Government
Office for Science

Artificial intelligence: opportunities and implications for the future of decision making

Government Office for Science

Foreword

We are in currently in the foothills of a new technological revolution. Artificial intelligence has the potential to be as transformative in our lifetimes as the steam-powered economy of the 19th century.

Already it's letting us to talk to our smartphones, recommending us music, describing photos for the visually impaired and flagging up fire risks in cities.

In the near future we could see it deployed in everything from driverless cars, to intelligent energy grids, to the eradication of infectious diseases.

In government too we are looking at the potential applications of this technology in the delivery of public services.

Our Government Data Programme is increasing the number of projects and data scientists in government, while playing a leading role in establishing the appropriate use of these powerful new tools.

As one the world's leading digital nations, artificial intelligence presents a huge opportunity for the UK.

Get this right, and we can create a more prosperous economy with better and more fulfilling jobs. We can protect our environment by using resources more efficiently. And we can make government smarter, using the power of data to improve our public services.

As we've seen already in many areas, much routine cognitive work - the filing, sifting and sorting - can increasingly be automated, freeing people up to focus on the more human aspects of any job.

The Prime Minister has announced an independent review of modern employment practices, so that the support we provide businesses and workers keeps pace with changes in the labour market and the economy.

Artificial intelligence also poses new questions about ethics and governance, the responsible use of data and strong cyber defences. To realise the full potential of this revolution, again we have to be ready with answers.

I am pleased that the Royal Society and the British Academy are conducting a review that will consider how best the UK might manage the use of artificial intelligence.

This note sets out where the science is heading, describes some of the implications for society and government, and shows how we can responsibly use this technology to improve the lives and living standards of everyone in Britain.

It is a timely and important piece of work from the Government Chief Scientific Adviser.

Matt Hancock
Minister for Digital and Culture

Contents

Introduction

Artificial intelligence has arrived. In the online world it is already a part of everyday life, sitting invisibly behind a wide range of search engines and online commerce sites. It offers huge potential to enable more efficient and effective business and government but the use of artificial intelligence brings with it important questions about governance, accountability and ethics. Realising the full potential of artificial intelligence and avoiding possible adverse consequences requires societies to find satisfactory answers to these questions. This report sets out some possible approaches, and describes some of the ways government is already engaging with these issues.

Artificial intelligence is not a distinct technology. It depends for its power on a number of prerequisites: computing power, bandwidth, and large-scale data sets, all of which are elements of 'big data', the potential of which will only be realised using artificial intelligence. If data is the fuel, artificial intelligence is the engine of the digital revolution.

Much has already been written about the use of artificial intelligence and big data. This paper does not attempt to survey the whole field. Its origins lie in a seminar held at the British Academy in February 2016, chaired by Mark Walport, Government Chief Scientific Adviser and Mark Sedwill, Permanent Secretary at the Home Office, that discussed some of the legal and ethical issues around the use of artificial intelligence. The issues discussed there provide the core of this report, with additional material drawn from the views of a wide range of scientific and legal experts in the field, although we have sought to minimise detailed discussion of technical aspects in order to concentrate on the practical aspects of the debate. We hope that it serves as an introduction to the topic.

The report considers the following questions:

- What is artificial intelligence and how is it being employed?
- What benefits is it likely to bring for productivity?
- How do we best manage any ethical and legal risks arising from its use?

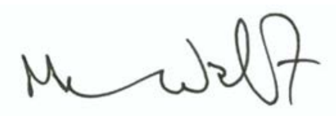

Sir Mark Walport
Government Chief Scientific Adviser

Mark Sedwill
Permanent Secretary, Home Office

What is artificial intelligence?

Artificial intelligence is more than the simple automation of existing processes: it involves, to greater or lesser degrees, setting an outcome and letting a computer program find its own way there. It is this creative capacity that gives artificial intelligence its power. But it also challenges some of our assumptions about the role of computers and our relationship to them.

Artificial intelligence is particularly useful for sorting data, finding patterns and making predictions. Current examples in everyday life are widespread: they include translation and speech recognition services that learn from language online, search engines that rank websites on their relevance to the user, and filters for email spam that recognise junk mail based on previous examples (see box for more uses). This list of applications is growing rapidly: artificial intelligence is enabling a new wave of innovation across every sector of the UK economy.

Artificial intelligence is a broad term (see box). More generally it refers to the analysis of data to model some aspect of the world. Inferences from these models are then used to predict and anticipate possible future events.

Some uses of artificial intelligence

Product recommendations from services such as Netflix and Amazon that evolve through users' web experiences are powered by machine learning.

The UK's 'smart motorways' use feedback on road conditions from embedded sensors and neural network systems to anticipate and manage traffic flow.

In financial markets, 'high-frequency trading' algorithms use pre-determined decision criteria to respond to market conditions many times faster than human traders are able to. Similar algorithms are being used by some financial advisors to automatically spot investment opportunities for clients.

Cornell University worked with machine learning specialists to identify the calls of right whales more accurately, making it easier to track individual whales.

Digital images from millions of satellite observations can be analysed for environmental or socio-economic trends using machine learning to identify patterns of change and development.

Statistical models are created using series of algorithms, or step-by-step instructions that computers can follow to perform a particular task. Computer algorithms are powerful tools for automating many aspects of life today, taking the step-by-step routines that underpin the administrative and operational tasks of organisations and digitising them, making them faster and more consistent. One approach to automation is to choose a series of rules to apply to inputs, leading a particular output. Most current medical self-diagnosis systems, both in books and online, use this logic. Certain combinations of answers to questions are deterministically linked to certain individual outputs. If you provide the same answers again, the algorithm will show the same result.

In recent years, however, available data and computing power have reached the point where it has become practical to develop **machine learning**: algorithms that change in response to their own output, or "computer programs that automatically improve with experience"[1].
Machine learning systems have often been shown to pick up difficult-to-spot relationships in data that may otherwise have been missed.

Most machine learning approaches are not restricted to producing a single prediction from given inputs.
Many algorithms produce probabilistic outputs, offering a range of likely predictions with associated estimates of uncertainty. The algorithms producing these probabilistic outputs are capable of being understood by humans. However, in the case of more complex machine learning systems (such as **deep learning**: see box), there are many layers of statistical operations between the input and output data. These operations have been defined by an algorithm, rather than a person. Because of this, not only is the output probabilistic, as with simpler algorithms, but the process that led to it cannot be displayed in human-understandable terms. This makes machine learning fundamentally different to the kinds of algorithms used to automate standard organisational routines.

Terminology

The range of different statistical techniques referred to here with the general term 'artificial intelligence' have emerged over a long time from many different research fields within statistics, computer science and cognitive psychology.

Consequently, authors from different disciplines tend to make different distinctions between terms like 'machine learning' and 'machine intelligence', using them to refer to related but distinct ideas.

As these techniques have been applied in different business areas they've become relevant to other tasks – so they're likely to feature also in discussions about 'data mining' and 'predictive analytics'.

While this paper looks ahead to a time when machine learning is more widespread than at present, many of the opportunities and challenges it discusses arise in other contexts too. So rather than be distracted with an academic discussion about terminology, we've chosen to use the umbrella term **artificial intelligence**.

There are many different kinds of algorithm used in machine learning. The key distinction between them is whether their learning is '**unsupervised**' or '**supervised**'.
Unsupervised learning presents a learning algorithm with an unlabelled set of data – that is, with no 'right' or 'wrong' answers – and asks it find structure in the data, perhaps by clustering elements together – for example, examining a batch of photographs of faces and learning how to say how many different people there are. Google's News service[2] uses this technique to group similar news stories together, as do researchers in genomics looking for differences in the degree to which a gene might be expressed in a given population, or marketers segmenting a target audience.

Supervised learning involves using a labelled data set to train a model, which can then be used to classify or sort a new, unseen set of data (for example, learning how to spot a particular person in a batch of photographs). This is useful for identifying elements in data (perhaps key phrases or physical attributes), predicting likely outcomes, or spotting anomalies and outliers. Essentially this approach presents the computer with a set of 'right answers' and asks it to find more of the same.

---

[1] Mitchell, T. (1997), *Machine Learning*
[2] http://news.google.com/

Current interest in machine learning is focused on ‘**deep learning**’, a supervised learning technique combining layers of neural networks to automatically identify the features of a data set that are relevant to decision-making. Deep learning is a powerful addition to the machine learning repertoire: however, it requires very large amounts of data to be effective.
The London-based firm DeepMind (owned by Google) is a world leader in this technique.

Deep learning

Deep learning is a subset of machine learning that depends on using layers of non-linear algorithmic processes to find patterns or classify data. There are many different techniques within this general approach – but the key feature is that they each use a layered or staged design, in which outputs from the previous layer are used as inputs for the next.

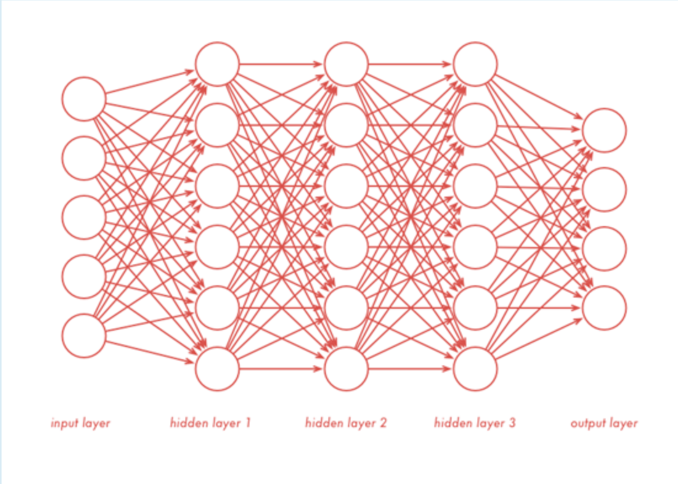

Central to the interest around machine learning is the potential it offers for autonomous decision-making.
Many algorithmic processes can be used to make decisions without human input (such as when mortgage providers automatically approve lending to individuals based on their credit score).
But real autonomy comes when a system is able to learn continuously and make deductions about the world without human input. For example, self-driving cars are able to make real-time decisions about speed and direction without input from a human driver, using many interlinked machine learning systems to make sense of information about their surroundings. They are not following pre-programmed decisions but responding to changes around them.

Using autonomous decision-making in other areas of society would be a significant change in the way we use data and make choices, bringing with it important questions around accountability and trust. This is particularly true for its use by government, in light of the compact that exists between it and citizens, and its responsibility for their wellbeing and security. At the present time there is generally agreement amongst experts and policy-makers that important decisions must involve a ‘human in the loop’ – but the exact nature of their role, or the degree to which they influence the outcome, is something that is likely to evolve as the technology develops over time.

Artificial intelligence for innovation and productivity

Artificial intelligence holds great potential for increasing productivity, most obviously by helping firms and people use resources more efficiently, and by streamlining the way we interact with large sets of data. For example, firms like Ocado and Amazon are making use of artificial intelligence to optimise their storage and distribution networks, planning the most efficient routes for delivery and making best use of their warehousing capacity. Artificial intelligence can help firms do familiar tasks in more efficient ways. Importantly, it can also enable entirely new business models and new approaches to old problems. For example, in healthcare, data from smart phones and fitness trackers that is analysed using new machine learning techniques can improve management of chronic conditions as well as predicting and preventing acute episodes of illness.

Artificial intelligence can help both companies and individual employees to be more productive. Routine administrative and operational jobs can be learned by software agents ('bots'), which can then prioritise tasks, manage routine interactions with colleagues (or other bots), and plan schedules. Email software like Google's Smart Reply can draft messages to respondents based on previous responses to similar messages. Newsrooms are increasingly using machine learning to write sports reports and to draft articles: in the office, similar technology can produce financial reports and executive briefings.

Artificial intelligence can reduce the burden of searching large sets of data. In the legal sector, groups like ROSS, Lex Machina and CaseText are using artificial intelligence to sift court documents and legal records for case-relevant information. Other firms are using similar techniques as part of due diligence. Artificial intelligence can also offer a way of interacting with these datasets, with platforms such as IBM's Watson able to support expert systems that can answer factual natural language questions. For cybersecurity firms, artificial intelligence offers a way of recognising unusual patterns of behaviour in a network.

These examples focus on using software to do the same thing as humans but, in many cases, analysing data of volume or complexity that is beyond the analytical capability of individual humans. Indeed, artificial intelligence is not a replacement, or substitute for human intelligence. It is an entirely different way of reaching conclusions. Artificial intelligence can complement or exceed our own abilities: it can work alongside us, and even teach us, as shown by Lee Sedol's unbroken string of victories since playing AlphaGo[3]. This offers new opportunities for creativity and innovation. Perhaps the real productivity gain from artificial intelligence will be in showing us new ways to think.

The UK is a world leader in the science underpinning this technology, with a rich ecosystem of investors, employers, developers and clients, and a network of supporting bodies such as the Alan Turing Institute. Innovations developed first in universities such as Cambridge, Oxford, Imperial and University College London have found their way into tools used by millions around

---

[3] "Lee Sedol has won every single game he has played since the #AlphaGo match inc. using some new AG-like strategies - truly inspiring to see!" Demis Hassabis, CEO of DeepMind, 4th May 2016 (https://twitter.com/demishassabis/status/728020177992945664, http://gokifu.com/player/Lee+Sedol)

the globe, with increasing numbers of UK startups electing to remain in the UK, further strengthening expertise and capability in this country.

The potential is driving a rapid take-up of artificial intelligence across a range of sectors[4]. In the words of the technology pundit Kevin Kelly: "the business plans of the next 10,000 startups are easy to forecast: Take X and add AI"[5]. Estimating the size of this growth is challenging given the different definitions of 'artificial intelligence', 'machine learning' and related terms: it is also hard to specify where sectors like 'big data' end and 'machine learning' begins. But a 2015 US report[6] suggests that the global market in 'robots and artificial intelligence-based systems will grow from $58bn in 2014 to $153bn in 2020.

Looking at big data more widely, a report earlier this year on the value of big data and the internet of things estimated £240 billion in cumulative benefits to the UK between 2015-20; manufacturing should derive the greatest benefits, with the greatest gains across all sectors set to come from efficiency savings[7]. Another report from 2015 predicts major operational savings for European governments by using big data, in addition to opportunities to increase tax revenues and reduce fraud and error[8]. A 2014 study of 500 UK businesses, meanwhile, concluded that those who make better use of customer and consumer data are between 8 and 13 per cent more productive than firms who don't[9]. More broadly-defined forecasts for the impact of systems combining robotics, data and artificial intelligence – sometimes labelled "Industry 4.0"[10] – also promise substantial gains.

According to McKinsey, companies themselves anticipate Industry 4.0 to increase revenues by 23 per cent and productivity by 26 per cent[11]. Artificial intelligence has a central role in enabling all of this growth.

---

[4] An overview of current work is offered by Shivon Zillis of Bloomberg BETA at: www.shivonzilis.com/machineintelligence.

[5] www.wired.com/2014/10/future-of-artificial-intelligence/

[6] www.ft.com/fastft/2015/11/05/robotics-ai-become-153bn-market-20-bofa/, www.bofaml.com/content/dam/boamlimages/documents/PDFs/robotics_and_ai_condensed_primer.pdf

[7] *The Value of Big Data and the Internet of Things to the UK Economy*, CEBR for SAS, February 2016.

[8] *Big data: The next frontier for innovation, competition, and productivity*, McKinsey Global Institute, 2011.

[9] Bakhshi, Bravo-Biosca and Mateos-Garcia (2014) – from *Data Driven Innovation* (OECD).

[10] Bundesministerium für Bildung und Forschung (2013), *Zukunftsbild "Industrie 4.0"*, PwC, *Industry 4.0: Building the digital enterprise*.

[11] McKinsey Industry 4.0 Global Expert Survey 2015.

The use of artificial intelligence by government

Government is already using data science techniques such as machine learning, and through the work of the Government Data Programme their use is growing[12]. These techniques are providing insights into a range of data, from feedback on digital service delivery to agricultural land use through the analysis satellite images. As their sophistication improves more benefits may be realised. For example, we might:

- Make existing services – such as health, social care, emergency services – more efficient by anticipating demand and tailoring services more exactly, enabling resources to be deployed to greatest effect.
- Make it easier for officials to use more data to inform decisions (through quickly accessing relevant information) and to reduce fraud and error.
- Make decisions more transparent (perhaps through capturing digital records of the process behind them, or by visualising the data that underpins a decision).
- Help departments better understand the groups they serve, in order to be sure that the right support and opportunity is offered to everyone.

Other applications will become apparent as the use of data and artificial intelligence becomes more mainstream.

Government is a special body, with unique obligations that do not fall on private organisations. It must be transparent about the way it acts, follow due process, and be accountable to its citizens. This means there are special responsibilities for government, beyond the general points outlined above, which follow from its use of artificial intelligence and big data. Recognising this, the government has published a guide to the ethical use of data science tools within government for government analysts[13]. This first iteration of a code of practice was developed with extensive external input and iterated with data scientists inside government to make it as practical and useful as possible.

Two uses with particular relevance for government are highlighted here: the use of artificial intelligence to advise, and possible legal implications of the use of artificial intelligence.

Advice

Artificial intelligence has a clear advisory role to play beyond traditional ‘decision-support systems’, supporting decision-making through assembling relevant data, identifying pertinent questions and topics for the attention of policy-makers, and in helping to generate written advice. Already, government is beginning to find value in using artificial intelligence to assist public servants in the delivery of digital services. It is likely that many types of government decisions will be deemed unsuitable to be handed over entirely to artificial intelligence systems. There will always be a ‘human in the loop’. This person's role, however, is not straightforward. If they never question the advice of the machine, the decision has *de facto* become automatic and they offer no oversight. If they question the advice they receive, however, they may be thought reckless, more so if events show their decision to be poor.

---

[12] https://data.blog.gov.uk/category/data-science/
[13] www.gov.uk/government/uploads/system/uploads/attachment_data/file/524298/Data_science_ethics_framework_v1.0_for_publication__1_.pdf