Print/online: Rethink government with AI

Policy makers should harness data to deliver government and public services that are responsive, efficient and fair, urge Helen Margetts and Cosmina Dorobantu

People produce more than 2.5 quintillion bytes of data each day --- a thousand times more than the human brain can store. Business is harnessing these data using artificial intelligence (AI) to add trillions of dollars in value to goods and services each year (https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-applications-and-value-of-deep-learning). Amazon dispatches items it anticipates customers will buy to regional hubs before they are purchased. Thanks to the vast extractive might of Google and Facebook, every bakery and bicycle shop is the beneficiary of personalised targeted advertising.

But governments have been slow to apply AI to hone their policies and services. The reams of data governments collect about citizens could one day be used to tailor education to the needs of each child and fit medical treatment to the genetics and lifestyle of each patient. They could help predict and prevent traffic deaths, street crime and the necessity of taking children into care. Huge costs of floods, disease outbreaks and financial crises could be alleviated using state-of-the-art modelling. All these services could be cheaper and more effective than today.

This halcyon future seems rather a distant dream. Governments have long struggled with earlier, much simpler technologies. Flagship policies that rely on information technology (IT) regularly flounder. Former US president Obama's Affordable Care Act nearly crumbled in 2013 when HealthCare.gov, the website enabling Americans to enrol in health insurance plans, kept crashing. Universal Credit, the biggest reform to the UK welfare state since the 1940s, is widely regarded as a disaster due to failing to pay claimants properly and wasting £837 million on developing one digital component, which was swiftly decommissioned (https://www.nao.org.uk/report/rolling-out-universal-credit). Canada's Phoenix payroll system, introduced in 2016, has remunerated 62% of employees incorrectly in each fiscal year since launch (http://www.oag-bvg.gc.ca/internet/English/parl_oag_201810_00_e_43161.html). And MyHealth Record, Australia's digital health records system, saw more than 2.5 million people opt out by the end

of January 2019 over privacy, security and efficacy concerns (https://www.abc.net.au/news/2019-02-20/my-health-record-opt-outs-top-2.5-million/10830220).

Failures like these matter. Technological innovation is essential for the state to maintain its position of authority in a data-intensive world. The digital realm is where citizens live and work, shop and play, meet and fight. Prices for goods are increasingly set by software. Work is mediated through online platforms such as Uber and Deliveroo. Voters receive targeted information via social media. Core tasks of governments, like enforcing regulation, setting employment rights and ensuring fair elections thus require understanding data and algorithms.

At [subs: authors point out that T in The should be capitalized>The Alan Turing Institute, the UK's national institute for data science and AI, we are working with governments to identify how they might use digital technologies to improve policy-making and services. Here we highlight the main lessons.

Responsive governance

Policy-making processes were designed in an era before citizens generated digital trails of data. Governments rely on 'official statistics' and other custom-built data, collected through traditional channels such as national statistical offices or surveys. They have no tradition of using transactional data about citizens' actual behaviour to improve policy or services.

But now governments' interactions with citizens generate vast amounts of digital data. For example, vehicle licencing authorities have databases containing information about the cars we own, how often we get stopped by the police, how many accidents we have, whether we pay our road taxes on time and when we obtained (or lost) our licenses.

AI could harness data about citizens' behaviour to enable government in three ways.

First, personalized public services can be developed, adapted to individual circumstances. Just as data can be used to target advertising to fine grained groups, they can help resources be targeted efficiently. For example, 'precision medicine' tailors treatments to an individual's genetic make-up or condition, while personal educational plans could be drawn up for each

child entering secondary school. These would be based on the kind of teaching that produced the best results for 'children like them' in the past --- in terms of educational performance, learning style and even personality.

Second, AI enables governments to make more accurate predictions and forecasts, helping them to plan. Machine learning algorithms 'learn' by identifying patterns in data and use them to predict future trends or events. Some UK local authorities are using predictive analytics to anticipate future needs in areas such as homelessness, special needs education, emergency services, health and safety inspections, adult social care and children's social services.[i] Similarly, schools could be singled out for inspection (as they are in the Netherlands); health organisations could predict where surgery for joint replacements would have the greatest impact.[ii]

More controversially, forecasts can be applied to individuals. Machine learning algorithms might pinpoint which children will drop out of school, or be deemed 'at risk,' allowing authorities to target scarce resources. Such an 'early warning system' is already in use in the US, New Zealand and Canada and is under consideration in the UK[iii].

Third, governments could simulate complex systems, from economies[iv] to military operations. This would enable them to experiment with different policy options and spot unintended consequences before committing to a measure.

Agent computing models, in combination with big data, can capture the complexities of the real world more ably than before and are beginning to be used for testing policies and interventions. For example, the Bank of England in the UK is modelling the housing market and simulating the effects of policy measures aimed at mitigating financial risk. The US federal government is assessing the impacts of potential disasters, such as a nuclear bomb exploding in the heart of Washington DC (https://www.sciencemag.org/news/2018/04/what-if-nuke-goes-washington-dc-simulations-artificial-societies-help-planners-cope). And policy advisors in Mexico are using agent computing models to identify the policy areas that the federal government needs to prioritise in order to reach the UN sustainable development goals (https://www.turing.ac.uk/research/research-projects/policy-priority-inference).

AI challenges

Making AI mainstream in government still has far to go, as a recent trial illustrates.
In 2017, the London Metropolitan Police tested AI-based facial recognition technology at a carnival to identify people on their 'wanted' list. The technology examined 1 million visitors and found 35 matches. Human reviewers ruled out 30. Police stopped the remaining 5 but only one match turned out to be the expected person. To make matters worse, the police then realized that the list was out of date, so this individual was no longer 'wanted'. (https://bigbrotherwatch.org.uk/wp-content/uploads/2018/03/Big-Brother-Watch-briefing-on-Facial-Recognition-for-Short-Debate-in-the-House-of-Lords-1-March-2018.pdf).

This failure illustrates the challenges of using AI in government. First, the AI technology used by the Metropolitan Police had worryingly low accuracy rates. Police forces -- and policy makers more generally -- will struggle to build high-performing, top of the line machine learning and AI applications for the same reasons that they struggled with earlier digital systems. These include lack of in-house expertise, inability to pay salaries that match the private sector, difficulties in evaluating the work contracted out to private providers, and cultural barriers amplified by past IT disasters[v].

Second, the stakes are high for government. When Facebook misses the target with a personalised ad, there are few ramifications. But public trust is eroded when government technological projects fail, limiting the ability to govern effectively in the future. For example, the use of patient data to improve healthcare generally has wide support. But in the wake of government failures, patients in countries from UK to Australia are withdrawing consent for their health data to be collected, with serious ramifications for medical research.

Third, the use of AI by public bodies brings calls for transparency. The Metropolitan Police did not inform the public that facial recognition was in operation and the data were collected and stored without consent. Transparency is crucial for assuring public trust. Processes such as citizens' juries are being used to understand public attitudes to AI. These bring in a cross-section of the public to consider questions like: 'would you like to be given an explanation of how the computer reaches its diagnosis even if that requirement makes the diagnosis less accurate?' (https://gmpstrc.wordpress.com/2019/02/27/should-artificial-intelligence-give-reasons-for-decisions-even-if-it-affects-accuracy-citizens-juries-deliberate/).

Fourth, policy makers need to decide when is it appropriate to use AI-based predictions to make decisions about individuals. Targeting large crowds of law-abiding citizens with facial recognition software to pick out a handful of criminals may be inappropriate as well as costly and labour-intensive for such a speculative task. As policy makers roll out similar technologies across sectors, new moral dilemmas will appear. What should a school do, for example, with a statistical probability of 60% --- or even 98% --- that a pupil will drop out of formal education? Should the school invest more resources in that child, or less?

Fifth, when the Metropolitan Police trialled the AI-based facial recognition technology in 2017, they had not tested it for racial bias (https://www.london.gov.uk/questions/2017/3632). This was despite clear indications, even at that time, that facial recognition algorithms were less accurate for blacks and ethnic minorities (https://oversight.house.gov/sites/democrats.oversight.house.gov/files/documents/032217%20EEC%20Opening%20Statement%20Facial%20Recogniton%20Technology.pdf). Ignoring biases when designing AI-based applications risks perpetuating them. Centuries of over-policing marginalised communities means that ethnic minorities are disproportionately represented on police forces' wanted lists. Facial recognition technologies' reliance on these lists to identify suspects, combined with the algorithms' lower accuracy when analysing non-white faces, will reinforce past biases.

So there are challenges for government to overcome. But there is reason for hope. Data-intensive technologies lay biases bare. This may force policy makers to start facing up to some deeply embedded societal issues that have nothing to do with the technology itself – such as the systematic bias shown in judicial decision-making over many decades before AI was on the scene[vi]. The data needed to measure such bias has not habitually been collected. For example, the UK courts system has tended not to record personal characteristics such as age, ethnicity, gender and disability. But in January 2019, in response to a Public Accounts Committee review of its £1 billion programme to modernize the courts, the Ministry of Justice pledged to 'do more to collect data on the protected characteristics of those who use the courts and tribunals in a way that will make it far easier to identify and tackle disproportionalities' (page 2 para 6 of Ministry of Justice, Evaluating our Reforms. Response to PAC recommendation 4, January 2019 https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/775588/Public_Accounts_Committee_Recommendation_4_31_Jan_2019pdf.pdf).

Next steps

While tech giants like Google, Amazon or Facebook are at the forefront of AI development in the public eye, independent academic researchers are better placed to help government maximise the potential of these technologies. Institutions developing AI across the world (such as the £1 billion Ethics and Governance of Artificial Intelligence Fund, led by MIT and Harvard University, or Stanford's Human-Centred AI Initiative) should introduce policy makers to the latest research and work with them to solve long-running policy problems.

At The Alan Turing Institute, we are using machine learning to identify offenders and victims of crime, from modern slavery to hate speech and radicalisation, to help policy makers to measure the scale and scope of the problems and to build counter-measures. We are using agent computing to simulate different levels of demand for police services and to tailor resources accordingly. And we are running citizens' juries, together with the Information Commissioner's Office, to develop guidance for explaining algorithmic decision-making.

Governments need to develop ethical frameworks for using AI, concurrently with the technology; ethics is part of the science in data science. Institutional development is essential here – after all, there are positive precedents, such as the Nuffield Council for Bio-ethics and the Human Fertilisation and Embryology Authority, which have engendered greater trust in these technologies that might otherwise have been the case. This is the rationale behind the creation of the UK government's Centre for Data Ethics and Innovation and the Ada Lovelace Institute, as well as private bodies such as the US-based Partnership on AI.

The payoffs to policy makers using data science and AI go well beyond cutting costs and making government more citizen-focused. The biases that machine learning technologies have revealed have existed for centuries within our governance systems. By revealing them, data-intensive technologies may also offer a way to overcome them. We hold technologies to a higher standard of account than we do humans; we expect driverless cars to be safer than those driven by people, for example. As a society, we might accept less bias in a system of government that uses AI. In this way, a data-driven government might actually be more fair, transparent and responsive than the human face of officialdom.

AUTHOR BIOGRAPHIC DETAILS:

Professor Helen Margetts, Programme Director for Public Policy at The Alan Turing Institute and Professor of Society and the Internet at the Oxford Internet Institute, University of Oxford

Dr Cosmina Dorobantu, Deputy Programme Director for Public Policy at The Alan Turing Institute and Research Associate at the Oxford Internet Institute, University of Oxford

The Alan Turing Institute, British Library, 96 Euston Road, London NW1 2DB, United Kingdom

REFERENCES:

[i] Bright, J, Ganesh, B, Vogl, T & Seidelin, C. (2019). *Data Science for Local Government: Challenges and Opportunities. Oxford Internet Institute.*

[ii] Kleinberg, J., Ludwig, J., Mullainathan, S. & Obermeyer, Z. *American Economic Review* **105(5)**, 491-495 (2015).

[iii] Cuccaro-Alamin, S., Foust, R., Vaithianathan, R. and Putnam-Hornstein, E., 2017. Risk assessment and decision making in child protective services: Predictive risk modeling in context. *Children and Youth Services Review*, *79*, pp.291-298.

[iv] Axtell, R. L. in *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems* (IFAAMAS, 2016).

[v] Margetts, H. *Information Technology in Government: Britain and America* (Routledge, 1999).

[vi] Spohn, C. 'Thirty years of sentencing reform: The quest for a racially neutral sentencing process', *Criminal justice* 3 (2000): 427-501; Spohn, C., Gruhl, J. and Welch, S. 'The effect of race on sentencing: A re-examination of an unsettled question.' *Law & Soc'y Rev.* 16 (1981): 71.