THE ROBOTS ARE COMING!

Are you ready?

"The first ultra-intelligent machine is the last invention that man need ever make."

I. J. GOOD, BRITISH MATHEMATICIAN & CRYPTOLOGIST

FOREWORD

We are on the brink of a dramatic transformation that will have a major effect on the global economy, on business models and on our working lives. Already, robots and increasingly intelligent and autonomous software systems are taking on ever more sophisticated jobs. But where is this heading?

We are already used to automated production lines and self-service machines in place of assembly line workers and check-out staff, but are we sufficiently prepared for robots and intelligent machines taking on the knowledge based tasks that we thought only humans could do?

There are real benefits to be gained from working in collaboration with sophisticated machines and artificially intelligent systems, but these seismic developments also present real challenges for business, policy makers and wider society:

- What changes will this mean for the global economy and levels of employment?
- What value is placed on human behaviours, skills and understanding?
- Where is the development of Artificial Intelligence (AI) all heading?

In exploring how these technology developments have the potential to reshape the world of work, we also consider how organisations can begin to prepare for a future that looks very different from today.

Rob Gear
Futurist, PA Consulting Group

THE SECOND MACHINE AGE IS ALREADY HERE

"20 years from now, labour demand for lots of skill sets will be substantially lower. I don't think people have that in their mental model."

BILL GATES. AMERICAN ENTERPRISE INSTITUTE, OCTOBER 2014.

The idea that technology continually creates and destroys specific types of work is not new. The world has already witnessed a number of technology-led transitions, such as the transformation from an agrarian to an industrial economy initiated by the Industrial Revolution that began in Great Britain in the 18th century. Jobs in agriculture dramatically reduced and people migrated to the cities to work in manufacturing as the emphasis of work changed from independent production to mass production.

As we move into the second machine age, characterised by hyper-connectivity, we stand on the edge of another shift in the nature and economics of work. As Bill Gates remarked during a speech in October 2014[1], "technology over time will reduce demand for jobs, particularly at the lower end of skill set… 20 years from now, labour demand for lots of skill sets will be substantially lower. I don't think people have that in their mental model."

The scale of this transformation is significant, and growing, as progress in artificial and machine intelligence accelerates. There are multiple approaches and dimensions of the technology, some of which already dramatically exceed human capabilities.

We are already seeing the effect on administrative work at the lower end of the skill set. Forbes reported in February 2011 that some 300,000 administrative jobs became surplus to requirements between 2004 and 2009 as technology such as voicemail replaced the need for clerical and secretarial work in the US. However, it is not just clerical work that is being automated.

1. http://www.aei.org/events/from-poverty-to-prosperity-a-conversation-with-bill-gates.

As levels of machine intelligence advance, work that has traditionally been regarded as the domain of the professional knowledge worker is under threat.

Some legal firms are now using software tools to scan and read mountains of legal documents which previously required the billable time of a solicitor.[2] In the world of journalism, algorithms that produce well-styled and grammatically correct reports are appearing. These bots can be far more productive than any reporter and can write articles in seconds, even about events that no journalist attended.[3] Elsewhere, medical diagnoses are being made by expert systems either in support, or independently, of professional doctors. For example, IBM's Watson system, famously lauded for beating the most skilled human players on the game show 'jeopardy', is now helping oncologists with cancer diagnoses.

In all these examples, the march of technology represents a double-edged sword. On the one hand, there is the potential for improvements in productivity and reliability. On the other hand, services that previously seemed entirely dependent on human interaction now look increasingly fragile. As MIT's Eric Brynjolfsson and Andrew McAfee put it, "the AI revolution is doing to white-collar jobs what robotics did to blue-collar jobs."[4]

In this context, the economic impact of such changes becomes an important question: will all this automation make society poorer?

1. As reported in Newsweek, 25 July 2011.
2. MIT Technology Review, October 2012.
3. Erik Brynjolfsson and Andrew McAfee, interviewed in Christian Science Monitor, 17 September 2012.

IS THE FUTURE ONE OF MASS UNEMPLOYMENT?

In previous waves of automation, the resulting increases in productivity have, in the end, generated wealth. New jobs have been created to replace those that have been lost. But who's to say that this will be the case this time around?

In their 2011 publication, *Race Against the Machine*[5], Brynjolfsson and McAfee warned that although automation leads to increased productivity, the fact that it cannibalises human activity results in wealth being removed from the economy. Unless new jobs are created in sufficient volumes to absorb the increase in unemployment, the gap between rich and poor gets exacerbated.

In the past, new technologies replaced manual work in certain sectors and environments, such as manufacturing production lines. This new wave of technology advancement is different because it has the potential to automate the work of skilled knowledge workers as well; financial traders replaced by algorithms and medical staff by digital diagnostic tools.

Moreover, the emergence of a model of work based entirely on interactions between machines – the 'autonomous economy'[6] – risks rendering swathes of jobs obsolete as digital processes communicate with other digital processes to create instructions and take decisions.

5. Race Against the Machine, Brynjolfsson and McAfee, October 2011, Digital Frontier Press.

6. As identified by W. Brian Arthur, visiting researcher at the Xerox Palo Alto Research Center's intelligence systems lab and a former economics professor at Stanford University.

Will automation make us poorer?

The potential scale of job losses would have significant implications for the economy and wider society. Brynjolfsson and McAfee highlight that the decade from 2000 to 2010 was the first since the Great Depression to end with a net loss in jobs in the US. This is despite the fact that per capita gross domestic product (GDP) is one-third higher than it was 20 years ago, and the country produces 75 per cent more goods than it did at that time.

Looking ahead, Oxford University's Carl Frey and Michael Osbourne suggest that 54% of jobs in the European Union are at risk from automation. Figure 1 gives an overview of the kinds of work forecast to be at risk, with administrative support, sales, service, transport and manufacturing industries predicted to see the highest levels of unemployment[7].

According to Okun's Law[8] – an accepted rule of thumb describing the relationship between productivity and employment – when the economy grows rapidly, unemployment should decline, often by 1% for every 3% increase in GDP. According to this formula, the US should have almost full employment by now. Instead, even before the financial crisis, few additional jobs had been created, although productivity had surged at the fastest pace since World War II.

Race Against the Machine outlines a challenge of achieving a balance between 'bounty' and 'spread'. Bounty is the increase in volume, variety, and quality of the many offerings brought by modern digital technologies, all enabled at a decreased cost. Spread is the negative aspect of this progress, leading to increasing wealth inequality, progressive unemployment, and reduction in social mobility. Without policy intervention, this spread looks likely to accelerate in the second machine age.

Figure 1: The impact of automation on employment by industry sector

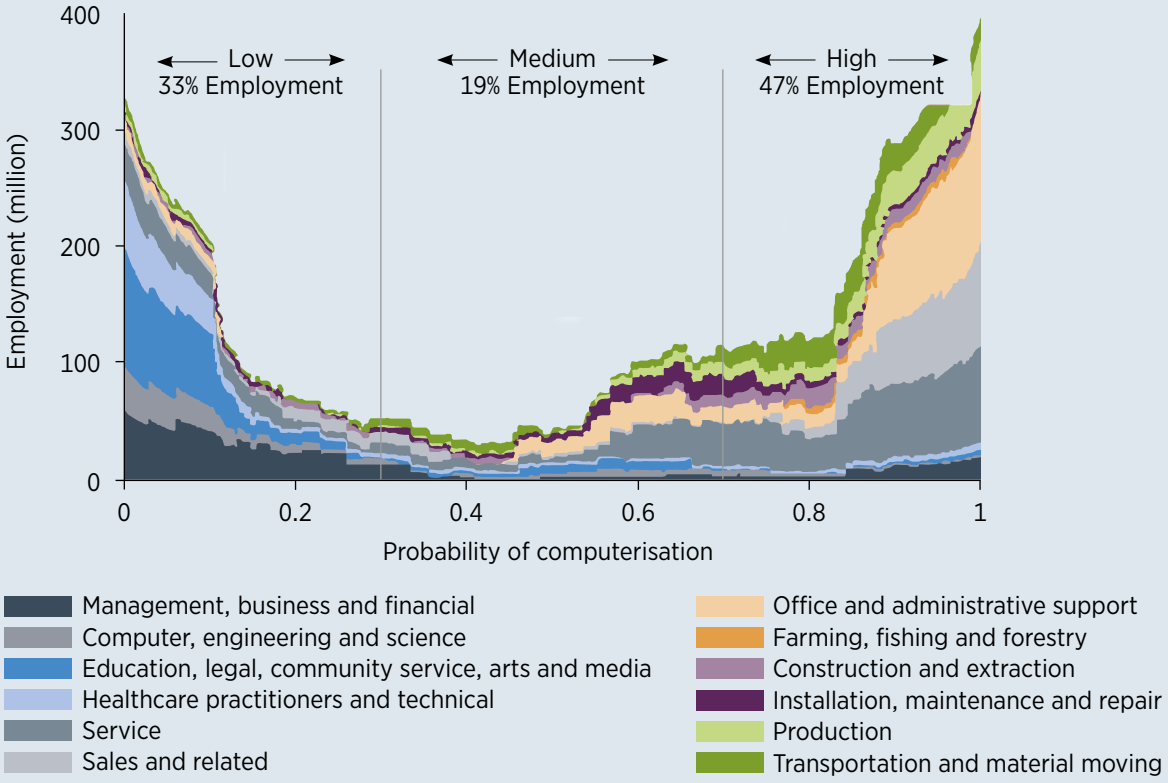

Source: ***The Future of Employment***. Carl Frey and Michael Osborne, Oxford University.

7. *The Future of Employment*, Carl Frey and Michael Osbourne, Oxford University, 2013.

8. Okun noted that, because of ongoing increases in the size of the labour force and in the level of productivity, real GDP growth close to the rate of growth of its potential is normally required, just to hold the unemployment rate steady. To reduce the unemployment rate, therefore, the economy must grow at a pace above its potential. More specifically, according to currently accepted versions of Okun's law, to achieve a 1 percentage point decline in the unemployment rate in the course of a year, real GDP must grow approximately 2 percentage points faster than the rate of growth of potential GDP over that period. So, for illustration, if the potential rate of GDP growth is 2%, Okun's law says that GDP must grow at about a 4% rate for one year to achieve a 1 percentage point reduction in the rate of unemployment.

A HUMAN ADVANTAGE?

Despite the increasing sophistication of technology, some argue that there are still many human qualities that machines cannot easily replicate. In 2012, the business author Don Peppers highlighted, via LinkedIn, two ways to 'beat the clock' of automation: "One way is to become very good at dealing with interpersonal issues – people skills. The other way is not to focus on solving problems but on discovering them."[9] In other words, it is perhaps better for humans not to try to compete where the sheer computing power of a machine will be faster or more reliable than the human brain. Rather, we should look at where the mind surpasses the machine. The creativity associated with discovering problems is one such example, but there are others.

In this context, it is helpful to explore the different dimensions and capabilities of AI and consider how these compare to human capabilities. To explain this further, we have looked at the three categories of sensing, thinking and interacting as defined in Figure 2 on the right-hand side.

Across these different dimensions, we can already see significant progress in machine intelligence. In terms of sensing, image recognition has achieved, and in some ways exceeds, human parity. Voice recognition has entered the mainstream via services such as Siri, Google Now and Cortana, and is improving all the time. With an explosion in sensing capabilities offered by the Internet of Things we will see a rapid expansion in the abilities of machines to be able sense what is going on in the physical environment.

In terms of thinking, machine memory so radically exceeds that of human beings in terms of raw capacity and longevity. It could also be argued that it is a radically different form of memory that is not bound like human memory to affective qualities such as nostalgia and indifferent to personal history and the passage of time.

9. Don Peppers commenting on LinkedIn, 24 May 2013.

Reasoning, planning, the creation of new ideas and learning may currently be regarded as operating at subhuman levels, but even here we see signals that these capabilities will also improve in the future, albeit expressed quite differently from their human analogs. Evolutionary algorithms, for example, can create quite radical new designs through a process of rapid iteration, based on a set of Darwinian evolutionary principles. The sheer number of iterations involved can make it very difficult for human engineers to understand exactly how a particular design has emerged through a very unfamiliar process of machine creativity.

In interacting, it is obvious that machines are presently far better at communicating with each other than with humans. In June 2014, there was a flurry of interest in 'Eugene Goostman', an intelligent chatbot, that was claimed to have passed the famous Turing Test – that "a computer would deserve to be called intelligent if it could deceive a human into believing that it was human".[10]

Figure 2: AI high level capabilities

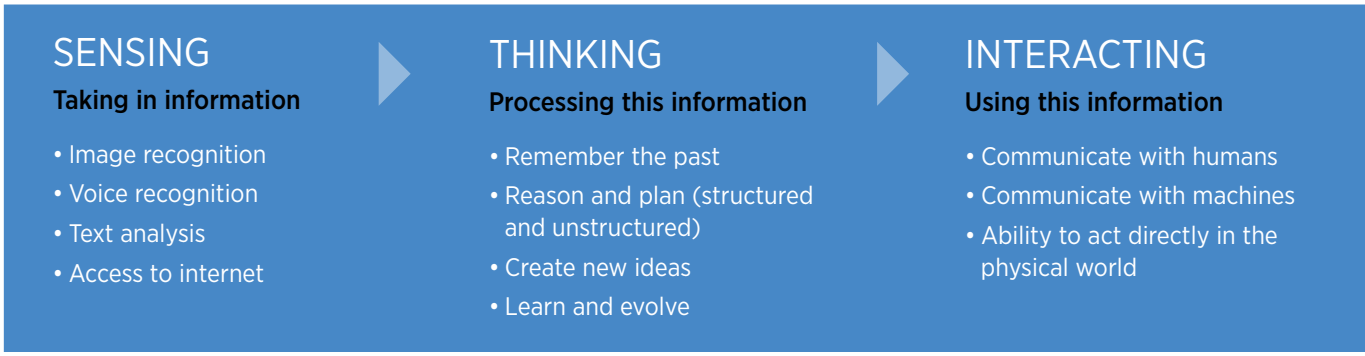

In the end, the consensus was that Eugene did not really manage it. He convinced just 10 of 30 judges from the Royal Society that he was human whilst pretending to be a 13-year-old Ukrainian boy for whom English was a second language. This is probably not what Turing had in mind when he conceived the test.

Nevertheless, the ability of machines to interact with humans is improving all the time and most experts believe that as machine learning algorithms improve, and process more and more of the wealth of conversational data that exists online, it is only a matter of time before the Turing Test is conclusively passed.

The ability of machines to act is also increasing. When machine intelligence combines with robotics it is given form and the ability to act in the physical world. On a more subtle level, machines are already acting and taking many routine decisions that impact on our daily lives, such as deciding on whether we are eligible for a loan or adjusting the room temperature when our house is deemed to be too warm or too cold. Machine agency will become far more controversial in the future and will require very careful ethical consideration and governance to ensure that machines are acting fairly. There will be a need for greater transparency in how business and governments are using AI to make decisions and act to ensure that flawed assumptions, prejudices, and other kinds of bias are not 'baked into' decision making algorithms.

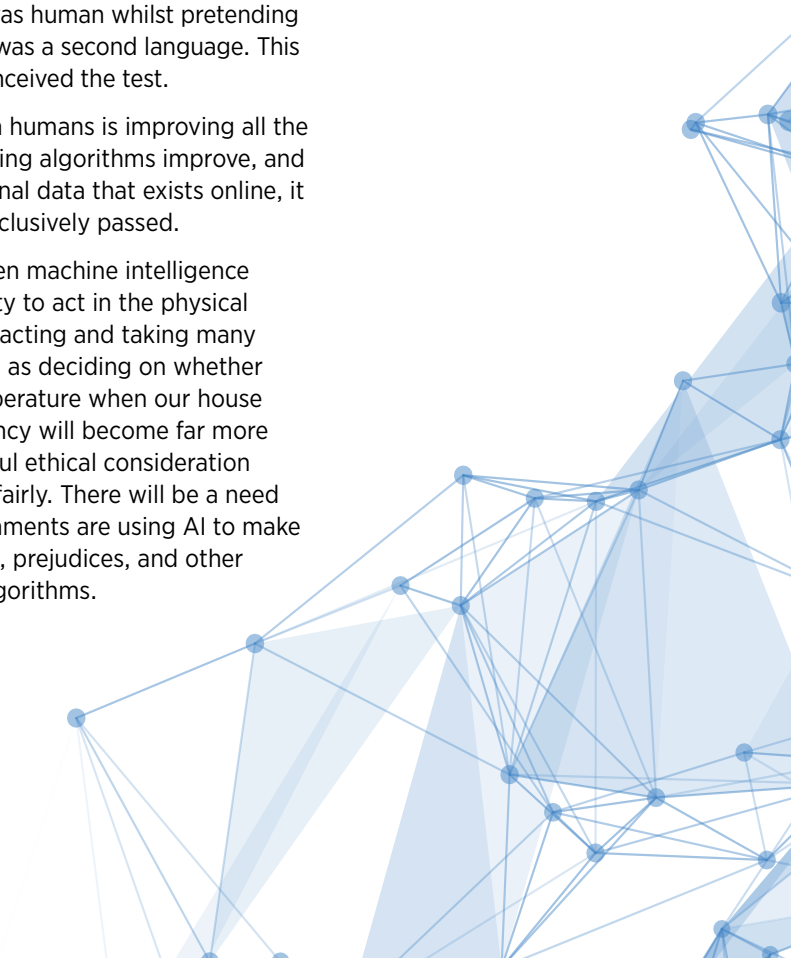

10. http://www.bbc.co.uk/news/technology-27762088.

Emphasising the human difference

Essentially, the kinds of skills that machines cannot acquire are those which are not easy to standardise or codify into an algorithm. Alongside creativity, these include leadership, motivation, intuition, empathy, abstract reasoning and lateral thinking; they also include the skills associated with craft.

The recent resurgence in the popularity of craft and hand made products has been linked both to the global economic crisis and the inherent positive and therapeutic benefits of cultivating and using physical skills. However, it could also be developed as a response to automation. It is possible to envisage a scenario where there is an increase in small shops, project-based work on the construction industry model and greater participation by households in more informal economies, characterised by production sharing, bartering and gifting.

Figure 3: An overview of current AI capabilities against those of humans

Source: PA Analysis

Alternatives to automation

Whilst the evidence is perhaps anecdotal at the moment, there has been a noticeable increase in crafted products over the past few years. From artisan bread, chocolate and beer to handcrafted bicycles, bags and belts, crafts have been celebrated in books, documentaries and design fairs. Artists like Joana Vasconcelos have adopted craft, such as crochet, for their chosen media and even graffiti has taken a crafted turn with new techniques in stencilling and knit-bombing.

With significant developments in progressive crafts such as digicraft, industrial designers are also turning to craft values instead of the traditional industrial ones. In the UK, a 2014 study by The Crafts Council[11] reports that 11,620 businesses involved in craft industries generate nearly £1.9 billion in turnover, contributing £746 million to the UK economy in Gross Value Added (GVA). Whilst somewhat insignificant at the moment, it will be interesting to observe how the fortunes of craft industries play out in the coming years. Might we see more people return to craft skills as a means of resisting increasing automation?

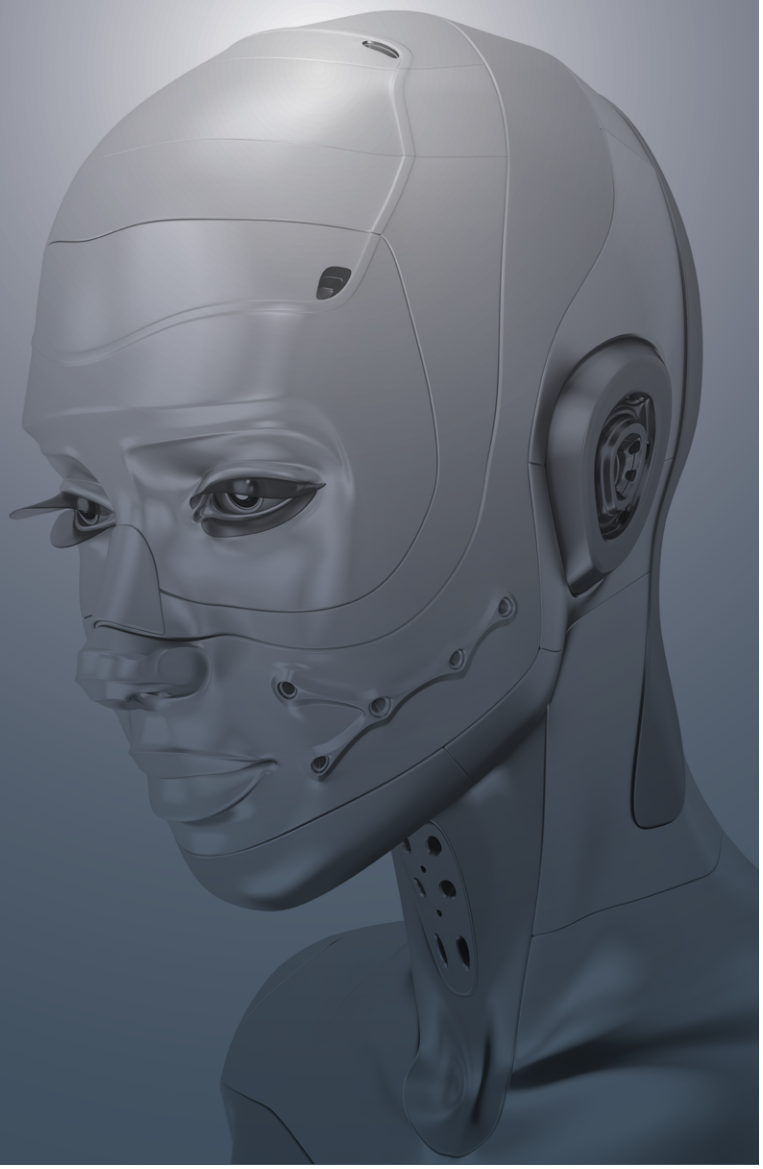

11. Measuring the Craft Economy, The Crafts Council, October 2014.

WHAT NEXT FOR AI?

As AI evolves in the future, meeting and subsequently exceeding the capabilities of humans, more significant questions will emerge around the risks and rewards associated with the technology. AI promises an unmatched ability to understand and predict behaviours that could offer huge benefits in everything from medical research and more efficient use of energy and resources to better fraud detection and crime prevention.

At the same time, machines with superhuman intelligence are unpredictable and could be harnessed for criminal intent, or simply behave in ways that create unexpected social, economic, or environmental consequences. In 2014 PayPal founder Elon Musk even cautioned that AI represents "our biggest existential threat".[12]

At a very high level we could consider the uncertainties around the impact of AI as being dependant on the motivations behind their development, as outlined in Figure 4.

Figure 4: The impact of AI development is uncertain

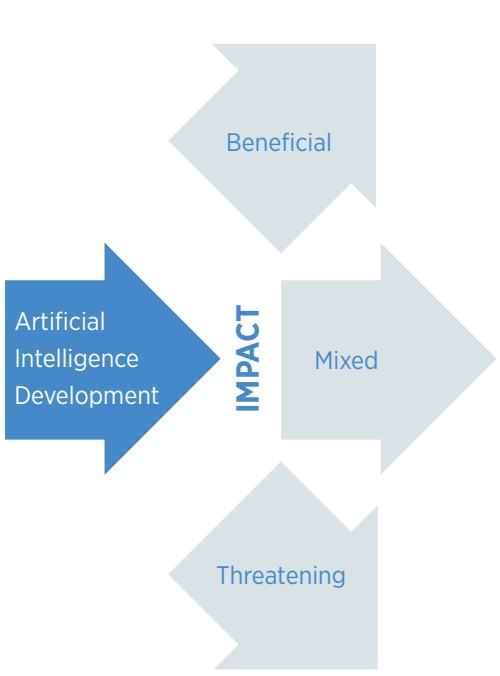

Intelligent technologies
Unmatched ability to understand and predict behaviours

"Anything that could give rise to smarter-than-human intelligence – in the form of AI ... wins hands down beyond contest as doing the most to change the world. Nothing else is even in the same league" ELIEZER YUDKOWSKY

- Better fraud detection
- Automated CCTV analysis
- Efficient home and commercial use
- Medical research

Increased productivity
Delivering growth but creating unemployment and instability

"Software substitution, whether it's for drivers or waiters or nurses ... it's progressing ... technology over time will reduce demand for jobs, particularly at the lower end of skill set ... 20 years from now, labour demand for lots of skill sets will be substantially lower. I don't think people have that in their mental model." BILL GATES

- Automated vehicles and supply chains
- Financial processes
- AI replacement for services that can be commoditised

Unpredictable machines
Super-human intelligence with unpredictable behaviours

"In the field of AI there is no more iconic and controversial milestone than the Turing Test ... having a computer that can trick a human into thinking that someone, or even something, is a person we trust is a wake-up call to cybercrime." KEVIN WARWICK

- Technologies used for the benefit of cyber crime
- Social engineering
- AI weapon technology

12. http://www.theguardian.com/technology/2014/oct/27/elon-musk-artificial-intelligence-ai-biggest-existential-threat.

It is important however to recognise that no outcome (positive or negative) is predetermined. Technology is a tool and we can shape its use through our control of our economic, social and political systems. The choices we make in reshaping these systems will ultimately determine the impact of automation on the world of work rather than technology itself.

There is a great deal of speculation about when we could see the creation of an Artificial General Intelligence (AGI) – an ultra-intelligent machine that some hypothesise might be 'our last invention' due to its abilities to further innovate and solve problems far more effectively than any human.

As I. J. Good, the British mathematician and cryptography contemporary of Alan Turing, put it:

"Let an ultra-intelligent machine be defined as a machine that can far surpass all the intellectual activities of any man however clever. Since the design of machines is one of these intellectual activities, an ultra-intelligent machine could design even better machines; there would then unquestionably be an 'intelligence explosion,' and the intelligence of man would be left far behind. Thus the first ultra-intelligent machine is the last invention that man need ever make."

There is a great deal of uncertainty as to when such an AGI might be created. Commentators such as Ray Kurzweil suggest that it might happen around 2045, triggering a Technological Singularity where the AGI's rapid and recursive self-improvements lead to an explosion of machine intelligence, the impacts of which will be hard to imagine with today's mindset. Others propose more conservative timelines.[13] Despite this uncertainty, it seems likely that any future machine intelligence will be non-humanlike and behave in ways that it will be hard to predict.

> "I think we should be very careful about AI. If I had to guess at what our biggest existential threat is, it's probably that."
>
> ELON MUSK, PAYPAL FOUNDER, OCTOBER 2014.

13. Based on expert consensus presented by Murray Shanahan, Professor of Cognitive Robotics at Imperial College London at public presentation at CSER (Cambridge), 20 February 2015.

EVOLUTIONARY FORCES

The evolutionary path of AI may be determined by a variety of forces, as illustrated in Figure 5.

Deliberate effort can be considered all of the formal programmes and initiatives focused on creating machine intelligence across the public and private sectors, and academia. There have been many notable investments in AI made in recent years by big Silicon Valley players such as Google, Facebook, IBM and Microsoft and we also need to consider the impact of national investment programmes.

Emergence suggests that intelligence might be created in a more chaotic and unplanned manner, arising as a consequence of the complexity inherent in our hyper-connected world. Perhaps some kind of distributed intelligence that emerges through the ad hoc combination of multiple narrow AI's and the embedded intelligence that will increase as more objects become connected in the Internet of Things, forecast by Cisco to reach 50 billion by 2020.

The learning and evolution route systems evolving greater intelligence through the creation of ever more sophisticated learning algorithms and access to an ever growing amount of data. For example, the Pentagon's research agency, Darpa, is leading efforts to build machines that can teach themselves using 'probabilistic programming'.

Whilst the future path of AI is far from certain, it is clear that we need to start the debate now to help us make the right choices and adapt to a very different world.

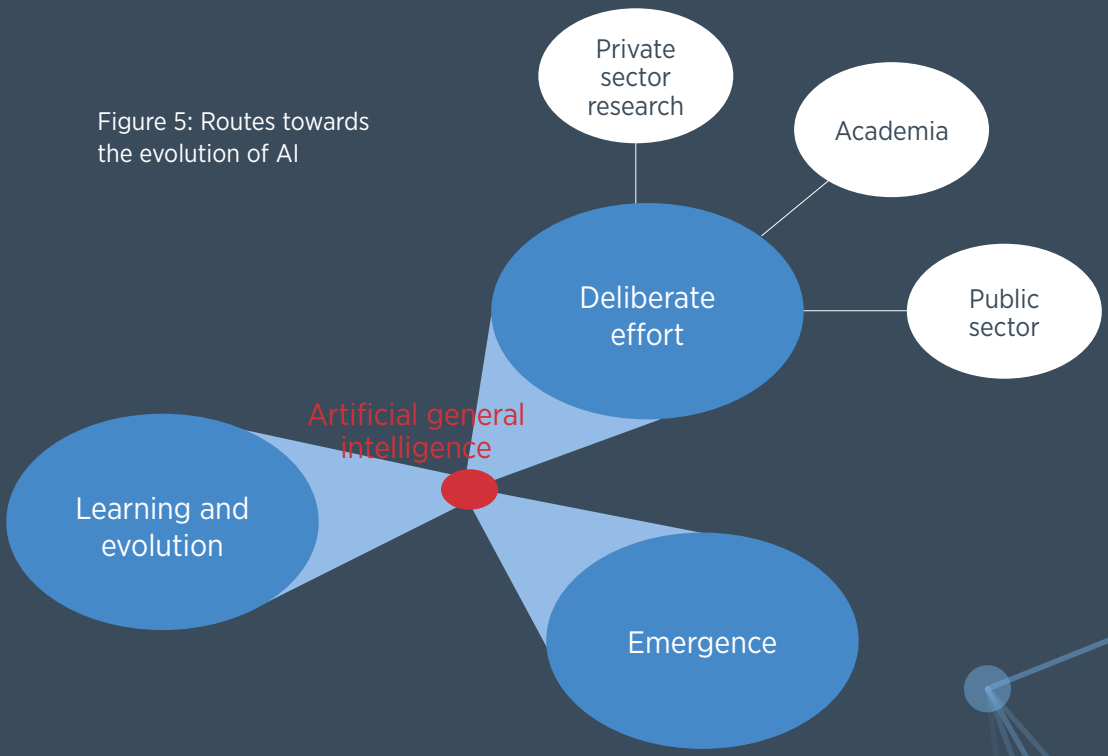

Figure 5: Routes towards the evolution of AI

TOWARDS A NEW KIND OF COLLABORATION

While the challenges presented by automation and AI can be worrying, it is important not to lose sight of the advantages they can bring. New opportunities for collaboration with machines could add significant value because there are many things a machine can do more efficiently than a person.

There are examples of how these kinds of collaborations can work. Baxter, a relatively inexpensive robot launched in 2012 by Rethink Robotics, works alongside people on production lines. Unlike many large-scale manufacturing robots, Baxter is capable of working safely with humans, and can be easily taught to adapt to its tasks and environment without complex programming.

Collaboration can also occur through intelligent software providing contextual support for human decision making, such as assisting with medical diagnosis. The nascent field of augmented cognition is seeking to develop technology that monitors an individuals' cognitive capacity and adapts information flows accordingly, combating problems such as information overload and so improving performance.

Ultimately we could achieve a greater degree of human-computer symbiosis. This would mean people and machines cooperate in making decisions and controlling complex situations without inflexible dependence on predetermined programs. In this kind of partnership, people will set the goals, formulate the hypotheses, determine the criteria, and perform the evaluations. Computers will perform the routine work to prepare the way for human insights and decisions in technical and scientific thinking.

This is not science fiction but is already happening today. A Carnegie Mellon University robot and an engineer partner built a space frame in 10 hours for $1,150. A group of human experts built the same frame in 10 hours at a cost $7,075. The robot's computer automatically ordered supplies, planned the sequence to perform the welding operations, and specified the optimal way to hold parts so they were secure. It told the engineer how to set up the complex construction using a projector displaying images and text,[14] and by working together they proved to be the most efficient team.

14. Scientific American, May 2015.

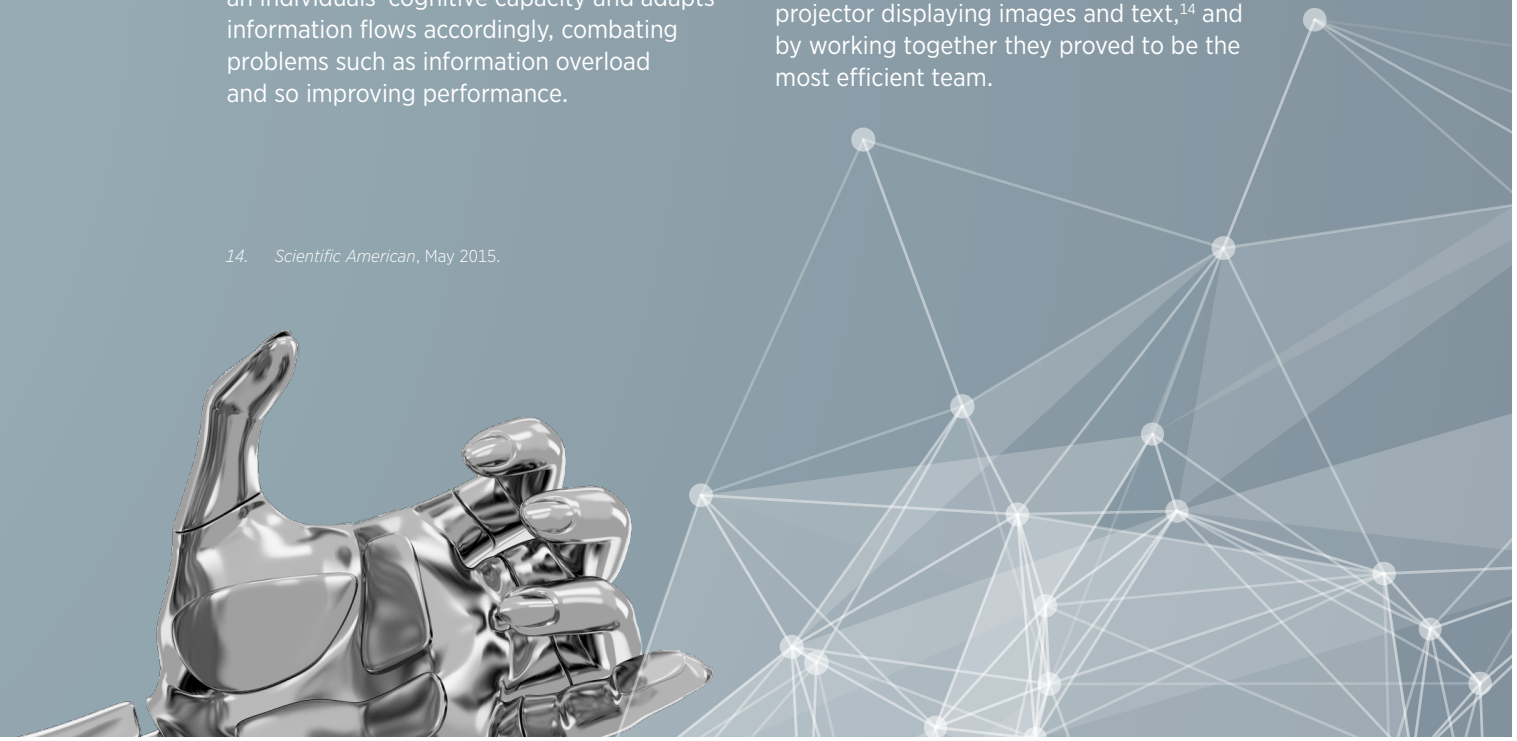

PREPARING FOR A DIFFERENT WORLD

It is important that business leaders are aware of the opportunities automation and AI technologies offer. Indeed, there are some pragmatic steps that any organisation can take now to help to better understand and prepare for the future possibilities of automation and AI:

1. Establishing a horizon scanning capability will ensure that their organisation keeps abreast of the latest developments and is well-placed to capitalise on early opportunities that it may offer.
2. Create scenarios to explore alternative possible futures and develop actions based on their implications. Seek to understand how your market and business model will look in a more automated world.
3. Establish a team to explore the impact of automation on the business. Focus on staff development and training and how to derive the maximum value from the human qualities, skills, and discretion of employees, and where necessary how to transition valued staff into new roles.

At PA Consulting Group, we help our clients look to the future and design an organisation or business model that is fit for the future rather than built to succeed in the past.

To speak further to a member of the PA team about establishing a horizon scanning workshop, scenario planning or exploring the impacts of automation on your organisation, please contact rob.gear@paconsulting.com.

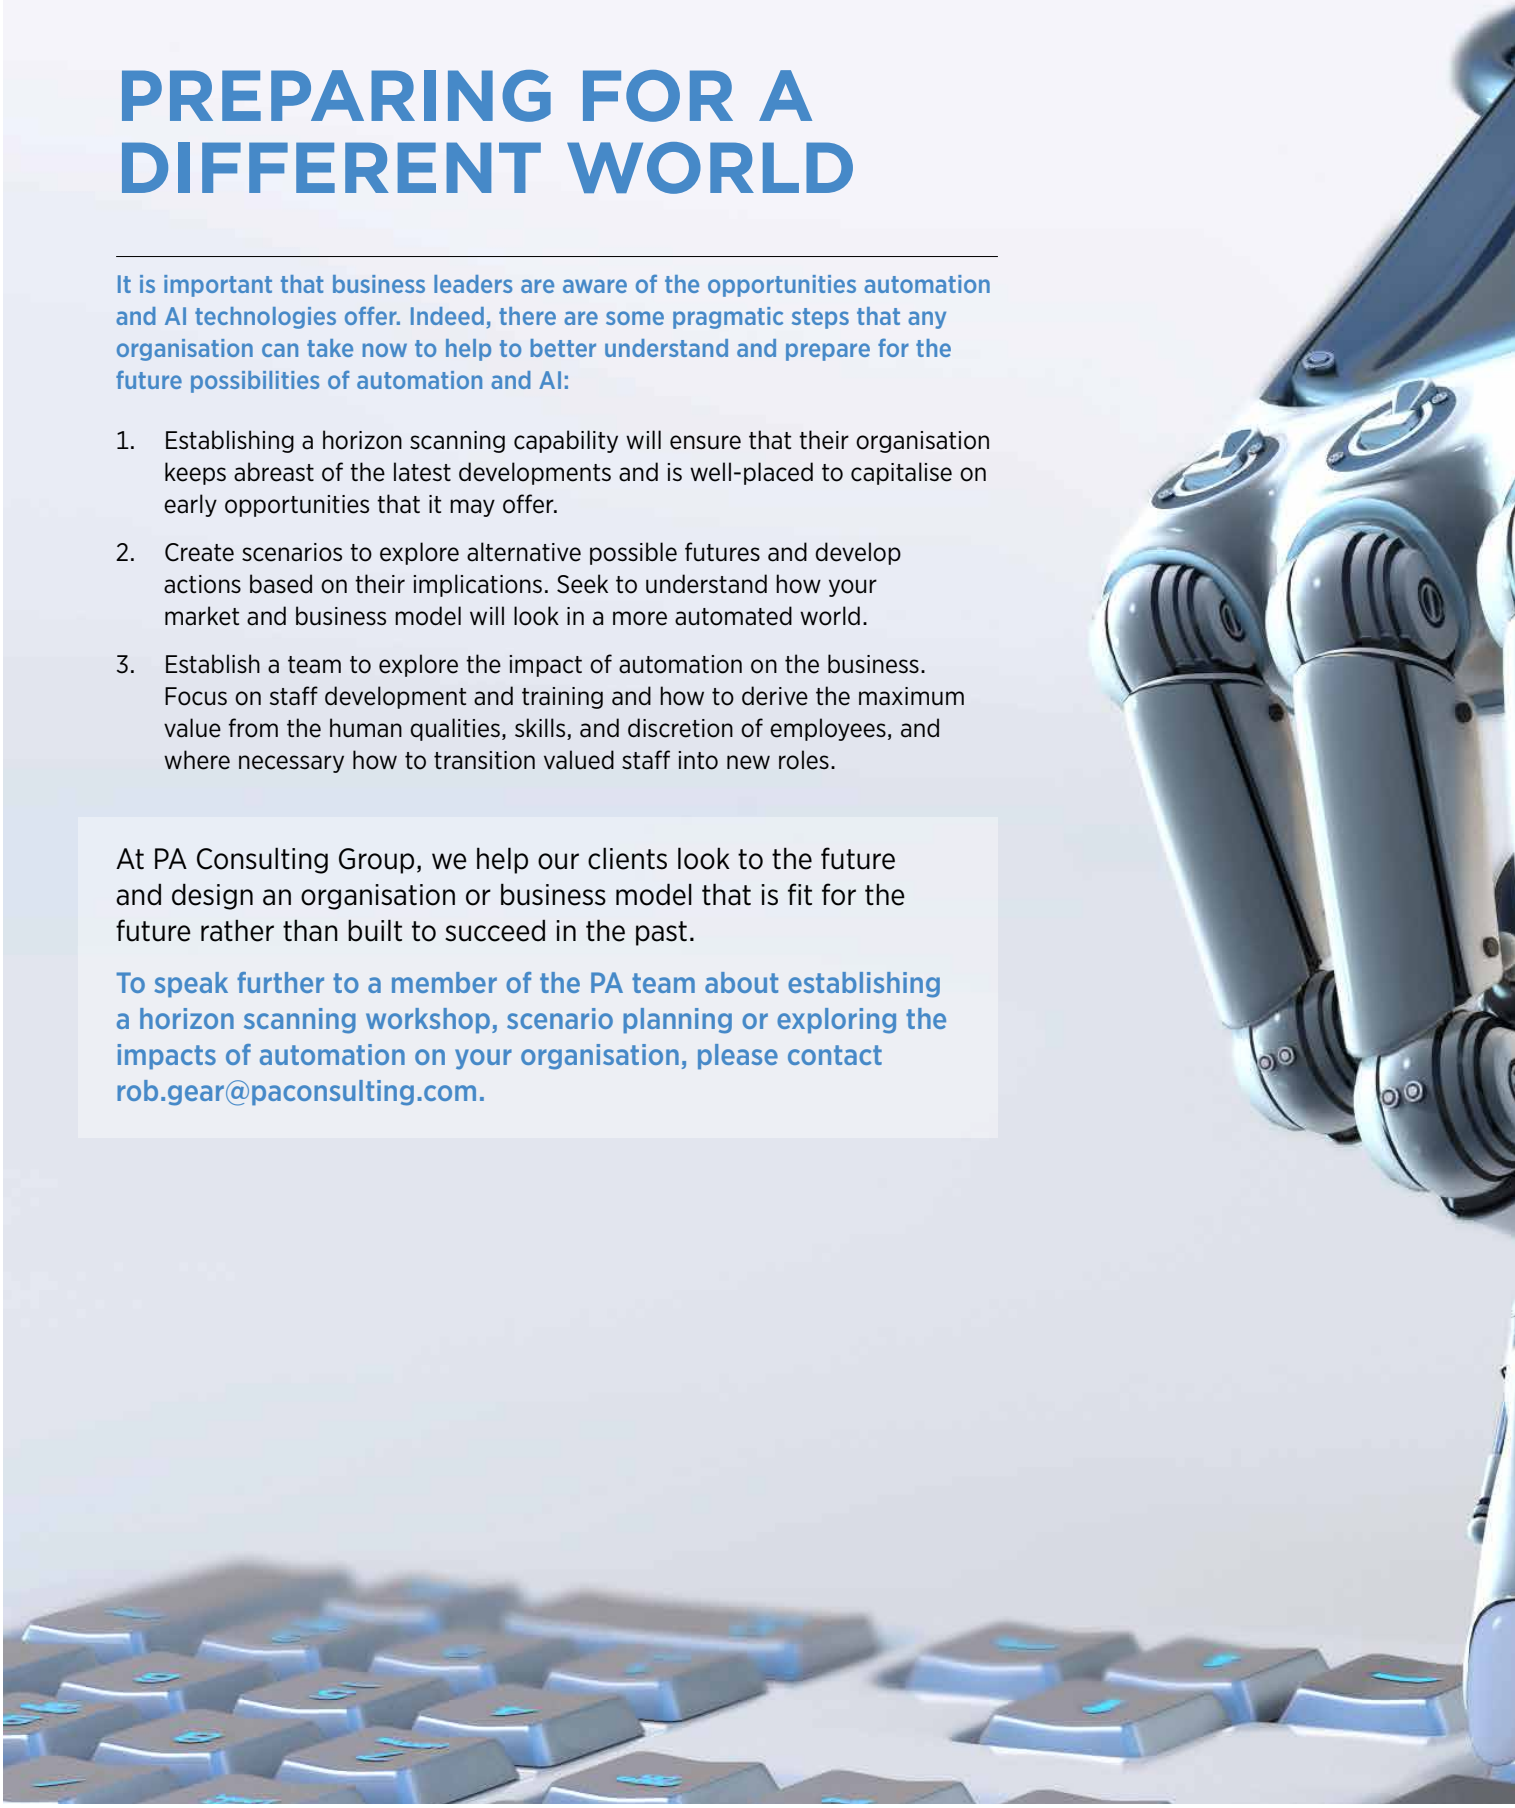

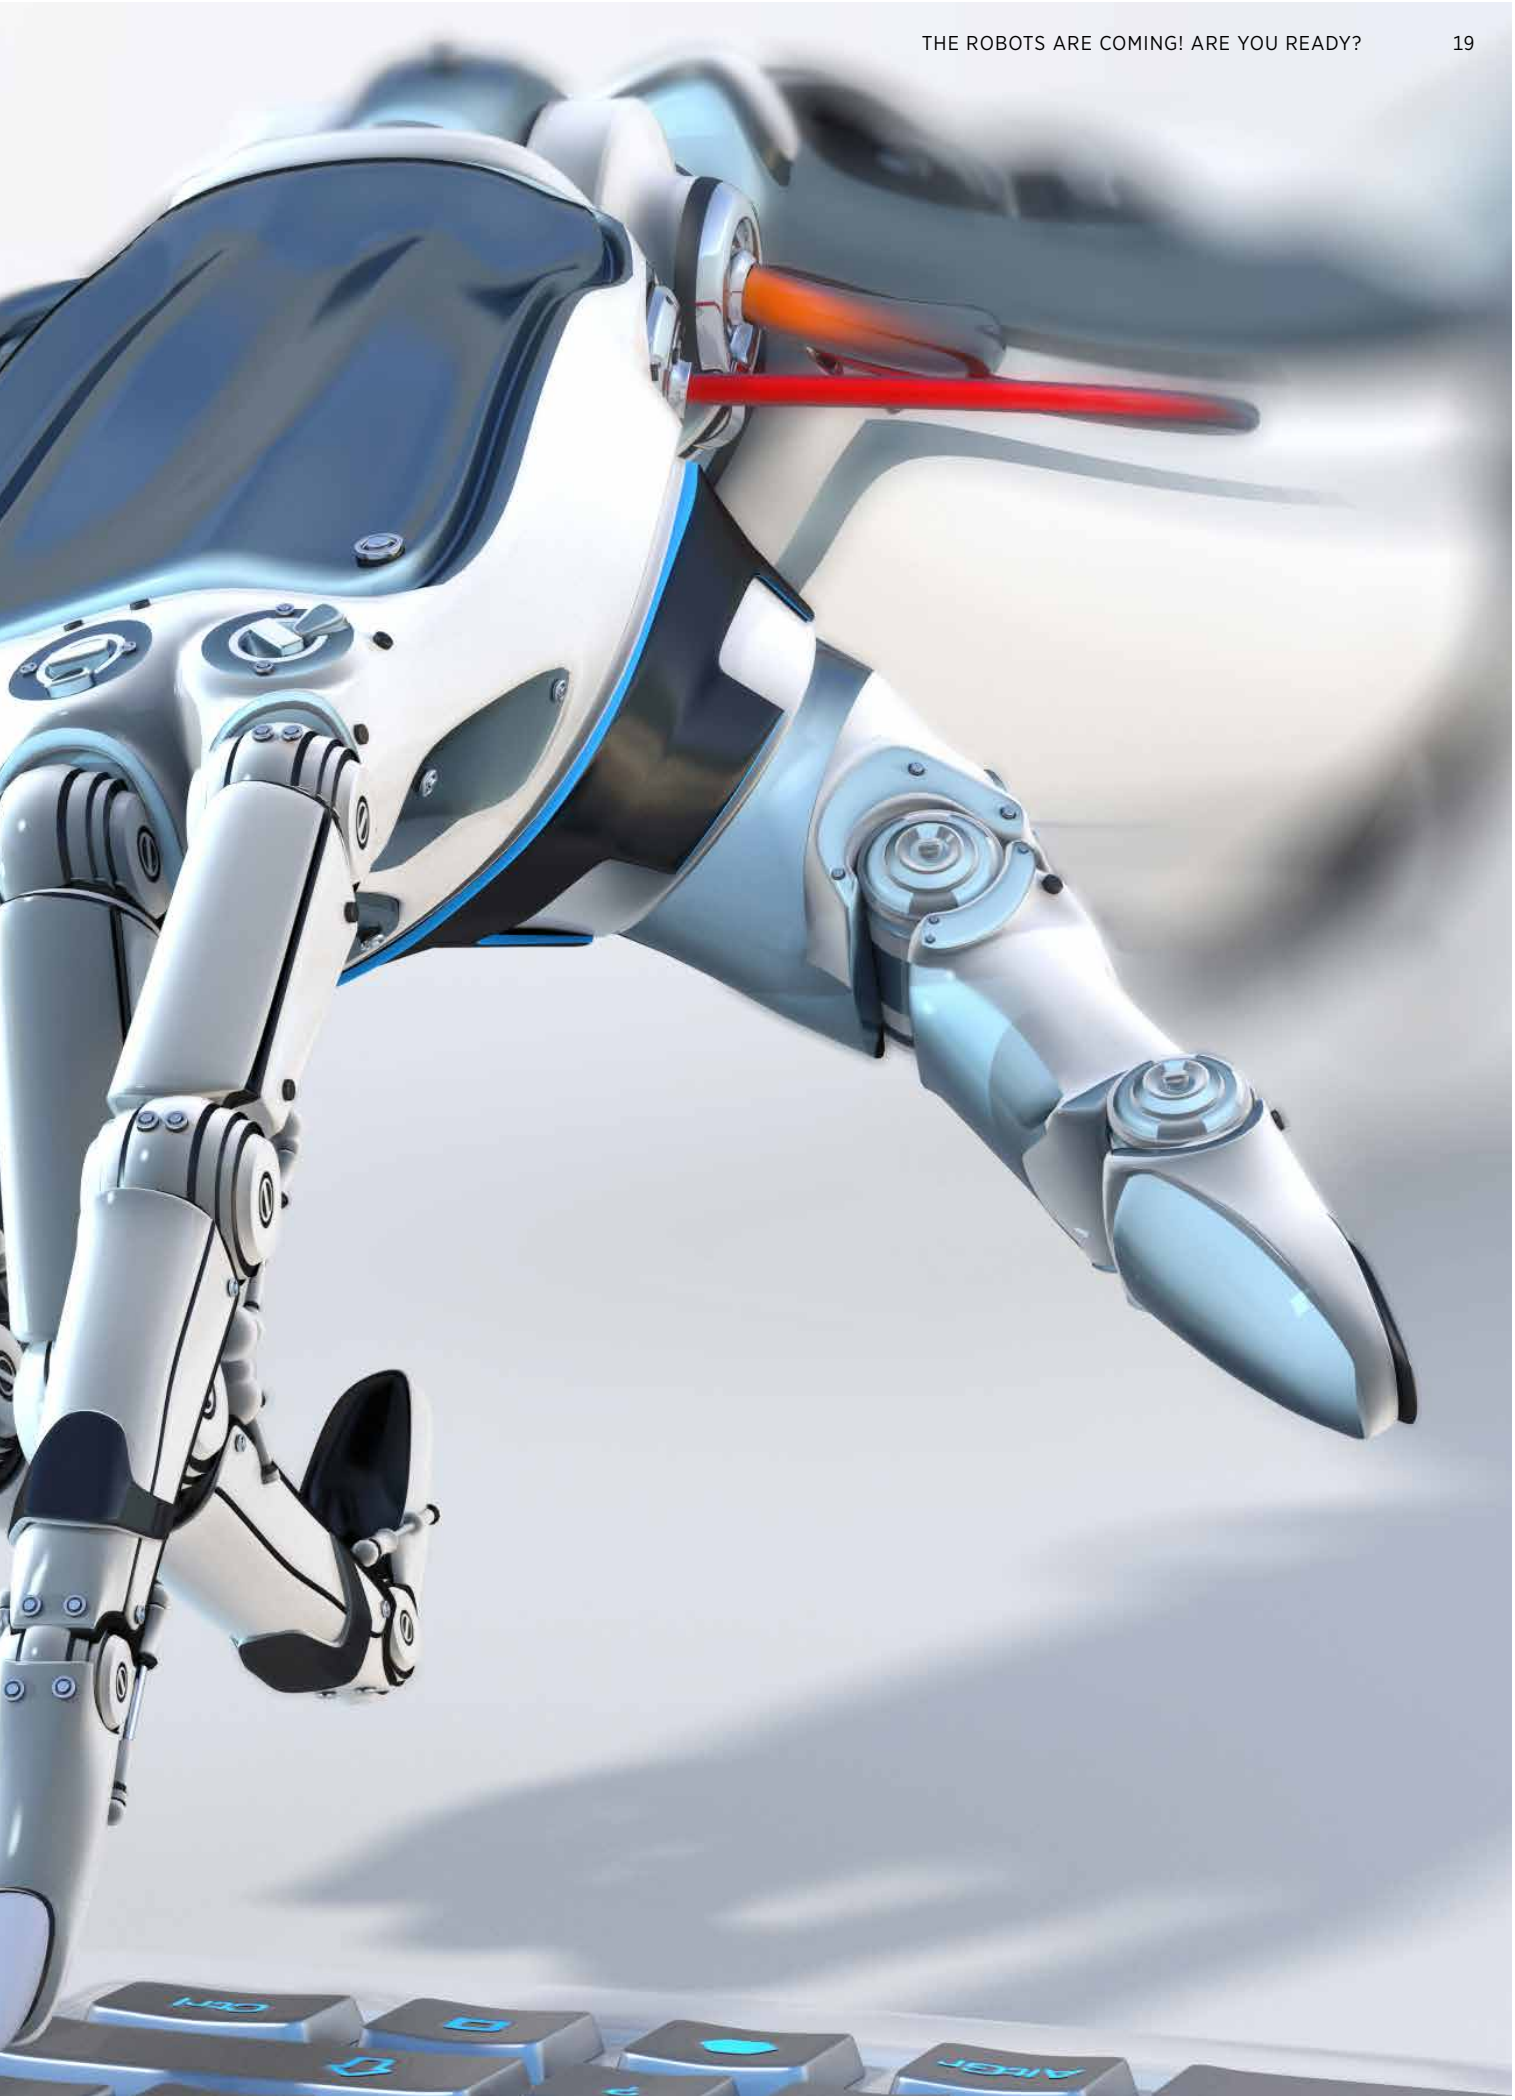

We are an employee-owned firm of over 2,500 people, operating globally from offices across North America, Europe, the Nordics, the Gulf and Asia Pacific.

We are experts in energy, financial services, life sciences and healthcare, manufacturing, government and public services, defence and security, telecommunications, transport and logistics.

Our deep industry knowledge together with skills in management consulting, technology and innovation allows us to challenge conventional thinking and deliver exceptional results with lasting impact.

1919-54

Corporate headquarters

123 Buckingham Palace Road
London SW1W 9SR
United Kingdom
+44 20 7730 9000

paconsulting.com