Places Catapult. As a civil servant, he worked on the 2017 review of Artificial Intelligence, the national Internet of Things Programme IoT.UK, the 2011 review of Intellectual Property and the Digital TV Switchover plan, among other technology policy programmes.

Dame Wendy Hall is Regius Professor of Computer Science, Associate Vice President (International Engagement) and is Director of the Web Science Institute at the University of Southampton.
Wendy is an advisor to the UK government and other governments and companies around the world. In 2023 Wendy was appointed to the United Nations high-level advisory body on AI. Her latest book, Four Internets, co-written with Kieron O'Hara, was published in 2021.

Public policy on digital technology has not been one of the main political battlegrounds in this general election, but the next administration will need to make decisions about it. It is possible for governments to avoid addressing these questions, but that is getting harder, and technology policy challenges emerge faster than they did in the past.

We hope that the new administration engages very actively with technology, because technology-driven change is going to keep impacting the economy and society. The new administration will need to keep building policy and delivery capability if they want technology to play a full part in delivering their policy aims.

Artificial intelligence is going to be one of the important areas to focus on. It has matured as a field of public policy. Compared to several years ago, far more government resources are now being devoted to AI. The government launched a new institute and hosted an international summit on AI safety in November, which has initiated an ongoing series of events. More recently it launched an incubator for public sector AI.

However, AI keeps on developing and generating new questions. Some of the more familiar public policy questions around it have become more pressing but not easier to answer. Large Language Models have surprised many people, including some technology experts and many policymakers, with their performance and ongoing improvement. The speed of development and uptake of related new applications, the range of potential impacts from them, and fear of missing out on potential benefits, all put pressure on governments. Collective UK AI capability is increasingly seen as a key kind of national asset that needs to be built up.

AI probably isn't accelerating towards an imminent super-intelligence as fast as some people say, but change is happening fast, and there is no time to waste, to get the best results from AI for the UK.

The incoming administration will be able to get advice if it wants. Expert organisations, policy thinktanks and bodies representing the interests of different groups have put forward technology manifestos. The tech sector representative body TechUK published A UK Tech Plan: How the next Government can use technology to build a better Britain, and Seven Tech Priorities for the next government. TechUK also helpfully summarised the implications for the tech sector of the manifestos published by Labour, the Conservatives and the Liberal Democrats. Open Rights Group invited all parliamentary candidates to support its manifesto to protect digital rights. BCS, the Chartered Institute for IT, set out aspirations for the next government to "transform society with ethics, education and equity in technology". TechSheCan's manifesto is a ten year vision to close the gender gap in technology sectors to deliver a more equal future.

On AI in particular, the Open Data Institute started a series of policy papers on the Future of Data and AI , the first two on data for training AI models and on intellectual property. The Ada Lovelace regularly publishes analysis and recommendations on social and political questions around AI and data.

Many of the policy issues relating to AI in particular are versions of ones already familiar in technology policy: more specific and more urgent, but not substantially new in kind. These include: gaps in skills and in representation; social and economic adaptation to increasing automation; the market power of technology companies; national competitiveness; achieving successful uptake by the public sector.

Some issues, notably bias and accountability, are more acute in relation to AI. Some questions are less familiar and relate specifically to how AI models are improved, in particular questions around training data and intellectual property. AI safety sits somewhere in between: there is an ongoing dispute as to whether the risks that deserve most attention are new forms of known ones (bias, inequality, lack of accountability) or the more novel threat of human extinction caused by superintelligent AI. It should be possible to pay critical attention to both.

Governance of AI: the view from Parliament

The next administration can also get advice about public policy on AI from Parliament. Just before Parliament dissolved for the election, the House of Commons Science and Technology Committee published its third and final report of session from its inquiry into the Governance of Artificial Intelligence.

Getting the report out at that point gave the Committee the last word in one conversation, and the first in its dialogue with the next administration. The Committee published its interim report in August 2023, and the Government's response to that last November. The final report is in part a response to that response, recognizing the Government's actions and positions to date, and suggesting where it should do more, in the opinion of the Committee. While addressed to the current government, it gives the next government (new or returned) ideas to consider, and perhaps more helpfully, it sets out the categories of challenge that the government needs to focus on.

Unlike many other governments and the EU in particular, the current government largely determined not to create new legislation specifically on AI, preferring to work to ensure that AI conforms with and supports existing rights, laws and regulations. The Committee does not quite oppose that position. Instead, it recommends that Government be better prepared to legislate if other measures are insufficient, explain what would trigger legislation, and report regularly on whether those thresholds seem likely to be crossed.

On regulation, the Committee also recommends improvements to the delivery of the Government's approach, rather than substantial changes of direction. It proposes a swift gap analysis of regulation, measures to manage potential conflict and overlap, and increased resourcing of regulators.

The report is brief about public sector application of AI, mostly recommending improvements to reporting. It does propose additional responsibilities i.AI, the new AI incubator for the public sector, which might be combine with its core objectives.

Aside from the Committee's report, public sector AI is the one area of electioneering where AI is getting mentions, for its potential to help close gaps between straitened resources and growing demand for provision of public sector services. Sometimes this has been more vague aspiration than actionable plan, but a few organisations have put forward practical suggestions. The Tony Blair Institute for Global Change and AI company Faculty published Governing in the Age of AI: A New Model to Transform the State. The report's recommendations on the leadership structure for government AI delivery will be, and should be, a matter for debate, but it is convincing in its description of points where change could deliver results.

Back to the Committee's report, while supporting the AI Safety Institute, the Committee asks Government to clarify whether it is successful in achieving access to test AI models as it is intended to do, and what the outcomes are. It broadly supports the current UK approach as compared to the EU and US, but suggests the value of continuing to learn from others internationally. This leads into overall support for the Governments policy direction on AI, as long as the Committee's Twelve Challenges of AI Governance (first set out in its interim report) are continually addressed.

While any of us might prefer to refine some, add a couple, or combine others, this list works well as a high level description of challenges that government needs to attend to, mostly on an ongoing basis, on AI. On some the Committee makes specific recommendations, on others it summarises the challenges and identifies the organisations which need to deal with them. Below I have summarised the Committee's latest recommendations, with some additional comments.

1) The Bias challenge. The Committee proposes regulatory requirements to submit AI models and tools to independent testing, and to report on measures taken to counter bias in datasets and bias in outputs. They use an interesting analogy for normalising this: "This data should be routinely disclosed in a similar way to company pay gap reporting.". In practice, thresholds and definitions for this reporting could prove hard to establish.

2) The Privacy challenge. The Committee follows the Government approach of putting responsibility onto sector regulators and outlines what that should mean: providing sectoral guidance on balancing privacy with benefits and issuing judgements where that has not been achieved, leading to sanctions and prohibitions. Again, the thresholds will take work to develop.

3) The Misrepresentation challenge. The Committee welcomes the amendment to the Criminal Justice Bill to address AI-assisted misrepresentation, including deepfake pornography, and recommends that if the Bill is not be passed before the election, the next government should take this up. The Bill was not passed, and it is likely that similar provisions on misrepresentation will reappear.

In relation to potential use of AI-enabled misrepresentation in election campaigning, the Committee proposes A cross-Government public awareness campaign. This has

not happened, in spite of a laudable initiative by a group of civil society organisations led by Demos to achieve “a cross-party agreement on their responsible use of generative AI ahead of the election.”

4) The Access to Data challenge. The Committee treats this first as a market dominance issue to be addressed by the Competition and Markets Authority.

It then adds a broad recommendation for a right to challenge automated decision-making: “particularly in areas that could affect the rights and standing of the individuals or entities concerned, such as insurance decisions or recruitment.”

The Committee then recommends development of a public data resource for startups to work with: “This could involve facilitating access to anonymised public data from data.gov.uk, the NHS and BBC via a National Data Bank, subject to appropriate safeguards.” Any action on this from the next administration would be conditioned by broader policies on national AI assets and on AI in the public sector.

The Open Data Institute has published an excellent introduction to the case for measures to deliver transparency about data used to train AI models.

5) The Access to Compute challenge. The Committee recommends addressing this as a national capability question, recommending “a feasibility study into the establishment of a National Compute Cluster that could be made available to researchers and startups.”

6) The Black Box challenge. Again, the Committee underlines the role of the regulators: in their approach to these models and tools, prioritise testing and verifying their outputs, as well seeking to establish—whilst accepting the difficulty of doing so with absolute certainty—how they arrived at them.”

7) The Open-Source challenge. The Committee takes a positive view of the mixed environment, evidenced by investment into open and closed models.

8) The Intellectual Property and Copyright Challenge. This has proved a difficult area to make progress in, with the proposed voluntary AI Copyright cod of practice being abandoned earlier this year. Recommending that the Government take up those difficult discussions again and conclude them, the Committee predicts a settlement of past claims and a licensing system for the future, and emphasises the need for that to operate internationally.

Again, the Open Data Institute has set out the background, key issues and steps that government could take in a paper.

9) The Liability challenge. The Committee recommends positive cross-government and cross-regulator action to develop guidance “to establish liability via statute rather than simply relying on jurisprudence”.

10) The Employment challenge. What action governments take to assess and prepare for the impacts of AI automation on jobs is likely to depend in part on their broader default assumptions about the role of government in changing markets. The Committee recommends “a review into the possible future skills and employment consequences of AI, along the lines of the 2017 Taylor Review of modern working

practices." This could bring in the work already being done to identify when and where in the UK disruption might be acute.

11) The International Coordination challenge. AI is a global technology, and the development of governance frameworks to regulate its uses must be an international undertaking. While supporting international action on AI safety, the Committee regards a "degree of distinction between different regulatory regimes" as inevitable, reflecting geopolitical differences and competition. It advocates against action towards a global AI governance regime, "unconvinced that such a prospect is either realistic or desirable."

12) The Existential challenge. The Committee leans towards the view expressed by many who have been critical of the emphasis on this, recommending that regulatory activity should be focused on here-and-now impacts, and leaving longterm safety to national security bodies, the AI Safety Institute and international fora, such as AI Safety Summits.

Looking ahead

As the dust settles, ministers in the new administration will receive a summary of the state of play on AI policy and legislation from their officials. The Science and Technology Select Committee's reports, developed with input from many experts, should be a good part of that.

There is every reason for the new administration to be optimistic. AI could be applied to serve public interests in very many ways. But that will not simply happen without appetite to adapt, and readiness to engage with everyone developing AI and everyone affected by it. The greatest challenge may be in the management of institutions: how to prepare many different areas of the public sector at the same time, to ensure AI can deliver for everyone.