

Artificial intelligence and government

Emma Martinho-Truswell and Sabrina Martin

Artificial intelligence offers a daunting challenge for government: a potent mix of high expectations, big budgets, a competitive international “AI arms race”,¹ and elevated public fears. Some governments, such as China, France and Canada, are determined that their countries be leaders in AI; others have been slower to make public policy statements but may be watching to see where their country might best fit in a field that is moving quickly.

An important question in considering how governments should respond to the acceleration of AI technologies is to ask whether they need to do anything at all. After all, many technological developments have useful social impacts without needing government intervention, and their adoption and safe use are largely managed using market mechanisms. This can be true even for technologies that dramatically change the way we live and work. While tools like washing machines, air conditioners and elevators transformed the lives of women, the comfort of homes and offices and the shape of our cities, none required substantial government involvement to manage their widespread use.

Some technologies, on the other hand, have such a dramatic impact on the way we live and the way we work that they create challenges that governments are best placed to solve. Many of these are infrastructure needs well served by

governments: railway lines for trains, roads for cars, and public wires for electricity and telephones. For other technologies, the role of government has been in helping to smooth changes, or to adopt technologies itself. As the use of personal computers increased rapidly in homes and offices, many governments not only invested in computers to improve their own operations but paid for schemes to ensure classrooms would have computers to teach basic word processing skills to a new generation of workers.

Artificial intelligence technologies are not washing machines: they will not go smoothly into people's lives while governments watch with interest. To ensure that the power of AI technologies is used to help make lives better for citizens, governments need to act. Decisions facing governments include:

1. What encouragement does a government want to provide to incentivise AI innovation to occur within its country, region, or city (rather than in other places)?
2. How can AI be used by governments to improve services?
3. What infrastructure needs to be in place for AI to work well?
4. How can governments help ensure that the technology is applied in an ethical way?
5. What does government need to do to ensure their workforce is ready for the changes ahead?

These questions fit into two overlapping categories, which might be considered government's "offensive" and its "defensive" roles related to the rise of AI.

Of these, many governments with AI strategies have so far focused on the "offensive" aspects of policy, hoping to ensure

their country or region becomes an early adopter of AI. Policies promoting AI research and development, supporting startups, providing faster internet, or giving funding for students to learn AI-related skills are all examples of this kind of approach. It is helpful in developing an important set of technologies, and those governments adopting these approaches are savvy. But promoting AI is not an essential task for governments.

On the other hand, every government will need to work out a “defensive” approach to challenges caused by AI. Governments should minimise harmful uses of AI, and be prepared to support citizens who lose their jobs or need retraining.

The five questions listed above are among those that most governments consider in their AI strategies. The earlier questions offer opportunities, but the later questions are unavoidable. Governments need to answer them before other actors, or market forces, answer them instead.

Components of a government AI strategy

While government AI strategies differ in content and approach, most strategies include five main areas. These are: government and public services; data and digital infrastructure; ethics; capacity, skills and education; and research and development. Below, we consider the kinds of questions governments should be addressing as part of these broad topic areas.

Government and public services

Governments may use AI tools to improve public services, with two broad categories of benefits: delivering services more quickly and/or cheaply, and delivering higher quality services.

Delivering services more quickly or cheaply

Many applications of AI in this category are areas of work that are currently done by humans but which can be automated using AI tools. These include mundane tasks such as tools that identify text from human handwriting in order to speed up postal delivery, through to algorithms providing answers to more complex questions such as where taxation fraud might be taking place, or whether someone should be granted bail.²

Citizens expect their government to respond to them quickly and fairly, and automation can help make this happen. And by automating some of the more monotonous, less creative parts of a public servant's job, they can free up more space for other parts of their job that might include, for example, developing new ideas or speaking with citizens.

Speeding up government decisions can come at a cost. Bail decisions provide a useful example, as an algorithm can reduce the amount of time people wait in jail before learning whether they will receive bail. In systems stressed by large numbers of accused and not enough judges, this can be a tempting solution for those overseeing the system, as well as for the accused facing long delays. Judges who take different approaches to granting bail can also create inconsistencies and unfairness in the system. But algorithms that help to inform bail decisions are based on historical data, which can contain racial and other biases. In the United States, statistical analysis by ProPublica of cases from one county in Florida showed that an algorithm used to inform bail decisions was two times more likely to falsely categorise black defendants than white defendants as future criminals.³ This helps to demonstrate the

importance of testing algorithms for bias. Even then, it is important that governments retain a role for a human decision-maker, and that this person or group has a clear understanding of how a recommendation from a machine learning algorithm has been generated.

Higher quality services

One of the most exciting possibilities created by AI is the ability for governments to better tailor services for citizens. Many of these tools will depend on accurate and complete data — a very high bar that many governments cannot yet reach in priority policy areas — but we are nonetheless starting to see early examples.

In education, programs driven by machine learning can adapt quickly to the learning style and knowledge gaps of a child. China-based Squirrel AI, for example, enables a student to practise mathematics at the same time as other classmates but in exercises designed just for them.⁴ Personalisation of services based on user data could also take place in sectors such as health and social care and give governments the opportunity to help provide better support for citizens.

Data and digital infrastructure

Data, any amount of it, is unusable without tools to store, structure, and analyse it. The term “data infrastructure” can be ambiguous but generally it refers to the means by which data is kept and processed — both physical, such as hardware and servers, and digital, such as software or cloud storage.

Governments’ effective use of data and digital infrastructure is essential to benefit from technological advances. For

example, different data systems or formats across departments (or even within them) can hinder data sharing, so careful planning around what systems are put in place, as well as the regulations for collecting, storing, sharing data are imperative.

As data usage becomes more prevalent in the delivery of public services, it will also be important for governments to consider how they want to collaborate with the private sector. As the majority of technological innovation occurs in the private sector, public/private partnerships are an essential part of developing a strong government digital infrastructure. One way to foster this collaboration is for the government to set up open data sources, where useful information is made publicly available, or data trusts, where the government guarantees accountability in storage of and access to data.

This then raises the question of who owns the data that the government collects: especially when it pertains to citizens and their private lives. In addition to the infrastructure itself, there must be rules around collecting, sharing and storing data, so that citizens' privacy is protected and the government maintains legitimacy and trust. This transitions us to the next topic: ethics.

Ethics

The growing importance of data in recent years has led to public discussions about whether digitised systems adhere to ethical standards, but so far no definitive principles have emerged for evaluating their ethical implications. 'Data ethics' is both a burgeoning academic field, as well as a policy area where various public and private entities have begun to develop broad principles and regulations that aim to guide the

use of data. The field of “AI ethics” is far behind both academically and in the public sphere, and most guidelines and regulations currently pertain primarily to data and how it is used to inform and program AI systems.

Sometimes, aspects of these principles get codified into law; for example, the EU’s GDPR, which is mostly aimed at providing data protection and privacy. However, many aspects surrounding data and its usage have been left out of the legislative space. As Luciano Floridi, a leading data ethicist, puts it: “... legislation is necessary but insufficient. It does not cover everything (nor should it)”.⁵ Data ethics can help us decide in what ways we want to use and be responsible for data and its outcomes. Furthermore, it is hard for laws to keep up with rapid technological developments resulting in potential ethical issues.

The field of data ethics thus takes up these extra-legal concerns related to the sharing and use of data.

This section discusses both academic data ethics, as well as public sector data ethics frameworks. The ethical principles that the frameworks recommend do not align neatly with the academic categorisation, meaning that at the current time, there is no clear agreement on a path forward.

Academic data ethics

Academic ethicists break data ethics down into three related subfields: “the ethics of data, the ethics of algorithms and the ethics of practices”.⁶ In other words, we have to look from an ethical perspective at the relation between (1) the data collected, (2) the processing of data, and (3) the outcomes generated by the data. Each of these phases will have different

but related ethical principles that apply. Moreover, data ethics sheds light not only on negative responsibilities — what we can't do with data — but also positive responsibilities — what we should be doing.⁷

An important gap in the academic literature is a discussion on the process of data sharing or linking. While private companies can often easily store their data in one place, differing data infrastructure across government services and departments mean that data is often kept in several different places. Therefore, this step needs to be factored in to government data ethics policies.

Data ethics frameworks

Data ethics frameworks vary in their emphasis. For instance, public sector frameworks tend to be oriented towards risk assessment, whereas frameworks aimed at commercial usage of data tend to try to balance monetary profit with public responsibility. These frameworks take various shapes, from online scoring systems⁸ to printable checklists⁹ and discussion papers¹⁰ to manifestos.¹¹

While there are many differing principles in these frameworks, there are three core values that each seem to have in common: fairness, transparency, and accountability. Fairness ensures that when working with diverse constituencies, governments do not treat any group inequitably. This is particularly relevant for the collection and use of data. Transparency enables relevant stakeholders to monitor the collection and processing of data, as well as the robustness of the technology involved. It also allows affected parties to seek redress if they have been adversely affected by the outputs of an algorithm.

Finally, accountability mechanisms ensure that governments are responsive to feedback and criticism, and facilitate the improvement of data-related practices.

One (non-comprehensive) example of how to marry the academic and public sector data ethics frameworks can be found below. It breaks down the stages of data usage in government in the way the academics do, then looks at what ethical principles might apply at what stage.

Capacity, skills and education

Governments play a critical role in developing the skills of their population: in terms of school curricula, the quality and accessibility of tertiary education, and ensuring that adults already in jobs have the means to retrain or to develop new skills.

At the school and university level, the most common policy response to AI is to place an increased emphasis on STEM (science, technology, engineering and mathematics) subjects. Ensuring that students have a good grounding in these areas helps to create a pipeline of future technologists, as well as increasing the level of data literacy among citizens. As people will work closely with technologies driven by machine learning in many areas of life, understanding the relationship between the data used to train an algorithm, and the outputs produced by that algorithm, will be increasingly important.

Along with instilling the knowledge that helps citizens become intelligent consumers (or, ultimately, producers) of technology, schools and universities should help citizens to develop “human specialties”. People have skills that machines do not in areas such as empathy, creativity, and collaboration. As it is difficult to effectively train an algorithm to display these

skills, so too they are the skills that will help citizens find employment as the economy changes. While a curriculum developed in response to the rise of AI should help prepare citizens to understand the technologies around them, it should also be rich with opportunities to learn subjects and ways of thinking that will inspire growth in emotional aptitude and creative thinking.

The trend towards learning new skills over the course of a career, and using old skills in new ways, will continue to accelerate as the capability of AI develops. In some cases, the private sector can provide reskilling for those who need it, either through paid professional development programs or by companies training their own teams as technology changes. But this is not likely to be enough for all those citizens who need help moving from one job to another, especially if they need to move between industries. The role for government here is to help those who are most at risk to access training early; to make that training affordable and accessible; and to ensure it is high-quality. Sometimes, even free training is not enough: citizens may need help to get to training, subsidised uniforms or equipment, or child care while they are in the classroom. Most educational institutions still assume they are teaching young people who have free time and limited responsibilities, and governments can help support the transition towards greater adult learning. This supports flexibility in the economy, and helps to look after those who will need help navigating changes to their livelihoods as more tasks are automated.

Research and development

As technology changes and adapts, so do citizens' expectations of how public services should be delivered. In order for governments both to keep up with technological advances, as well as provide the best services possible to their citizens, they need a clear and forward-looking research and development plan for AI. Research and development (R&D) includes incentives for companies to carry out research or test new products, as well as research spending, including PhD funding and the creation of AI research centres, thinktanks, grants, and so forth.

Public sector AI R&D is necessary above and beyond that in the private sector for several reasons. First, government R&D helps to support private sector innovation through funding, training and research. Second, governments' size means that they can help to oversee, coordinate, and spread emerging technologies throughout the economy.

Countries are now spending upwards of tens of billions of dollars on R&D investment in AI, and many countries around the world have published national AI strategies or investment plans¹² detailing their intentions for implementing new technologies and roadmapping development strategies. In 2017, China, for example, released its "New Generation Artificial Intelligence Development Plan" and plans to spend USD150bn through 2030 to make it the world leader in AI.¹³ This type of strategic planning is essential because in addition to keeping up with service provision, governments also need to keep up with emerging technologies so that economic growth and international competitiveness do not suffer.

How are governments responding?

The global landscape

Research shows that currently, governments in the Global North are overall better placed to use AI in the delivery of public services than their counterparts in the Global South.¹⁴ In the 2019 Government AI Readiness Index, Singapore was ranked as the top government in the world. It was followed by the United Kingdom, Germany, the United States, and Finland as the top five countries. India is the first developing country on the list at number 17, followed by China at number 20 (though the methodology of this ranking means that China's ranking is probably lower than it would be with better data availability). The highest ranking Latin American country is Mexico at 32nd, and the highest ranking African country is Kenya at 52nd.

Several countries have instituted AI spending plans, budgets, or even more comprehensive national AI strategies. The existence of these national plans demonstrates a commitment to develop policies focused on AI, and bolstered countries' rankings in the Government AI Readiness Index. As part of their USD150bn investment plan, for example, China has selected national corporations to invest in to form their AI "national team".¹⁵ This strategy can be a double-edged sword: while investing heavily in a few major corporations can help secure their status as world-leaders in AI, a few years down the line, having a few state-owned giants may make it harder for other companies to break into the market and for the Chinese technology sector to grow.

The United States, on the other hand, is taking a more laissez-faire approach to their AI strategy, leaving research and development mostly up to US-based corporations and the Department of Defense,¹⁶ which has a department-level AI strategy.¹⁷ Instead of a single, unitary AI policy, the US has several smaller policies directing the use AI in various sectors. In 2019, the White House launched AI.gov which details each of the federal-level initiatives.¹⁸

Singapore and across the EU, the UK included, governments are focusing efforts on the ethical use of AI and data and legal structures to support AI development. A large part of the European Commission's AI strategy is built around preparing citizens for any socioeconomic changes that occur because of AI. The European Commission is also investing in building ethical and legal frameworks to protect individual rights.¹⁹ Singapore's national AI strategies centre on the social impacts of AI, hosting bimonthly AI for everyone ("AI4E") workshops to spread knowledge about the use and benefits of AI,²⁰ establishing three new ethics frameworks, and investing heavily in AI education at the secondary and post-secondary levels, to ensure that its population is prepared for the future of work.

Singapore's and the EU's social responsibility-type models of AI strategies are important for the ethical and responsible use of AI. Domestically, there is a risk that governments using AI for their public services without adequate preparation may end up worsening problems that AI is supposed to fix. For example, governments that do not have accurate or complete data about their citizens or public services may accidentally deepen inequalities within their countries.

Globally, the best performing region in the Government AI Readiness Index, on average, is North America, while the worst performing regions are Africa and the Asia-Pacific. Part of the problem with these disparities point not only to readiness in implementing AI solutions, but also in accessing the technology behind them in the first place.²¹ There is a risk that countries in the Global South could be left behind by the so-called “fourth industrial revolution”. Not only will they not reap the potential benefits of AI, but there is also the danger that unequal implementation widens domestic and global inequalities.

Different levels of government

To date, most government action on AI has been at the national level. Indeed, efforts by some countries to accelerate investment in AI, especially for military uses, has led to concerns about a global “AI Arms Race”. China, Russia and the USA are most commonly cited as participants in the race, and Russian President Vladimir Putin has been the most explicit about his goals: speaking about AI in 2017, he said “whoever becomes the leader in this sphere will become the ruler of the world”.²²

National attention for AI initiatives has helped to provide political attention and large funds for research. This is true especially of those countries that expect to be leaders in AI, with public funds playing an important role in helping to develop new technologies.

At the time of writing, there has been far less attention given to AI by subnational governments. This is partly because decisions about many of the policy areas associated with AI tend to be made at the national level, such as research and

development incentives or funding for doctoral programmes. But many areas of policy are important for state and city governments, who are able to tailor their policies to smaller populations, and may also be closest to the impacts of job disruptions due to automation. At the national level, for example, a government may be confident that technological change will lead to as much job creation as job loss. For a state, a province or a town, this picture may look different, as the effects of automation will not be felt evenly.

In 2018 and 2019, we worked with the city of Stockton in northern California to develop an AI strategy, making it among the first cities in the world to do so. Aspects of Stockton's demography make it especially susceptible to the challenges of automation, as its population has relatively low skills levels and high unemployment. Large numbers of its workers are employed in job groups more susceptible to automation, such as construction, manufacturing and transportation.

Focusing on "defensive" priorities, cities like Stockton will benefit from preparing early for the changes ahead, including by identifying those groups who may be most at risk to offer tailored opportunities for retraining or skills development. City and regional governments can also take advantage of opportunities by preparing early, including testing new government uses for AI technology and working to attract entrepreneurs.

Conclusion

There is nothing inevitable about the impact of technology. Technology is designed by people, and it is people who decide how it will be used. Governments have the power to encourage the development of technology, regulate its use, and moderate

its impacts. This is true for AI, as it has been for generations of technologies. The decisions governments are making now — about areas as diverse as their data collection, their welfare policies, and what children learn at school — will help determine whether the impacts of AI should be feared or should be welcomed.

Endnotes

- 1 <https://foreignpolicy.com/2019/03/05/whoever-predicts-the-future-correctly-will-win-the-ai-arms-race-russia-china-united-states-artificial-intelligence-defense/>
- 2 <https://engineering.stanford.edu/magazine/article/can-ai-help-judges-make-bail-system-fairer-and-safer>
- 3 Angwin J et al. (2016). Machine bias. *ProPublica*, <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- 4 Hao K (2019). China has started a grand experiment in AI education. It could reshape how the world learns. *MIT Technology Review*, August.
- 5 Floridi L (2018). Soft ethics: Its application to the General Data Protection Regulation and its dual advantage. <https://link.springer.com/article/10.1007/s13347-018-0315-5>
- 6 Floridi L & Taddeo M (2016). What is data ethics? <https://royalsocietypublishing.org/doi/pdf/10.1098/rsta.2016.0360>
- 7 Nesta has produced a useful collection of case studies highlighting the benefits of predictive analytics across the UK. Copeland E et al. (2019). Public Sector Data Analytics. https://media.nesta.org.uk/documents/Public_Sector_Data_Analytics_-_A_Nesta_Guide_byCwKTI.pdf
- 8 Government of Canada (2019). Algorithmic impact assessment. <https://www.canada.ca/en/treasury-board-secretariat/corporate/news/algorithmic-impact-assessment.html>
- 9 UK Government (2018). Data Ethics Framework. <https://www.gov.uk/government/publications/data-ethics-framework/data-ethics-framework>
- 10 Dawson D et al. (2019). Artificial intelligence: Australia's Ethics Framework. <https://consult.industry.gov.au/strategic-policy/artifi->

- cial-intelligence-ethics-framework/supporting_documents/Artificial
Intelligenceethicsframeworkdiscussionpaper.pdf
- 11 Future of Life Institute (2017). Asilomar AI Principles. <https://futureoflife.org/ai-principles>; Accenture (2019). Responsible AI and robotics: an ethical framework. <https://www.accenture.com/gb-en/company-responsible-ai-robotics>.
 - 12 Dutton T (2018). An overview of national AI strategies. *Medium*. <https://medium.com/politics-ai/an-overview-of-national-ai-strategies-2a70ec6edfd>
 - 13 <https://futureoflife.org/ai-policy-china/>
 - 14 <https://www.oxfordinsights.com/ai-readiness2019>
 - 15 <https://www.scmp.com/tech/china-tech/article/2120913/china-recruits-baidu-alibaba-and-tencent-ai-national-team>
 - 16 <https://medium.com/politics-ai/an-overview-of-national-ai-strategies-2a70ec6edfd>
 - 17 <https://media.defense.gov/2019/Feb/12/2002088963/-1/-1/1/summary-of-dod-ai-strategy.pdf>
 - 18 <https://www.whitehouse.gov/ai/ai-american-innovation/>
 - 19 <https://ec.europa.eu/digital-single-market/en/artificial-intelligence>
 - 20 <https://www.aisingapore.org/industryinnovation/ai4e/>
 - 21 Another problem reinforcing these inequalities is also availability of data. Countries with stronger economies, such as those in the Global North, tend to more data availability.
 - 22 Vincent J (2017). Putin says the nation that leads in AI “will be the ruler of the world”. *The Verge*. <https://www.theverge.com/2017/9/4/16251226/russia-ai-putin-rule-the-world>