

**ARTIFICIAL
INTELLIGENCE**

WELCOME

Welcome to the debut edition of *Antenna*, our glance into the middle distance and the life-changing developments the near future will bring. In this edition, we will be looking at the impact of artificial intelligence (AI) on our lives at home, at work and for society in general. We'll also look at the investment implications which, as ever, are far from intuitive. Nobody can deny the impact of technology on our lives, but investors in Blackberry (formerly RIM) or lastminute.com know that not all boats float on the technology tide. The internet revolution has moved beyond technology stocks and into retail and media. The AI revolution stands to become even more pervasive. Headlines are being made by autonomous vehicles and robotics, but more than any preceding theme, this is one that infiltrates all sectors of the economy. AI is set to transform mundane clerical industries where productivity has stagnated over recent decades and the enhancements are already coming. The revolution has started.

Guy Foster
Head of Research, Brewin Dolphin

BREWIN DOLPHIN

ANTENNA is published for Brewin Dolphin by Wardour

01 Artificial Intelligence published June 2018

Brewin Dolphin editorial panel

Charles Ferry
Guy Foster
Nick Gardner
Gregory Thorpe

Wardour

www.wardour.co.uk
Editor Jane Douglas
Editorial contributors Jules Gray, Emily Reynolds,
Matthew Burgess, Stephen Armstrong
Creative Director Ben Barrett
Senior Designer Johan Shufiyan
Senior Account Director David Poulton
Production Director Angela Derbyshire
Managing Director Claire Oldfield
CEO Martin MacConnol

All rights reserved. Reproduction in whole or in part without written permission is strictly prohibited.

This magazine is printed by a certified carbon neutral printer, using paper sourced from sustainable sources.



IMPORTANT NOTE: The value of your investments may go down as well as up. Past performance is not a guide to future performance. Any tax allowances or thresholds mentioned are based on personal circumstances and current legislation, which are subject to change. Some products or services may be affected by changes in currency exchange rates. If you invest in currencies other than your own, the value of your investment may move independently of the underlying asset.

All information within this publication is for illustrative purposes only and is not intended as investment advice; no investment is suitable in all cases and if you have any doubts as to an investment's suitability then you should contact us or your financial adviser.

We or a connected person may have positions in or options on the securities mentioned herein or may buy, sell or offer to make a purchase or sale of such securities from time to time. In addition we reserve the right to act as principal or agent with regard to the sale or purchase of any security mentioned in this document. For further information, please refer to our conflicts policy, which is available on request or can be accessed via our website at www.brewin.co.uk. The opinions expressed in this publication are not necessarily the views held throughout the Brewin Dolphin Group. No Director, representative or employee of the Brewin Dolphin Group accepts liability for any direct or consequential loss arising from the use of this document or its contents. The information contained in this publication is believed to be reliable and accurate, but without further investigation cannot be warranted as to accuracy or completeness.



WHAT IS IT?

A layperson's guide
p.4



MACHINE 'ART'

If machines create original visual works, should we call them art? And does this mean that even creative jobs are in the sights of algorithms?

p.12

A BRIEF HISTORY



How AI grew from the imaginations of the ancient Greeks to machines that can think for themselves, via Alan Turing and autonomous Mars rovers

p.6

AT WORK: LABOUR PAINS

There is no denying that the world of work is entering a highly unstable phase, but there are some things humans just do better than algorithms

p.8



SOCIETY: FLUX CAPACITY

If we learn from the past and remain open to big thinking, AI will benefit society. But if we do not question the impact of the technology, the future could take a very different course

p.14



AT HOME: CLOSE TO HOME

Mundane tasks such as shopping and ordering a takeaway could become a thing of the past as AI gets to know us

p.22



THE ETHICS OF AI

The scope of influence this technology has over our lives raises serious ethical questions. Expert Joanna Bryson walks us through the state of play

p.18

ONE QUESTION

What is the single most transformational impact AI will have on our lives? Three experts give their perspectives

p.21



SCIENCE AND FICTION

Sci-fi has shaped the public's perception of AI, but has it also influenced the technology's development?

p.26



FURTHER READING

p.27



SNAPSHOT:



“THERE ARE MORE POSSIBLE MOVES IN GO THAN THERE ARE ATOMS IN THE OBSERVABLE UNIVERSE”

1 WHAT IS AI?

Scratch the surface of this question and you will find layers of different technologies and fields of study. But fundamentally, it is the development of machines that behave like human brains, performing tasks that would otherwise require human intelligence, such as learning, planning, reasoning and problem-solving. Although that might sound futuristic, AI is already here. It powers your internet searches, enables your voice-activated personal assistants, recommends products online and detects fraud on bank accounts, among many other things. Combined with other technologies, it also enables robots and driverless cars.

All AI that exists today is classified as **narrow** (sometimes called ‘weak’), which means it can perform only the specific task it is designed for. Technologists are working towards artificial **general** intelligence (AGI, or ‘strong AI’), which would perform any intellectual task a human can.

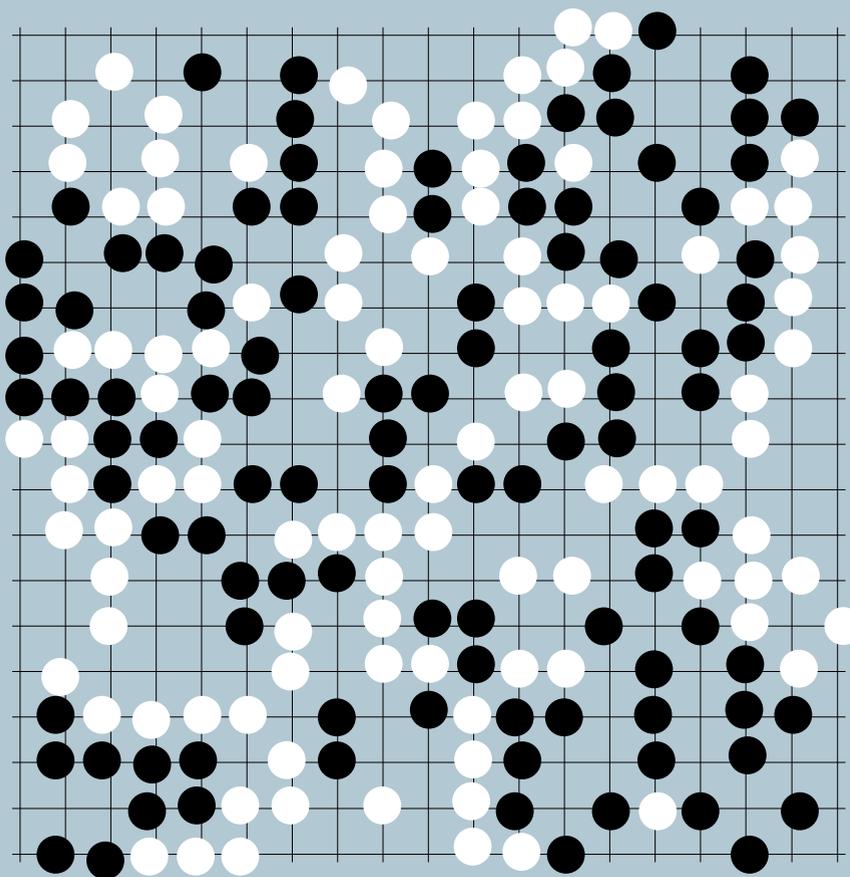
2 HOW ADVANCED IS THE TECHNOLOGY?

We are still a long way from creating AGI, but we are getting closer and the journey is being driven by computer games such as chess and Go. UK-based AI company DeepMind, owned by Google, seems to be leading the way with its ‘Alpha’ series of programs. AlphaGo Zero effectively became the world champion at the Chinese game Go in 2017. Two things make this remarkable. First, it is impossible to win the game simply using the force of processing power because there are more possible moves in Go than there are atoms in the observable universe. Second, unlike previous models, the program trained without examples of human play, using just the rules of the game.

When machines teach themselves in this way, they are not constrained by human knowledge and think in what DeepMind co-founder Demis Hassabis has described as an “alien” way. At a conference shortly after the release of the company’s new headline-grabbing

program, AlphaZero – which builds on AlphaGo Zero with the ability to learn not just Go, but also chess and Shogi (learning more than one task being a major hurdle in developing AGI) – Hassabis said: “It doesn’t play like a human, and it doesn’t play like a program.” He described maverick chess plays such as moving the queen into the corner of the board and reflected: “Maybe our conception of chess has been too limited.” The teacher, then, may become the pupil, as we are able to learn new methods from AI.

The ‘Alphas’ mentioned above are examples of **machine learning** (ML), or **deep learning** (a more advanced form), which is a subset of AI whereby systems learn from data without being explicitly programmed and make predictions. ML is thought to be the best path to AGI and is already in common use in things like internet searches. It is enabled by technology called **neural networks**. These comprise thousands of processing nodes, or layers of computation (think spreadsheets), that mimic the human brain by spotting



patterns in images, sound and text. That ability to interpret the world is enabled by advances in **natural language processing** (NLP, see below) and visual reasoning.

So AI is a discipline that comprises many layers of research, a lot of data and an ever-increasing amount of processing power. When those ingredients are brought together with a pinch of human imagination and other technologies, practical applications are born. While it is impossible to predict all the uses humans will find for AI, the problems researchers are working to solve at the moment provide clues as to the direction we are heading in.

3 WHAT WILL AI BE ABLE TO DO IN THE FUTURE?

● UNDERSTAND HUMANS

NLP aims to reduce the gap between what a human says and what a machine understands. Siri and Alexa are sparkling examples of how far this technology has come. But NLP systems are in their infancy. They are improving

steadily in their understanding of context and analogy, which is implicit to humans – for example, they are starting to analyse the form text takes, not just its constituent words, which reveals extra meaning (an abrupt email could signify an angry author). One area where NLP is already excelling is unlocking meaning from unstructured data, such as vast volumes of emails. High-street bank RBS uses NLP to spot trends in customer feedback by analysing emails and telephone calls. A related development is **affective computing**, which aims to create AI that can recognise and interpret human emotions and, ultimately, simulate empathy. Affective computing is applied to speech and video and is being trained in cultural context.

● REASON DEEPLY

While AI may have better memory than humans, it is a long way behind us in its ability to reason – a core facet of intelligence. A human, for example, knows that if you stack too many boxes on top of each other,

the bottom one will collapse – this is implicit. A machine needs to be told. Understanding the relationship between things and reasoning without specific previous examples is the cutting edge of research. Advances in this field could get AI systems thinking creatively about problems in the world around us, such as climate change, without the limitations of human understanding.

● STEP OUT INTO THE WORLD

The most obvious examples are robots that run AI programs, enabling them to traverse uneven surfaces, learning to balance as they go, and autonomous cars that interpret the world around them and take decisions accordingly. But our imaginations are the limit for how AI might get out into the world and transform our lives. One cutting-edge example that comes from Massachusetts Institute of Technology is an ingestible robot.

● MULTITASK

As the technology advances, the hope is that AI will become capable of adapting to different tasks and situations, rather than performing just the task it is designed for. Intuitive interfaces will enable users to be instructive with AI, so that eventually they will effectively be able to reprogram systems to do different jobs simply by talking to them.

4 WILL THE ROBOTS TAKE OVER THE WORLD?

With researchers determined to build AGI, there are concerns that, at some point, machine intelligence could surpass our own. This is called **superintelligence** and its advent is referred to as the **singularity**, which, it is hypothesised, would trigger runaway technological growth and unfathomable changes to civilisation. Although we are unlikely to create a superintelligence, as we do not even understand how the human mind works, leading thinkers are keen for AI to be developed within a framework that will limit the possibility: they want **safe AI**. ●

A BRIEF HISTORY

The story of AI, from the imaginations of the Ancient Greeks to machines that can develop their own scientific theories

1970s

The 'AI winter' – a period when progress stalls, partly because of a lack of computing power and data. Funders pull out, but research does not stop. Advances are made in visual perception and computer reasoning. By 1979, INTERNIST-I has been developed by Pittsburgh University academics to give medical diagnoses based on pre-programmed knowledge, and a program from Carnegie Mellon University defeats the world backgammon champion.

1960s

In 1965, Massachusetts Institute of Technology's Joseph Weizenbaum builds ELIZA, a computer program that can have conversations with humans based on scripts and pattern matching. In the same year, Stanford University academics begin work on the Dendral system, which will identify unknown organic molecules, automating the decision-making and problem-solving processes of chemists.

1950s

Alan Turing publishes a paper considering the question 'Can machines think?' and develops the 'Turing test' to gauge a machine's ability to exhibit intelligent behaviour. By 1955, Logic Theorist, arguably the first AI program, is engineered by researchers at RAND Corporation to mimic human problem-solving, and the term 'artificial intelligence' is coined by Dartmouth Assistant Professor John McCarthy. Four years later, IBM's Arthur Samuel uses the term 'machine learning'.

1940s

Mathematician and codebreaker Alan Turing starts to discuss the practicalities of creating intelligent machines with neurologist William Grey Walter. In 1946, one of the earliest electronic general computers, the ENIAC, is revealed to the world.

1830s-40s

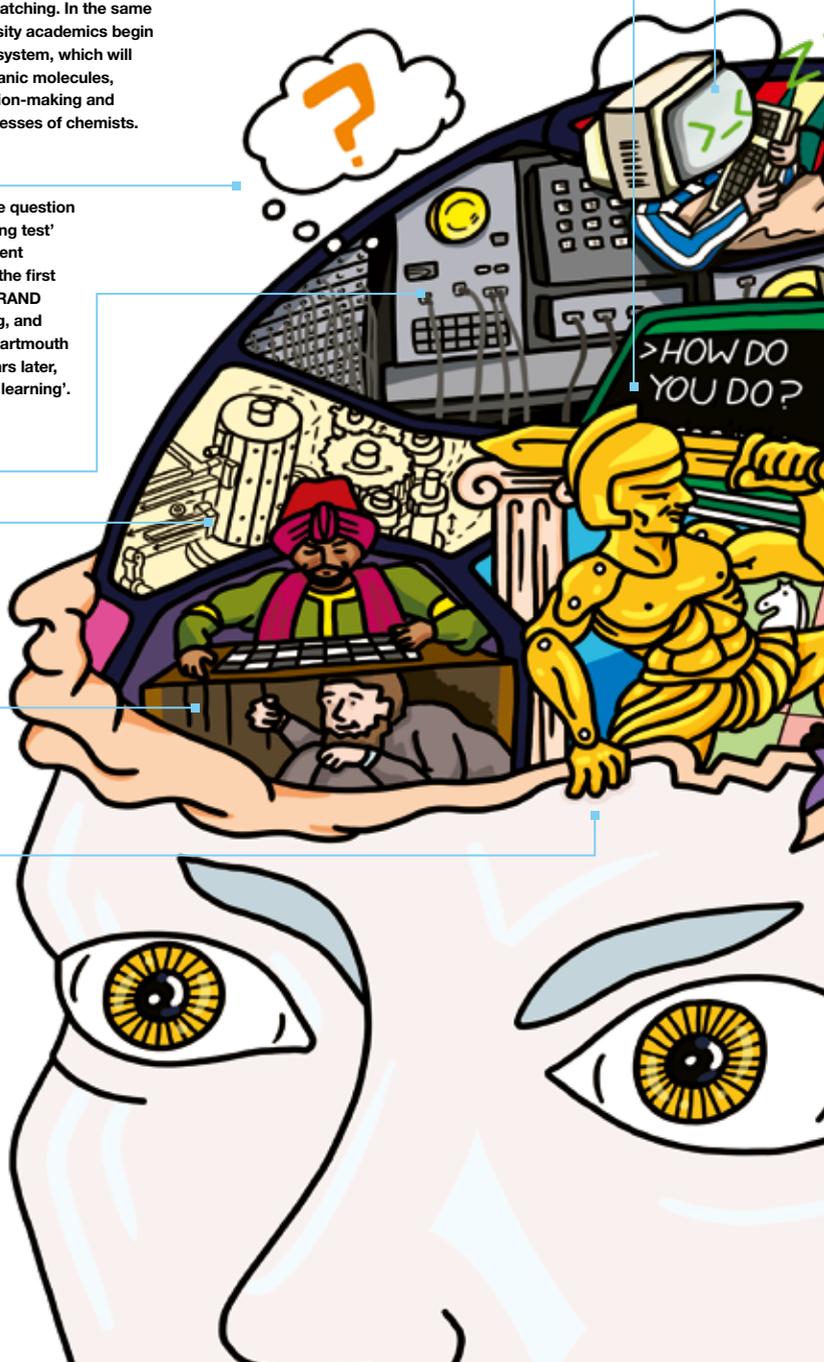
Ada Lovelace and Charles Babbage invent an early form of programming, and Babbage's Analytical Engine design lays the groundwork for computers.

1770

The chess-playing 'automaton' The Mechanical Turk makes its first appearance in Austria. It captures the imagination of the public, but is later revealed to be a hoax.

c 400 BC

The Ancient Greeks are the first to record the idea of a mechanical man. In Greek mythology, Talos – a giant automaton – protects Crete from pirates and invaders.



**“WHEN THE CHILDREN WHO
FILL CLASSROOMS TODAY
GROW UP AND START
TO WORK, MANY OF THEM
WILL BE DOING JOBS THAT
DO NOT YET EXIST”**

1980s

Expert systems that mirror the decision-making capability of human professionals begin to proliferate and the commercial potential of AI is starting to be recognised. The focus shifts to producing effective 'narrow' rather than 'general' AI (see page 4). By the mid-'80s, neural networks are in wide use and dynamic computer vision pioneer Ernst Dickmanns and his team at Bundeswehr University demonstrate an advanced early form of driverless car.

1990s

Advances in machine learning accelerate rapidly, becoming more data-driven as the internet develops. The DART AI program is deployed during the Gulf War to optimise logistics and is said to repay the US Government's entire investment in AI since the '50s. Then, in 1997, IBM's Deep Blue beats chess champion Garry Kasparov, providing a boost to the publicity and potential of AI.

2000s

AI starts to proliferate in our everyday lives, through internet searches, product recommendations based on our digital footprints and credit scoring by banks, to name a few examples. In 2004, NASA rovers Spirit and Opportunity autonomously traverse the surface of Mars.

2010s

Deep learning capabilities advance quickly, benefiting from great strides in cloud computing power and the availability of big data. In 2011, IBM's Watson wins the TV game show *Jeopardy!* and Apple launches the voice-activated assistant Siri on the iPhone, showing how far natural language processing has come. In 2017, DeepMind's AlphaZero system is able to teach itself multiple complex games using just the rules, taking a step towards general AI, and Saudi Arabia grants citizenship to Sophia the robot. But concerns start to mount over the safety of AI, with scientist Stephen Hawking and Tesla CEO Elon Musk both saying the technology is a risk to human civilisation.

2018

Chinese tech firm Alibaba produces a neural network that performs better than humans at Stanford University's reading and comprehension test, demonstrating China's stake in the development of cutting-edge AI. And Cambridge University researchers announce that AI-enabled drug-discovery system Eve has helped identify an ingredient in toothpaste that can fight drug-resistant malaria. Eve is able to hypothesise and then test scientific theories in an automated lab.

THE FUTURE...

The changes brought about by the march of technology will rewrite life as we know it. When the children who fill classrooms today grow up and start to work, many of them will be doing jobs that do not yet exist, a challenge we must prepare them for.



Illustration by Mimi Leung



LABOUR PAINS

The world of work is becoming turbulent as AI scales the list of skills previously thought to be the preserve of humans. But when the dust settles, AI should enable us to fulfil our real potential

“WHERE PHYSICAL WORK WAS THE TARGET OF AUTOMATION IN THE PAST, AI MAKES IT POSSIBLE TO AUTOMATE BRAINPOWER”

Fear is proliferating through almost every profession as the realisation dawns that the aim of AI - to create machines that can perform any cognitive task a human can - seemingly puts every job in the sights of technologists. In reality, while there will be many casualties as the technology improves, AI should ultimately enhance what we can achieve and allow us to concentrate on the most human skills.

Research from Oxford University suggests that 47% of all jobs in the US could be automated within 20 years. “If your job is essentially routine and predictable, you’re at risk, while skilled manual trades like electricians, plumbers and nurses can’t be replaced,” says futurist Martin Ford, author of *Rise of the Robots: Technology and the Threat of a Jobless Future*. The latter set of jobs is protected by the difficulty and expense of building dexterous robots to give AI physical form, and the possibly insurmountable problem of creating machines capable of empathy.

Repetitive jobs have often been the first to be disrupted through history. But where physical work - be it on the production line or the telephone switchboard - was the target of automation in the past, AI makes it possible to automate brainpower. “AI broadens the range of jobs that machines can take over into almost everything involving data,” explains Tom Hulme, General Partner at Google’s venture capital arm, known as GV.

The kinds of jobs that will be affected first are rules-based and non-subjective, such as data collection and synthesis. One face of that change is robotic process automation (RPA): the software slots into a company’s existing IT system, logging on as if from a desk and mimicking the behaviour of a human employee. “With RPA you can take, say, 60% of your company’s rules-based repetitive tasks and get the robot to take the repetitive work away from employees,” explains David Wright, Director in Robotic and Cognitive Automation at Deloitte. “Software robots can open email attachments,

complete e-forms, record and re-key data, and replace many middle- and back-office processing centres. Our analysis suggests the cost of a digital worker is about one-third the cost of an offshore worker.”

RPA-based systems also run non-stop, so it is easy to see the productivity potential. According to Raja Subramanian, Group Enterprise Architect at technology services company Xchanging, robots processing insurance-related tasks can clear over 30,000 cases per month, reducing processing time from five minutes to under 10 seconds.

Advances in natural language processing mean AI is even steadily encroaching on jobs involving human interaction. Already, Barclays is using conversational AI chatbots to deal with unbanked populations, including the illiterate. In education, AI is starting to replace routine jobs - for example, language learning platform Duolingo teaches 200 million people new languages by getting students to listen and repeat. And taxi drivers are starting to grow concerned for their jobs.

So tasks based on set rules and processes, even those that may seem on the surface very human, are already feeling the effect of AI. But the programs processing and analysing data still need human oversight and with the opportunities companies are seeing in unstructured data blooming, they are starting to employ more data scientists and advanced analytics and machine-learning experts.

TRANSFORMING THE PROFESSIONS

It is not just low-paid routine work that is changing. The great strides that have been taken in deep learning (see page 4) mean that knowledge-based professions such as law and medicine are also set to be transformed. The work that UK-based AI company DeepMind has done with the NHS, building on lessons learnt from a computer-game playing algorithm, is at the forefront of this trend. In 2015, the company unveiled a deep learning system that could increase the score in Atari console games 

**“PEOPLE TEND TO
OVERESTIMATE HOW
MUCH OF THEIR
WORK DEMANDS
JUDGEMENT,
CREATIVITY
AND EMPATHY”**



without information on the rules. DeepMind recently started applying this technology to develop diagnostic tools for head and neck cancer with University College London Hospitals, acute kidney injury with the Royal Free Hospital and age-related blindness with Moorfields. In an interview with *The BMJ*, DeepMind Co-founder Mustafa Suleyman explained that diagnosis is “actually very similar” to the Atari research. When preparing for radiotherapy, a patient’s CT scan is marked up by a radiotherapist to identify healthy and cancerous tissues for accurate targeting. “This takes about four hours for each radiographer to do, so there’s a very long delay from deciding on treatment to being treated,” he said. “We can train a model to do this faster and more accurately.”

In the surgeries and hospitals of the near future, routine check-ups, mammograms, lung cancer scans and even dementia diagnosis will be done by AI. In some cases, such as a pneumonia-detecting program developed by researchers at Stanford University, the software is even better than humans at spotting problems. And as machine learning advances, programs are starting to not just detect, but predict diseases. It is not hard to imagine a point when – as AI starts to reason for itself, unrestrained by human understanding – disease prediction will be vastly improved by the technology.

For now, these developments do not mean specialist doctors, such as radiotherapists, ophthalmologists and nephrologists,

will all lose their jobs. Indeed, it should make them more precise and free them up to concentrate on elements of their work that require judgement, creativity and empathy.

However, it is likely there will be some losses. Government adviser Richard Susskind, who looked at the impact of AI across law, education, audit, tax, consulting and architecture, points out that people tend to overestimate how much of their work demands judgement, creativity and empathy. When broken down into its constituent parts, the work of many professionals is routine and process-based. Taking a macro view, a recent OECD study looked at the susceptibility of tasks within jobs to automation and found that, across 21 countries, 9% of jobs are at risk. Compared with the 47% figure cited by the Oxford study, this may seem small, and shows how much debate there is around the topic, but this amounts to millions of jobs.

In the longer term, advances such as AlphaGo Zero (see page 4) show how AI might one day be capable of problem solving in ways we could not imagine. And with teams of researchers plugging away at the challenge of affective computing (understanding human emotions), natural language processing and reasoning, and machines already creating original artworks (see page 12), it is possible that many more tasks could be automated in the future. Even those plumbers and electricians might be out of a job once

AI can cheaply be given humanlike form through advances in robotics. That future, of course, is very distant and reliant on technical breakthroughs that may not be possible.

REWRITING THE WORK LANDSCAPE

It is likely that people will try to resist progress, but ultimately, the advantages of AI, as with steam and electrification before it, will make a necessity of change. Technological revolutions completely rewrite the landscape so that it is not just jobs that are redefined, but the fundamentals of the organisations creating them. As the Massachusetts Institute of Technology's Andrew McAfee and Erik Brynjolfsson have highlighted, in retrospect, it is clear electric power was so transformational to industry because it enabled the redesign of the production line in ways previously inconceivable, such as the introduction of conveyor belts. At first those advantages were not at all obvious, as the 'curse of knowledge' blinded incumbents to the possibilities of the technology, but eventually those who did not forget what they knew about business and go back to the drawing board, failed. Almost every industry will go through this kind of process as AI picks up pace.

It is impossible, standing where we are, to see what is beyond the horizon. But given the sheer scale and pace of change that is to come, everyone will need to be adaptable, from the individual worker to the businesses that employ them. But, while we are headed into turbulent times, AI should ultimately make our work more enjoyable. Waves of technological progress in the past have relieved humans from the drudgery of hard manual labour and created more fulfilling, knowledge-based roles. For now, we vastly outstrip machines at perception, reasoning, communication, creativity (be it coming up with theories or musical compositions) and anything involving empathy or social skills (including things like negotiation, persuasion or care). As Google's Hulme says: "Sixty-five per cent of children entering primary school today will ultimately end up working in completely new job types that don't yet exist. Those who truly excel will be able to harness technology for repetitive tasks and save their efforts for more creative work... it could be the best thing that's ever happened to us." ■

THREE TAKEAWAYS

- ① Rules-based, repetitive jobs will be the first to be replaced by AI, from back-office data processing to dealing with basic customer enquiries.
- ② Highly skilled work will also be hit by the coming wave of AI, as the technology automates functions within medicine, law, audit, architecture, education, tax, consulting and more.
- ③ New technical jobs are being created to make the most of AI, and humans remain better at perception, reasoning, communication, creativity and social skills.

OPPORTUNITIES: SPOTTING THE WINNERS

There are many disruptive technologies, but who will ultimately gain from them?

The value chain for AI-related assets is long and complex and many of the players, such as Alphabet (Google), feature at almost every stage; first as a proprietary data source, but also through the modelling and training stages, where they provide infrastructure, and finally into applications, such as those found in self-driving vehicles.

The workhorses of AI are chips. They are at the heart of the data centres where AI training, the interrogation of large datasets to facilitate deep learning, and inference (the use of that learning) take place. Again, some of the big chipmakers, such as Intel, work with AI companies from the design and production phases, while elsewhere the process is separated, with firms such as ARM specialising in design but outsourcing the actual manufacturing to a foundry. Chipmakers come with the requisite risks, however, as the benefits of AI will be shared among many firms within the technology sector, such as those providing customer service and process automation.

To sidestep the perils of a complex value chain, the best play on robotic cars may well be tyre companies. They will benefit from the anticipated increase in miles driven regardless of which technology provider wins the race to perfect the technology.





WOULD YOU CALL THIS ART?

Your gut instinct might be to say 'yes', but what if you knew it was created by a machine?

Researchers at Rutgers University in the US gave an algorithm 80,000 images of western paintings, then asked it to generate original artworks that did not copy existing styles. "When you look at the results, you see that the machine chose colour in an aesthetically meaningful way, that its compositions followed established rules, and it balanced warm and cold colours. All these things the machine figured out for itself, by looking at art. But, at the same time, it developed new ideas," says Ahmed Elgammal, Professor of Computer Science at Rutgers University, who led the project. And the results were convincing – in a survey of the general public, 75% of respondents thought the pictures were by a human artist.

So does that make it art? Boris Čučković Berger, Associate Lecturer at The Courtauld Institute of Art, does not think so. "The machine is reproducing a graphic style as an abstraction of what is the norm. Art happens in dialogue with the social and aesthetic features of our environment; it produces culturally or historically significant references, and potentially shifts or undermines the norms."

But given its aesthetic appeal, does this technology mean that creatives will be competing with algorithms for work? "I don't see it as competing," says Elgammal, "more as AI providing creative tools. When the camera was invented, anyone could push a button and create a picture, so artists had to change the way they made art. Art became more of an abstraction of the artists' ideas about the world. Artists will find a way to use this new creative tool to make forms of art for the 21st century."

Čučković Berger agrees that it is human agency that is ultimately the creative force here: "The machine itself is not occupying the role of the artist any more or less than a paint brush did in the past. The question remains what can be done with it: what kind of a world does it portray?"



FLUX CAPACITY

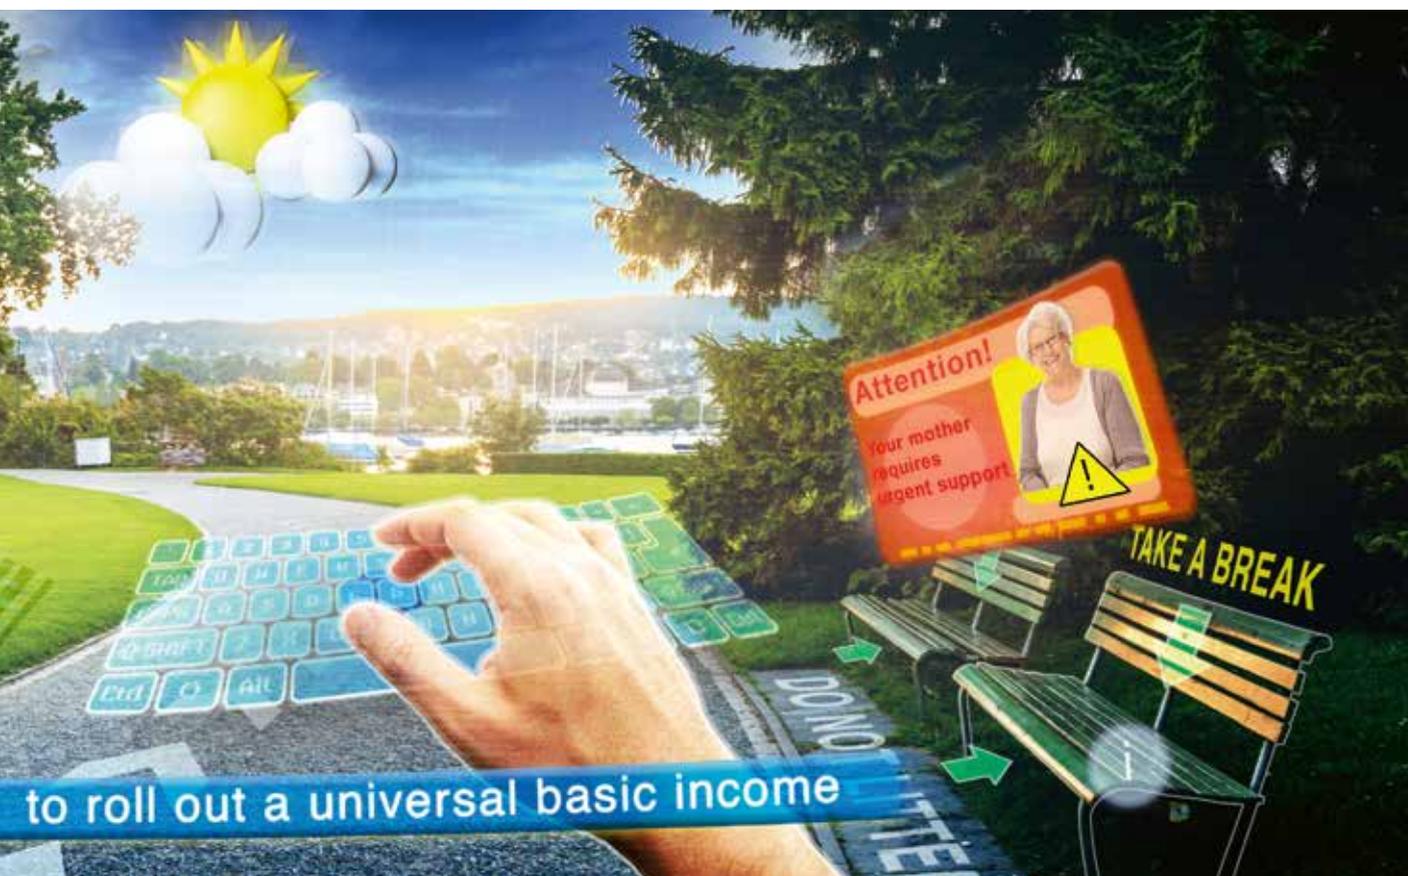
Technological change in the past benefited society in the long run, but it wreaked havoc along the way. History does not have to repeat itself

Writer and professor Isaac Asimov lamented in 1988: “The saddest aspect of life right now is that science gathers knowledge faster than society gathers wisdom.” His words echo down the years as AI places us on the cusp of a new technological revolution. But science does not have to outstrip wisdom. We are more organised and better connected than ever before in human history, which means we have the tools needed to adapt quickly and, in the long run, to harness AI for positive social change.

Given the complex, interdependent systems that make up our society, it is impossible to predict how individual changes might add up to macro effects, just as we could not have foreseen how social media would polarise political debate. But there are some broad themes, such as unstable work, changes in human interaction, and even a new phase in warfare that we can already see emerging, and for which we can be prepared, if we are willing to think big.

SOCIETY IN TRANSITION

As we have seen, AI is unlikely to replace work entirely any time soon, but we are heading into a period of instability, and the shift from knowledge-based jobs to creative and social roles will be tumultuous.



“There are a lot of important parallels with the industrial revolution,” says Ryan Avent, author of *The Wealth of Humans*. “The most significant is that the relationship between workers and the tools they use is changing.” Just as with the industrial revolution, automation today means machines have the potential to do work that previously required many people, with fewer individuals required to oversee them. In the long run, this trend towards the concentration of wealth could see workers further organising to “compel” their employers to share the benefits of AI, the author says.

One solution to large-scale automation is universal basic income (UBI). The concept is simple: the state pays money to every citizen. The amount can vary, but in one country it could be enough for a person to live without working; in another it could be a supplement to working, or a lump sum to allow for retraining. While this might seem like a radical idea, trials are already underway in Barcelona, Ontario and Stockton, California, with more governments considering pilots and four Scottish cities due to start experiments this year. Although, the Finnish Government recently refused a request by its social security agency to extend a two-year trial that paid a fixed, unconditional amount only to the unemployed, opting instead to trial other systems. But the UBI trial will run its course and results will be out in 2019.

It is clear to see how a full UBI system would help individuals adapt to the changing work landscape: enabling through-life retraining as the job market morphs. But for the scheme to be adopted to the degree that someone could live on a UBI because there are not enough jobs to go around, a big shift in social norms and politics would be needed.

“You have this huge social question of what do people do with themselves?” says Avent. To avoid a world where people spend all their time on the sofa, and prevent friction with those still working, society would need to create norms that drive people to be active. “It’s much easier for us to imagine coming up with jobs that don’t really need to be done than to imagine us coming up with a social system that makes it acceptable for people to not work,” he explains. Those still in work would, after all, want people who are not economically productive to make some kind of social contribution.

Our education system, too, will need to reflect the change in work, with more emphasis on creative and social skills. Without knowledge, of course, creative, social and critical thinking skills are redundant, but we are already seeing some schools develop curricula that provide children with a more expansive education. School 21 in east London, for example, focuses not just on the academic success of its pupils, but also on character, problem-solving and



“FOR THE ELDERLY, AI COULD PROVIDE A LIFELINE”

idea generation. It teaches children to be adaptable, eloquent and knowledgeable so they are equipped for the challenges of the 21st century. Scientific, mathematical and technical skills will also be valuable as new jobs are created in this space.

HUMAN BONDS

On an individual level, AI will affect the way we interact with each other. This process has already begun, with social networks such as Facebook and Twitter using algorithms to identify who we should become friends with and suggesting groups of like-minded strangers. “It can counteract loneliness produced by social, geographic, physical and, thanks to Google Translate, linguistic isolation,” says Mark Sprevak, Senior Lecturer in Philosophy at the University of Edinburgh. But there are concerns too. “There is a more insidious worry that, as we spend more time interacting with ‘friends’ over AI-mediated connections, we have less time to spend on genuine friends and even the idea of a ‘genuine’ friend may lose its meaning and value,” he says.

For the elderly, typically the focus of concerns about loneliness, AI could provide a lifeline. In Japan, robotic home assistants using elements of conversational AI and facial recognition systems are being deployed to help care for senior citizens. When the technology is more advanced,

it could even help people maintain independence for longer. Japan, like the UK, has an ageing population and there are worries about the overall cost of healthcare. Effy Vayena, bioethics and digital health expert at the university ETH Zurich, believes AI systems can “offer more cost-effective approaches to elderly care”.

The opportunity exists, then, to solve one of the biggest challenges facing our society today: how to care for the elderly. But, again, we need to be attuned to the long-term implications of the technology across society, as the loss of wellbeing and the sense of belonging that comes with human interaction could ultimately compound loneliness.

THIRD REVOLUTION IN WARFARE

Security is another area where we can be proactive in managing the effects of AI. Machine learning has a role in national security and crime fighting, according to Chris Hankin, Co-Director at Imperial College London’s Institute for Security Science and Technology. “In physical security, the possibility of using intelligent robots in hostile environments, thus avoiding putting humans at risk, is an important application,” he says.

Such applications will need to be managed carefully, though, as the technology could be used to build autonomous

lethal weapons, which 116 robotics and AI experts (including CEOs of companies leading research in these fields) described in a letter to the United Nations (UN) last year as a “Pandora’s Box”. Cautioning against the development of such technology, which could stem from the same advances that enable driverless cars to avoid pedestrians, the experts cautioned: “Once developed, they will permit armed conflict to be fought at a scale greater than ever, and at timescales faster than humans can comprehend.”

At the present level of sophistication, the scope for error when machines encounter a situation they are not familiar with is also very high, so this kind of technology would be risky on the battlefield for several reasons. But it is not inconceivable, as another group of high-profile AI experts, including Tesla’s Elon Musk and Apple co-founder Steve Wozniak, have cautioned, that an arms race could begin. It is reassuring, then, that the UN has already started talking about setting limits on weapons that can kill without human intervention. But again, this is an area where momentum needs to build, as Russia has already voiced opposition to a formal ban.

On a more immediate issue, Microsoft released a policy paper last year, in the context of proliferating cyber-attacks, calling for a digital Geneva Convention between states to prevent them from launching such attacks if civilians could be harmed in the process; for example, when hospitals or critical infrastructure are targeted.

We do have the opportunity, then, to limit the dangerous consequences of weaponised AI, and it is interesting to see pressure coming not just from human rights groups, but from some big technology companies that might benefit from unfettered development of AI.

Ultimately, the AI we create and its impact on society will mirror those who build it. Given the scope of influence the technology will have over society, it is the responsibility of everyone – business, government and civil society – to proactively consider the potential butterfly effect of each new development, and react wisely. If it is mismanaged, AI could exacerbate inequality, amplify human pathologies and create a third revolution in warfare. If managed well, the technology could benefit our working lives, welfare and security. ■

THREE TAKEAWAYS

- ① Social and political norms may need to be dramatically reshaped to enable people to deal with the tumultuous world of work.
- ② AI can be deployed to solve big social questions, such as how to care for an ageing population, but we need to keep an eye on the side effects.
- ③ In the realm of security, too, we need to be mindful of the consequences of new technologies to avoid a dystopian future of warfare.

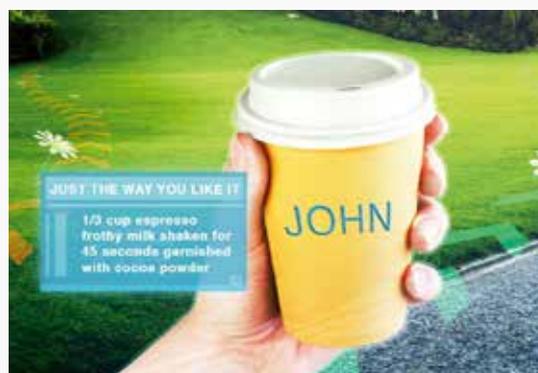
OPPORTUNITIES: CAPITAL GAINS

Although AI is steeped in uncertainty, there is some clarity about who will benefit

Plus ça change, plus c’est la même chose, as the old saying goes. Changes to the employment landscape brought about by AI may well be huge, but forecasters find it difficult to determine whether this will result in a big increase in unemployment or the creation of a huge number of new roles. We may well see both.

Productivity gains should enable big increases in wages for the people at the top of the AI revolution, and unskilled workers may share in this to a lesser extent. The owners of capital stand to gain the most through returns on their investments, while the economic gains should increase the amount potentially available to fund a social safety net (which could be a universal basic income) if one is required.

Generally rising living standards are unlikely to flatten the wealth and income inequality we have seen in recent years, which means wealth brackets will still be identifiable. This creates opportunities for brands to target different segments, facilitated by logistics, platform and data companies and the right real estate (AI-enabled retail experiences are starting to seek out physical premises). While in the past owners of brands held the advantage by monopolising distribution, with shelf space at most major shops (Unilever, Colgate-Palmolive), those with the greatest intellectual property are likely to have the edge from now on (LVMH, Apple, L’Oréal, for example). Companies that can use the data they hold to personalise and customise their products may well have the advantage online.





The
ethics
of
AI

Joanna Bryson is a leading expert in AI who divides her time between the universities of Bath and Princeton. Speaking to *Antenna*, she explains that AI is forcing decisions on age-old ethical quandaries

Humans arrive at ethical decisions based on complex cultural learning that spans a lifetime. In the age of machine learning, algorithms are starting to take those decisions on our behalf: should a driverless car hit a pedestrian or endanger the life of the passenger if it has to choose one? The idea of machines taking very human decisions is sending a shiver down the spines of many. But, explains Joanna Bryson, who is an Associate Professor in the Department of Computing at Bath University, these are actually very human problems: “The great ethical questions people are asking about AI are actually just great ethical questions about humanity and about our society. For whatever reason, we were afraid to rock the boat when we were just thinking about ourselves, but when we put a machine there, we can ask these questions.”

In an accident scenario, each human reacts differently – some would hit a pedestrian rather than drive into a river and endanger their own life, others would not. “The only thing that’s different now is that the decision is mass produced so we have to make up our minds about what is the ‘right’ course of action,” says Bryson, comparing the process to law-making, which already codifies human values.

So AI is forcing society to confront ethical questions that were previously in the personal domain. But who is in charge of this process? Are these decisions being made in the labs and boardrooms of companies, or in parliaments and citizen forums? For the most part, how AI reacts in any situation is determined by the parameters set by the technology developers. Some companies, such as DeepMind, are acting responsibly by convening ethics panels that include voices from civil society to consider questions such as how values designed into AI systems can be truly representative of society. DeepMind – alongside Amazon, Apple, Facebook, Google, IBM and Microsoft – is also part of the Partnership on AI, which provides an open platform for discussion and engagement around AI’s impact on society. As Google co-founder Sergey Brin said in a letter to Alphabet shareholders in April, the current era of AI requires “tremendous thoughtfulness and responsibility”, citing the impact on employment, fairness and manipulation as concerns. Whether deeds will follow words remains to be seen.

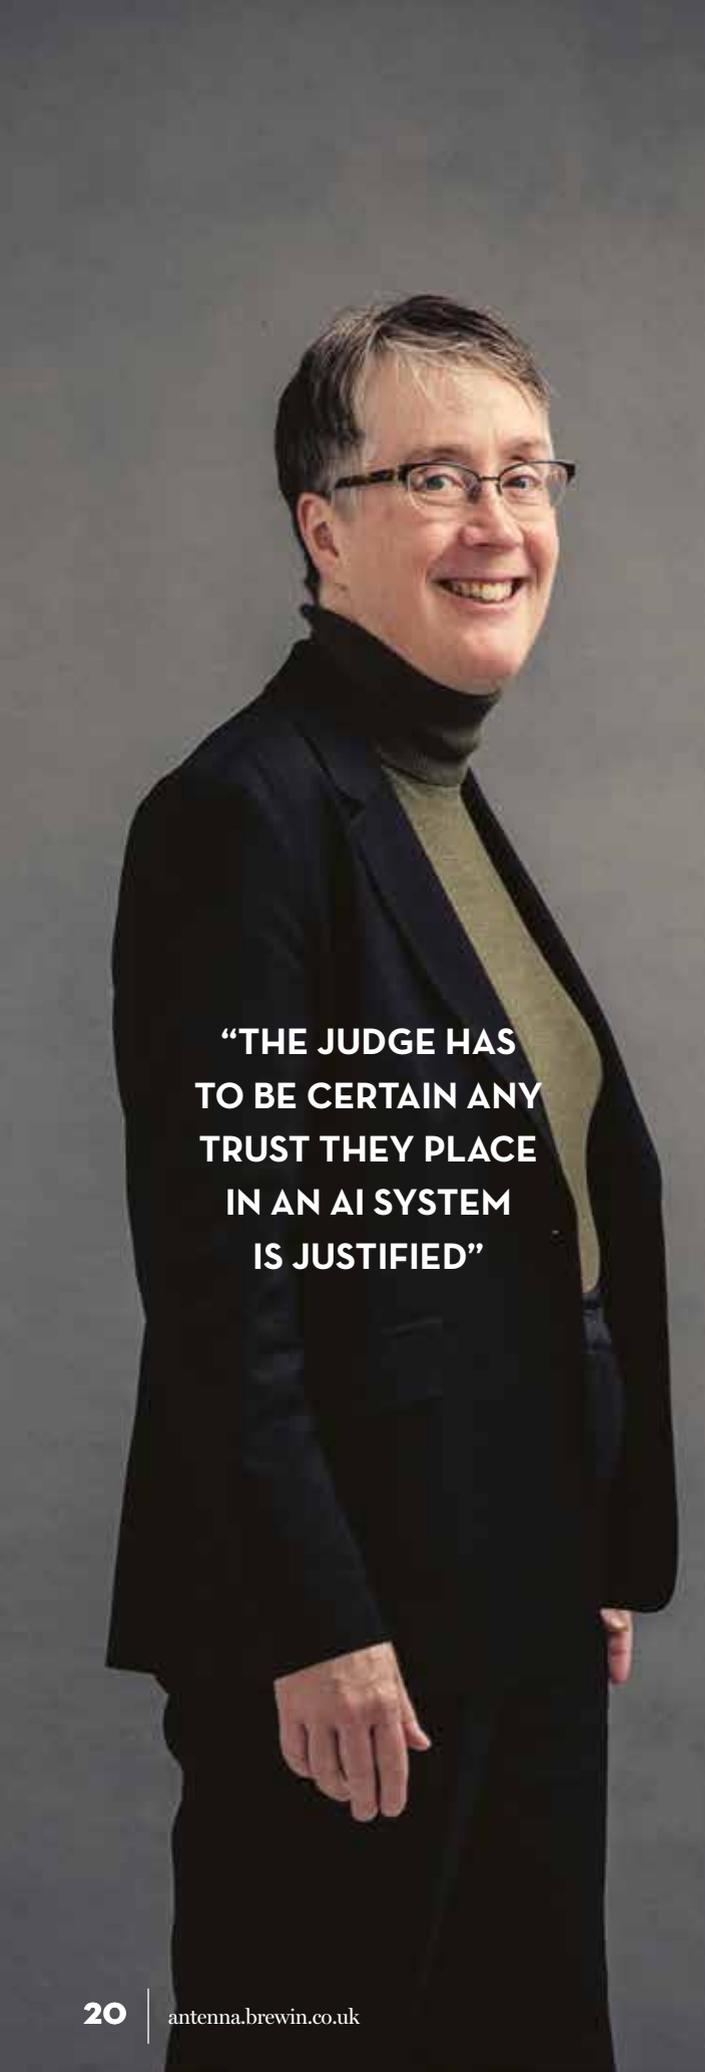
Governments have been slow to update regulation, probably for fear of stifling innovation. But one country that is on the front foot is Germany, where the Government has developed guidelines on driverless car ethics. In the event of unavoidable accidents, they read, any distinction between individuals based on personal features (age, gender, physical or mental constitution) is impermissible. This makes Bryson happy, as

“the car itself is not responsible for that decision because that decision was already made”. Society has recognised that it is responsible rather than the machines.

But the ethics of AI is about more than morality puzzles in any given situation. Bryson flags the broader, indirect effects of AI as one of her top concerns: “We don’t understand what we’re doing to our democracy, what we’re doing to our economy, and with AI the pace of change is faster.” She explains: “Think about London taxi drivers – anyone can be a London taxi driver now, all they need is an iPhone and Uber. In a way, that’s incredibly cool, but on the other hand it means one of the ways we used to differentiate wages has gone away,” she says, referring to the changes to the economy brought about by new, platform-based businesses. These often concentrate wealth with the technology developers, while those powering the businesses are termed self-employed rather than workers, which means they are not entitled to holiday or sick pay and the platforms do not pay employment taxes. “That’s part of the reason you get inequality, and when you have high inequality, you get high political polarisation, violence and social chaos.” Indeed Uber recently bowed to public and legal pressure and offered access to medical cover, including sick pay and parental leave, to its European drivers.

She has concerns over the direct consequences of the technology, too, and says there is a limit to how far individuals can protect against malicious applications of AI, even if they hold back their personal data. “Even at this point, if you decide to have no computers in your house, no mobile phones, you are still exposed. The more data we have about people, the better models we can build, and the better models we have built, the less data we need to predict what any one person is going to do.” Bryson uses the example of an AI model created by two Stanford University academics to determine whether people are gay (with a middling degree of accuracy) to exemplify how this can be dangerous. “Anyone with your picture could have an idea of who you’d like to date and that’s scary.” Her words ring true in a world where same-sex relationships are illegal in 72 countries, according to a 2017 report by the International Lesbian, Gay, Bisexual, Trans and Intersex Association.

Even when AI models are in just hands, they can lead to unethical outcomes. In courtrooms across the US, algorithms are used to predict the likelihood of a criminal reoffending. Judges use that information to decide on bail terms and sentences, but in some cases the software could be biased. Non-profit news site ProPublica analysed AI-generated risk assessments for 7,000 arrests in one US county. They found the program, used by judges to assess recidivism, was almost twice as likely to wrongly label black



**“THE JUDGE HAS
TO BE CERTAIN ANY
TRUST THEY PLACE
IN AN AI SYSTEM
IS JUSTIFIED”**

defendants as future criminals than their white counterparts, while white people were more likely to be mislabelled low risk. Of course, the system was not programmed to be biased, and it may be that the factors by which it came to its conclusion are commonly used to assess recidivism, but if ProPublica's analysis is correct, the results are nonetheless concerning.

Bryson emphasises the importance of remembering that humans are in the driving seat here. “If you say AI is making the decision about who should go to jail, you're thinking about it wrong. The judges are making the decision. If they use an AI system for advice, the question is what was that AI device telling them? The judge has to take responsibility for their own decisions, and therefore be certain that any trust they place in an AI system is justified.”

When machines teach themselves, we sometimes do not know why they have reached certain conclusions – the ‘black box’ problem – which has led to concerns over transparency. Bryson believes that, so long as the developers of such systems are held accountable for their products, they will find the means to make systems adequately transparent so that they can be used safely. This may slow the rate at which new technologies are released, but it will ensure they are more usable and therefore more sustainable in the long run.

HAVE A WORD

Bryson's research has found that machine bias is sometimes learnt from our language. “People don't usually think that implicit biases are a part of what a word means or how we use words, but our research shows they are. This tells us all kinds of things about how we learn prejudice. The fact that humans don't always act on our implicit biases shows how important our explicit knowledge and beliefs are. We're able as a society to come together and negotiate new and better ways to be, and then act on those negotiations. Similarly, in AI, we can use implicit learning to automatically absorb information from the world and culture, but we can use explicit programming to ensure that AI acts in ways we consider acceptable, and to make sure that everyone can see and understand what rules AI is programmed to use.”

For Bryson, the most interesting question in the ethics of AI right now is what the sudden access to so much information will do to society and human agency. “If we have enough information about people, such as personality type and what sort of day they've had at work, and we can predict with some certainty that they're going to go home and commit domestic abuse, are we obliged to intervene?” she asks. “I think once you know then yes, you have an obligation to act, but that reduces human agency if suddenly we're all watching over each other.” She sees that as the critical ethical question: “As we get to know more about ourselves and each other, how is that going to change what it is to be human? How does that change what a community is, how does that change our relationship to our country, our city, our family, our church and our job, and the corporations who know all about us?” ■

Photography by Will Bremridge

ONE QUESTION

WHAT IS THE SINGLE MOST TRANSFORMATIONAL IMPACT AI WILL HAVE ON OUR LIVES?



OLIVER RATZESBERGER,
Chief Operating Officer, Teradata

“Ubiquitous AI, not the AI of sci-fi fame, but AI built into everything around us, will automate repetitive or predictable actions in every aspect of our lives.

Entry-level functions of all forms are being automated. Skilled roles like salespeople, law clerks, drivers and hedge fund analysts are being replaced by AI-based software. Repair roles are reduced by [AI] guiding maintenance to intervene before predicted failures. The replaced roles are part of an ecosystem: when drivers are not required, neither are services to support them. AI-based automation will widen the gulf between the 1% and the rest of humanity without significant change at the policy level. It is up to all of us to retrain and retool ourselves, teach our children to continuously develop new skills and advocate for education and skills training for those whose skills are obsolete by AI and automation.”



PROFESSOR NICK JENNINGS CB, FEng,
Vice-Provost (Research and Enterprise),
Imperial College London

“AI will transform all aspects of our society that rely on information for decision-making and action. Since this is almost everything we do, it will have profound effects on our lives at home, work and play. I see AI systems as working in partnership with humans: suggesting options we may not have thought of, challenging the decisions we make when they appear unusual, and dealing with the routine. The AI systems will need to learn to work alongside us and learn how to be good collaborators and team players. Humans will need to adapt to working alongside smart machines and to trust them. By working together, we can raise both the humans and the machines.”



RAJ BIYANI served as the
Managing Director for Microsoft
IT India, Hyderabad from 2010-16

“There’s a popular time management theory based on the ‘rock, pebbles, and sand in the jar’ story. If the imaginary jar in the story represents the 24 hours per day in one’s life, the rocks represent the most important priorities (family, health, spiritual development), the pebbles represent the second set of priorities (professional, financial), and the sand represents the filler activities. However, the dilemma confronting most of us is that so much of our time is consumed by the filler activities (responding to routine emails/WhatsApp/text messages, managing one’s calendar, etc) that we don’t have enough hours in a day left to focus on the big rocks. AI can provide each one of us with a personal digital assistant to help address these filler activities (coordinating with your friend’s digital assistant to schedule the best time and location for you to meet for dinner; or adding the appropriate travel buffer time between appointments to your calendar to account for the traffic; or advising you on weather-appropriate clothing... the list goes on).”

CLOSE TO HOME

In our personal lives, AI will take over mundane tasks, making some very familiar objects redundant and improving the consumer experience, but only if providers build a relationship of trust with users

Technological progress has created an always-on world, where the boundaries between work and home seem forever under assault and our to-do lists overflow. But AI could be the antidote. It is already saving us precious spare time: ever asked Siri or Alexa for directions, called an Uber, or had inappropriate emails filtered into a spam folder? Then AI has made you more efficient.

This trend is set to continue as the applications of AI multiply. But as it steps up to the next level, AI will not just advise us, it will start to take decisions on our behalf and, as it gets to know us almost better than we know ourselves, everything will become more personalised.

So how will this play out in our day-to-day home lives? “Sadly, I don’t think we’ll see robot butlers any time soon!” laughs Professor Michael Wooldridge, head of the University of Oxford’s Computer Science unit and author of the *Ladybird Expert Guide to Artificial Intelligence*. He says that near-future applications of AI are going to be “a bit less obvious”, pointing to AI-driven “smart heating and power-management systems” as one of his major predictions for the future. Such systems already exist, of course, with companies including Nest, Hive and Honeywell selling products that allow you to control your heating remotely, adjust temperatures in separate rooms of your home, and thus help you save money.

SMART HOMES

But such systems will only get smarter. Dave Coplin, Microsoft’s former Chief Envisioning Officer and an expert on future technologies, believes their influence will expand to include

other areas of the home in ways that could significantly change our lives. “It’s yet to reach widespread adoption, but the use of AI technologies to connect a bunch of ‘dumb’ devices, creating a truly smart home, is coming fast,” he says. “If every aspect of your home is connected – the doors, windows, fridge, TV, heating, cooling, lighting, curtains, garage door and so on – then AI can provide the intelligence to connect and manipulate those devices in a way that is completely tailored to you.”

What this means in reality, Coplin says, is having a home that not only responds to your commands, but actively learns your habits, and performs tasks it predicts you may want before you even think to ask for them. “When you arrive home, the bath has been run, the temperature is perfect and the drinks have been chilled to perfection; your favourite TV boxset is primed and ready, and, just as you finish laying the table, the doorbell rings with your food delivery,” he says. “You haven’t actually ‘asked’ for anything, but everything has been done for you based on what you’d normally do in such a scenario.” This may seem fantastical, but, as with the smart heating systems already mentioned, it is starting to happen, admittedly with human intervention. Futuristic devices such as smart fridges can be found in homes around the world, although inventories have to be updated by owners, rather than devices automatically re-ordering food. If Coplin’s predictions are right, however, it may not be too long before our fridges are re-stocking themselves with our favourite drinks.

To arrive at the ideal smart home, we will need to upgrade the physical infrastructure in our homes. This process has not happened as quickly as some predicted, 





possibly because of the cost and security fears. But in the US, the market has seen steady growth over the past few years, with the number of connected homes rising from 17 million in 2015, to 22 million in 2016 and 29 million in 2017, according to consultancy McKinsey.

“The home is a setting where there is a natural repurposing of space and function over time. This is nothing new,” explains designer Kim Colin. She was part of the London-based Future Facility design team who conceived the Amazon Apartment at the Design Museum, a “conceptual home of the future” for older people that automatically re-ordered household products and self-maintained appliances.

Colin believes that, in the transitional period between “invention and total adoption”, we will likely “retrofit” smart home architecture. “Before we change the infrastructure in the built environment – which is a very, very slow process compared with the rate of change of technology – we will likely work around existing architecture,” she explains. But innovators, she says, will make a totally “fresh start”. We have seen tech companies “unburdened by the past” make innovations that change not only industries but also the way we live our lives. Why would architecture be any different?

PEOPLE CARRIER

If the home is our biggest personal expense, then travel is the second, and another area where AI will take over mundane tasks is on the roads. Driverless cars may seem more like a

utopian vision than the life-changing but fairly quotidian smart heating systems. In actual fact, driverless cars are much closer than you might think. Professor Wooldridge believes that real driverless car technologies will hit the roads in about five years’ time, suggesting there will be a lane on the motorway designated for driverless cars: “In 20 years? I think it will be common.”

Companies are already working on this. Bea Longworth, head of automotive PR at technology company NVIDIA, says the technology behind driverless cars will change more than just how we drive. She explains: “They’ll change our whole relationship with transportation. Car-sharing schemes already offer many people the benefits of having a car without the downsides – insurance, maintenance, road tax. Now imagine if you could combine that with the convenience of a taxi. Rather than owning a car, you could subscribe to one.”

NVIDIA is working with companies including Mercedes and Toyota to develop such a future, where “an on-demand autonomous vehicle automatically personalises itself to your preferences or even provides different kinds of environment depending on what you want to do during your journey”. Cars could act as an office, cinema or even a sleeping compartment.

“Driverless cars seem an inevitability to me,” Professor Wooldridge says. “The only question is when it will happen.” Whatever experts’ answers were to that question at the start of 2018, they are likely to have changed since March, when the first fatal crash involving a pedestrian and an autonomous

vehicle occurred in the US. Elaine Herzberg was tragically struck and killed by an autonomous Uber while crossing a road in Arizona. Neither the vehicle's AI-powered systems nor the human driver reacted in time, calling into question the sophistication of such systems.

PRIVATE LIVES?

As the crash reminded us, it can be easy to overlook the negative consequences of technology among the excitement of innovation. One of the most widespread drawbacks of AI for individuals is the potential lack of privacy. In 2017, iRobot, the maker of the robotic Roomba vacuum cleaner, faced controversy after it announced plans to sell 'mapped floor plans' of customers' homes, ostensibly to help the operation of other smart devices. Many, very reasonably, believed that this was an invasion of privacy. Others argue that giving up such data is not only par for the course when it comes to AI and smart tech, but a necessary part of achieving the utopian, tech-enabled homes that many of us want.

The real challenge, says Coplin, is for tech companies, advertisers and brands to earn the trust of the people that use their services. "Ultimately, the opportunity provided by AI is simply too great for us to ignore, and as a result we will need to drive a conversation with consumers, tech companies and governments to create an environment that is respectful of the privacy of individuals, but is able to ensure that data can be used in a way that can drive the innovation and opportunity to deliver a transformational change in how we might live, work and play." Wooldridge describes himself as a "pessimist" about privacy, adding that younger people are particularly at risk due to their propensity to "routinely share information about themselves that their elders would not dream of". The recent scandal around Facebook data being used to manipulate election results shows how seemingly innocuous data can be manipulated. Being careful and thoughtful about what data is shared is key here.

How technology will develop in the future is still, to some extent, a mystery. From AI and machine learning to more tangible pieces of consumer tech, the field is so rapidly expanding, research so vastly accelerating, that there is no real way of saying how our lives or homes might look in 10 or 20 years' time. What we do know, though, is that our lives are going to change – hopefully for the better. ■

THREE TAKEAWAYS

- 1 AI will embed everywhere decisions are made. Smart homes will use our data to learn our habits and take decisions on our behalf.
- 2 Car ownership will be a thing of the past when driverless cars hit the roads. Instead, we will subscribe to vehicles.
- 3 We will only reap the full benefits of AI if we are willing to give up our data, but consumers need to be careful about what they share.

OPPORTUNITIES: VALUE FOR MONEY

Make sure you invest in products that offer real benefits to users

The important thing when looking at AI investments is not to be overawed by the technology but to focus instead on the value it brings; there are many things that are possible with AI but not all of them are valuable. The development of the internet provides some good examples of this principle. Online shopping was unheard of a couple of decades ago, but is ubiquitous today because it is convenient. But Google's smart glasses that enabled users to see internet search results right in front of their eyes proved a flop.

So while the US National Retail Federation's 2018 Big Show boasted of robotic shopping carts that can lead shoppers on the most efficient route through the store based upon their list, such technology smacks of desperation.

The investable benefits of AI come from data, both the services that are able to harvest it and the platforms that are able to use it to serve customers better, such as Amazon. Behind the virtual counter they are able to exploit that data. Better customer analysis means better customer targeting and inventory selection, enabling retailers to suggest products to shoppers with increasing degrees of success. Meanwhile, automation of supply chains, including the use of autonomous vehicles, reduces lead times and waste and capital requirements, helping companies to be more profitable.

These sectors are developing fast, but we are watching them closely.





Science fiction has coloured our understanding of AI, but has it also influenced its development?

There have been many depictions of AI in film, television, novels and comic books, many wide of the mark, some remarkably accurate.

One example of the latter is the virtual assistants seen in *Star Trek*, which have recently become reality in products such as Alexa and Siri. In this case, both Amazon and Apple engineers name-checked Gene Roddenberry's series as a big influence.

Aside from inspiring specific technologies, many works of fiction have been prescient of questions that we are now asking for real. American writer Isaac Asimov's book *I, Robot*, published in 1950, asked what a world where humans and intelligent robots lived alongside one another would look like. It is perhaps most famous for its 'three laws of robotics', designed to protect humans from harm. Today, real-world discussions about the ethics of AI have referenced those laws.

One of the earliest examples of AI in cinema is the supercomputer HAL 9000 in Stanley Kubrick's 1968 film *2001: A Space Odyssey*, an adaptation of Arthur C Clarke's *Space Odyssey* novels. The sentient computer is the antagonist, operating the Discovery One spaceship for a team of astronauts and speaking with an eerily calm and conversational voice. HAL changes from a reliable member of the ship's crew to one bent on killing the astronauts after they decide to shut it down. Fears of AI turning against humans have been a consistent theme in film since then. James Cameron's 1984 convoluted thriller *The Terminator* shows Arnold Schwarzenegger as a murderous robot sent from the future to kill the mother of a man fighting to defeat the machines, which have taken over the world. Recently, fears over the existential threat posed by

AI have been echoed by leading thinkers, including the late Professor Stephen Hawking and technologist Elon Musk. While their concerns are rooted in rather more academic thought, they are made more plausible in the minds of the public because malign robots are a familiar trope.

Closer to home, in Charlie Brooker's dystopian television series *Black Mirror*, AI is used by a woman to overcome the grief of losing her boyfriend. A virtual version of him is created from his past online and social media activity. Today, tech giants such as Facebook have already created extremely accurate profiles of users based on their online activity. Brooker's stories are so compelling because the ideas within them sit on our periphery, close enough to become reality if someone develops existing technology. Could Brooker's human replicas become reality, or, given that his plots often serve as a warning, will he act as a canary in the mine for technology developers? The Chinese Government clearly has not been watching his series, as they have plans to use big data to introduce a 'social credit' score for every citizen in a move that echoes a particularly Orwellian episode of *Black Mirror* in which people's real-world prospects are heightened or diminished based on how they are rated online. The move in China will see people assigned a value based on school, employment, criminal and other records.

While many of these depictions are fanciful – psychotic robots from the future – science fiction has clearly helped shape the debate around AI. It is also continuing to influence innovators, with Mark Zuckerberg determined to recreate the Jarvis computer assistant from the *Iron Man* films in his own home – voiced by none other than Morgan Freeman. ■

ARTIFICIAL INTELLIGENCE

FURTHER READING

NICK BOSTROM SUPERINTELLIGENCE

This is the book that inspired Elon Musk to fear AI. Oxford University philosophy professor Nick Bostrom asks how humans can stay on top of the food chain. While full of big ideas, this influential work is not for the casual reader as it has an academic tone.

RYAN AVENT THE WEALTH OF HUMANS

Economist Ryan Avent argues that the new digital phase we are facing is similar in profile to the Industrial Revolution. Just like in the past, society must go through wrenching political change before we can adjust to the new economic realities.

OLIVER RATZESBERGER AND MOHANBIR SAWHNEY THE SENTIENT ENTERPRISE

A methodology for using data and analytics to transform business practices. Everyone knows that data is the key to success in an AI-enabled world, but not everyone knows how to use it for competitive advantage.

ANDREW MCAFEE AND ERIK BRYNJOLFSSON MACHINE, PLATFORM, CROWD

Following their acclaimed book, *The Second Machine Age*, which explored how digital technologies will affect our lives, these two MIT academics explain how businesses can navigate the oncoming disruption. Get ready to tear up your rule book as this highly accessible read makes a convincing argument for fresh thinking.



ANTENNA 01
ARTIFICIAL INTELLIGENCE



antenna.brewin.co.uk