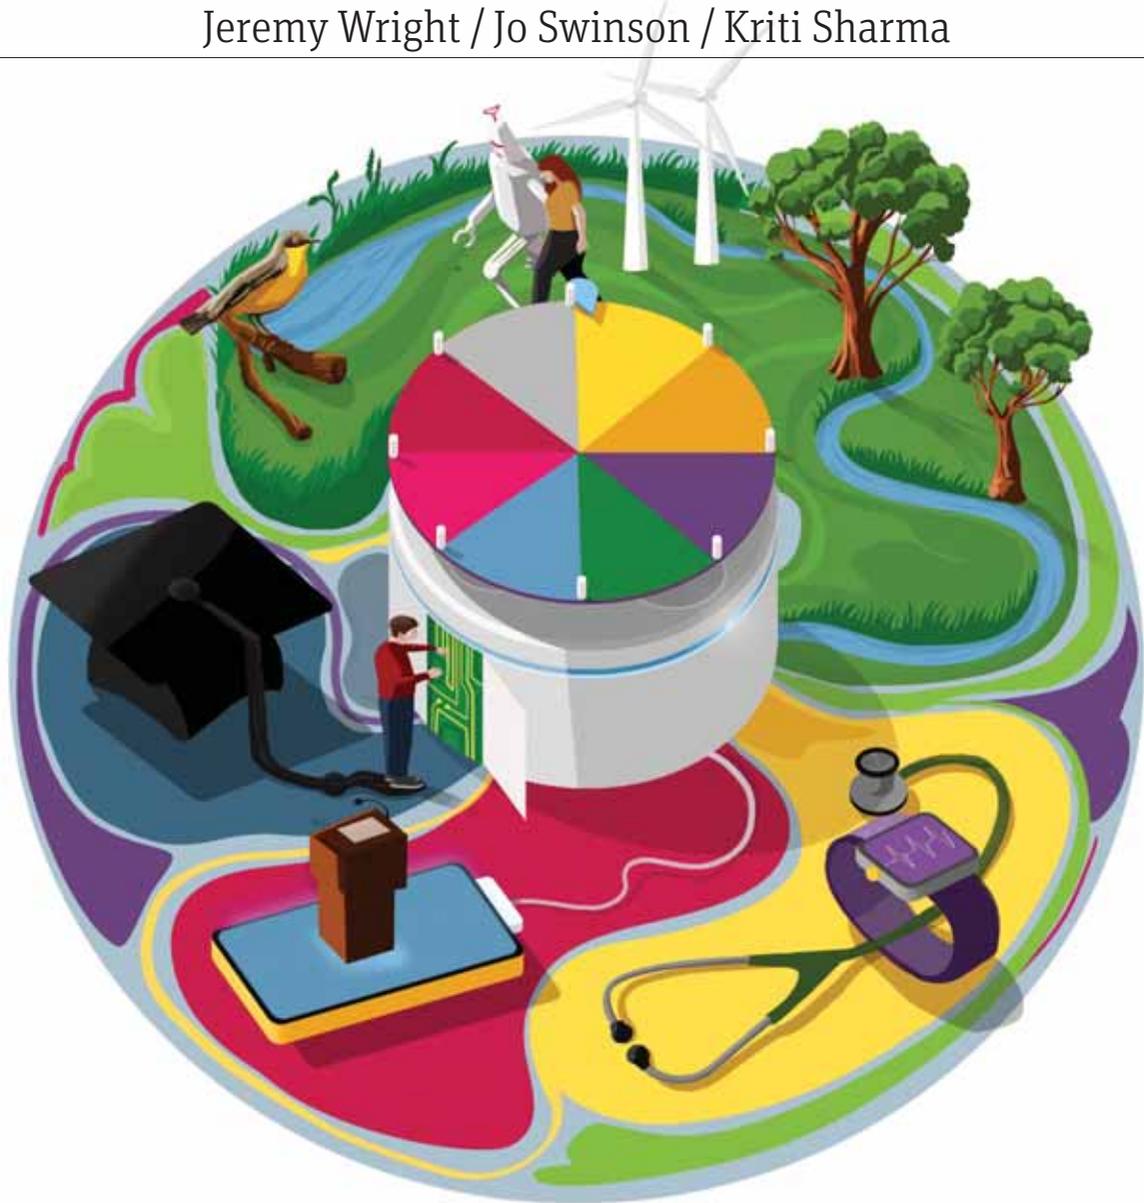


Spotlight

ARTIFICIAL INTELLIGENCE: THE FIRST LAWS OF ROBOTICS

Jeremy Wright / Jo Swinson / Kriti Sharma



Spotlight



Read more in-depth interviews
and features and download
full policy reports at:
newstatesman.com/spotlight

The danger of the black box



There are many examples of fiction for previous decades that, with hindsight, seems to have predicted the modern age, but few appear as prescient – and from such a distance in time – as EM Forster’s *The Machine Stops*, written in 1909. The story begins with a scene that many have argued presaged internet use by a century. A woman sits alone in a room, surrounded by technology. She can summon food or entertainment at the touch of a button, but she is bothered by constant notifications from an unseen community – “she knew several thousand people” – but she never sees anyone in person. In a line every internet user will relate to, she says: “Here I am in the dark, wasting my time”.

But Forster’s story may come to seem even more prescient in the years to come. We already live in a time in which most people are not able to explain more than a superficial amount about the technology by which they are surrounded every day. This is a new situation for humanity, and a confusing one, but it is still possible to learn how things work. As artificial intelligence becomes more sophisticated and more widespread, people will begin to use technology that no human is able to explain.

In fact, this is already happening. It is already the case that engineers and researchers are unable to explain the ways in which the AI algorithms they’ve created reach the conclusions that they do. There are engineers developing self-driving cars who are not able to get the information about why, in a given instant of a journey, the car decided to act in a certain way. Deep learning techniques applied to medical science appear to be very good at predicting the onset of some diseases, but because the software has written itself, researchers are not able to say how they work.

Few technologists worry about a robot uprising. But there is increasing unease that if the technologies that run our world become inexplicable, we may stumble into the situation Forster predicted: “No one confessed the Machine was out of hand. Year by year, it was served with increased efficiency and decreased intelligence. The better a man knew his own duties upon it, the less he understood the duties of his neighbour, and in all the world there is not one who understood the monster as a whole.”

6 / Jeremy Wright

The digital and culture secretary on AI optimism

8 / Buying the debate

Spotlight investigates claims that AI ethics are subject to conflicts of interest

12 / Innovation and impact

When AI brings 50,000 new patients into a GP’s surgery, can NHS funding structures cope?

21 / Code and conduct

UCL’s Rainer Kattel on the beginnings of AI legislation

28 / Jo Swinson

The Lib Dems’ deputy leader on the ethics of AI

31 / Kriti Sharma

On eliminating bias in AI

NewStatesman

Standard House
12-13 Essex Street
London, WC2R 3AA
Subscription inquiries:
digital.subscriptions@
newstatesman.co.uk

Account Managers

Bishara Abdi
Jugal Lalsodagar
Justin Payne
Dom Rae
Cyrus Ramezani

Commercial Director

Peter Coombs

Special Projects Editor

Will Dunn
Special Projects Writers
Jonny Ball
Rohan Banerjee

Design and Production

Leon Parks
Cover illustration
Sam Falconer



WINNER

First published as a supplement to the *New Statesman* of 7th June 2019. ©New Statesman Ltd. All rights reserved. Registered as a newspaper in the UK and US. The paper in this magazine is sourced from sustainable forests, responsibly managed to strict environmental, social and economic standards. The manufacturing mills have both FSC and PEFC certification and also ISO9001 and ISO14001 accreditation.

This supplement can be downloaded from:
newstatesman.com/page/supplements

News



SHUTTERSTOCK

OECD adopts AI principles

Will Dunn

Forty-two countries, including all 36 members of the Organisation for Economic Co-operation and Development (OECD), adopted a set of principles for the uses, operation and accountability of AI at the OECD's recent annual meeting in Paris.

The principles state that AI should benefit “people and the planet”, respect laws, rights, “democratic values and diversity”, and that they must be transparent and safe. Furthermore it places accountability for the decisions of AI on those who develop or deploy the technology.

The OECD's principles have been warmly endorsed by the United States, and also have the backing of the European Commission.

While OECD Principles are not legally binding, they have in the past informed national and international law. The organisation's Principles of Corporate Governance, first published in 1999, have guided the development of legislation in many countries, while individual organisations such as the London Metals Exchange have used principles developed by the OECD to set their trading standards.

Fake news created by AI algorithm

Jonny Ball

Scientists at the University of Washington have developed a new algorithm, Grover, which is able to generate fake news stories that are, according to research, more believable than fake news written by humans.

The Grover software mimics the writing style of various publications, including the *New York Times*, the *Washington Post* and *Wired*, and can generate article names, author names and article leads based on just an inputted headline. After being read by human subjects, the AI algorithm, which has created fake news pieces on the links between autism and vaccination, as well as on Donald Trump's impeachment by Congress, was found to be more believable than man-made fake news.

The developers at the university's Allen Institute for Artificial Intelligence intend Grover to be used to detect, rather than create, fake news online, and the system can spot misinformation with a high degree of accuracy. The tool will be released to the public and will remain open source, which the makers say will give the public “recourse against adversarial attacks”.

AI generates faces from voices

Rohan Banerjee

A new artificial intelligence technology, Speech 2 Face, developed by researchers at the Massachusetts Institute of Technology, can generate images of people based on an analysis of their voice. The team at MIT trained the AI using videos from YouTube and elsewhere online to recognise various vocal characteristics in order to use them to construct a “rough” picture of what it

understood the person speaking might look like.

The vocal samples used to train Speech 2 Face were mainly those of celebrities. The level of accuracy with which the images reflected their corresponding video clips was mixed, but with more successes than failures.

The researchers noted the potential privacy concerns which may arise from this project in a statement. “We feel that it is important to explicitly discuss in the paper a set of ethical considerations due to the potential sensitivity of facial information.” The team insisted that any “practical use of this technology will be carefully tested to ensure that the training data is representative of the intended user population.”



David Cameron joins AI firm

Will Dunn

The former Prime Minister has been appointed chair of the advisory board of Afiniti, where he joins Elizabeth Murdoch, the former French PM Francois Fillon, and the former CEOs of BP, Verizon, Thomson Reuters and Sony. Princess Beatrice of York also works for Afiniti as vice president of partnerships and strategy. Afiniti is an American company that uses artificial intelligence to gather information about customers

– including from social media - when they speak to call centres, and to use that information to improve sales or customer service.

Since his abrupt resignation following the result of the UK’s 2016 Brexit referendum, Cameron has taken a number of senior advisory positions at companies, charities and other organisations, including consultancy positions at the American financial services firm First Data and Illumina, a biotechnology company also from the United States, both of which have developed artificial intelligence solutions in their businesses.

Estonia develops “robot judge”

Rohan Banerjee

The Republic of Estonia is developing a “robot judge” to handle small claims court disputes, with the aim of freeing up human legal professionals to work on more serious cases.

The Baltic nation already has a reputation for technical prowess, with a well-known “e-residency” programme that allows non-residents to set up businesses from abroad, and to access all of the country’s government services online.

Estonia’s Ministry of Defence asked Ott Velsberg, the country’s chief data officer, to lead on the robot judge project, in which legal databases are used to train algorithms to produce opinions on certain disputes. The verdict, however, will still be reviewed and formalised by a human judge.

The project is in its infancy, with no set date announced for a pilot model, but Velsberg is optimistic and enthusiastic about artificial intelligence’s potential in this area. He told *Wired*: “We want the government to be as lean as possible. People worry that if we lower the number of civil employees, the quality of service will suffer. But the AI agent will help us.”



Chicago, Detroit and San Francisco in AI face-off

Rohan Banerjee

Facial recognition technologies which use artificial intelligence to analyse video footage are dividing opinion in the United States. While police departments in Chicago and Detroit have opted to use software that can identify people from images captured on surveillance cameras, as revealed in a report from Georgetown University, San Francisco has issued a ban on its authorities acting in this way.

Georgetown’s investigation found that law enforcement agencies in Chicago and Detroit have purchased real-time facial recognition systems from the company DataWorks Plus. A description on the company’s website says the technology, called FaceWatch Plus, “provides continuous screening and monitoring of live video streams.”

Facial recognition has been used by police on static images to identify suspects for a long time, but the introduction of real-time capabilities has coincided with advancements in AI.

Critics of real-time facial recognition suggest that it could infringe on civil liberties, especially if attached to cameras in public spaces. Studies from Microsoft and IBM, meanwhile, have shown that many facial recognition systems struggle to accurately identify ethnic minorities and transgender people.

Innovation must be guided by social responsibility, writes **Jeremy Wright**, Secretary of State for Digital, Culture, Media and Sport

Technology and the common good



When we look back on this remarkable era of technological change, we will not judge the progress we have made by the numbers on balance sheets. We will judge it by the impact it has had on humanity as a whole, and whether it has made us safer, happier and more fulfilled. Governments have an important role in supporting the mission-driven organisations that want to use technology for the greater good, especially in fields such as artificial intelligence and automation.

The United Kingdom is committed to being open and optimistic about the world beyond us and the opportunities provided by new technologies. A recent report from Tech Nation named the UK as a “critical hub in the global tech ecosystem”, with a £2.3bn valuation of UK-based “tech for good” organisations.

From financial inclusion to sanitation, healthcare to housing, technology has the power to enable and improve a wide

range of sectors. Earlier this year, I announced the UK government will invest £1m to drive social tech innovation in civil society, to help develop solutions to tackle social isolation and bring communities together. Successful participants will be rewarded with a cash incentive and ongoing business support.

And we have also backed the UK’s leading dedicated supporter in social tech ventures – the Social Tech Trust – to set up a new investment fund. This fund will provide ventures with the access to capital that they need at the right time. The aim is to raise up to £30m for this investment fund, to help ventures focused in three key areas of social transformation: health, wealth and communities.

To continue this momentum, Microsoft is also working in partnership with the Social Tech Trust to deliver an immersive accelerator programme,



focused on “AI for good”. Ventures have access to resources to help them scale solutions in the areas of AI for accessibility and AI for environmental sustainability. The UK represents a natural home for tech pioneers such as these, who have big ideas and big ambitions to make the world a much better place.

Part of this is about having an immigration system that welcomes the world’s top tech talent, which is why we have also doubled the number of Tier 1 Exceptional Talent visas. And we have

Innovation comes with a duty of care

SHUTTERSTOCK/MIKEDOTTI

opened two new visa routes for tech specialists. The startup route for those starting a business for the first time and the innovator route for more experienced business people with funds to invest in their firm. We are open to those with the skills and the determination to join this tech for good movement. But the other way we can show our commitment to this agenda is by playing our part in the big debates surrounding the development of new technologies.

Over the past few months we have seen a number of major questions coming to the fore about the development of new technologies. Issues such as content moderation, online disinformation and data ethics are now the subject of mainstream debate, as digital technology plays an increasingly important part in our everyday life. It is vital we work with industry to provide coherent answers and solutions to these issues to foster a level of optimism and enthusiasm towards the digital revolution.

I believe we find ourselves at a crossroads here, particularly around online harms. The tragic events in New Zealand in March showed just how quickly illegal terrorist and extremist content can spread. We are also seeing increasing numbers of people closing their social media accounts after unacceptable online abuse. We cannot allow this behaviour to undermine the very real benefits that the digital revolution can bring.

The government is setting out, therefore, how to create a new duty of care, establishing that tech companies have a responsibility for the safety of their users, and must take reasonable steps to tackle harmful content and activity. Compliance will be overseen and enforced by an independent regulator, with significant penalties available to it.

But government and regulators cannot solve these problems on their own. Technology itself has a crucial role to play. I have seen some groundbreaking work already, including AI-enabled software that identifies terrorist content

and apps to help young people monitor the time they spend online.

To encourage this, the new regulator will also have broader responsibilities to promote the development and adoption of technology to tackle online harms. The government will also work with industry and civil society to develop a framework for safety by design – setting out clear principles on how to include online safety features in new applications.

From online shopping and suggesting songs for our playlists to diagnosing illnesses, algorithms and artificial intelligence are playing a greater role in our everyday lives. It is important the use of AI and data is fair and accurate. To help address this, the Centre for Data Ethics and Innovation – a world-leading advisory body designed to make sure data-driven technology delivers the best possible outcomes for society – launched a call for evidence to support its reviews into online targeting of content, products and services to different individuals and bias in algorithmic decision making.

And I know a number of countries have already been in touch with the centre to find out more about their approach. The centre will be looking internationally to work with partners from around the world, and to learn from best practice wherever it can be found. With this in mind, I am also pleased to have appointed leading experts to our new AI Council, chaired by Tabitha Goldstaub. This is an independent committee bringing together experts from a range of sectors to promote growth in the artificial intelligence sector and encourage its responsible adoption across the economy.

The aim is to realise the full potential of AI to the economy in a way that works for everyone, to help us put in place the right skills, data and ethics so we can all make the most of AI technologies. We remain committed to boosting the UK’s global position as a leader in these technologies, and I am looking forward to meeting leaders from across the sector during London Tech Week to continue this important conversation.

How Big Tech funds British research into AI ethics



The Silicon Valley giants are spending billions of pounds developing AI. But they are also funding the ethical debate set to shape the technology's future, reports Oscar Williams

Most people are unaware of the extent to which algorithms already make life-changing decisions on their behalf. Software can diagnose illnesses, shortlist you for a job interview or assess whether you should be granted parole. As machine learning and artificial intelligence become more sophisticated, automated decision-making will become ever more influential, and the rules that underpin these decisions may one day have as much effect on daily life as the laws of nation states. But while elected governments draw up legislation in public, the technology industry is spending millions of pounds quietly attempting to shape the ethical debate about the decisions machines make.

Spotlight and *NS Tech* have spent the past two months investigating the role that the industry plays in supporting British research into AI ethics. Data obtained under freedom of information

laws reveals that Google has spent millions of pounds funding research at British universities over the last five years. Oxford University alone has received at least £17m from Google.

The search giant, whose parent company was worth more than \$750bn at the time of going to print, has been fined more than €6bn by the European Commission for its business practices. The French data protection authority fined the company a further €50m in January, and in the UK the Information Commissioner's Office is currently investigating claims that the company has breached EU data laws.

While much of the funding provided to Oxford goes towards technical research, Google and DeepMind – a British AI company wholly owned by Google's parent company, Alphabet – have also supported work at the Oxford Internet Institute (OII) exploring the ethics of AI, the civic responsibilities of



tech firms and research into the auditing and transparency of automated decision-making (ADM).

A number of academics at the OII, most notably Professor Luciano Floridi, Associate Professor Sandra Wachter and Dr Brent Mittelstadt, are prolific public commentators on ethical AI and ADM. Floridi is an advisor to the Information Commissioner's Office, a board member for the UK government's Centre for Data Ethics and Innovation and a representative of the European Commission's High Level Expert Group (HLEG) on AI. He also served on Google's short-lived AI ethics board earlier this year.

In June 2017, the academic journal *International Data Privacy Law* published an article written by Floridi, Wachter and Mittelstadt which argued that European data law does not give EU citizens the right to an explanation of decisions made about them by machines.

Other experts expressed concern at what this might mean when software made decisions that were contentious. According to a page on the OII's website, the same researchers then received a grant from Google covering work on the "underlying ethical and legal principles, as well as latent policies shaping the auditing [of automated decision-making] debate". A spokesperson said the funding enabled the delivery of an internal workshop at the company.

Later that year, DeepMind also committed to funding a series of work at the OII on a similar subject. A page on DeepMind's website states that the funding would support a "research project on explainable and accountable algorithms and automated decision-making in Europe". The project started in January last year and is led by Wachter, but also supports some of Mittelstadt's work.

When *Spotlight* asked the OII for evidence of papers that acknowledged funding from either Google or DeepMind, a spokesperson presented four articles, each of which were written or published in or after April this year.

However, two papers identified by *Spotlight*, which were authored by Wachter and Mittelstadt, published since January 2018 and address approaches to explaining algorithms, do not acknowledge that funding had been received from DeepMind. Other academics consulted for this feature said it would be normal practice to disclose any corporate funding in the realm of a paper's topic, and if necessary to disclaim it for the sake of transparency.

An OII spokesperson said that only a small proportion of its funding comes from business, that it is committed "to the highest standards of academic independence" and that all projects are carried out "in line with Oxford's culture of academic freedom", meaning "external funding does not compromise the value or integrity of the research output". The institute also noted that Wachter and Mittelstadt have recently written papers which acknowledge DeepMind's funding, are critical of the industry and

call for more stringent regulation.

"In the interests of transparency," the spokesperson said, "we publish the details of all our corporate supporters on our website and the University of Oxford publishes a consolidated breakdown of its annual income and expenditure." However, unlike many other organisations contacted for this piece, the OII refused to reveal how much the individual grants it had received from Google were worth. Another spokesperson for the university said that "the publication of research results, for example in journal articles or conference proceedings, will always acknowledge the source of any funding that has supported the work". DeepMind said it was passionate about supporting the research community and proud to have supported Oxford through several "unrestricted donations".

Google has been funding research into ethics of technology for more than a decade. In 2017, the Campaign for Accountability (CfA) claimed to have identified 329 research papers published from 2005 onwards about "public policy matters of interest to Google that were in some way funded by the company". The company has also partnered with a number of European organisations to launch policy research units, including Readie - a joint initiative with the innovation agency, Nesta - in London.

"Google uses its immense wealth and power to attempt to influence policymakers at every level," said Daniel Stevens, a campaigner and executive director of CfA, at the time. Addressing the research identified by his organisation as having been funded by the company, he said: "At a minimum, regulators should be aware that the allegedly independent legal and academic work on which they rely has been brought to them by Google."

However, it later emerged that CfA had in turn received funding from Oracle, which had been fighting Google in court. A recent call for papers on corporate funding published by the University of Amsterdam observed that the incident "perfectly captures the complexity the

Funding can create a “level of dependency”

situation [industry funding] creates around science”.

A Google spokesperson said that “academic institutions must properly disclose our funding” and that grants are only offered “for discrete pieces of research, not to shape academics’ subsequent scholarship. The researchers and institutions to whom we award research grants will often publish research with which we disagree.”

The company is not alone in funding this sort of research. Facebook announced in January that it was set to donate \$7.5m to the Technical University of Munich, to fund the launch of a new AI ethics research centre. Two months later, Amazon announced it would partner with the US National Science Foundation (NSF) “to commit up to \$10m each in research grants over the next three years focused on fairness in AI”.

Academics are divided over the ethics of such funding. While some say businesses play an important role in supporting research, others question whether this work can ever be entirely independent. “It creates a level of dependency, especially in an age where there’s such pressure to get external funding,” says Lina Dencik, the co-director of the Data Justice Lab at Cardiff University and a member of the Funding Matters initiative. “In that context, it’s even more pressing that there are more discussions around who universities should and shouldn’t accept money from.” Oxford confirmed that it had not rejected any offers of funding from Google, Amazon, Facebook, Apple or Microsoft over the last five years.

“A university abdicates its central role when it accepts funding from a firm to study the moral, political and legal implications of practices that are core to the business model of that firm,” writes Yochai Benkler, a Harvard Law School professor, in a recent article for *Nature*. “So too do governments that delegate policy frameworks to industry-dominated panels.”

Last year Thomas Metzinger, a German philosophy professor, was named as one of the select group of ethicists who

had been appointed to the European Union’s AI expert group. Around half of the 52 members of the HLEG were industry representatives. Google, Facebook, Amazon, Microsoft and Apple were represented by an organisation called DigitalEurope, which has since denied that the group was tilted in the industry’s favour.

Metzinger and a machine learning expert were tasked with drawing up the “red lines” for the use of AI in Europe. “We did this for six months,” he says. “We talked to many other experts and extracted their proposals, formulated it and put it in a document of non-negotiable normative principles.” Two of the “red lines” Metzinger proposed were that AI should never be used to build either autonomous lethal weapons or social credit score systems akin to China’s. But after he submitted his proposals, some members of the group pushed for a softer stance.

“Suddenly there was this categorical sentiment by various industry representatives and others that this phrase ‘red line’ cannot appear in that document anymore,” he recalls. The final version of the text, which has now been shared with the European Commission, instead referred to these issues as “critical concerns”.

The Commission aspires for Europe to invest €20bn a year on AI in the coming years, and some experts have proposed the creation of a new dedicated centre of research similar to CERN. One way, says Metzinger, to ensure regulation keeps pace with the technology would be to establish an allied Ethics Centre of Excellence with a series of rotating fellowships, as well as a wider European network of 720 professorships in the applied ethics of AI. “For the Centre, the idea would be to have the young ethicists right there [working with the developers] and for the network to educate new generations of students and to keep the wider public informed and to structure public debates,” says Metzinger. “Europe is the last grown-up in the room between China and [President] Trump. We have to guide the debate.”

How patterns can power better treatment

Zia Chishti,
chairman of the
board and CEO at
Afiniti, discusses
the role that artificial
intelligence can
play in healthcare

How can healthcare harness artificial intelligence?

The goal of healthcare is to improve the quality and increase the length of people's lives. AI has some great opportunities to help in this field – but many are misunderstood or mis-sold. There is a lot of hype around diagnostics, for instance. Today, with only a small handful of exceptions, many applications of AI in diagnostics have fallen flat.

Where AI can really help generate better health outcomes is by matching patients to clinicians more efficiently. Home visits, for example, can be critical to a patient's care but are all too often neglected. This is because educating patients that home visits are good for them is difficult, even though they are key to reducing morbidity and mortality rates. The education process is important – and who is trying to do the educating? Other humans; the insurer talking to the insured. Creating an interaction between those two people that increases the likelihood that the patient accepts a home visit can be an incredibly powerful thing. It helps ensure patients adhere to their treatments and often flags serious health issues before they become any worse.

Afiniti's AI platform identifies subtle and valuable patterns in data and in human interaction. We match individual agents working for healthcare providers such as United Healthcare in the United States to individual patients, creating interactions that increase the number

of home visits, which help to deliver better health outcomes.

How does Afiniti's AI work?

Afiniti's AI already has a great track record in business, working closely with leading British companies such as Sky and Virgin Media to help improve the interactions they have with millions of customers. Rather than customer calls being automatically assigned to agents in the order they come in, the algorithm analyses the behavioral traits of both agents and callers based on data such as prior interactions each party has had. What does this mean in practice? A better conversation that increases the likelihood that the customer receives the best TV and internet package for them and their family.

It is a similar principle in healthcare, albeit with more at stake. We can predict which pairings between insurers and insureds will generate the best dialogue that is more likely to mean that the patient will agree to a home visit from a clinician.

What is the role of humans in rolling out more AI-enabled technologies?

Firstly, we need to distinguish between data analytics and AI. Data collection, usage and analysis need to be regulated, particularly around issues of privacy. But there is a train of thought which says that government should regulate AI. This is wrong and represents a fundamental misunderstanding of what AI is. It is maths – and you cannot regulate maths.

Secondly, we need to recognise that AI is oriented around specific use cases. The era of the automaton is not upon us, despite the hyperbole surrounding the technology. Businesses, including healthcare providers, need to understand how specific applications of AI will help directly help them and their customers. Don't believe all the hype!

For more information, please visit:
www.afiniti.com

IN ASSOCIATION WITH

afiniti

The arrival of AI-powered health services in one constituency has caused concern amongst local NHS groups, reports **Jonny Ball**

The NHS meets disruptive AI



On a busy road in Fulham, West London, between the Pizza Go-Go and the A-Z Wine Mart, stands England's fastest-growing GP surgery. In April 2017, the Lillie Road Medical Centre had 2,500 patients – well below the average for an English GP practice. The surgery now caters for almost 50,000 patients, a 1,900 per cent increase in little over two years. But the rise in patient numbers isn't down to mass migration into the local area. Lillie Road Medical Centre has become the hub for an AI-powered app, Babylon GP at Hand, and anyone working in central London or within 40 minutes' travel of a GP at Hand centre can register, making them technically a patient of the Lillie Road surgery.

While the app serves patients remotely, the NHS allocates funding to GP practices based on the number of patients registered. So, while the vast majority of GP at Hand patients do not visit the surgery, this huge influx of patients has led to funding pressure for

which the local Clinical Commissioning Group (CCG) has not been prepared. The local MP, Andy Slaughter, claims that the arrival of GP at Hand has created a funding deficit for the Hammersmith and Fulham CCG.

The GP at Hand app offers users an AI-powered chatbot service, which uses machine learning to give medical advice, as well as connecting patients to real doctors through video and telephone consultations. It has won plaudits from the Health Secretary, Matt Hancock, who is himself an enthusiastic user of the service. Hancock told the *Telegraph* last year that he wanted GP at Hand to be made “available to all, not based on their postcode”.

But for Hammersmith and Fulham's CCG, which encompasses the Lillie Road Medical Centre, the influx of the GP at Hand patient list – less than ten per cent of whom live in the CCG's catchment area – has created a significant accounting problem. A letter sent by Slaughter to the Health and Social Care Select Committee

last month put the funding deficit created by the sudden spike in patient numbers at “at least £26 million”.

The funding deficit has been created by an unwillingness on the part of other NHS CCGs to pay for GP at Hand patients to “move” to the Lillie Road practice. GPs are funded according to the number of patients they have registered. Under normal circumstances, a shift in the number of patients from one CCG to another would result in a change in the funding allocated to each. But the skyrocketing number of app users listed at Lillie Road Medical Centre has led other CCGs to question whether they should transfer 50,000 patients' worth of funding out of their own budgets for patients who still live within their catchment areas.

“That money should eventually come to Hammersmith,” says Slaughter, who has been the Labour MP for Hammersmith since 2010. “But at the moment,” he continues, “the other CCGs are refusing to passport money



A huge surge in patients put the CCG £26m in deficit

GETTY IMAGES/CARL COURT

because they're saying they're not sure it is all being done fairly."

This sense of unfairness comes from the type of patients that are using GP at Hand. A recent independent evaluation of the app, prepared by Ipsos MORI for NHS England, reported that users of the app were "younger and potentially more affluent than patients at the average practice in London and nationally." Some 94 per cent of Babylon GP at Hand patients were found to be under the age of 45 and were described as "predominantly city sophisticates" and "career climbers". Patients using the app were "healthier than those at other practices in the CCG, even after adjusting for age", and the evaluation concluded that the service was "not being used by large numbers of older people, or large numbers of people with more complex health needs".

Separate data from the Institute for Fiscal Studies shows that an average 65-year-old costs the NHS two and a half times more than the average 30-year-old. An 85-year-old costs, on average, more than five times as much to treat.

Slaughter says that services aimed at affluent, technologically adept younger patients are "creaming off a particular part of the public who, in terms of their health needs, are the least demanding". As he sees it, the GP practices these users leave when they join GP at Hand are being deprived of funding that was effectively subsidising the higher costs of caring for older patients with more complex health needs in the same CCG.

"You're segregating the market," Slaughter says, "telling GPs to deal with people with very difficult and complex cases that are very demanding of your time and service, and all the people who are on your books, who give you an income, that are relatively straightforward to deal with, they'll all be taken by a private company who'll make a profit out of it." He adds: "It looks like a license to print money."

A Babylon spokesperson responded that "GP at Hand does not exclude anyone – it is designed to provide high-quality NHS primary medical

services for people of all ages, whether in good health or not. If our model became mainstream then a large segment of patients with complex needs could benefit by having access to a GP 24 hours a day." Furthermore, the spokesperson responded that "Babylon is not taking money from other practices by having younger patients. There is a Carr-Hill formula that is the basis for how NHS GP clinics are funded. It weights payments by age and gender, as well as other factors so that practices are paid depending on who registers with them." Overall, they argued, Babylon expects "GP at Hand to reduce the costs on the overall NHS because fast access every hour of every day means that fewer people go to A&E and don't face delays to diagnosis and treatment."

What is certain is that, with a deficit running into the tens of millions, Hammersmith and Fulham CCG cannot afford the sudden change in funding allocation that the arrival of AI has brought with it. The group is already struggling with cuts to frontline services, including out of hours GP visits, radiology, cardiology, gynaecology and dementia support, following budget cuts of £8m last year and £11m this year, according to Slaughter. The group has applied directly to NHS England for financial help, but no guarantees have yet been given.

At an event at Babylon's offices last September, attended by Matt Hancock, the company's CEO, Ali Parsa, announced a \$100m expansion of its UK operation. Slaughter worries that if Babylon expands its services without more consideration of existing NHS structures, more CCGs will find themselves struggling to account for the changes AI brings. "Babylon could have paired up with lots of GPs around the country and had lots of local Babylon GP at Hand stations," he reflects. "But it's all run from one particular address in Fulham, for accounting and management purposes. They've been asked about this and they just say it's what suits them, it's within the rules, and therefore the NHS will have to adapt."

Connecting and protecting the digital economy

Artificial intelligence can help human analysts fight back faster against cyber threats, explains Professor **Tim Whitley**, BT's managing director of applied research and of Adastral Park



Security professionals at BT typically protect our network from around 125,000 cyberattacks a month. And it's an area where AI can be harnessed to great effect. We've done exactly that, developing BT's SATURN as part of BT's Cyber Security Platform, to help cyber experts predict and detect attacks at an early stage.

By combining patented machine learning and advanced data visualisation it allows analysts in our security operations centres to identify individual threats amongst billions of network events. We've deployed this technology so it now protects British businesses, institutions and critical national infrastructure alike.

SATURN follows decades of investment and innovation in AI by BT. It started principally with a predecessor in my role, the late John Alvey who inspired the UK government's first structured funding programme for AI in 1983. In its wake followed BT's

first full application of AI – to the deployment of BT field engineers servicing the network. The goal was to match the right person to the right job at the right time.

Using data sets from our field force, networks, weather feeds and roads, we've continued to improve the system. This has helped us speed up the rate at which we connect homes and maintain the network. To illustrate the scale involved, in the last year alone, Openreach engineers have deployed around 2.6m kilometres of fibre. The efficiency gains therefore equate to savings of many millions of pounds.

Overall, our field force AI has been good for customers, our people and our network. Ultimately, it has been good for the UK. In short, AI can serve many of BT's priorities.

Take how we've harnessed AI to benefit our customers, who are our first priority. One example is "Call

IN ASSOCIATION WITH





Protect” which uses algorithms based on BT data science to screen out illegal cold calls which are such a nuisance to customers and often a menace to the vulnerable.

Or look at some ways in which we support our colleagues using AI. Engineers providing or fixing a connection in the local cabinet can face an intimidating spaghetti of cables that has built up as more and more premises have taken up broadband. That’s why we use image recognition tools to help identify the right connection types and spaces within a cabinet to make the engineer’s job easier and less prone to faults.

AI is also helping us build not just the safest but the best networks. For example, we’ve supported our network resilience by streamlining supply chains with an award-winning tool to optimise the distribution of spare parts, developed in collaboration with the EBTIC research centre in Abu

Dhabi. Additionally, we are using AI algorithms to plan our fibre networks and help accelerate rollout across the UK.

These AI-powered programmes all add to BT’s overall contribution to the country, supporting the UK’s ambitions for the world-class digital connectivity which will drive productivity and underpin its Industrial Strategy.

Yet while BT has been the third highest investor in R&D in the UK over the past decade, we – and other UK companies – face intense global competition for ideas, talent and scale. The UK has benefitted from an AI Sector Deal, Grand Challenge and the government’s skills investment. It punches well above its weight. But it can only continue to do so via an open innovation model of partnerships between multinationals, start-ups, universities and government bodies all working together to ensure the diffusion and penetration of the best new technologies from across the globe, wherever they can serve UK economy and society.

Among the changes we expect in the next decade is the democratisation of AI, with non-specialists able to use increasingly accessible AI tools for new and currently untapped business or personal uses. We also anticipate more data-driven innovation. While in the past, data has been collected and curated as the “exhaust” of applications, it will increasingly itself drive the creation of the applications of the future.

Another expected trend is a growing interest in privacy-enhancing technologies, including Edge AI. In other words, using AI to process consumers or businesses’ data locally (or on hardware), providing potential privacy and security benefits.

While there are undoubtedly many more as yet unimagined opportunities, the risks of AI are perhaps better documented. Around a third of the UK population think they may lose their job to automation and those that do may well not benefit from the new

opportunities and jobs AI creates. Navigating this disruption will require government and companies to adapt. For example, by reshaping the skills we provide to future and current workers, and how we teach our children.

BT has already made digital skills a core part of its commitment to the country. Via BT’s Barefoot programme alone, 2m primary school children have already benefitted from training. Last month we announced a new ambition to train 10m people in the UK with digital skills.

Alongside job losses, questions about AI ethics are at the forefront of the thinking on AI by public authorities and regulators. These include the ICO’s consultation on an AI auditing framework and the Centre for Data Ethics and Innovation’s review of bias in algorithmic decision making. These are just two examples of current initiatives, but there are many, and all set against the backdrop of government rightly considering how the UK’s regulatory model might need to be modernised so as to address digital competition. We want to be part of these discussions, sharing our experience to date and working with others to design the right governance framework and regulatory principles to support an ethical approach to AI in the UK.

AI is a complex topic and this month’s London Tech Week is a good place to debate what risks to watch out for. But it’s also a great place for the debate about the exciting opportunities ahead – how AI can make the country more productive, safer and better connected. We’ve come a long way already. Investment in UK-based AI startups alone has increased fivefold in the five years since the inaugural London Tech Week and, more importantly, the overall ecosystem has thrived. We look forward to making our contribution over the next five years to catalyse even further growth.

For more information, please visit: www.btplc.com/Innovation

How AI will deliver cost savings and create new jobs

Robotic process automation is the key to a more competitive economy, writes **Bruno Ferreira**, area vice president for UK and Ireland at UiPath



Last year, the government launched the AI Sector Deal along with the Centre for Data Ethics and Government Office for AI. Tim Clement-Jones, the chair of the House of Lords' first AI Select Committee, warned that the country was falling behind South Korea, Canada and Germany in its AI-sector endeavours.

His speech at the AI Expo coincided with the release of a 181-page House of Lords report that called for a more structured framework for AI development and addressed the importance of grasping its economic, ethical and social implications.

Fast-forward to 2019. In May, the UK announced the first members of its new AI Council. Representatives from the private sector, data privacy organisations and academia were among the appointees. "[The Council] will represent the UK's AI Sector on the international stage, and help us put in place the right skills, data and ethics so we can make the most of AI technologies," Jeremy Wright, Secretary

of State for Digital, Culture, Media and Sport, said in his keynote address at the Viva Tech conference.

So far, the UK has pledged £115m for 16 new AI Centres for Doctoral Training (CDTs), which will fund PhDs for 1,000 students. The AI Council represents another milestone in encouraging tech leadership, educating the general public and promoting ethical and thoughtful AI adoption in both the private and public sectors.

But some of the most concrete examples of the UK's AI agenda have been less pronounced, particularly in public sector departments. In 2017, the Department for Work and Pensions (DWP), the government's largest department, completed a highly successful robotic process automation (RPA) project. RPA refers to "software robots" that function over existing IT infrastructure to automate repetitive, rules-driven tasks.

One of the DWP's functions is to provision welfare and pension services to about 20m people every year. This

IN ASSOCIATION WITH





includes processing pension claims, which prior to RPA implementation, had been completed manually, resulting in a backlog of 30,000 claims.

Shaun Williamson, senior product manager at DWP, estimated that catching up would require a labour force in the thousands working over many thousands of hours. But by using only 12 UiPath software robots, the DWP cleared that entire backlog in just two weeks. “Our experience suggests the return on investment is around 15:1,” Williamson said.

Despite this public sector success story, RPA has arguably achieved greater public recognition outside of the UK. For example, the European Commission released its Coordinated Plan on Artificial Intelligence in March that, “acknowledges the use of robotic process automation [RPA] and the impact it has had on improving public sector processes.” European Parliament members also emphasised AI and robotics’ transformative potential in the health, energy, agriculture and

transportation spheres.

Meanwhile, as the UK moves towards Brexit, taking the lead on AI is no doubt a chance to assert its tech acumen as it confirms an EU departure. In his announcement of the AI Council, Wright acknowledged that leaving the EU doesn’t mean the country is “turning inwards”. He added: “We are still committed to being open, outward looking and optimistic about the world beyond us and opportunities provided by new technologies. And we still care about the positive development of these technologies, and making them work for the benefit of humankind.”

So should “these technologies” include RPA? RPA has clearly demonstrated its potential to serve the public by boosting departmental productivity in the DWP. Meanwhile, the body that the UK hopes to distinguish its identity from has already said robotics are an asset in social and economic digital transformation. Additionally, RPA can function with legacy public sector technologies, and is quick and cost-effective to deploy.

Most importantly, RPA is a gateway to AI innovation in both the public and the private sectors. Unlike AI, RPA is limited to rules-based processes that usually involve structured data. This curbs its ability to complete cognitive processes such as predicting outcomes or processing conversational language.

But when RPA is enhanced with AI, robotic software can continuously improve its efficiency and performance. For instance, it could:

- | Process conversational language.
- | Prioritise certain automated tasks over others.
- | Visually interpret user interfaces to interact with software more fluidly.
- | Contextualise unstructured data by understanding document contents.

UiPath, the same vendor that helped the DWP clear its pension claims backlog, is already investing in these capabilities. As the UK stimulates AI development, it must keep all possible paths to AI top of mind. This will be

important for educating the general public, but also for leading by example.

According to a recent UiPath survey, 61 per cent of public sector employees in the UK spend 30 per cent of their time each week on repetitive chores that can be automated by RPA. The majority of respondents (78 per cent) also said that they’re not afraid that automation will take their jobs.

Research from Deloitte supports this sentiment. In 2017, it estimated that automation contributed to the loss of 800,000 low-skilled jobs in the UK, but that it created 3.5m higher-skilled, and higher-paying, new jobs. The World Economic Forum paints a similar picture. It predicts that by 2022, AI and robotics will create almost 60m more jobs than they will destroy.

In the UK’s public sector, automation has predominantly been well-received; 74 per cent of government employees said their experiences with RPA have been positive.

Nevertheless, misconceptions about AI, RPA and other transformative technologies remain. Many employees, especially outside of the 18-24 age range, are sceptical about automation’s ability to make them more productive, despite admitting that they spend a significant portion of their time on mundane, administrative tasks.

This likely stems from a general lack of awareness about AI and robotic capabilities, one that the UK’s AI-forward initiatives will hopefully address in the coming years. The onus is also on the private sector and big tech to contribute thought leadership that helps inform the responsible and ethical progression of the technology.

For now, though, all eyes will be on the public sector to facilitate important conversations about new technologies. Just as importantly, the UK government will need to take a pragmatic approach to the internal implementation of AI and robotics. And technically speaking, RPA looks like the most pragmatic approach of all.

For more information, please visit: www.uipath.com

How AI will be a driver for economic growth

Artificial intelligence and machine learning are the great enablers for the office of the future, according to **Darragh Jones**, CTO and co-founder at **deputi.ai**

Last year, a report from the Department for Business, Energy and Industrial Strategy (BEIS) projected that by 2035 artificial intelligence (AI) could add up to £630bn to the United Kingdom's economy. As AI drives growth across the whole economy, small and medium-sized enterprises (SMEs) – the so-called backbone of the British economy – must not be left behind by the pace and scale of change and, as with larger firms, must make the most of opportunities that new technology presents.

Preparing SMEs for change

We need to prepare SMEs for the radical changes that AI will bring to the workplace. The playing fields are not level and large enterprises have the resources to leverage a competitive advantage. Helping SMEs to apply the right technologies at the right time to support productivity and drive cost savings should be paramount.

It is in everyone's interests that SMEs are equipped with the resources and capability to compete and fuel growth. Powered by the right mix of AI rocket fuel one of today's startups might be best placed to compete on the global stage.

We are getting there; initiatives such as the Digital Catapult, The Department of AI and the Innovation Fund are helping break down the financial, technological and risk barriers for SMEs to leverage AI, but these need to be expanded and promoted further.

Building a corporate memory

Business operations are primed for automation, leveraging AI. Common to all businesses, startups through to multinationals, are a core set of administrative principles including financial control, performance reporting and record keeping.

Primed for automation, these functions all contribute to a collective corporate memory, enabling further automation, greater transparency and partner integrations. Together with other enablers such as the API economy, open banking and digital tax it could fast-track uptake of artificial intelligence in the workplace.

AI can help SMEs to achieve more with less

The workplace of the future

Fast-forward ten years and we envisage an AI-powered workplace as being the norm in everyday business operations, removing the burden of business administration and letting organisations focus on their most valuable customers, assets or strategies. AI bots or agents could be seen as a valuable team member, supporting decision making and a catalyst for change in how SMEs can deliver their products and services.

Cutting back on the time and money spent on administrative tasks frees up new funds to invest in other areas of the business, enabling staff to be retrained and redeployed more effectively. For SMEs who perhaps do not have the luxury of large teams, AI could be the key to achieving more with less.

For more information, please visit:
www.deputi.ai

IN ASSOCIATION WITH

deputi.ai

Building Better Businesses

Navigating a new era of intelligent automation

Neil Murphy,
vice president, head
of global business
development at
Abbyy, considers
the steps businesses
can take to get the
most out of digital
transformation

Digital transformation can be a complex undertaking, especially at the outset, where large buy-in and investment is needed. However, the tools and technologies to make new solutions a reality are readily available. In fact, the core technologies for digital transformation may already be in use in certain areas of your organisation. In most companies, intelligent capture, robotic process automation (RPA) and artificial intelligence (AI) are the central building blocks of digital transformation. Here are steps your business can take to prepare for the new era of intelligent automation...

Leverage RPA as the driving force for positive change

The volume and complexity of content is growing exponentially. To leverage the value in that information, businesses need to do more than just scan and archive documents. They need to automatically extract all relevant data and integrate that information into their operation.

RPA is key to accomplishing this. It is sophisticated software designed to mimic what a user does in any given application. It can connect legacy systems, apply business logic, and connect software and processes to eliminate the repetitive work that humans often perform.

It works via software “robots” which follow a pre-determined process to automatically initiate actions, perform calculations, move or populate information between

pre-identified locations and/or launch downstream activities.

Move beyond traditional capture

In the past, paper-based documents were essential to a functioning business. However, with the advent of scanners and mobile devices, organisations quickly realised that digitising paper documents and data could improve their operations. Today, that scanning process known as “capture” is driving a shift in how organisations use and consume data. Capture has evolved to connect new digital technologies such as RPA to handle invoicing, new account opening, claims processing and more.

However, to grow and expand the use of RPA within an enterprise, robots must become smarter to be able to interpret and understand unstructured content (documents, images and text) and turn into actionable structured information.

Consume AI to deliver content IQ

Content IQ is defined as a class of enabling technologies that help digital workforces understand and create meaning from enterprise content. Content IQ provides the ability to automatically extract all relevant information from documents and breaks down processing of content into easy to use and consume technology that can be leveraged directly within an automation solution such as RPA.

Content IQ is helping transform business at all levels of RPA, starting with the most basic automation robots all the way to designing robots that automate tasks involving more intuition, judgement and problem solving. By incorporating these steps towards Content IQ, businesses will gain new meaning and opportunity to apply intelligence to their content and connect it to their business processes.

For more information, please visit:
www.abbyy.com/en-gb/

IN ASSOCIATION WITH

ABBYY®

Helping humans reach their full potential

WorkDone improves your life by watching you work and then doing your work for you, explains **Joseph T. Rogers**, founder and CEO of WorkDone.AI

WorkDone offers a machine learning-based technology that transparently acquires and retains knowledge worker expertise, which allows the storage of operational best practices, or Work Heuristics™, in a persistent corporate memory. Organisations can avoid the costs associated with institutional knowledge loss when valued employees retire or leave the company by leveraging WorkDone's patent-pending Expertise Capture technology. With roughly 76m American baby boomers set to retire (part of the 1.6bn people worldwide aged 46-65), WorkDone will play a crucial role in helping the global economy dodge this costly bullet.

Work heuristics

WorkDone saves companies time, labour and money by enabling them to seamlessly automate repetitive human-intensive tasks between major SaaS platforms through work heuristics, with no client-side training or programming required. The work heuristics in turn direct virtual processors (WorkDone Agents), which free up knowledge workers to do higher value, more strategic work. Learned work heuristics also enable the process of self-documentation and systems of intelligence creation.

Focusing on the back office, WorkDone is first targeting ERP use cases such as AP and moving on to adjacent use cases from there. WorkDone's frictionless technology ensures rapid customer

adoption among both SMEs and larger businesses. WorkDone's distribution channels of only top-tier SaaS platform partners/resellers with existing customer bases create a flywheel effect and sustainable advantage in the field of work heuristics.

Social impact

WorkDone takes corporate social responsibility seriously by supporting the transition and reinvention of workers displaced by automation. For this reason, WorkDone was created as a Public Benefit Corporation whose bylaws commit a percentage of revenue to assist displaced workers through funding of the Work Forward Foundation. By providing a soft landing and a supportive community, individuals who may never have seriously considered entrepreneurship are encouraged to pursue more creative endeavours.

Ethical AI

From a technical and commercial perspective, the sky is the limit for the AI industry. More concerning, however, are the ethical and socio-economic considerations that must be proactively addressed around the issues of job automation and displaced labour. In the drive to cut costs through automation, businesses may be bankrupting their own current and future customers. Case in point, even Henry Ford paid his workers enough to afford to purchase a Model T.

Historically in the United States, companies have avoided addressing the negative side-effects of their products (think environmental pollution and social media manipulation) until it was too late. In the case of AI, it is imperative that any unintended detrimental consequences must be anticipated and mitigated or avoided completely. The AI industry needs to mature into the benevolent stewards of the various sectors of society it impacts.

For more information, please visit:
www.WorkDone.AI

IN ASSOCIATION WITH



Do coders need a code of conduct?

Rainer Kattel, professor of innovation and public governance at the UCL Institute for Innovation and Public Purpose, discusses the need to regulate and monitor AI technologies



Are public organisations ready for artificial intelligence? Recent tragic accidents involving the Boeing 737 MAX suggest they are not. While modern aviation, including rules-based software it relies on, is overwhelmingly safe, the slippery slope of self-regulation by industry shows the limits of human-machine interactions.

The step towards probabilistic software, such as that used in self-driving vehicles, is not just highly demanding technologically speaking, and uncertain at this point, it also puts public regulators into a uniquely complicated situation. Namely, there are things we may want to regulate – for example, search algorithms that exclude competitors, social media feeds inciting violence – that will change as they are being monitored. The object of regulation is dynamic. And the complexity only increases as we are applying AI to less technological areas, such as health or education, while trying to understand how these complexities affect people, systems and society.

The European Commission has fined Google €8.2bn over two years for abusing its monopoly power in online advertising, shopping and in its Android operating system. This is a highly commendable action. Yet, it is unlikely that will change Google's behaviour, mainly because such rulings misunderstand the source of Google's ability to dominate. The market power of big data companies does not rely on their means to pressure websites into using their advertising tools. Instead, power comes from the combination of the sheer endless amount of data about its users

and clients, and its code and algorithms that continuously learn about the users and clients.

Code is not only law, but code is also learning. Historically, learning and in particular tacit aspects of it, such as the ability of teams to work well together, have been fundamental to innovation. Thus, the question is how competition authorities should curtail the ability to learn within big data companies. Fining them for external anticompetitive behaviour forces them to come up with better internal processes – better code – to circumvent not only competitors but also regulators.

Breaking up companies like Facebook, as recently suggested by one of its co-founders, Chris Hughes, would probably not change the underlying dynamics. Radical open source solutions like the data sovereignty approach by cities such as Barcelona is a much more promising alternative.

Put bluntly, AI makes code and algorithms into economic and political agents, and our economic and policy frameworks have limited tools to deal with such non-human actors whose primary goal is to circumvent human agency. This poses the critical question for AI and public policy: what is the purpose of public policy, such as competition policy or seamless public services, and who and how defines this purpose?

In truth, 20th-century public organisations are not supposed to have the capabilities to check, to question, to redefine the purpose of public policy. In this age of super-wicked problems, the temptation will only increase to cut out the human from the decision-making processes.

AI will be, eventually, taught to learn from millions of cases of purpose, of policy choices – and it will decide. Thus, the question we really need to be asking is not how to create seamless public services, or how to diminish big tech's market power, but rather who will teach machines about what is innovation and what is political and policy choice?

AI will lead to more engaging employment

AI and automation are poised to inject fresh enthusiasm and impetus into modern workforces, writes **Nisha Deo**, head of communications at Rainbird



What does the future of work look like? Some believe it's already arrived with the gig economy and globalisation. But I believe we can do so much better. Automation will clearly be a central player. But while we hear about a boost in efficiencies across multiple industries, we rarely hear enough about the positive transformative impact that automation can have on a business and the way its staff operate. This is arguably down to the fact that most cases of artificial intelligence implementation have kept people at arm's length – whether via learning algorithms mining data in the background, or robotic process automation labouring through isolated workflows. The fact that working people have rarely been involved in the Fourth Industrial Revolution should be a cause for concern.

Playing to our strengths

A good craftsman never blames his tools, we're led to believe. But what if

the nature of the technology really has been the problem in the lack of tech uptake? According to a recent EY survey, a lack of AI talent in the marketplace is a core barrier to AI adoption, citing the fact that so many platforms are targeted at data scientists and AI specialists. But herein lies the problem: people shouldn't have to gain what we commonly term as "AI skills". The technology should accommodate us, rather than the other way around. We have to encourage technology that can be built and used in the most intuitive, accessible way, empowering people to integrate it into the transformation of their working lives.

What kinds of skills do we need to kickstart job transformation? Two Accenture heads tackled the question in Reimagining Work in the Age of AI – and found that elite STEM skills were pretty low down the list. Nuanced human reasoning, creativity, social and emotional intelligence were the traits thought to be most integral – skills that

IN ASSOCIATION WITH



SHUTTERSTOCK / 4 PM PRODUCTION



we all possess.

So it's important that we debunk the notion that the evolving tech-boosted workplace is set to leave ordinary people behind, when in fact it can play to all of our strengths.

Breaking barriers

There are, however, undoubtedly tough cultural barriers to tech adoption in the workplace. Innovate UK's Stephen Browning admitted that mass adoption is often held back by "human aspects" around the technology rather than the technology itself. Many C-Suite executives are still far from convinced about how the technology works and whether it can be trusted. But even where business leaders are convinced, the technology still requires buy-in from employees. This can be hard to come by. A landmark EY survey found that employee trust (33 per cent) is one of the greatest barriers to AI adoption, despite the fact that 87 per cent of CEOs and business leaders claimed to completely

or somewhat trust the technology.

If the blocker is widespread trust and understanding, as the surveys suggest, the way to breakthrough is with technology that is transparent and usable from a business perspective. Fear of the unknown is the "oldest and strongest emotion of mankind", as HP Lovecraft once wrote. Our first – and biggest – task with implementing AI is familiarising the unknown.

Putting people first

Businesses should be creating a culture of change, and this begins with human-centric technology and implementation. AI development is at a crossroads: either we continue down a road that posits humans as mere supervisors while uber-algorithms go to work, or we turn the tables and put the people in the driving seat. "Can we reframe the education, the research and dialogue of AI and technology in a human-centric way?", asked Fei-Fei Li at a recent Stanford discussion.

That road is paved with interpretable, business-friendly AI-powered platforms. This means systems that can be modelled without the need for data scientists or coders, and into which business people can accurately model the logic that will combine with data to automate large-scale, transactional, complex decisions.

AI-empowered decisions

If an insurance company on-boards an automation platform to assist in risk assessment, those underwriters will be asking themselves: what will this technology do for me? Will it enhance my work, or will it consign me to the sidelines of my own job?

In terms of performance, automation can be very much a friend. McKinsey surveyed 1,200 managers across the globe and found strong signs of growing frustration with broken decision-making processes. Fewer than half of the survey respondents said that decisions are timely enough, while other factors cited included an overload of disruptive "noise", and company cultures that lack

staff empowerment.

Arming employees with real-time, best-practice recommendations to inform their decisions and interactions can quickly upgrade the way they go about their business, while improving the quality of their organisation's high-volume decisions.

Have we been too fixated on job numbers to give enough thought to job quality? Let's put to one side the World Economic Forum's prediction that technology will create 58 million more jobs than it displaces by 2022, and focus on the fact that the quality of jobs that already exist today can improve markedly with automation.

"Automation is coming for the part of our jobs that we don't like," said Graeme Codrington, futurist and global speaker. It's also coming to improve the parts of our jobs that we do like: the soft skills, the nuanced thinking, the things we're best at.

A tech-powered marketplace

For businesses themselves, a more empowered, engaged, excelling workforce can be the edge in competitive, fast-changing markets. Take the legal sector as an example. Countless conventional law firms are faced with an emerging \$10bn market of alternative legal service providers, and half of clients and small businesses going online for legal providers. But by placing the emphasis on staff understanding and collaborative implementation strategy, any law firm can effectively become a tech company.

At Rainbird, we saw a huge appetite for embracing technological change from Taylor Wessing, where dozens of employees underwent our certified training program. The end result is a firm that can self-sufficiently identify and implement its own AI-led solutions. If such an example can be replicated across our industries, we could be about to see an extended period of business and market transformation, all driven by an AI-powered workforce.

For more information, please visit:
www.rainbird.com

Transforming care delivery with AI

AI represents a game-changer in the field of diagnostics and will be a critical tool in the move to value-based care. Siemens Healthineers' strategic focus on AI continues to support the digitisation of healthcare



In recent years artificial intelligence has become part of our everyday lives. Intelligent computer algorithms are utilised in internet search engines, smart phones, financial transactions and self-driving cars. In 2017, the machine learning-based computer programme “AlphaGo” defeated the world champion of “Go”, the strategic board game, once again demonstrating the potential of AI¹.

Meanwhile, in the healthcare sector, and specifically in the field of diagnostics, a fundamental transformation is on the horizon. According to a recent poll by *The Economist*, more than 50 per cent of global healthcare leaders expect an expanding role of AI in monitoring and diagnosis².

In the United Kingdom, an ageing population, shortage of radiologists, strain on pathology, as well as economic austerity has underpinned a difficult workforce shortage. Demand for CT and MRI scans has grown by more than 30 per cent between 2013 and 2016,

while growth in the number of radiologists over the same period staggered at eight per cent. Additionally, the number of pathologists reporting that they worked well beyond a 39-hour week is also increasing, which is directly linked to adverse events⁴.

With the above in mind, and these 12-hour plus workdays becoming more common³, what role could AI play in managing and standardising workload, while mitigating against the risk of reading fatigue and clinician burnout?

We at Siemens Healthineers believe that AI is poised to play an increasingly important role in medical imaging and pathology. We believe the role of the radiologist and pathologist will be augmented by these new tools, and we believe these tools will improve the patient experience while also helping to meet the rising demand for diagnostic imaging and faster turnaround times in the lab. Digitisation will make way for more efficient and data-driven decision making.

At the core of our approach to

IN ASSOCIATION WITH

SIEMENS
Healthineers

The future of chest readings with AI-RAD Companion Chest CT



healthcare digitisation is a drive to not only cut costs but to impact real and positive change on both the working lives of clinicians as well as patient outcomes, and to always ensure patient safety and privacy is kept paramount. It is no secret that reading and reporting on medical images or pathology results can be labour intensive and repetitive at times. With that in mind, we set ourselves the goal of giving clinicians an accurate, reliable aid to help minimise

AI has the potential to benefit both patients and clinicians

repetitive tasks such as manual labelling and measurement of organs, to give them more time that could be spent reporting on more severe cases.

Enter the AI RAD companion; after being “trained” on more than three million images, the AI RAD companion is a cloud-based software that will automatically analyse CT Chest scans to report on metrics such as aortic diameter, heart calcification, emphysema staging, and lung nodule size and position. This is all done within minutes of a patient’s scan being completed. While the AI RAD companion is a classic example of AI in action within healthcare, the applications of machine learning within healthcare can sometimes be subtle and often reach beyond the limits of one hospital department or discipline.

An example of such an application is our AI Pathway companion, another AI-driven software that Siemens Healthineers is developing with a key focus on the management of patients through their journey along an often complicated, painful and life-altering clinical pathway.

The Pathway Companion currently focuses on some of the main killers of our time; namely cancers of the prostate, breast and lung, as well as heart disease. It connects to data sources and extracts insights using machine learning and natural language processing from pathology and radiology results, patient records and notes to locate and analyse the most relevant data and facilitate a comprehensive clinical picture of the patient.

Imagine you’re a patient with suspected prostate cancer, and in a single appointment your oncologist can see your latest blood results as well as a history of those results over the last few years. They can also see the results of your latest biopsy and radiological exams, and lastly, they see how all this data would fit into the clinical pathway, and what the recommended next step would be for you as an individual. Data from three or four different sources and/or departments in one place, with a

recommendation for action to be considered by your doctor – or doctors in a multidisciplinary team setting.

Value-based care – leveraging AI to deliver data-driven precision medicine

The digitisation of diagnostic results and shortening the time required by clinicians to interpret what those results mean is a key output of artificial intelligence; but what it really delivers is value-based care, where the patient is at the centre of diagnostic decision making and success is determined by patient outcomes, not just operational efficiency.

While certain technical metrics such as specificity, sensitivity, and reliability are crucial to the acceptance of any AI solution into clinical use, equally important is how it fits within the clinical workflow and the value it adds to the clinical decision-making process. Far from replacing any clinicians, an AI system is only worth as much as the value it produces for clinicians.

Artificial intelligence is not the answer to every issue being faced in the NHS, nor will it be a perfect and infallible solution. However, the potential for AI to introduce highly practical and life-changing implications for both patients and clinicians is high enough that it must be given a real chance to prove its worth in our NHS.

By Matt Gibson, head of sales and digital services – UK&I, and Hasan Jouni, business development – digital services UK&I, at Siemens Healthineers

1. Mozur, Paul (May 23, 2017) “Google’s AlphaGo Defeats Chinese Go Master in Win for A.I.”
2. *The Economist* (2017) *The future of healthcare – Realities or science fiction?*
3. *The Royal College of Radiologists* (2017): *UK workforce census 2016 report.*
4. <https://thepathologist.com/outside-the-lab/all-in-a-days-work>
5. <https://www.siemens-healthineers.com/en-uk/press-room/press-features/lunghealthchecks.html>,
6. <https://www.siemens-healthineers.com/press-room/press-releases/pr-20181125047shs.html>

The ethical and transparent path to automation

Auditable and people-first technology can achieve the transparency demanded of a modern business, writes **Ben Taylor**, chief executive officer at Rainbird



Transparency is a necessity in today's marketplace. Any forward-thinking business is leveraging the business value of artificial intelligence; but any business leveraging AI will soon find that their regulators, customers, staff and compliance teams will all be demanding "explainability". This means that no matter how vast or complex the decisions being automated across a variety of sectors, the guidance of a clear and detailed audit trail should be sacrosanct.

This is not just a technical feature to ward off regulators; it makes all the difference in consumer interactions. If a fraud agent can explain the rationale behind a transaction risk and have a better informed conversation, it'll garner more customer trust and loyalty in their provider. To seal transparency's integral place in the business landscape, the concept is now being enshrined by governments, with the European Union placing it front and centre in its requirements for ethical AI. By this

point, the "why" for explainability hardly needs to be established any more.

Still, we can't take the idea of transparency for granted. The explosion in the number of algorithms used by employers, banks, police forces and many more in recent years has proven one thing: not all AI and automated decision-making systems are created equal. In regard to highly complex AI technologies such as deep learning and neural networks, it's almost impossible for developers, who possess an understanding of the technology on an algorithmic level, to explain all of the many processes in a language that consumers, regulators and business people can understand.

As we've seen, some can, and do, demonstrate bias and make bad decisions that detrimentally impact people's lives, whether handing out credit loans or selecting job applicants. But even where automation is beneficial, being able to scrutinise the decision-making process that impacts your life

IN ASSOCIATION WITH





should be easily available.

Transparent implementation

If vendors and businesses are to fully appreciate this, the fundamentals of technology design and implementation also need to evolve. In recent years, AI implementation has seen business expertise get “lost in translation” when developers are left to try and represent their clients’ thinking in code. The preparation and validation of data, and the selection and testing of models and their hyper-parameters, have been the preserve of data scientists. This disconnect between technology and real-life expertise increases the likelihood of baked-in biases.

With knowledge-based automation platforms, however, the experts can be on hand to make sure their logic is accurately and ethically represented, and learn how to marshal and modify the technology according to their business challenges. This last point has proved crucial from our perspective as

technology providers. When Taylor Wessing recently underwent Rainbird implementation, integral to successful technology adoption, in the words of its IT director Kevin Harris, it was “getting the people on the front lines comfortable and understanding how it works”.

Closing this knowledge gap requires a collaborative implementation process, helmed by specialists who can translate business logic into technology. Some call them analytics translators – at Rainbird we have knowledge engineers, and they’re a crucial way of capturing the nuanced, explainable logic of a client’s subject matter experts and scaling it into automation software.

Auditability eases compliance

Whether it’s ISFR 17 in the insurance sector, IR35 in the accountancy space, the continued refinement of EU customer data policies, or any number of logistical challenges associated with Brexit, businesses are increasingly having to scramble for the right tools to react to the impact of legislation. Accountancy firms could quickly find themselves making panicked, overly risk-averse IR35 decisions because they lack the ability to easily grasp and reconfigure their large-scale approach to different cases while remaining compliant. Automation with audit trails and transparent logic will grant businesses this crucial capability, enabling them to strengthen consumer and regulator trust while preparing themselves for the changing tides of legislation.

From an internal compliance standpoint, built-in explainability can give IT teams the assurances they need that AI can mitigate risk rather than introduce it. This is currently hampering attempts at AI adoption, with KPMG finding that 80 per cent of IT internal audit teams said they were not confident of AI governance in their organisation.

Towards data transparency

Strategies that heavily focus only on data could see some compromise as

legislation and regulation are demanding transparency. The industry has enjoyed a laissez-faire approach to data in recent times, but governments and regulators are finally getting up to speed. The Center for Data Innovation in Washington warns that legislation like the EU’s General Data Protection Regulation (GDPR) will limit the possibilities of data-hungry AI, with the risk of breach fines ever more likely.

With as much as 70 per cent of organisations not data-ready according to IBM, explainability for heavy data-mining cannot be anything less than immaculate in order to be conducive to a culture in which automation can be accountable and responsible, let alone trusted by consumers. Where there are no comprehensive audit trails, and with questionable data hygiene plaguing legacy industries, organisations and consumers are both in their right minds to be hesitant about trusting AI.

Interpretable automation

We underestimate the importance of understanding the things that impact us. The mission-creep of unaccountable automation into our lives has limited our autonomy and agency as citizens - but interpretable automation is how we claim it back. Business leaders are beginning to appreciate this notion. In an IBM Institute of Business Value study last year, about 60 per cent of 5,000 executives expressed concern “about being able to explain how AI is using data”.

Humanist writer Yuval Noah Harari recently discussed how our financial systems, from fraud and credit decisioning to risk modelling, have spiralled into ever-more complex, obtuse workings. When we take a wide view of the issue, explainability is even more than a question of compliance, or legality, or customer care. It’s about empowering people with understanding about how the world around them operates.

For more information, please visit:
www.rainbird.com

Jo Swinson, deputy leader of the Liberal Democrats, discusses the importance of transparency and ethics alongside new innovations

Why all AI must be ethical AI



Diagnosing rare genetic disorders is difficult. Because cases are few and far between, it makes it harder to train medical professionals in what to look for. This is precisely the kind of activity that artificial intelligence can make easier.

A new app called Face2Gene is giving doctors a second opinion on their diagnoses, using machine learning and neural networks. It looks for certain tell-tale facial features and presents doctors with a likely list of congenital and neurodevelopmental disorders. And as healthcare AI continues to develop, it will lead to more accurate and faster diagnoses, while freeing up doctors and nurses to spend more time with their patients, providing the care, advice and empathy that algorithms can't.

Advances in AI have the potential to change our lives in ways we previously thought impossible. Whether it's solving

the climate emergency, improving healthcare and care for the elderly, or making transport more efficient, the technological revolution ahead of us should be, and can be, harnessed for good. But, as well as the awesome examples of artificial intelligence in action, we hear worrying reports of how this technology is being abused too. Facial recognition can lead to better medical outcomes, but it can also be used for monitoring.

Just last month, a man was fined £90 because he refused to show his face to police trialling new facial recognition software in east London, in a modern-day equivalent of Harry Willcock refusing to show his ID card. As a liberal, I empathise with that man's attempt to protect himself from a blatant invasion of his privacy, with no consent or understanding how that data would be used.



We are also at risk of imbuing this technology with divine-like abilities to predict the future and assume that it will make better decisions than a lowly human being. The danger is we view AI as a crystal ball that can look into the future. In reality, all it does is reflect the information we feed it. It is more like a mirror that we are holding up to society.

And society isn't all that pretty, so we shouldn't be surprised that the data we are using to train these systems lead to entrenching the biases that we are still

Data should be used but not abused

trying to eradicate: sexism, racism, and all other forms of discrimination. Last year, Amazon had to stop using an algorithm in its recruitment process when it realised that, in an attempt to reflect the current male-heavy workforce at the company, it was penalising women.

As certain as I am that these technologies can make our lives infinitely better, I am also certain that they will throw up fundamental questions about the values we want to live by. And as these technologies transcend national borders, there will be a clash between western, liberal values that centre on individual freedoms and cultures where the collective good is placed above individual rights.

Given how heavily China is investing in AI, that is a battle liberalism stands to lose unless we create a global alliance committed to the ethical application of AI, which protects privacy and other fundamental freedoms and rights. The UK should be at the heart of those efforts and we should be leading by example.

The UK can and must become a world leader in ethical AI. We are already off to a good start. There is a plethora of organisations which are looking at how we ensure AI systems are transparent, accountable, fair and secure. And from the conversations I have been having in recent months with world-leading experts and businesses, through the Lib Dem Technology Commission, there is significant co-operation on tackling the challenges ahead.

The government has also recently set up its Centre for Data Ethics and Innovation. Though it's a welcome step, the centre's funding is woefully low and it's too early to judge if it will have the powers it needs to steer the development and deployment of AI.

There is more we can do. We should start with a frank and informed public debate on how we develop and deploy new technologies. In Montreal, for example, there is an institute that organises public meet-ups to inform policy development on how AI can be ethical, safe and inclusive. That's

something the Centre for Data Ethics and Innovation in the UK should consider as a model.

Just because AI can change every aspect of our life and how we interact with each other, with business and with government, it doesn't mean it should. If, for example, someone's freedom is at stake, do we want that decision to be made by a human being who can be held accountable or by an algorithm?

That's a debate we need to have, and the answer won't be clear cut. I can definitely see the merit in an algorithm that can produce fast and accurate analysis of case law in a trial, but it's a different matter if the algorithm is in charge of making the decision.

We should also appreciate just how fast technology is evolving and the pressures that places on our regulators. All our regulatory bodies need to have a solid understanding of where technology is headed and how it affects them. But we should also create the structures for these bodies to be able to share intelligence and use consistent definitions – that will help close the loopholes and business will appreciate the clarity.

And finally, the government needs to set an example for the ethical use of AI across public services. Examples like the use of facial recognition by the police that I mentioned at the start will create unhelpful public backlash and create mistrust. Government and public services need to be transparent about how they are already using AI and involve the public in determining how it should be used in the future.

The opportunities that AI offers are immense. We are so lucky to be living in this age of technological wonder that hopefully holds the answers to the biggest challenges humanity faces. And the UK is in prime position to capitalise on these developments. To succeed, we need to bring the public along with us and we need to join forces with other countries, like Canada and our allies in the European Union, to lead the charge for ethical AI that makes our lives better without sacrificing our liberal values of fairness, privacy and transparency.

Democratising AI is the best way to improve business

Innovation must be accessible for the masses, writes **Elliot Wellsted-Crook**, head of media and communications at The AI Summit

So far, the conversation around artificial intelligence (AI) for enterprise has been dominated by the opportunity it presents for large businesses who can afford to reap the benefits (and mitigate potential losses) of adoption. But with small and medium-sized enterprises (SMEs) representing 99 per cent of all UK businesses, the dialogue around the impact of AI needs to be made accessible for the whole business community. The UK government recently released the AI Sector Deal which sings to the tune of making the UK the “world’s most innovative economy” but doesn’t include how SMEs can be better prepared for the opportunities and challenges AI presents.

AI could help smaller businesses punch well above their weight. This view is maintained by AI advocates who paint a picture of the accessibility and utility of AI for all. As well as the immediate benefits AI presents for large-scale enterprises, AI can glean insights from data faster than ever before, drive operational productivity by automating routine tasks, and enable greater personalisation without the need for large staffing increases.

All of this sounds dreamy for smaller businesses with limited resources. But cutting through the buzz and understanding the bottom line ROI of AI isn’t easy, and you would be hard-pressed to find a traditional small business owner prepared to invest the time and eye-watering money into an AI application when the potential

benefits and use cases are unclear.

Indeed, when it comes to innovation more broadly, research by the Federation of Small Businesses (FSB) reported that 80 per cent of smaller business owners said that a lack of time or skills were the main barriers to innovation. The report also found that the uptake of “disruptive tech” is low, with less than half of smaller businesses having used even basic cloud services, let alone anything else.

There’s a gung-ho attitude from the AI industry that the revolutionary benefits of AI for business are open to everyone. But the radical language of AI and its promise of “transformation” and “disruption” (while true at a top industry level) does not resonate with SMEs. The jargonistic, technocratic narratives around AI have alienated smaller business owners from understanding the AI opportunity. The focus for SMEs should not be on market-changing innovations, but how AI applications can lead to incremental business improvement – not radical innovation.

This brings into focus a shortcoming of the growing commentaries around “democratising AI”. While noble in intent, the dialogue on democratising AI has focused on the potential of AI to address humanity’s greatest challenges. Meanwhile, tech giants, such as Google, have committed to open source research via the cloud to ensure resources are shared and that the power of AI doesn’t become concentrated in the hands of an elite few. The dialogue on AI’s impact needs to be made accessible for the entire business community. Repackaging “transformation” as “business improvement”, while speaking about AI more understandably, would go a long way towards democratising AI for the masses.

The AI Summit will take place at The ExCel Centre, London, on 12-13 June. The New Statesman is an official media partner for the event.

For more information, please visit:
www.theaisummit.com

IN ASSOCIATION WITH



AI risks replicating the prejudices of the past

The biggest issue that artificial intelligence faces at the moment is not a problem of technical advancement – there are leaps being made all the time. It is about designing systems and products that humans can trust. And trust comes from transparency, responsibility, and ethical design. If an algorithm tells you to do something, you won't do it if you don't have confidence in its motives.

Perhaps the most important hurdle that AI needs to get over is the issue of bias. The algorithms in AI systems are trained using large datasets, and if those underlying datasets are biased, the output is likely to be biased as well.

This creates problems if you're using AI in systems such as making credit decisions about who gets a mortgage or a credit card, or who gets invited to a job interview. If there are historic patterns, such as a higher concentration of men in senior leadership roles, then the AI is going to make its decision based on those patterns.

Bias can also be introduced to AI by the people working on it. I don't believe it's malicious, but if teams are not diverse then bias emerges, at a very low level, during the design process, and this affects the end product. The statistics on gender equality in AI are fairly depressing – women make up about 12 per cent of the workforce.

Take the growing number of voice assistant devices powered by AI. Voice assistants such as Alexa, Siri, Cortana and Google Assistant have female voices



Trust and responsibility are vital to the success of AI, writes Kriti Sharma, founder of AI for Good and technology advisor to the UK government and the UN

or feminine personalities. And they do mundane tasks, such as switching your lights on and off, or ordering your shopping, and playing your favourite music. The “male” AIs, such as IBM Watson and Salesforce Einstein, are the ones designed to make important business decisions.

It's not just a question of gender, either. There is evidence that facial recognition systems are biased against ethnic minorities and women, because the algorithms were trained on certain kinds of faces.

And background, too, makes a difference. People's social mobility and education level impact the kinds of problems that they are interested in solving with AI, and this affects the question of who the technology is being designed for. It would certainly be a shame if the greatest technological advancement of our times wasn't used for social good – improving healthcare for all, providing high-quality education, reducing inequality and so on.

I'm optimistic, because policymakers and legislators are now deeply interested in this topic. I find that very encouraging and refreshing. But I do feel there needs to be more responsibility taken by businesses. I genuinely believe that the future of our society should not be designed just by geeks like me. We need to allow a wider combination of people – people who are concerned with law, ethics, anthropology and the humanities – to take part in the AI movement, and not just be people it happens to.

Hello, we're UiPath.

We make robots so people don't have to *be* robots.

UiPath® does something called Robotic Process Automation, or RPA. That's where robots – software robots – work for you, doing the dull, repetitive tasks people loathe. This frees people to do the work they love. And that's what's helped make UiPath one of the world's most highly valued AI enterprise software companies.

"Tedium is
our medium."

"I'm your virtual
robotic assistant.
But I do real work."

"We are maestros
of monotony."



Digital transformation starts with RPA. Learn more at uipath.com