

COMMENT

PHYSICS Lee Smolin, no fan of quantum physics, offers some alternatives **p.166**


HISTORY Graphic novel of Humboldt's adventures **p.168**



CANCER An immunologist's memoir recalls decades of hope and disappointment **p.169**

NEUROSCIENCE Sex differences in brain diseases need to be studied **p.171**

SUZANNE KREITER/THE BOSTON GLOBE/GETTY



Artificial intelligence could one day be used to tailor education to the needs of each individual child.

Rethink government with AI

Policymakers should harness data to deliver 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. Businesses are harnessing these riches using artificial intelligence (AI) to add trillions of dollars in value to goods and services each year. 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 personalized targeted advertising.

But governments have been slow to apply AI to hone their policies and services. The reams of data that governments collect about citizens could, in theory, be used to tailor education to the needs of each child or to

fit health care to the genetics and lifestyle of each patient. They could help to predict and prevent traffic deaths, street crime or 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 of these services could become cheaper and more effective. ▶

► This dream seems rather distant. Governments have long struggled with much simpler technologies. Flagship policies that rely on information technology (IT) regularly flounder. The Affordable Care Act of former US president Barack Obama 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 because of its failure to pay claimants properly. It has also wasted £837 million (US\$1.1 billion) on developing one component of its digital system that was swiftly decommissioned. Canada's Phoenix pay system, introduced in 2016 to overhaul the federal government's payroll process, has remunerated 62% of employees incorrectly in each fiscal year since its launch. And My Health Record, Australia's digital health-records system, saw more than 2.5 million people opt out by the end of January this year over privacy, security and efficacy concerns — roughly 1 in 10 of those who were eligible.

Such failures 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 — and disinformation — through social media.

Thus the core tasks of governments, such as enforcing regulation, setting employment rights and ensuring fair elections require an understanding of data and algorithms. Here we highlight the main priorities, drawn from our experience of working with policymakers at The Alan Turing Institute in London.

RESPONSIVE GOVERNANCE

Policymaking processes were designed in very different times. Governments rely on custom-built data, collected through national statistical offices or surveys. They have no tradition of using transactional data about people's actual behaviour to improve policy or services.

Today, governments' interactions with citizens generate trails of digital data. For example, vehicle-licensing authorities have databases containing information about our cars, 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 driving licences.

AI could harness data about citizens' behaviour to enable government in three ways. First, personalized public services can be developed and adapted to individual circumstances. Just as data are used to target advertising in a fine-grained way, similar methods can help resources to be



Police officers watch the crowds from the rooftops at the Notting Hill Carnival in London in 2017.

targeted efficiently. For example, a government platform might personalize services according to your personal details and past interactions with the state, as is happening in Queensland, Australia. And in New Zealand, the mobile app SmartStart provides personalized information from across all government agencies to expectant parents, allowing them to fill out forms and apply for a birth certificate from their mobile phones.

Second, AI enables governments to make forecasts that are more accurate, helping them to plan. Machine-learning algorithms identify patterns in data and then use them to predict future trends or events. Some UK local authorities are experimenting with the use of analytics to anticipate future needs in areas such as homelessness, emergency services and social services¹. For example, machine-learning models can simulate future demand for special-needs education, and how that varies if policy or other external factors change.

AI can also be used to target health and safety inspections rather than using randomization. The health department of Las Vegas, Nevada, working with the University of Rochester in New York, has used social-media data and machine learning to identify restaurants associated with food poisoning. The researchers estimate that their system could prevent more than 9,000 cases of food poisoning and almost 560 hospitalizations in Las Vegas each year².

More controversially, forecasts can be applied to individuals. Machine-learning algorithms might pinpoint which children are likely to drop out of school or be

deemed at risk on the basis of data about their previous interactions with public-sector agencies. This would enable authorities to target scarce resources. Such an early-warning system is already in use in the United States and New Zealand, and one is under consideration in the United Kingdom³.

Third, governments could simulate complex systems, from military operations to the private sectors of entire countries⁴. This would enable governments to experiment with different policy options and to spot unintended consequences before committing to a measure.

Agent computing models, in combination with large-scale 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 is modelling the UK 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. And advisers in Mexico are using an agent computing model⁵ to identify what the federal government needs to prioritize to reach the United Nations Sustainable Development Goals.

AI CHALLENGES

Making AI mainstream in government still has far to go, as this recent trial shows.

In 2017, the London Metropolitan Police tested a facial-recognition algorithm at a carnival that attracts 1 million visitors

to identify people on its ‘wanted’ list. The technology flagged 35 ‘matches’. Human reviewers ruled out 30. Police stopped the remaining five. Just one turned out to be the expected person. To make matters worse, the police then realized that the list was out of date and this individual was no longer wanted in relation to a crime.

This failure illustrates five challenges. First, the technology used had worryingly low accuracy. Police forces — and policymakers more generally — will struggle to build top of the line machine-learning and AI applications for the same reasons that they struggled with earlier digital systems. These include a 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 disasters⁶.

Second, the stakes are high for governments. When Netflix misses the target with a film recommendation, there are few ramifications. But trust is eroded when public-sector projects fail, limiting the ability to govern effectively in the future. For example, the use of data from individuals to improve health care generally has wide support. But in the wake of government failures, such as the UK National Health Service agreeing in 2017 to hand over patient information to the immigration authorities⁷, individuals in many countries are withdrawing consent for their health data to be collected. This has serious consequences for medical research.

Third, the use of AI by public bodies brings calls for transparency. The Metropolitan Police did not release information about how many carnival-goers were aware that facial recognition was in operation, nor did they release details about how the data were collected and stored. Transparency is crucial for assuring public trust. Processes such as citizens’ juries are being used to understand attitudes to AI. These bring in a cross-section of the public to consider questions such as: ‘Would you like to be given an explanation of how the computer reaches its diagnosis, even if that requirement makes the diagnosis less accurate?’

Fourth, policymakers need to decide when it is 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 might be inappropriate, as well as costly and labour-intensive for such a speculative task. When policymakers roll out similar technologies across sectors, new moral dilemmas will arise. For example, what should a school do with a statistical probability of 60% — or even 98% — that a pupil will drop out of formal education? Should the state invest more resources in that child, or less?

Fifth, when the Metropolitan Police trialed the facial-recognition tool in 2017, it had not tested it for racial bias. This was despite clear indications, even at that time, that such algorithms were less accurate for black people and individuals from minority ethnic groups. Ignoring biases when designing AI applications increases the risk of perpetuating those biases. Centuries of over-policing in marginalized communities also means that some groups are disproportionately represented in policing data. The technology’s reliance on such lists to identify suspects or to target patrols⁸, combined with the lower accuracy of algorithms when analysing the faces of people of colour, is likely to reinforce over-policing of black and minority ethnic groups.

Nonetheless, there is reason for hope. Better use of data could force policymakers to start facing up to some entrenched societal issues that have nothing to do with technology. One such is the systematic bias shown in judicial decision-making over many decades before AI was on the scene^{9,10}. The data needed to track such bias have not habitually been collected. For example, the UK courts system has tended not to record protected characteristics among users of courts and tribunals, such as age, ethnicity, sexual orientation or

“We hold some technologies to a higher standard than we do humans.”

disability. In January, in response to a 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.”

NEXT STEPS

Although tech giants such as Google, Amazon and Facebook are at the forefront of AI development in the public eye, independent academic researchers are better placed to help governments to maximize the potential of these technologies. Institutions developing AI across the world should introduce policymakers to the latest research and work with them to solve long-running policy problems. Examples of these include The Alan Turing Institute in London, UK; the Stanford Institute for Human-Centered Artificial Intelligence in Stanford, California; and the Ethics and Governance of Artificial Intelligence Initiative, led by the Massachusetts Institute of Technology and Harvard University, both in Cambridge, Massachusetts.

At The Alan Turing Institute, we are using machine learning to identify offenders and victims of crime in areas ranging from modern slavery to hate speech and radicalization.

We aim to help policymakers to measure the scale and scope of these problems, and to build countermeasures. 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 UK Information Commissioner’s Office, to develop guidance for explaining algorithmic decision making.

Governments need to develop ethical frameworks for using AI. Institutional development is essential. There are positive precedents — in the United Kingdom, examples include the Nuffield Council on Bioethics and the Human Fertilisation and Embryology Authority, which have built trust in technologies such as stem-cell therapy. This is the rationale behind the creation of the UK government’s Centre for Data Ethics and Innovation, the Ada Lovelace Institute in London and private bodies such as the US-based Partnership on AI.

The pay-offs for policymakers using data science and AI go well beyond cutting costs and making government more citizen-focused. The biases revealed by machine-learning technologies have existed for centuries in governance systems. By laying them bare, data-intensive technologies could offer a way to overcome them. We hold some technologies to a higher standard 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. ■

Helen Margetts is programme director for public policy at The Alan Turing Institute in London, and professor of society and the Internet at the Oxford Internet Institute, University of Oxford, UK. **Cosmina**

Dorobantu is deputy programme director for public policy at The Alan Turing Institute in London, and research associate at the Oxford Internet Institute, University of Oxford, UK.

e-mail: hmargetts@turing.ac.uk

- Bright, J., Ganesh, B., Seidelin, C. & Vogl, T. *Data Science for Local Government* (Oxford Internet Institute, 2019).
- Sadilek, A. et al. *AI Magazine* **38**(1), 37–48 (2017).
- Cuccaro-Alamin, S., Foust, R., Vaithianathan, R. & Putnam-Hornstein, E. *Child. Youth Serv. Rev.* **79**, 291–298 (2017).
- Axtell, R. *Handb. Comput. Econ.* **4**, 157–213 (2018).
- Castañeda, G., Chávez-Juárez, F. & Guerrero, O. A. *J. Econ. Behav. Organ.* **154**, 335–361 (2018).
- Margetts, H. *Information Technology in Government: Britain and America* (Routledge, 1999).
- Gulland, A. *Br. Med. J.* **356**, j911 (2017).
- Shapiro, A. *Nature* **541**, 458–460 (2017).
- Spohn, C. C. *Crim. Justice* **3**, 427–501 (2000).
- Spohn, C., Gruhl, J. & Welch, S. *Law Soc. Rev.* **16**, 71–88 (1981).