



TONY BLAIR
INSTITUTE
FOR GLOBAL
CHANGE

Public Artificial Intelligence: A Crash Course for Politicians and Policy Professionals

LIMOR GULTCHIN

RENEWING
THE CENTRE

Contents

Executive Summary	3
Introduction	8
Public-Service Delivery	9
How to Realise the Potential of AI	25
Recommendations	5

EXECUTIVE SUMMARY

Artificial intelligence (AI) promises a radical transformation of public-service delivery, allowing governments to meet each citizen's needs, freeing up time for civil servants and front-line staff, and putting data at the heart of decision-making. Understanding the potential of these technologies must be central to any serious modernisation of the state.

Policymakers should take steps to embrace the promise of AI while preparing for its impact and complications. Doing so would radically transform public-service delivery, allowing governments to better meet citizens' needs while putting data at the heart of decision-making.

This briefing articulates a vision of how AI can improve the public sector, by explaining what AI can do, where its application might be most impactful and feasible, and how policymakers can harness its potential in the public interest.

WHAT AI CAN DO

AI and related techniques—machine learning (ML) and deep learning (DL)—allow for a paradigm shift in how governments and services operate and interact with citizens. Their capabilities can be classified into five broad categories:

- **Personalising the citizen experience:** To meet citizens' expectations of public services, chatbots can answer questions quickly and effectively; recommendations and information feeds can be personalised for businesses; and services can be tailored to users' needs.
- **Monitoring services in real time:** Inefficient or delayed data collection can undermine service delivery. Analysing data collected from sensors with ML techniques can help services to track and analyse progress in real time.
- **Classifying cases more effectively:** Categorising events, users or information correctly can make the difference when it comes to a patient seeing a doctor or a citizen getting welfare support. ML-based classification can help services to prioritise cases like these.
- **Predicting outcomes more accurately:** Large data sets and ever-improving computing power can help people to learn from the past by predicting the likelihood of events recurring, such as pressure points on hospitals, and allowing services to prepare in advance.
- **Modelling complex systems:** It can be difficult for governments to track and measure the myriad consequences of delivery programmes. New simulation techniques can inform the research of long-term effects and improve decision-making.

WHERE AI CAN HELP

Mapping AI's capabilities onto domains of government delivery, and assessing their impact and current technical feasibility, reveals four categories of application for policymakers to consider:

- **Priority areas: healthcare, welfare, and energy and the environment.** These areas have high feasibility and would have high impact. Examples include personalising medical care, simplifying and automatically adjusting benefit payments when needed, and cutting energy bills by modelling demand better.
- **No-regrets moves: tax, pensions and licensing.** These domains have high feasibility but would have low impact. Examples include simplifying tax reform and bypassing inflexible tax codes,

predicting future pension needs and deficits, and reducing licensing duplication and redundancies.

- **Investment areas: transport, education and foreign policy.** These domains have low feasibility but would have high impact. Examples include dynamic transport scheduling and on-demand mobility, personalised educational materials and informing geopolitical strategy by modelling global actors.
- **Low-priority areas: housing, emergency services and regulation.** These areas have low feasibility and would have low impact. There is some narrow applicability of ML to these areas—for example, parametric design in social housing, diagnosing patterns of emergency services call-outs and real-time economic data—but all are currently unfeasible and there is no broad application in each domain.

RECOMMENDATIONS

To realise the potential of AI while preparing for its potentially disruptive economic and social impact, policymakers must focus on the twin priorities of resources and responsibilities.

Resources

Policymakers should:

- **Sort out government data.** AI models are only as good as the underlying data, yet government data are often siloed, poorly labelled and unutilised. Proper labelling and merging of citizen and municipal data sets will allow governments to use data to deliver improved services, just as many tech companies already do. However, governments have a greater duty of care, so they should be transparent about how data are used and provide opt-outs for data collection where appropriate.
- **Invest in skills by linking with public service.** To plug the ever-growing AI skills gap and compete with industry, governments should massively scale up PhD funding, but this should include a commitment to several years of work in civil society or the public sector.
- **Ready procurement for ML.** Advances in cloud computing mean

governments should avoid building bespoke, unwieldy IT solutions themselves. However, beyond computing infrastructure, governments should be wary of procuring ready-made ML applications as these may not meet the higher standards of risk, bias and fairness assessment required of public services.

- **Create AI-ready institutions.** A vibrant ecosystem of research, industry and governmental bodies is required to shape AI development in the public interest and provide anchors around which clusters of start-ups, scale-ups and civil-society actors can operate. Creating this environment must be a central part of any AI development and governance strategy.

Responsibilities

Policymakers should also:

- **Establish feedback mechanisms involving the public and civil society.** Given the serious ethical challenges involved in AI, policymakers should incorporate feedback from the public, civil servants and anyone affected by AI-driven services into continual, iterated AI development. This can help individuals to trust in governments' application of AI.
- **Harness automation to reduce the burden on civil servants.** As many civil servants' workloads become unmanageable, AI can automate many administrative tasks and free up time for government workers to think creatively and improve interactions with colleagues and citizens.
- **Take responsibility for ethics, fairness and transparency.** Algorithmic bias and fairness are serious concerns when it comes to deploying AI for social causes. Governments must provide further support for research in this area, help coordinate industry ethics codes, be transparent about how data are used so users understand these technologies, and help develop standard dispute-resolution procedures for citizens and businesses that need protection from AI.
- **Start small but think big.** AI can have a huge impact in both the short and the long term, but to build trust with citizens, governments should focus on delivering quick wins and ensuring

the foundations of a reliable data infrastructure. However, this must not come at the expense of AI's bigger promise, so investment in skills, infrastructure and computing power will help accelerate those later advances.

- **Set realistic expectations.** Without a doubt, AI has huge potential, but policymakers must cut through the hype and be clear about expectations, both internally and publicly. This requires careful, ongoing deliberation about thorny ethical issues as well as encouraging start-ups, scale-ups and established industry players to develop AI responsibly.

INTRODUCTION

Governments have a lot of responsibilities, and lots of ways of meeting them. It is up to politicians and stakeholders to make sure governments not only remain vital and relevant for the benefit of citizens but also harness the potential to take public service delivery into the artificial intelligence (AI)– and machine learning (ML)–infused era.

The benefits of AI and ML for government services can be massive. A 2017 Deloitte report estimated that AI applications could free up 30 per cent of government workers' time, while McKinsey estimated that deep learning (DL) techniques would account for \$3.5–5.8 trillion of value creation annually in the private sector, the equivalent of 1–9 per cent in 2016 revenue.¹

This briefing presents the promise of AI, ML and DL for public-service delivery and emphasises that this terrain is unique and cannot be dealt with like any other digital technology. The aim is not to enumerate every application but to articulate the scope of what is possible and highlight where policymakers should focus their attention to realise the potential of these technologies.

¹ Peter Viechnicki and William D. Eggers, "How much time and money can AI save government?", Deloitte, 26 April 2017, <https://www2.deloitte.com/insights/us/en/focus/cognitive-technologies/artificial-intelligence-government-analysis.html>; Michael Chui, James Manyika, Mehdi Miremadi, Nicolaus Henke, Rita Chung, Pieter Nel, Sankalp Malhotra, "Notes From the AI Frontier: Insights From Hundreds of Use Cases", McKinsey Global Institute, discussion paper, April 2018, https://www.mckinsey.com/~media/mckinsey/featured%20insights/artificial%20intelligence/notes%20from%20the%20ai%20frontier%20applications%20and%20value%20of%20deep%20learning/mgi_notes-from-ai-frontier_discussion-paper.ashx.

PUBLIC-SERVICE DELIVERY

Reform of public-service delivery encompasses organisational change as much as the technological, but there is no doubt that AI and ML allow for a paradigm shift in how government and services both operate and interact with citizens. Whether it is automation of back-end services, pattern recognition applied to transport data or personalising the interface between public services and citizens, understanding the potential of these technologies must be central to any serious modernisation of the state.

This briefing starts by explaining the different capabilities provided by these technologies, before mapping these capabilities onto domains of government delivery.

WHAT AI CAN DO

What follows is an introduction to five broad capabilities of AI and ML technologies. This is an approximate classification intended only as an illustration of how their application can radically improve the delivery of public services, both on the front line and in the back office. It is followed by an analysis of more concrete case studies.

Each category is accompanied by a toolbox of AI techniques used for each application.

Personalising the Citizen Experience

Every individual is different, and every interaction with government has its own drivers, circumstances and constraints. AI can help in areas like answering peoples' questions quickly and effectively, personalising recommendations and information feeds, and tailoring services to users' needs. In a world where personalised products are increasingly the standard, governments can and should offer such services to enable greater trust and satisfaction with public service.

Toolbox: clustering,² natural language processing,³ collaborative filtering.⁴

Monitoring Services in Real Time

Currently, most government services and data are delivered, stored or analysed after the fact, based on forms (such as tax returns), self-assessments (like expenditure surveys and censuses) or citizens' reports (such as emergency calls). This often leads to lags, inaccuracies and inefficiencies.⁵ However, in the age of big data and machine learning this no longer has to be the case. Systems can now be built to track data in real time and analyse them instantly, making decisions, predictions or recommendations quickly and efficiently.

2 Jure Leskovec and Anand Rajaraman, "Clustering Algorithms", Stanford University, presentation, <https://web.stanford.edu/class/cs345a/slides/12-clustering.pdf>.

3 Natural language processing is the branch of ML that focuses on the analysis, understanding and reproducing of human language in its natural form—that is, as it usually occurs in human interaction, and not in a formal or constrained form, such as through code. For more information, see Karen Sparck Jones, "Natural language processing: a historical review", University of Cambridge, October 2001, <https://www.cl.cam.ac.uk/archive/ksj21/histdw4.pdf>.

4 A common approach to recommendation problems, such as that to which the Netflix Prize was dedicated. The prize encourages research teams to enter a competition to help improve iteratively on the Netflix film recommendation algorithms, from 2006 to 2010. See "The Netflix Prize Rules", Netflix Prize, accessed 30 October 2018, <https://www.netflixprize.com/rules.html>, and Prince Grover, "Various Implementations of Collaborative Filtering", Towards Data Science, accessed 30 October 2018, <https://towardsdatascience.com/various-implementations-of-collaborative-filtering-100385c6dfe0>.

5 Torsten Bell and Adam Corlett, "A history lesson wouldn't hurt – at least when it comes to child poverty", Resolution Foundation, 24 July 2018, <https://www.resolutionfoundation.org/media/blog/a-history-lesson-wouldnt-hurt-at-least-when-it-comes-to-child-poverty/>.

Toolbox: classifiers,⁶ random decision trees,⁷ clustering,⁸ regression.⁹

Classifying Cases More Effectively

Many government services require classifying cases properly—placing events, citizens or information into the right category—whether it is to determine which citizen is entitled to government support or which patient should see a doctor most urgently. Classifiers can immensely help with these types of tasks: these machine-learning models are built to reason through cases, look for similarities and differences between points of data and consequently place them into appropriate categories.

Toolbox: regression,¹⁰ deep learning,¹¹ classifiers,¹² clustering.¹³

Predicting Outcomes More Accurately

ML algorithms are especially suitable for prediction questions—those that try to recognise trends and predict future behaviour. This is an area where massive amounts of data and ever-improving computer power can do an immensely better job than even human experts; humans are, after all, notoriously bad at making predictions and reasoning about the future.¹⁴ Therefore, a promising area in which ML can make a real difference is assessing the likelihood and frequency of events. People have been using

6 Mandeep Sidana, “Types of classification algorithms in Machine Learning”, Sifium Technologies, 28 February 2017, <https://medium.com/@sifium/machine-learning-types-of-classification-9497bd4f2e14>.

7 Ibid.

8 Jure Leskovec and Anand Rajaraman, “Clustering Algorithms”, Stanford University, accessed 30 October 2018, <https://web.stanford.edu/class/cs345a/slides/12-clustering.pdf>.

9 Statistics Solutions, “What Is Linear Regression?”, accessed 30 October 2018, <http://www.statisticssolutions.com/what-is-linear-regression/>.

10 Ibid.

11 Favio Vázquez, “A ‘weird’ introduction to Deep Learning”, Towards Data Science, accessed 30 October 2018, <https://towardsdatascience.com/a-weird-introduction-to-deep-learning-7828803693b0>.

12 Sidana, “Types of classification algorithms in Machine Learning”.

13 Leskovec and Rajaraman, “Clustering Algorithms”.

14 Caroline Beaton, “Humans Are Bad at Predicting Futures That Don’t Benefit Them”, Atlantic, 2 November 2017, <https://www.theatlantic.com/science/archive/2017/11/humans-are-bad-at-predicting-futures-that-dont-benefit-them/544709/>.

analytics and statistics to reason through such scenarios, but ML allows these assessments to be done more efficiently, quickly and cheaply, and on a much bigger scale.

Toolbox: regression,¹⁵ deep learning,¹⁶ agent-based models.¹⁷

Modelling Complex Systems

Government programmes and departments are, almost by definition, big, complex and convoluted. Each proposal, reform or delivery programme has far-reaching effects that are often not obvious from their onset, and reasoning through all possible outcomes of a policy programme is an almost disheartening task. However, new technologies such as agent-based simulations and deep learning can help study the effects of complex systems over time, and even test policies before implementing them fully to gain more insight into their long-term consequences.

Toolbox: agent-based models,¹⁸ simulation,¹⁹ reinforcement learning.²⁰

WHERE AI CAN HELP

With the broad capabilities of AI in mind, these features can be mapped onto various fields of application that comprise government services. Using a matrix of technical feasibility and impact allows four categories to be developed (see table 1):

15 Statistics Solutions, “What Is Linear Regression?”.

16 Vázquez, “A ‘weird’ introduction to Deep Learning”.

17 Eric Bonabeau, “Agent-based modeling: Methods and techniques for simulating human systems”, PNAS 99 (suppl 3) 7280–7287, 14 May 2002, http://www.pnas.org/content/99/suppl_3/7280.

18 “Agent Based Modelling: Introduction”, Geography Department, University of Leeds, accessed 30 October 2018, <http://www.geog.leeds.ac.uk/courses/other/crime/abm/general-modelling/index.html>.

19 “Integrating Artificial Intelligence and Simulation Modeling”, PwC Artificial Intelligence Accelerator, presentation at AnyLogic Conference, April 2018, <https://www.anylogic.com/upload/conference/2018/presentations/integrating-artificial-intelligence-with-simulation-modeling.pdf>.

20 Vishal Maini, “Machine Learning for Humans, Part 5: Reinforcement Learning”, Machine Learning for Humans, 19 August 2017, <https://medium.com/machine-learning-for-humans/reinforcement-learning-6eacf258b265>.

- priority areas (high feasibility, high impact);
- no-regrets moves (high feasibility, low impact);
- investment areas (low feasibility, high impact); and
- low-priority areas (low feasibility, low impact).

Table 1: Opportunities for AI in the Public Sector

Impact	Feasibility	
	High	Low
High	Priority Areas: Healthcare Welfare Energy and Environment	Investment Areas: Transport Education Foreign Policy
Low	No-regrets moves: Tax Pensions Licensing	Low-priority areas: Housing Emergency Services Regulation

This typology highlights that beyond straightforward high- and low-priority areas, there are many applications with low feasibility that could be highly impactful and therefore require investment, and others with low impact but high feasibility and are therefore worth pursuing as quick wins.

Below are examples of how each capability can be applied to different policy domains, beginning with the priority areas. Each section leads with the most impactful application, so the order of capabilities varies throughout. For instance, while simulating complexity is particularly impactful for transport, it is less so for licensing and registrations of services.

These examples are not comprehensive and are intended only as illustrations. Given the applicability of multiple ML capabilities to each domain, some examples across domains bear similarity to one another. The capabilities are colour coded as follows:

- █ Customer experience
- █ Real-time monitoring

■ Diagnostics or classification

■ Prediction or forecasting

■ Modelling complexity

Priority Areas

Applications in healthcare, welfare, and energy and the environment would have a high impact and already have high technical feasibility. Governments should therefore prioritise these domains when designing AI-driven service strategy.

Healthcare

Real-time medical data: Medical check-ups are rarely frequent, and even in hospital patients may be waiting longer than expected or symptoms may be missed. Wearable technologies allow for real-time monitoring of significant medical data, with devices now able to warn of irregularities and automatically notify doctors.

Automating diagnoses: Patients must often wait weeks to receive diagnoses for serious illnesses, and struggling health systems coupled with an increasing burden over time mean this situation is likely to worsen before it improves. Automating certain diagnoses—such as of cancer, blood tests and viruses—using pattern recognition and computer vision could radically speed this up.

Personalised healthcare: Demand for healthcare often outstrips supply. Even when a diagnosis is made, it can be hard to find a functional treatment, let alone optimise it. However, new technologies, powered by ML, can provide personalised, on-demand healthcare while reducing the burden on the health system. This can take the form of 24/7 accessibility, continuous notifications and recommendations based on similarity to other patients' needs, or even triage of symptoms to recommend basic treatments (based on doctor approval when needed).

Predicting hospital administration requirements: Although some periods (such as winter) are known for increasing pressure

on hospitals, granular change in healthcare demand is highly dynamic and therefore unpredictable. Using historical data to recognise patterns and real-time data on weather and people movements, ML systems can help predict equipment and staffing requirements based on likely patient inflow, as well as inform the decision-making process for locations of new hospitals based on geographical areas of need.

Analysing DNA to improve healthcare provision: Despite many advances in healthcare, optimising dosages or distinguishing effective drugs from others remains difficult. As the breadth of ailments will only increase, the intractable nature of this issue will only worsen. In response, by sequencing patients' genomes, modelling their complexity and applying a neural net, researchers can simulate possible drug designs and dosages to personalise treatment according to what is most effective. Combining this with real-time medical data could ignite a step change in treatment success.

Welfare

Simplifying and automating benefits: Welfare systems across the world are often unwieldy and bureaucratic, making it difficult for claimants to receive all the benefits they are entitled to. What is more, often they are not even aware of everything they qualify for. Using ML to automate this process—based on data the government either has or can request simply—can bridge the knowledge gap and ensure everyone receives what they need without undue stress or strain.

Automatically adjusting welfare: Job losses and other changes in personal circumstances leave individuals requiring support. However, it can often take a long time for welfare to kick in. Real-time monitoring of these circumstances can inform a more flexible benefit calculation and allow the system to react quickly, immediately paying out unemployment support.

Diagnosing system failures: From a citizen's perspective, welfare systems can often fail or make poor judgements about

individual cases while being unaware of others in their position. For welfare administrators, not linking up these cases means important insights may be missed. Using pattern recognition to find commonalities among these groups or discover those who are underserved can help find a more targeted solution that can address issues together.

Speeding up delivery capacity by predicting demand surges:

When the external environment changes rapidly, thousands of people may need help. However, crisis-response strategies may focus on immediate damage rather than long-term support, for example in the case of job-market disruption. Predicting demand surges and trends with greater accuracy can improve provision while reducing waste.

Improving long-term delivery: Diagnosing system failures can often struggle due to limited resources. The sheer complexity of all the different decisions and payment permutations makes auditing previous strategies particularly difficult. Allowing a simulator to tackle these questions could cut through the resource difficulty while still providing the insights, perhaps recommending action to avoid unexpected consequences.

Energy and Environment

Shortening power cuts by increasing response times:

Inefficient power distribution can waste vital resources, and when things go wrong it can be tricky to diagnose where the issue is. With real-time, online monitoring of energy outages, emissions and leaks, power plants can reduce reaction times, possibly conduct remote repairs and generally take a more preventive approach.

Better modelling of energy demands: As climate change worsens, an ever-better understanding of energy consumption will be necessary. And this will only increase in importance as populations continue to grow. Using simulation can tackle these challenges including for edge cases—problems that occur only at

extreme operating parameters—perhaps by comparing natural phenomena and iterating possible solutions.

Personalised energy plans: An increased need for sustainability demands that everyone on the planet works together to reduce energy use. However, it can be very difficult to coordinate this sort of action, with some advice being relevant for some but damaging for others. Designing personalised energy systems and recommendations, based on tracking an individual’s usage and comparing it with others’, can provide crucial guidance and automated energy management.

Overcoming siloes to improve policy design: Operating interdepartmentally can be challenging at the best of times, but during natural crises this situation can prove even more complicated. Using ML to categorise data events better and demonstrate cross-thematic insights can provide a stronger evidence basis for collaboration between departments such as those of agriculture, transport, infrastructure and the environment.

Predicting power outages: Power outages are often unexpected, disruptive and costly. As climate change leads to more natural disasters and severe weather patterns, the occurrence of outages is likely to increase. AI can help reduce this impact by predicting power shortages, planning inspections and repairs more accurately, and effectively managing supply and demand.

No-Regrets Moves

Applications in tax, pensions and licensing have high technical feasibility but would have a low impact.

Tax

Simplifying tax reform: Tax laws can be some of the most complex to reform, and as the nature of work changes with an increase in casual and freelance workers, this may grow more

complicated still. Using ML to analyse edge cases effectively can help cut through the complexity: who pays too much, who pays too little, who benefits the most after taxes and transfers, and is this intentional?

Reducing inefficiencies in tax collection: Changes in employment, including becoming self-employed, often illustrate the slow, inflexible nature of the tax code, with employers and employees alike having to wait months for issues to be corrected. Real-time monitoring of employment and financial circumstances can form the foundation for flexible tax collection that reduces these inefficiencies.

Planning tax audits less randomly, based on data-driven predictions: Because it is impractical to audit every individual and corporate taxpayer, most administrations rely on a random sample of accounts to check for mistakes or abuse. As predictive systems improve, it is becoming easier to target the accounts more likely to be at risk of abuse or with more money at stake, resulting in a better outcome for taxpayers.

Personalised tax schemes: Changes in circumstances can affect both earnings and benefits, but at present the two systems are treated separately. Self-employed earners may also have varied earnings each month, with little room to smooth out tax payments. With the right data, an intelligent system could combine important insights in both cases to design a personal tax scheme that overcomes these issues.

Optimising the tax system: In a single system with millions of taxpayers, it is difficult to expose unintended consequences or sort the outliers from the mainstream. Running simulations from historical and current data can aid with this.

Pensions

Helping defuse the pensions time bomb: As long as savings rates remain low, pressure on pensions will increase. With populations aging, this picture only becomes more difficult.

Using ML can help predict people's future needs and possible deficits, as well as tailor their pension schemes accordingly.

Satisfying different attitudes to risk: Pension funds are often pooled together to use greater capital to maximise returns, but this is an imprecise way of meeting people's different attitudes to risk or personal preferences regarding specific companies. Designing more flexible pension schemes, with personalised recommendations based on past data and similarities to other cases, can overcome this while maintaining return on investment.

Financial advice for all: Expert financial advice can be extremely expensive, while increasing demands on income can make saving particularly difficult. A virtual pensions assistant could keep track of long-time trends, people's appetite for risk and market, and their personal circumstances to make recommendations.

Helping providers to tailor schemes: When serving millions of customers, pension providers could benefit from a detailed classification of customers to tailor their offering. However, this complexity is often difficult to overcome. ML can automate this categorisation and keep it updated, producing insights to be reflected in pension schemes.

Future-proofing pensions: Understanding future demands and historical errors go hand in hand when it comes to business strategy, but detecting unintended consequences can be a particularly challenging task. Simulating events can reveal these cases, allowing responses to be more tailored in future.

Licensing

Automating licensing in real time: A legacy of the paper-based licensing system is that even current digital systems can replicate historically inefficient processes. When circumstances change rapidly, reflecting this change can take months or years. Including continuous, real-time assessment allows for registration to be updated instantly, removing inefficiency.

Making registrations easy: A change in personal circumstances can require changing many licences or registrations, while starting a new business might require an individual to apply for new ones. Using pattern recognition to recommend certain applications and auto-filling data based on previous interactions with the user can make the entire interaction much easier.

Reducing duplication and redundancies: Although regulation is important in many domains, some overspecific or overlapping applications, registrations or requirements can be too burdensome. In turn, this can risk disincentivising new entrants. Creating simpler categories of licences by grouping users and their needs, and understanding their commonalities, can help reduce these redundancies while maintaining standards.

Identifying and predicting fraud: By its very nature, fraud can be difficult to detect. The bigger the data set, the harder the task becomes. However, rather than treating each case as anomalous, grouping past cases together to detect future instances of fraud can improve planning of inspections and, in turn, conviction rates.

Detecting unintended consequences: Determining whether policies are exclusionary is already tricky, but it can be more difficult still to understand whether registrations genuinely improve access to services. Using simulation to study negative and unintended consequences such as these can inform future policy improvements.

Investment Areas

Applications in transport, education and foreign policy have the potential for high impact but currently have low technical feasibility.

Transport

Reducing costs and complications when building infrastructure: Large infrastructure projects are generally very costly and often overrun. As environmental and economic

sustainability become increasingly important, these issues can escalate, causing many complications during a project. Using simulation techniques to model new transport lines or economic and environmental impact can help people prepare in advance.

Dynamic transport scheduling: Many public transport systems are rigid and overcrowded. As cities become more densely populated, this situation will worsen. Using ML to monitor demand from sensor and Global Positioning System (GPS) data in real time can allow for more sophisticated online scheduling, dynamically adjusted to need and congestion.

On-demand mobility: The saying “You wait ages for a bus and then two come along at once” is a good illustration of the inefficiencies and frustrations of public transport. Turning to a taxi is not always an instant or affordable solution, either. However, if cities were to model user demand, they could provide on-demand public transport, following companies such as ViaVan or CityMapper. Think: a world of buses, but no bus stops.

New methods of transport: Although many improvements to existing transport infrastructure can be made, a long-term approach would benefit from a more detailed understanding of traveller needs and trends. Classifying similar groups could reveal a more efficient and novel way to provide transport, perhaps skipping current systems.

Improving transport administration: As populations increase and cities grow denser, transport systems will come under greater strain. Improvements can encompass upgrades of existing infrastructure or new systems altogether, but assessing which is most effective is not always clear. Predicting demand with greater accuracy and having a clearer forecast of environmental and economic impact can aid this decision and make cost-benefit analyses more representative.

Education

Improved diagnosis of educational needs: High teacher-student ratios can make it difficult to understand each student's educational needs. Young children may also find it hard to provide feedback and understand their needs themselves. Combining data sets to spot patterns can identify similar cases and expose the most effective teaching techniques, which can in turn inform more tailored teaching by matching classes to teachers and using other environmental cues to aid progress.

Personalised education design: Citizens and governments alike care about proper education, but delivery is occasionally imprecise, grouping children together simply because they are the same age and teaching a one-size-fits-all curriculum. In response, ML can be used to recognise areas of difficulty and generate personalised recommendations of relevant educational materials—or perhaps even create them automatically—for different types of learning, topic and difficulty.

Real-time assessment: Exams at the end of the school year, supported by homework and weekly tests, can be a highly imprecise way to measure attainment. In turn this can mean students miss out on crucial interventions in their learning and development. Building ongoing assessment into the system can help highlight edge cases and identify particular needs or areas where students are performing well far earlier, allowing the feedback loop to improve delivery earlier too.

Anticipating students' weaknesses: High-achieving students may do well all year until they meet a particular topic that they struggle with. These children may typically have less direct support, exacerbating the situation when things become tricky. Using historical data to spot patterns can produce predictions of where these cases may arise, in turn giving teachers an early warning to reallocate support.

Meeting the skills demand of the next century: As the nature of work, education and life changes, so will the demands on students and workers. However, education reforms are difficult to evaluate in advance. Simulating education reforms before

implementing them, be it curriculum changes or school restructuring, can help policymakers avoid unintended outcomes while maintaining standards and innovating education provision.

Foreign Policy

Automated database of diplomatic and military action:

International relations, given their complex, networked nature, are highly difficult to navigate. Secrecy of diplomatic and military action creates asymmetries of information between both partners and rivals. Using AI to trawl the Internet, verify unconfirmed reports and amass them in an automatically updating database can plug this gap and level the playing field.

Learning from history: It is often said that a student of history can avoid repeating mistakes in the future, but complex geopolitics and the fallibility of human decision-making can undermine this hope. Using ML to compare foreign policy situations with past examples and understand which are more similar can produce an interactive library to recommend relevant cases and insights from varied academic, governmental and military sources.

Modelling actors in international-relations settings: Many state and nonstate actors can be unpredictable in their actions, while understanding their incentives and broader geopolitical motivations or restrictions can be an incredibly complex task. Modelling their behaviour is a way to capture the totality of these issues and balance them against each other, in turn aiding strategic decision-making.

Simulating alternative outcomes of past crises: History is full of tipping points and sliding doors, but it is impossible to verify what would have happened had a given actor behaved differently. However, new technologies allow people to simulate these scenarios and examine the unintended consequences of international pacts or actions, helping to avoid them in future.

Low-Priority Areas

Applications in housing, emergency services and regulation would have low impact and have low technical feasibility. Governments should therefore deprioritise these domains when designing AI-driven service strategy. This is not to say that ML has no applicability here. For instance, parametric design in social housing, real-time economic indicator monitoring and diagnosing patterns from emergency service call-out data are all valuable innovations. However, there is no broad application for ML in these domains yet. This recommendation is supported by the fact that other areas clearly call for far more prioritisation.

HOW TO REALISE THE POTENTIAL OF AI

Policymakers who want to leverage the power of AI technologies need to be aware of a range of obstacles covering resources (data, computing power and skills) and responsibilities (disruption, privacy, data ownership and algorithmic fairness). Although complex issues in and of themselves, resource problems require more straightforward responses that can begin immediately. They should therefore become departmental priorities. Questions of responsibility generally require a broader and more deliberative approach, including engagement from a wide range of stakeholders and possibly the public. That does not make them any less important, but because appropriate actions are less clear-cut, they follow the analysis of resource problems.

RESOURCES

Resource problems can broadly be split into questions of data, skills and computing power. Each certainly requires investment to build capability, but just as important is an approach that embraces experimentation and entrepreneurship. This applies as much to privacy and data protection as to plugging the skills gap.

Data

The clear majority of common ML models require data. Lots and lots of data. The amount of data available can determine the level of accuracy a model can attain and the tasks that AI can or cannot be complete. Without enough data to train and validate a model, it simply falls apart. This means that to develop a serious model to deal with a given problem, it is first necessary to make sure there are enough data: data that are cleaned, well maintained, appropriately labelled and readable by a computer via an application programming interface (API). Ensuring all data, historical and current, meet these requirements should be the primary concern for departmental leaders.

Beyond that, paper forms should be abolished wherever possible, and a reliable digital form of storage should be adopted that could then be reused not only for lookup but also for the training of other

models. Connecting these data is also essential to allow departments across governments to access the same pool of data whenever possible, and to ensure accuracy, consistency and efficiency. Moreover, providing the regulatory room to use data beyond their initial consented purpose to support future service innovations requires a serious discussion about privacy and protection of personal data.

Educational Requirements and Expertise

The level of education required to be an expert in AI and ML-model development is ever increasing. As Professor Michael Woolridge of the University of Cambridge has put it, having the right skills is “not the same as having people that can program . . . It’s people with master’s degrees, PhDs and so on”.²¹

So while general software engineers in the technology industry often hold undergraduate degrees, the skill level of ML experts and data scientists is often higher.²² What is more, there is a dearth of trained people in the field compared with the number of industry jobs available, and future demand is expected to increase sharply in the next few years.²³ Academics with the right skill set are therefore in high demand, a situation that risks a brain drain from university faculties. While this could be a setback for AI research and the rolling out of ML models, it could also be a sign of just how valuable these technologies are and how quality education will continue to play an important role.

As a result, significant investment in postgraduate, doctoral and postdoctoral academic programmes will be required to plug the skills gap. Meanwhile, encouraging greater collaboration between industry and academia over outright poaching can help use expertise across the board without undermining the development

²¹ Artificial Intelligence Committee, Tuesday 10 October 2017, House of Lords, Parliament, accessed 30 October 2018, <https://parliamentlive.tv/Event/Index/073717ca-484b-4015-bd10-f847cea3f249>.

²² Will Markow, Soumya Braganza, and Bledi Taska, “The Quant Crunch: How The Demand For Data Science Skills Is Disrupting The Job Market”, Burning Glass Technologies, IBM and Business Higher Education Forum, 2017, accessed 30 October 2018,; <https://public.dhe.ibm.com/common/ssi/ecm/im/en/im14576usen/analytics-analytics-platform-im-analyst-paper-or-report-im14576usen-20171229.pdf>.

²³ Ibid.

of future computer scientists. The Chinese programme of starting AI education in secondary school may be a model for wider reform.²⁴

Computing Power

AI has known cycles of hype, anticipation and disappointment. The year 1990 was the lowest point of funding for AI, and the time of reduced funding and excitement became known as the “AI winter”. As inventor Ray Kurzweil wrote in *The Singularity Is Near*, this period does not mean that the original sense of promise was misguided:

*Many observers still think that the AI winter was the end of the story and that nothing since has come of the AI field. Yet today many thousands of AI applications are deeply embedded in the infrastructure of every industry . . . the AI winter is long since over.*²⁵

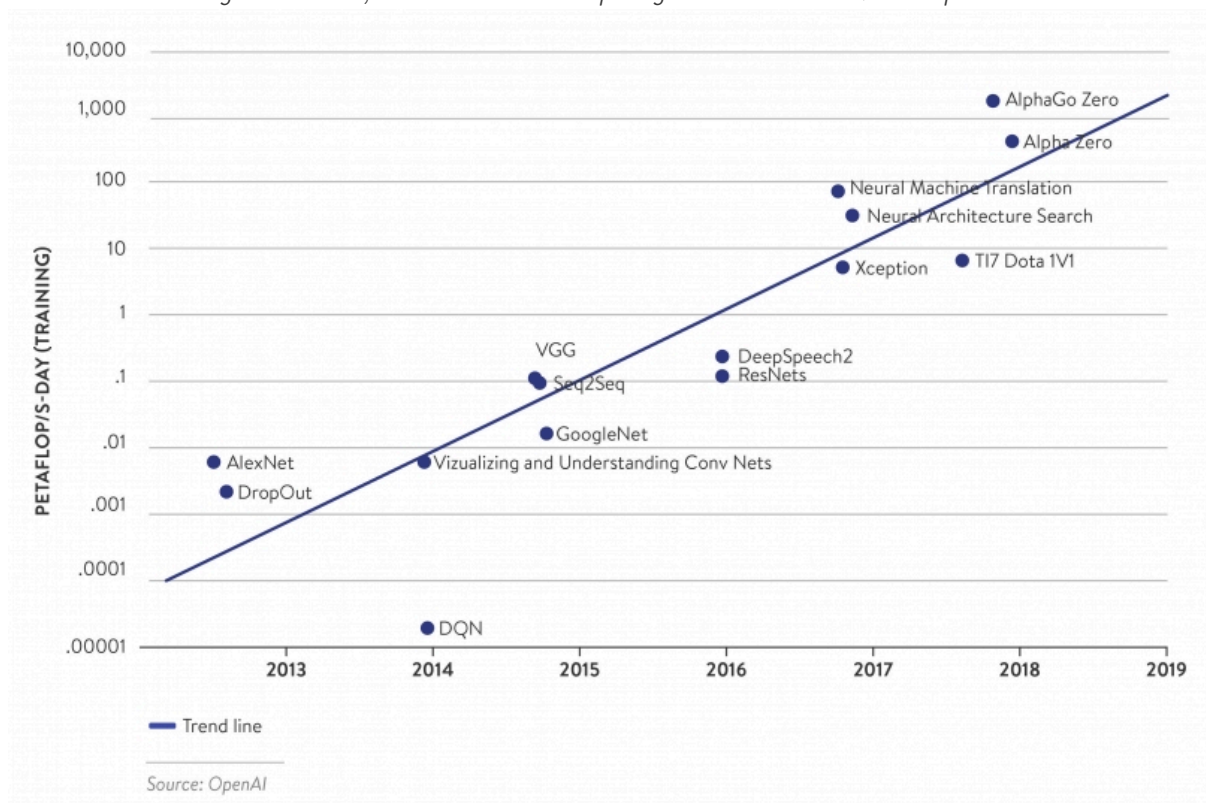
The three main factors leading to that period were a scarcity of computational power, data and overhype. Much has changed since in the landscape of computing power: many ML algorithms, even those trained on decently sized data sets, can be trained and run rather quickly and efficiently. However, computing power and hardware are not resolved issues and are still increasingly essential to developing advanced AI systems. The computing requirements in the most cutting-edge ML technology, as measured in petaflops (a unit of computing speed), has grown exponentially over the years (see figure 2).²⁶

24 “First AI textbook for high school students released”, China Daily, 11 June 2018, accessed 30 October 2018, <http://www.chinadaily.com.cn/a/201806/11/WS5b1de85fa31001b82571f4ca.html>.

25 Ray Kurzweil, *The Singularity Is Near* (London: Gerald Duckworth & Co Ltd, 2005), 263–264.

26 Graph via Open AI, “AI and Compute”, 16 May 2018, accessed 30 October 2018, <https://blog.openai.com/ai-and-compute/>.

Figure 2: A 300,000x Increase in Computing Power From AlexNet to AlphaGo Zero



Ensuring governments have access to this level of computing power is as important as having clean, well-labelled training data; without computing power, the models will struggle to get going. That said, governments need not build their own systems entirely by themselves. Procuring ML-capable cloud computing provides a way to benefit from the same computing power as leading private businesses. However, enumerating detailed requirements and analysing the different provider options should be near the top of departments' to-do lists.

RESPONSIBILITIES AND REGULATION

There are also several more opaque issues for policymakers to contend with. These cover economic and social disruption, privacy, ownership of data, and algorithmic fairness and ethics. This section highlights the key issues in each domain, emphasising that their resolution requires broad, ongoing engagement with a variety of industry, policy and civil-society actors, and perhaps even the wider public, given the political character of issues to do with inequality,

disruption and privacy. It should be noted that it is not possible to ‘do the ethics first’ and move on; this is a fast-moving field with ever-changing technologies as well as social and political consequences.

Disruption and Displacement

In most conversations about AI, questions about job losses, unemployment or retraining also appear. Although the precise effect of AI and automation on the workforce is unknown, when the current wave of AI promises so much in terms of efficiency and a step change in innovation, large-scale disruption and realignment of labour are inevitable.

A 2018 study by PricewaterhouseCoopers suggested that there will be multiple waves of automation, with a low risk of job displacement at first followed by up to 30 per cent of jobs being at risk by the mid-2030s, particularly among people with low education. However, the report also suggests that “AI, robotics and other forms of smart automation have the potential . . . to contribute up to \$15 trillion to global GDP by 2030”, which in turn will generate demand for many roles.²⁷

The headline figures may not be clear, but what seems certain is that individual tasks, rather than entire jobs, are more likely to be automated, while workers at risk will need to reskill to some degree to be more employable as the nature of tasks—and work generally—changes. A realignment of jobs, from old clusters of tasks to new ones, is therefore the more tractable issue for governments to consider. For those who are displaced, a 2013 Oxford University report suggested:

*As technology races ahead, low-skill workers will reallocate to tasks that are non-susceptible to computerisation – i.e., tasks requiring creative and social intelligence. For workers to win the race, however, they will have to acquire creative and social skills.*²⁸

27 “Will robots really steal our jobs? An international analysis of the potential long term impact of automation”, PwC Economics, accessed 30 October 2018, <https://www.pwc.co.uk/services/economics-policy/insights/the-impact-of-automation-on-jobs.html>.

28 Carl Benedikt Frey and Michael A. Osborne, “The future of employment: How susceptible are jobs to computerisation?”, *Technological Forecasting and Social Change* 114 (2017): 254–280.

The job of government, therefore, should not be to push back against technological progress; but neither is it to sit on the sidelines and watch as the forces of technology and big business march on. Retraining programmes will be needed to mitigate the disruption ahead, with a greater focus on creative thinking, social skills and care. This could not only soften the blow of displacement but also, perhaps, help more people find satisfying and fulfilling places in the workforce.

Privacy

ML is in many ways as good as the data that underlie it. And the data underneath are as available, granular and accurate as society allows. The need for personal data for providing personalised services and accurate classifications and predictions is clear. But this necessitates an honest discussion about the implications of data gathering on privacy. How aware should people be about the data collected and used to target them? How much data should governments have and use for ML purposes? How far should consent extend? These are difficult—and political—questions that policymakers cannot ignore.

These questions should not deter the use of ML, but they require a careful consideration of the issues. At a minimum this means an honest political debate; including such questions in civic education is also a step in the right direction.²⁹

Citizens' anxieties must be acknowledged and their rights guaranteed. Recent European Union legislation, the General Data Protection Regulation (GDPR), is a positive move towards protecting privacy and independence online as well as supporting the data-hungry world. But fast-moving technologies and social environments demand ongoing deliberation and revision.

Ownership of Data

The question of data ownership is nebulous, far reaching and difficult to keep tractable. On the one hand, regulation already does

²⁹ Catherine Miller, Rachel Coldicutt, Abbey Kos, "Only a third of people are aware that data they have not actively chosen to share has been collected. A quarter have no idea how internet companies make their money.", in "People, Power and Technology: The 2018 Digital Attitudes Report", Doteveryone, 2018, <http://attitudes.doteveryone.org.uk/>.

a good job of protecting raw personal data. On the other, as technologies have changed, concern is now focused on how companies process insights from myriad sources, often without transparency.

Moreover, some have sought to recharacterise data as labour, arguing that users should share in the value that companies generate from data. This can cover simple personal data or liking a page on social media: When users share or produce data on social media, do they have a right to share in the monetisation of those data? Can users organise collectively to withhold data en masse or take their insights elsewhere?

Here is a cruder example: Answering a CAPTCHA form involves labelling images that are then used as training data for ML models—supervised learning—so that algorithms can later recognise objects for themselves (so-called computer vision). Who should benefit from this work—the company and engineers who developed the AI algorithm alone or those who independently tagged the data over the years?

In the aftermath of the Cambridge Analytica episode, when the data analytics firm used a Facebook app to harvest the profile data of millions of users, which it then incorporated into aggressive political advertising campaigns—and amid a growing awareness of how data are monetised and a frustration with the scarcity of realistic competition in social media—the tech backlash has been getting louder and these questions have been raised increasingly frequently. These issues should be part of the debate about technology, data, ML and the surrounding topics. Policymakers should become familiar with them and help shape data and ML regulation together with the public, expert organisations such as Doteveryone and OpenAI and technology companies themselves.

Algorithmic Bias, Fairness and Ethics

The growing concern with the level of transparency of black box algorithms, computational models with reasoning that cannot be audited and assessed, is a subsection of a growing movement concerned more generally with the fairness and ethics of using ML or AI techniques in decision-making that can have grave social implications. Such is the case with algorithms involved in criminal

sentencing that came under fire for potential racial bias, or the mechanisms behind targeted advertising or health-insurance schemes offered based on the collection of personal data or available aggregate data.³⁰

On the horizon there are ethical considerations related to the form and amount of data collected as well as the design of fair systems that do not amplify existing gender and racial discrimination. However, just as governments should aim to be leaders in technical AI applications, so too should they aim to lead the moral space to build AI for the public good. As with the other issues in this section, the ethical and political consequences here are so fundamental that governments must engage with researchers and civil-society organisations to promote different disciplines and perspectives in the process of developing ML and AI-based techniques.

30 Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, “Machine Bias”, ProPublica, 23 May 2016, accessed 30 October 2018, <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>; Cathy O’Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (New York: Penguin Random House, 2016).

RECOMMENDATIONS

It is up to governments to realise AI's potential for public-service delivery. Policymakers should take the necessary steps to embrace AI's promise in government while preparing for its impact and complications. Disengaging will leave citizens dissatisfied and disaffected by government delivery, in a world of ever-smarter services and subsequently raised expectations. Policymakers should therefore focus on the twin priorities of resources and responsibilities.

RESOURCES

Policymakers should:

- **Sort out government data.** The development of high-quality data sets is crucial for the deployment of ML-based technologies. A model is only as good as the data that underlie it. Issues concerning data are sensitive but also extremely important. Given the amount of data tech companies already have on their customers, governments should be able to hold more data on their citizens to be able to deliver services of a similar quality. However, governments have even more responsibility to citizens and should lead the way in ensuring transparent data collection and use. Data should be labelled properly, and governments should inform citizens of the existence of such data and allow people to opt out of government data programmes. Governments can also introduce incentives for citizens to share data when possible.
- **Invest in skills by linking with public service.** The key to developing successful AI systems for the service of governments is investing in the right personnel. Expert, ML-specific skills take years to acquire and significant investment to encourage. As such, beginning this as early as possible is essential. Governments should not copy and paste technology solutions developed for industry—though it should build on them—but at the heart of good algorithm design must be human talent. This need for local knowledge in, or contracted by, government requires investment in committed, highly skilled workers. To

compete with industry, governments should offer sponsorship schemes for doctoral programmes in the fields of AI, ML and DL that include a commitment to several years of work in civil society or the public sector.

- **Ready procurement for ML.** As cloud computing has grown, so has the capacity for large organisations to scale their projects. Now, cloud providers are also making graphics processing units required for intensive ML applications available through the same infrastructure. Many of these cloud providers are the same companies that innovate in hardware, ensuring that customers can benefit from the latest speeds and capabilities as quickly as possible. As such, for governments to build hardware systems that suffice for the scale of all public-sector projects would be unwise and unwieldy. However, governments should be warier of procuring ready-made ML solutions as these may not reach the higher standards of risk, bias and fairness assessment required of public services. To that end, policymakers must begin preparations for a procurement strategy that incorporates these requirements, both technical and social.
- **Create AI-ready institutions.** The United Kingdom (UK)–based Nuffield Foundation has announced a new centre for data, AI and ethics, while the UK government funds several similar bodies.³¹ More countries could follow this path and encourage higher education institutions to open faculties that will train more dedicated researchers for the subject while populating the institutional ecosystem. These bodies would help shape strategy, research and governance of AI, as well as providing key anchors on which clusters of start-ups and civil-society actors can operate.

RESPONSIBILITIES

Policymakers should also:

³¹ Examples are the Ada Lovelace Institute (<https://www.adalovelaceinstitute.org>), the Turing Institute (<https://www.turing.ac.uk/>) and the Centre for Data Ethics and Innovation (<https://www.gov.uk/government/consultations/consultation-on-the-centre-for-data-ethics-and-innovation/centre-for-data-ethics-and-innovation-consultation>).

- **Establish feedback mechanisms involving the public and civil society.** The significant impact AI can have on public services and beyond also means much responsibility. AI can provide greater productivity, efficiency and satisfaction for all stakeholders, but only if they all get to participate in the process and flag issues as they arise. Policymakers' priority should be to give people better services, before simplifying work and easing the burden for governments. At the same time, civil servants and government workers are often the main experts on government service delivery. Their insights should be integral for developing systems that could aid their work and automate away the 'boring stuff', leaving more time for creative work, social interaction and meaningful engagement than for bureaucracy and administrative tasks. Governments should design the right feedback mechanisms by delivering smarter tools and continuously incorporating user input after such systems are deployed. Feedback should be incorporated into the design process as soon as possible to avoid vicious cycles that could arise from poorly executed systems.
- **Harness automation to reduce the burden on civil servants.** The conversation about AI and ML quickly lends itself to concerns about job losses and unemployment and can thus induce anxiety. Reports from the United States (US) and Europe show staff overworked and crashing under the burden of their current tasks. In the UK, a survey of children's social workers found that four out of five think their case load is unmanageable, while the US government recently reported that employees are overworked, a fact exacerbated by staffing gaps.³² However, this is not inevitable. AI can augment these workers, free them from easily automatable administrative tasks, and give them more time to think creatively and interact with colleagues and citizens. Automation's effect on the destruction or creation of jobs remains uncertain, but governments should nevertheless embrace AI: they can adapt existing tools to aid workers with

32 Luke Stevenson, "Four out of five social workers think their caseload isn't manageable", Community Care, 11 April 2018, accessed 30 October 2018, <https://www.communitycare.co.uk/2018/04/11/four-five-social-workers-think-caseload-isnt-manageable/>; Juana Summers, "The government is concerned about a lack of government workers", CNN, 9 February 2018, <https://edition.cnn.com/2018/02/09/politics/white-house-report-staffing-gaps-agencies/index.html>.

existing problems. Some staff changes are likely in the long run, but recognising that departments have influence over how they manage automation can change the conversation from one about redundancies to one about empowering employees.

- **Take responsibility for industry and public-sector ethics, fairness and transparency.** Algorithmic bias and fairness are serious concerns when it comes to deploying AI for social causes. Research in this area is taking its first steps, and governments should encourage and support it. Many companies and research groups dealing with AI have already introduced ethics and policy teams or advising committees, but governments must assume a coordinating role to ensure everyone is on the same page. Policymakers should make sure use of AI technologies is accompanied by adequate explanations for citizens who could be affected; transparency and data-usage guidelines should ensure algorithms will not harm citizens. Governments should develop standard procedures for citizens to demand information, raise concerns or protest outcomes. Coordination with insurance companies and the creation of arbitration and dispute-resolution mechanisms will be needed to protect citizens and businesses harmed by AI.³³
- **Start small; think big.** Some AI tools already exist and can make services significantly better even without transforming the entirety of government delivery. For now, policymakers should focus on measures with immediate impact, such as updating the use and maintenance of large data sets that can no longer be maintained manually; innovating citizen advice by introducing chatbots or AI-based assistance; using ML for fraud detection; automating form filling; and optimising resource use by drawing on past data to make future predictions.³⁴ However,

33 For more recommendations tailored to handle bias and fairness, see Alex Campolo, Madelyn Sanfilippo, Meredith Whittaker and Kate Crawford, “AI Now 2017 Report”, AI Now, 2017, https://ainowinstitute.org/AI_Now_2017_Report.pdf, and Eddie Copeland, “10 principles for public sector use of algorithmic decision making”, Nesta, 20 February 2018, <https://www.nesta.org.uk/blog/10-principles-for-public-sector-use-of-algorithmic-decision-making/>. While the author does not necessarily agree with all points made by these pieces, this is an intricate and essential issue that requires consideration of a diversity of opinions.

34 For more focused examples and recommendations for this level of AI, see Hila Mehr, “Artificial Intelligence for Citizen Services and Government”, Ash

governments should not lose sight of the bigger promise of AI, as exemplified by cases with potentially huge impact but low technical feasibility at present. Keeping future benefits in mind will help policymakers put the quicker wins into context, while investment in talent, data and computing power can help accelerate those oncoming advances. Establishing quality measures will ensure that more complex use cases can be managed safely and consistently.

- **Set realistic expectations, internally and publicly.** AI is an area of much promise, but it is also moving incredibly fast. The public is constantly exposed to headlines announcing new AI capabilities, sometimes beating the average human performance. This could lead to overhyped promises about how quickly automation and greater productivity are coming, or it could stoke fears of an imminent robot takeover. The truth is far from both. Policymakers should engage with the political questions of AI and advocate responsible development of AI for the betterment of as many people as possible. Governments should work closely with research groups and companies working on AI development, and consider investing in AI, particularly in applications that could ensure better government services.

ACKNOWLEDGEMENTS

This briefing was written with assistance from Andrew Bennett and Chris Yiu.

Policymakers should take steps to embrace the promise of AI while preparing for its impact and complications. Doing so would radically transform public-service delivery, allowing governments to better meet citizens' needs while putting data at the heart of decision-making.

FOLLOW US

facebook.com/instituteglobal

twitter.com/instituteGC

instagram.com/institutegc

GENERAL ENQUIRIES

info@institute.global

FIND OUT MORE
INSTITUTE.GLOBAL

Copyright © November 2018 by the Tony Blair Institute for Global Change

All rights reserved. Citation, reproduction and or translation of this publication, in whole or in part, for educational or other non-commercial purposes is authorised provided the source is fully acknowledged. Tony Blair Institute, trading as Tony Blair Institute for Global Change, is a company limited by guarantee registered in England and Wales (registered company number: 10505963) whose registered office is 50 Broadway, London, SW1H 0BL.