


Tomorrow comes today: How policymakers should approach AI



Dr Joanna Bryson is Reader in the [Department of Computer Science](#) at the [University of Bath](#).

Artificial intelligence poses a range of challenges to policymakers. As a technology that is now pervasive, it is impacting on democracy, security and the global economy in ways that are not yet well-known to publics around the world – and, being covert, these impacts are generally not balanced against strong political will to shape them with effective policymaking. Equally, it is a field of technology beset by alarmist sentiments that have little bearing on the actual risks which it presents, or may yet present, to humanity.

Thoughtful policymaking will be required over the coming years and decades if AI is to be successfully and productively incorporated into human society – and it was with this in mind that I recently answered [the call by the UK's House of Lords](#) for evidence on AI. My answers to their questions are below, and I will add some concluding notes at the end.

The pace of technological change

1. What is the current state of AI and what factors have contributed to this? How is it likely to develop over the next 5, 10 and 20 years? What factors, technical or societal, will accelerate or hinder this development?

AI is now a pervasive technology. For clear thinking about AI policy it is best to take a very simple, straightforward **definition of AI as any technological artefact that generates action in response to its own perception of context**. With this definition we can see a clear continuous progress from the mechanical governors of the industrial revolution to the “self-learning” systems of the last few years. While machine learning has produced advances that stun us all with their capacity to capture human intelligence, it is important to realise that a)

there is a great deal of precedent for what happens each time technology advances our capacity to compute, and b) that computation is a physical process. This latter is important because it excludes one class of alarmist concerns about AI: that one nation, company, or even machine will suddenly create perfect omniscience and thus dominate the world. In fact, laws of computation are laws of nature, and it is provably intractable to know or foresee everything. Computation is not an abstraction like math; computation requires time, energy, and space for storage of intermediate results.

Having said that, AI is already super-human in many domains and in the next 5-20 years it is quite likely that we will be able to capture and express all of extant culturally-communicated human knowledge with it. Already we are far better at predicting individuals' behaviour than individuals are happy to know, and therefore than companies are happy to publicly reveal. Individuals and parties exploiting this are very likely compromising democracy globally, notably in the UK. There is an incredibly large project here for the social sciences and the humanities as we urgently address the political, economic, and existential (in the philosophical sense) challenges of massive improvements in communication, computation, and prediction.

Again, natural laws of biology tell us to anticipate an accelerated pace of change given the increased plasticity of increased intelligence. Therefore we need to ensure our societies are robust to this increase, with sufficient resilience built into the system to allow individuals to have periods out of work finding a new place in the economy. This requires adequate minimum wages, adequate individual savings, and an adequate civil safety net. The greatest decelerators of this process would be: 1) war – including cyber/stealth war inducing democracies to dismantle their own critical infrastructures and 2) cybersecurity. The government's present policy of outlawing adequate encryption is a severe threat to the UK on many levels, but particularly with respect to AI.

2. Is the current level of excitement which surrounds artificial intelligence warranted?

See above. Basically, yes, it is if anything belated given that AI is already the core technology of the richest corporations on both sides of the great firewall of China, and given the impact on individual security and on democracy. But no, AI itself is not itself a legal or moral actor and will not take over the world on its own, and there is no particular new threat beyond the damage already done and our increasing reliance on a more-easily-assaulted digital/electric infrastructure. I say again because I cannot understate the importance: having backdoors in our encryption is a substantial security error.

Impact on society

3. How can the general public best be prepared for more widespread use of artificial intelligence?

See first my answer to Question 1, which addresses retraining. The most important thing is that we reduce the Gini coefficient so that our population retains (or recovers) its social mobility, and those able to innovate have the freedom to do so and the ability to hire others. The productivity and invention intelligent technology should generate should be sufficient to solve the problems of society providing that the economic and political renovations necessary to handle the new redistribution challenges are made.

I am particularly concerned that we are again, as in the nineteenth through mid-twentieth centuries, in a context of increased inequality and its concomitant political polarisation. We need to remember, as we knew in 1945, that it is in the interest of the elite even more than the rest to have a society sufficiently stable to run nations and businesses. The redistribution we practiced from 1945-1978 was not a (successful) war on communism, but rather a necessary economic tactic to counter the technological innovations of petroleum and early ICT. Late (contemporary) ICT requires even greater innovations in shared transnational regulation; the treaties the EU has been experimenting with are not perfect but they need to be improved and extended globally, because the economy is now global.

4. Who in society is gaining the most from the development and use of artificial intelligence and data? Who is gaining the least? How can potential disparities be mitigated?

It is critical to realise that we have all gained immeasurably from having knowledge at our fingertips. Poor people now have a longer life expectancy than billionaires a century ago. Any talk of “wage stagnation” just tells us how impoverished prices are as an indication of economic value, and how poorly the discipline of economics is serving our society – we need to make massive investment to improve the social sciences. Having said that, and reiterating from Question 3, the current aggravation of the essential political problems of a high Gini coefficient economy, and also of sustainability, must necessarily be addressed because they threaten stability.

Public perception

5. Should efforts be made to improve the public’s understanding of, and engagement with, artificial intelligence? If so, how?

The UK is doing an outstanding job of this, a credit to universities, government, the BBC, *The Guardian*, and the Royal Society. We should maintain this level of investment, and probably offer more – particularly through digital university outreach.

Industry

6. What are the key sectors that stand to benefit from the development and use of artificial intelligence? Which sectors do not?

Artificial Intelligence affects every aspect of life and all sectors. It is essential that we research how to make AI a standard part of software engineering, and introduce software engineering earlier in education even than A-level.

7. How can the data-based monopolies of some large corporations, and the ‘winner-takes-all’ economies associated with them, be addressed? How can data be managed and safeguarded to ensure it contributes to the public good and a well-functioning economy?

Firstly, although data is very important, I believe that the “winner-take-all” nature of Internet commerce is not just about data, but rather about the relatively low (but by no means zero!) cost of transport of the outcomes of computation. Historically, the cost of travel has been a reliable pressure for wealth distribution – you would not go to the best bakery in the world or even in your country, you would have some individualised function of quality times the cost

of travel. New technologies challenge this, whether canals, rail, the exploitation of petroleum rather than coal or wood – each of these innovations required new countermeasures for redistribution.

In addition to this challenge, corporations have learned to evade taxes by bartering in non-denominated ways. Every interaction with Google or Facebook is a barter of information. With no money changing hands there is no tax revenue to support the needs of the global populations facilitating the created value. One of two things has to happen: either we need to find a way to denominate these transactions, or we need to abandon the policy of throttling income with taxation, and turn instead to taxing existing wealth. Although income may be becoming easier to hide, existing wealth is becoming harder than it has been historically, exactly because of the information age. Economic theory shows that it is far easier to design a stable economy through regulating wealth than through regulating income, but historically this has not been practical because of the power associated with wealth. That this is changing now may be one reason that our democracies have been under such risky assault by the extremely wealthy – perhaps they have good reason to fear that their relative advantage will soon be reduced. However, as I said earlier, creating a stabler society and economy by reducing the Gini coefficient and thus ensuring that individuals and corporations cannot destabilise states has the potential to benefit everyone.

Ethics

8. What are the ethical implications of the development and use of artificial intelligence? How can any negative implications be resolved?

Ethics is the set of behaviours a society uses to maintain itself – as such everything I've said above is relevant to ethics. However, I have above particularly focussed on aspects of safety related to economics and democracy, and only briefly mentioned aspects of safety related to privacy and diversity, so will go into more detail on that here. I will not address consent because I lack expertise in that.

The issues I described in answer to Question 1 concerning prediction are exactly the problems of privacy. It is not only that we do not wish others to know about us, we do not wish others to be able to use that knowledge, and for good reason: because they can then manipulate us. Humanity and human innovation have historically depended on individual diversity, which is part of the basis of our notion of dignity. Thus privacy and respect for diversity are both absolutely essential if our society is to prosper, as well as being essential to our individual mental health and wellbeing.

It is important to note that diversity is under assault not only from the misapplication of AI but also from other forms of algorithmisation. I am particularly concerned about the detailed legislation of teaching which reduces the autonomy of individual teachers. This has been generated by a combination of parents' fears of chance events compromising their children's opportunities, and governments' desire to control. In pursuit of equality of opportunity we have generated enforced mediocracy, exactly when what most benefits a citizen is a unique basket of skills, knowledge, and opportunities for insight.

9. In what situations is a relative lack of transparency in artificial intelligence systems (so-called 'black boxing') acceptable? When should it not be permissible?

If you read Frank Pasquale's excellent book *The Black Box Society*, the black boxes emerge not so much from AI (the algorithms or source code) as the unregulated gathering and diffusion of data about people. The current system is hopelessly complex in a way we would never permit for money and other legal obligations. There is no question in my mind that AI and ICT more generally have become sufficiently central to every aspect of our wellbeing that they require dedicated regulatory bodies just as we have for drugs or the environment. However, given that many of these issues have to do with impact on democracy, it is probably not a good idea to have governance only at the national level, since the party in power may well be a beneficiary of any irregularity. Thus I strongly recommend continuing to participate in the EU's world-leading efforts to govern both data and, independently of that, AI.

Please note that saying AI should be subject to regulation and audit is not the same as saying that AI cannot have proprietary intellectual property or must all be open source. Medicine is full of intellectual property, yet it is well regulated.

The role of the Government

10. What role should the Government take in the development and use of artificial intelligence in the United Kingdom? Should artificial intelligence be regulated? If so, how?

Yes, please see my answer to Question 9.

Citizens (or perhaps citizens' advocates, see next paragraph) should be able to trigger audits of software systems when they suspect conditions such as: a) the inappropriate or unauthorised use of data, or b) unfair or unlawful bias. With respect to data, I advocate for the position that data about a person is a part of the person and belongs to that person. It should be used only for purposes to which that person has consented. Government regulation and the possibility of audits should encourage companies to use clear, transparent methods to aggregate data and secure methods to store it. With respect to fairness, it should be possible to demonstrate that decisions execute lawful duties and do not disadvantage on the basis of protected characteristics, nor are they arbitrary. Note that the right to audit does not demand that all code is transparent, symbolic, or open source. What is at issue is effects, so demonstrating valid code is just one possible defence against an audit. Others include: showing that the intelligent system behaves appropriately against a relevant range of inputs, identifying what aspect of an individual's profile produces the contested result, or demonstrating a legitimate source of data that results in an output presumed to be based on inappropriately sourced data.

Ultimately it would be ideal if automated systems were in place to answer any individual's complaint or query, but at least initially it will probably be necessary to require citizens to aggregate some threshold number of examples of suspected misconduct before audit procedures are triggered. Again, automated systems might be used to find related filings, and to provide access to already established explanations. Both governments and NGOs should probably be expected to set up such systems.

Learning from others

11. What lessons can be learnt from other countries or international organisations (e.g. the European Union, the World Economic Forum) in their policy approach to artificial intelligence?

We should continue participating in EU efforts. As I mentioned in the introduction, I have also provided under separate cover my 30-page, fully referenced recommendations to the OECD; I hope their final white paper will also be useful.

In conclusion, AI is not a challenge that policymakers will face 10 or 20 years in the future; it is a technology that is already shaping how individuals lead their lives and how they interact with corporations, government entities and one another. On this basis, the measures that I set out above – be they new powers for citizens, principles for moving forward with the development of policy around AI or suggestions for how we might begin to think differently about this technology – must be read as advice for concerned policymakers today, not tomorrow.

The evidence presented here was originally published on Dr Bryson's personal blog [Adventures in Natural Intelligence](#).

This blog post is part of the [Future Policy Challenges](#) series, a new series of IPR Blogs with a focus on science, technology and innovation that highlights some of the crucial issues policymakers may face in the coming years. Subscribe to the IPR blog to get the latest blog posts, or to keep up to date with our activities, connect with us on [Twitter](#), [Facebook](#) or [LinkedIn](#). You can also follow the hashtag [#FuturePolicyChallenges](#) for more on this series.