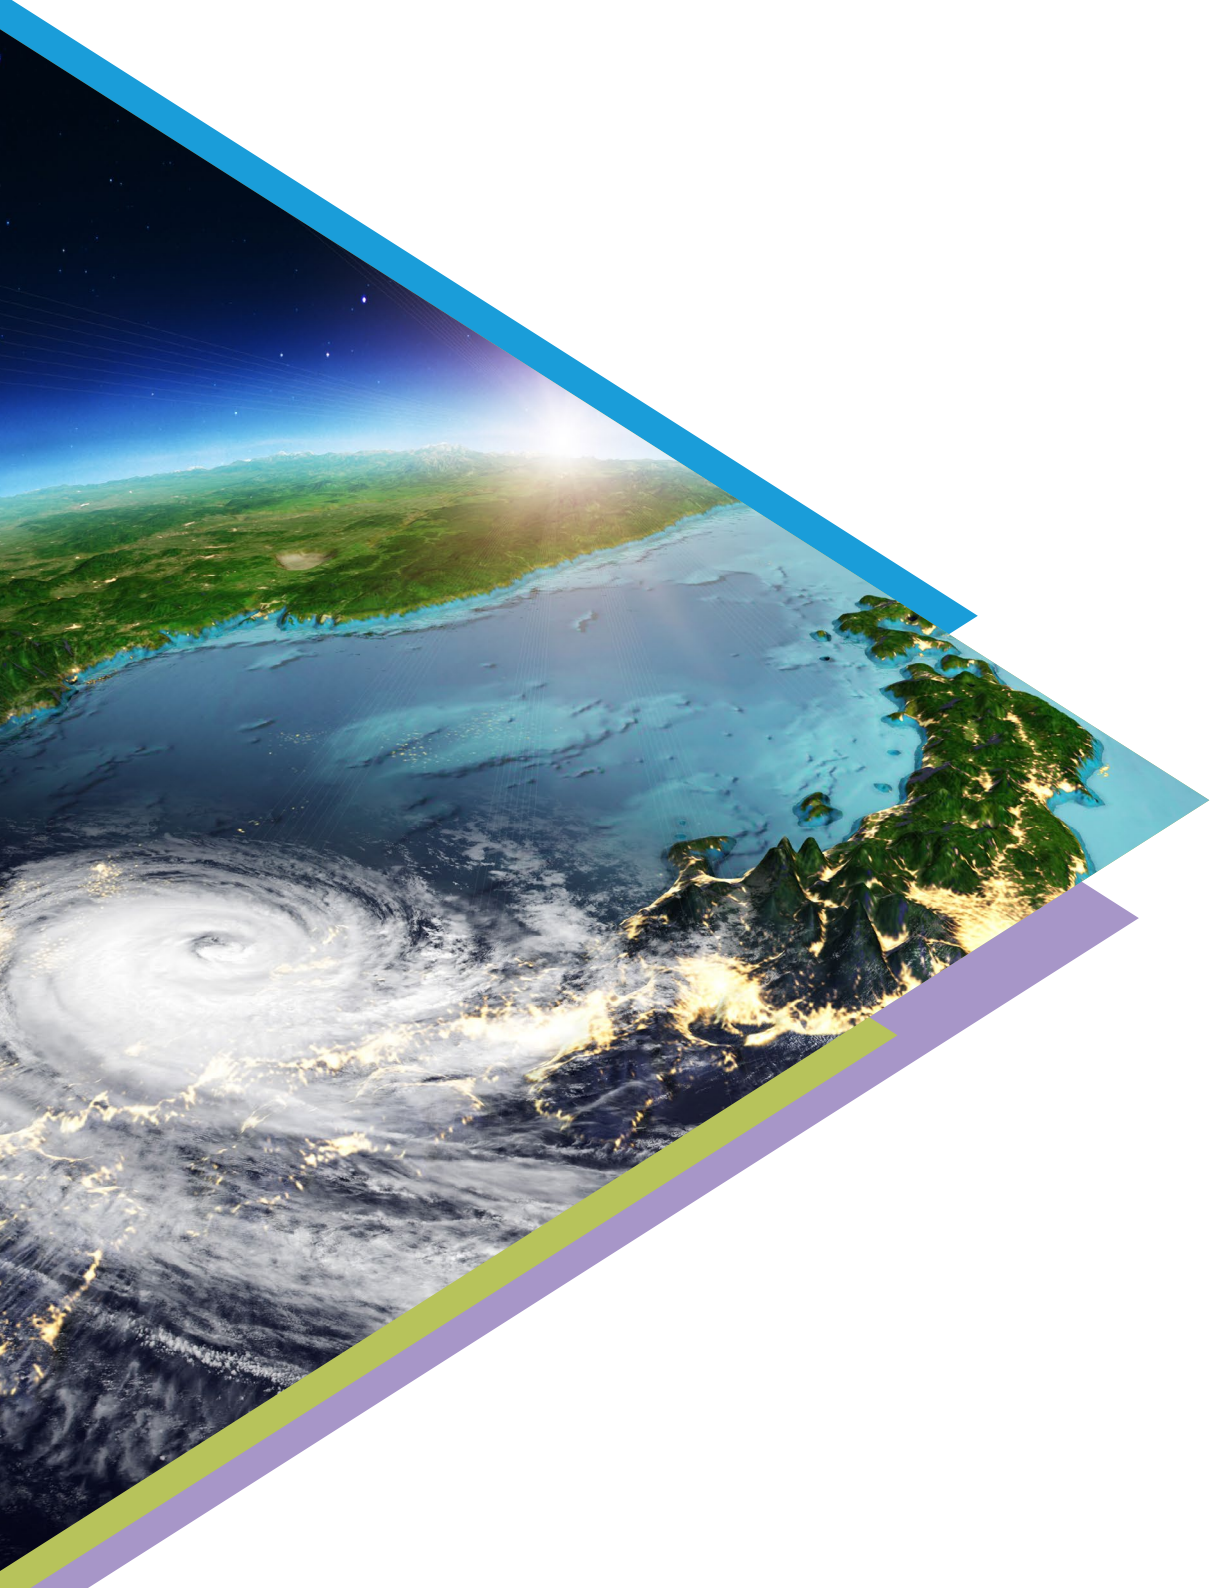

Future Matters

THE ADVANCE OF AI:

Challenging investors in a post-Covid world

COLUMBIA THREADNEEDLE INVESTMENTS

Your success. Our priority.

EXECUTIVE SUMMARY

- **The Covid-19 pandemic has hugely accelerated take-up of artificial intelligence (AI).** It is forcing companies to rethink their approach to intelligent automation, speeding the spread of AI-based technologies from the tech and internet sectors through the rest of the economy. Early adopters stand to gain huge competitive advantages.

- **While the holy grail of "artificial general intelligence" remains distant, the real-world applications of AI in its current form will be hugely economically significant.** AI works best today when carrying out limited and well-defined tasks, where large amounts of data are available to train the algorithms effectively.

- **AI is set to boost economic productivity as soon as three to five years from now.** Businesses are deploying AI to cut costs, grow revenues and enable disruption, although it remains in its early stages of development. Similarly to personal computers and other "general purpose technologies" such as electricity and the internet, AI will prove a powerful tool for creating economic value.

- **Active portfolio managers will be challenged to identify and invest in "superstar companies" that successfully use AI to sharpen their competitive edges and dominate their respective sectors.** All industries are now adopting AI; it is just a question of how quickly. In time, it will affect many stocks in an investment portfolio.

- **AI will create formidable "economic moats" enabling businesses to establish and maintain competitive advantages.** Leaders in AI are set to benefit from a virtuous circle: advances in AI lead to better products and services that attract more users, who provide more data, which leads to further advances in AI. This cycle will also benefit companies by helping them to attract the best people and grow profitability, fostering still more investment in AI. AI leaders will increase their returns on capital significantly. Meanwhile, some undervalued companies face major challenges to their business models and may not exist 10 years from now.

CONTENTS

An investment perspective

David M. Egan, CFA®
Senior Analyst, Semiconductors and Head of Thought Leadership, Research

Neil Robson
Head of Global Equities

James Waters
Senior Data Scientist

A business perspective

Anand Rao
PwC's Global AI Lead

An academic perspective

Professor David Barber
Director of the UCL Centre for Artificial Intelligence

Dr Ali Shafti
Senior Research Associate in Robotics and AI, at the Brain & Behaviour Lab, Imperial College London

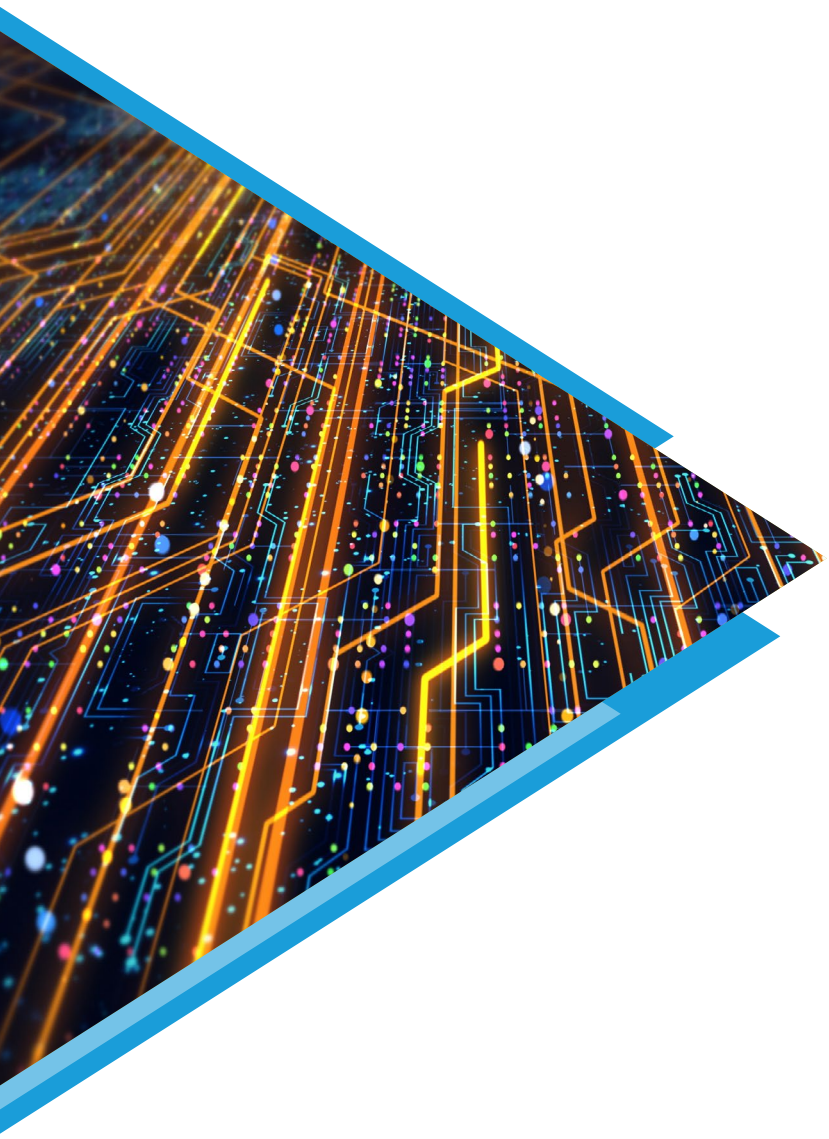

An investment perspective:

COVID-19 'MASSIVELY ACCELERATES' TAKE-UP OF AI

The pandemic has forced companies to rethink their approach to intelligent automation, speeding the spread of AI-based technologies from the tech and internet sectors through the rest of the economy. Early adopters stand to gain huge competitive advantages, according to David M. Egan, Senior Analyst, Semiconductors and Head of Thought Leadership, Research, at Columbia Threadneedle Investments.

QUICK READ

- **In healthcare and elsewhere, AI-related technologies have received a "battlefield promotion" as the crisis forced the pace of innovation and adoption.**
- **Companies that move fastest to adopt technology innovations such as AI and intelligent automation open up a productivity and growth advantage over those that move more slowly, which quickly becomes unassailable.**
- **It is reminiscent of the 1990's, when early adopters that installed industrial robots thrived, expanding their headcount by more than 50% over the following 18 years as their increased productivity delivered gains in market share. By contrast, non-adopters cut jobs by 20%[1].**

[1] Source: Robots and firms; VOX CEPR Policy Portal; 01 July, 2019. https://voxeu.org/article/robots-and-firms.

David M. Egan, CFA®
Senior Analyst, Semiconductors and Head of Thought Leadership, Research

Even before the Covid-19 pandemic struck in early 2020, artificial intelligence (AI) and related computing techniques were already spreading beyond their heartland in tech and internet companies to other industrial sectors. The Covid-19 crisis, however, has prompted a "massive acceleration" of the trend towards intelligent automation, says David M. Egan, Senior Analyst, Semiconductors, at Columbia Threadneedle Investments.

In some cases, this is directly related to the fight against the virus. Egan cites examples including Royal Bolton Hospital in the north of England, which fast-tracked the introduction of an AI-based system to triage the huge flow of patients presenting with suspected Covid-19. The algorithm read chest X-rays for signs of Covid-related lung infection, allowing medics to identify those patients most in need. Similarly, researchers seeking possible drug therapies as the crisis escalated relied on very computationally-intensive parallel processing techniques that are fundamental to AI. This allowed them to scan the massive body of published research and identify promising candidates.

Covid-19 hands AI a 'battlefield promotion'

Egan says that in healthcare and elsewhere, AI-related technologies have received a "battlefield promotion" as the crisis forced the pace of innovation and adoption, sweeping aside doubts over whether the time was right to start experimenting with new techniques. In sectors ranging from logistics to financial services and facilities management, companies are accelerating their take-up of intelligent automation.

The immediate impetus may be to combat the effects of the Covid-19 pandemic on their business operations, but these are investments they would have made eventually in any case – the arrival of Covid-19 simply pushed them to the top of the agenda. As a result, the adoption of technology, AI and data science will become ever more pervasive as companies that did not previously prioritise the 'digital revolution' rethink its potential impact on their businesses.

In examples such as hospitals tackling Covid-19, AI tools that prove their worth after being fast-tracked into service under crisis conditions for one purpose, will spread through the organisation as the crisis passes and further use cases are identified, says Egan. "Are you going to throw them away? No, you're going to say 'Oh, I tried this thing and it worked really well. Let's see what else we can

> "The Covid-19 crisis has prompted a "massive acceleration" of the trend towards intelligent automation."

> "This is an existential question of whether companies want to prosper and succeed or fail."

use it for.'" Healthcare has been a direct beneficiary of AI's accelerated adoption, but Egan says he is also seeing a surge of interest in robotics among companies that need to respond to the productivity hit caused by social distancing requirements in warehouses. "We're only now getting to the point where we can start doing pick and place, so the robot can look using sensors and pick things up and move them into another place autonomously," he says.

As before, robots were starting to be used to pick and move packages around warehouses before the crisis, but Covid-19 is speeding up adoption significantly. Although this automation trend is not yet showing the same momentum in factories, where the more complex production stage poses greater challenges, for the handling of finished products the shift is now well under way.

Accelerated uptake brings unassailable advantages

Our experience from previous technology adoption waves demonstrates the effects that an accelerated uptake of AI and related technologies is likely to have.

Data from Spain compiled by *Encuesta sobre Estrategias Empresariales*, an annual survey covering 1,900 manufacturers, showed the commanding competitive edge that early adopters of traditional industrial robots gained. Those that installed robots between 1990 and 1998 thrived, expanding their headcount by more than 50% over the following 18 years, from 1998 to 2016, as their increased productivity delivered gains in market share. By contrast, non-adopters cut jobs by 20%.[1]

"This is an existential question of whether companies want to prosper and succeed or fail," says Egan. He points to research by Accenture consultants Paul Daughtery and James Wilson, published in *MIT Sloan Management Review*, who studied 8,300 companies across 20 industries[2]. Daughtery and Wilson report: "We found that the top 10% of these companies, in terms of their levels of technology adoption, technology penetration and organisational change are achieving levels of revenue growth that are double those of the bottom 25% and grow revenues more than 50% faster than the middle 20% of the companies... At critical stages of systems evolution, the 10% of companies that lead the way boldly choose the most challenging, but most rewarding, of the technology options typically available. In contrast, laggards

[1] Robots and firms; VOX CEPR Policy Portal; 01 July, 2019. https://voxeu.org/article/robots-and-firms.
[2] https://sloanreview.mit.edu/article/how-leading-organizations-are-getting-the-most-value-from-it/ The slide deck explaining their research sets out a framework explaining how companies can move up the tech adoption curve. Egan believes this framework is also relevant for investors in evaluating companies. See: https://www.accenture.com/_acnmedia/Thought-Leadership-Assets/PDF-2/Accenture-Legacy-or-Legend-PDF-Report.pdf

fail to achieve full value from their investments in new technology because they make defensible but suboptimal decisions that inhibit their ability to share and scale technology-driven innovation across business units and processes."

Egan observes: "That growth advantage compounds every year – it's a devastating conclusion because the compounding curve is exponential. People don't think in exponentials very well, so they don't understand the magnitude of the difference. They don't get that there is something fundamentally different going on here, which is that those big tech companies, which they say are overvalued, can scale their businesses to a level that we've never seen before and they can do it with incredibly high margins and free cash flow."

He points out that, as of December 2020, the five largest US technology companies (Alphabet, Amazon, Apple, Facebook and Microsoft) account for 18.8% of the market capitalisation of the Russell 1000 index. During 2020 just two sectors – information technology and healthcare – accounted for 28% of revenues in the Russell 1000, 32% of EBITDA, 55% of free cash flow and 41% of the index's market capitalisation, he says.

His conclusion is stark: companies that move fastest to adopt technology innovations such as AI and intelligent automation open up a productivity and growth advantage over those that move more slowly, which quickly becomes unassailable.

The effects of this trend are already obvious in the US tech sector but as AI-related technologies spread further into a wider range of industries – for example, to automated loan decisioning in financial services – he expects a similar divergence in growth rates and profitability between the leaders and the rest. The effect of the Covid-19 pandemic has been to accelerate adoption of these technologies and therefore increase the pressure on investors to identify those companies that are best positioned to benefit.

"The laggards who think they're trying things out – they're going to have problems. It may not be tomorrow, but it will be in five, 10 or 20 years because they're going to be so far behind. The scale that you're getting from using technology and AI allows a massive shift in what goes on in the economy.

"The tech intensity of the world economy is growing and computational approaches will take a higher share of

> "Big tech companies can scale their businesses to a level that we've never seen before."

> "Autonomous vehicles, if we ever get them, are the ultimate example of edge computing."

economic output." he says. "Over time, the companies that are more open to trying these things are more likely to gain control of the economy than others."

The tool makers for AI

This raises the obvious question of how best to invest in this trend. The technology and internet sectors remain the obvious focuses for investors, but Egan also highlights opportunities among the tool makers for AI.

The same visual sensors used in Amazon Go's stores instead of checkouts, for instance, are used in automated vehicles and robots. The companies making these sensors stand to do well. Similarly, NVIDIA designs the graphical processing units (GPUs) for machine learning that Audi, Mercedes and Toyota are experimenting with in self-driving vehicles and healthcare companies use for purposes such as early detection of cancer. Based in Santa Clara, California, NVIDIA is the market leader in GPUs, which in effect are the brain of AI solutions.

Having designed the GPUs, NVIDIA also has NVIDIA Drive, a platform that enables car manufacturers to create their own algorithms for automated vehicles. In healthcare, its platform allows radiologists to build a data set and create algorithms that make predictions about the kind of surgery people need. AI is perfectly suited to making early and better diagnosis of disease, as there is a huge amount of data available.

However, complexities around handling data, such as confidential health information, are spurring innovation in other areas, he says.

In cases where data has to be processed locally rather than sent to the cloud for regulatory reasons, or because the latency involved in sending data to the cloud is too great, the processing will have to be carried out locally – at the edge, in the jargon. "Autonomous vehicles, if we ever get them, are the ultimate example of edge computing," he says. "You cannot guarantee that the vehicle will always have connectivity to a centralised cloud, so you're going to need edge computing there."

He says a range of service providers are trying to come up with an outsourced, cloud-type service that is available remotely, at the edge, for users such as hospitals handling confidential patient data. Innovations in this area are likely to open a further growth area for tool makers.

The pandemic brings data challenges for AI

While the Covid-19 pandemic has accelerated adoption of AI-related technologies across numerous sectors, its arrival has also highlighted the technology's limitations. AI operates by making predictions based on the patterns it detects within historical data sets. If patterns of human behaviour and demand for products and services change in fundamental ways, as they have done since the onset of global lockdowns, then historical data will no longer be relevant and predictions made using AI systems relying on historic data will no longer apply.

"That's what the pandemic is like for systems making supply chain predictions about how much you need here or there. Now your supply chain system is getting what looks like an anomalous set of data and it doesn't know how to react to it, because it's never seen this before." This is the point at which automated systems have to be augmented by human intelligence that is able to make judgements based on more than historical data.

Although the crisis may challenge existing AI systems in the short term, as they struggle to interpret data patterns they have never seen before, ultimately Covid-19 will widen the range of information that can be used to train AI systems and will accelerate their spread into all corners of the economy.

David M. Egan, CFA® biography

David Egan is a senior analyst in central research at Columbia Threadneedle Investments and is responsible for the semiconductor vertical. In addition, he leads the team's thought leadership effort.

Egan joined Columbia Threadneedle Investments in 2008. Previously, he was a research analyst at Lehman Brothers, where he covered semiconductor equipment. Prior to that, he worked at various internet and technology startups, including Lycos, in the San Francisco, California Bay area. Egan began his career doing actuarial work at the pension and benefits consulting firm Watson Wyatt.

Egan has been a member of the investment community since 2004. He received a BA from Duke University, a master's degree in Finance from CEMA University in Buenos Aires, Argentina and an MBA from the University of California at Berkeley. In addition, he holds the Chartered Financial Analyst® designation.

An investment perspective:

WINNERS TAKE ALL IN THE AI EVOLUTION

For active asset managers, artificial intelligence promises deep change that will touch many holdings in their portfolios. Interview with Neil Robson, Head of Global Equities at Columbia Threadneedle Investments.

QUICK READ

- **"Superstar companies" that successfully harness AI to gain formidable competitive advantages will grow their profitability significantly. Meanwhile, some undervalued companies face major challenges to their business models and may not exist 10 years from now.**
- **AI's progress is set to accelerate in the next few years and will affect many stocks in an investment portfolio.**
- **A key benefit of AI in many sectors will be the opportunity to capture efficiency and productivity gains. Relatively small productivity gains could lead to far larger increases in returns on invested capital.**
- **Other businesses will harness AI to accelerate revenue growth by developing new products and services, based on the insights provided by data that they own, generate and, in some cases, buy in.**
- **AI will entrench the dominance and superior profitability of a few leaders in major business sectors.**

Neil Robson
Head of Global Equities

In AI's accelerating evolution, the winner is surely likely to take a greater market share. Think of Amazon, Google and Microsoft in cloud computing, a vital form of infrastructure for AI. They dominate the sector that is poised to grow swiftly as AI supercharges its adoption, and such is their scale that new entrants will find it hard to compete.

In most other areas, the biggest winners of AI are harder to identify; even so their potential is great. AI has the power to transform productivity and supercharge revenues. On this view, adoption of AI will effectively lead to a wider dispersion of outcomes for companies and their shareholders, as successful companies compound their competitive advantages over time.

"As AI develops, the businesses that can adapt are more likely to accelerate this trend (of diverging performance) than to see any sort of mean reversion," says Neil Robson, Head of Global Equities at Columbia Threadneedle Investments. "In terms of growth versus value investing, the real message is that there's an underlying change and many businesses will see the reverse happen, where their business models face a massive challenge. An awful lot of value stocks have major issues – on a 10-year horizon, will they still exist?"

In terms of commercial applications, AI remains in its formative stages: many 'AI winners' have yet to emerge and some sectors that will ultimately be transformed by AI remain largely untouched. However, Robson believes that AI's progress will accelerate significantly over the next few years and that some, at least, of the dynamics that will characterise this process are already visible.

Changing the game in productivity

A key benefit of AI in many sectors will be the opportunity to capture efficiency and productivity gains by using AI rather than humans to make critical, real-time decisions. Much has been written already about the potential to automate many routine administrative tasks in sectors such as banking. But Robson also highlights the potential for AI to transform efficiency in capital-intensive industrial settings, where its effects could be significant.

He cites the example of semiconductor fabrication plants, which are already heavily automated, but where Intel has suggested that having decision-making by AI could lead to entirely automated chip production, giving incremental productivity gains of 2%-3%. "A 2%-3% gain in output from a fab probably represents the difference between

"AI's progress will accelerate significantly over the next few years and some of the dynamics that will characterise this process are already visible."

achieving a high-teens return on invested capital and achieving 30%. The gains could be that big," he says.

Similarly, Robson highlights the example of a small-cap company which produces 'quasi-AI' automated control systems for industrial applications including refineries. By aggregating real-time price data for the range of products the refinery can derive from a barrel of crude, it can optimise the refining process in order to maximise the value of the outputs from every barrel that goes through. Again, the steady progress of technology systems towards intelligent automation has the potential to manage complex processes and unlock significant efficiency gains. In capital-hungry sectors, even small productivity gains can transform the economics of the industry leaders.

As well as optimising key business operations, he suggests AI will fundamentally change near-universal corporate practices such as forecasting, which today frequently involves entering information manually in complex, error-prone spreadsheets. As that process is automated, it will change the way companies behave, Robson suggests. "AI enables any industry or business that does any sort of forecasting to do it better, quicker and cheaper. If you can do it better, quicker and cheaper, you'll probably do more of it. So I think the amount of forecasting and modelling that people will do is undoubtedly going to go up." If companies can reduce the need for human intervention, increase their forecasting capacity and improve their results, the benefits should be significant.

Forecast rise in labour productivity with AI by 2035

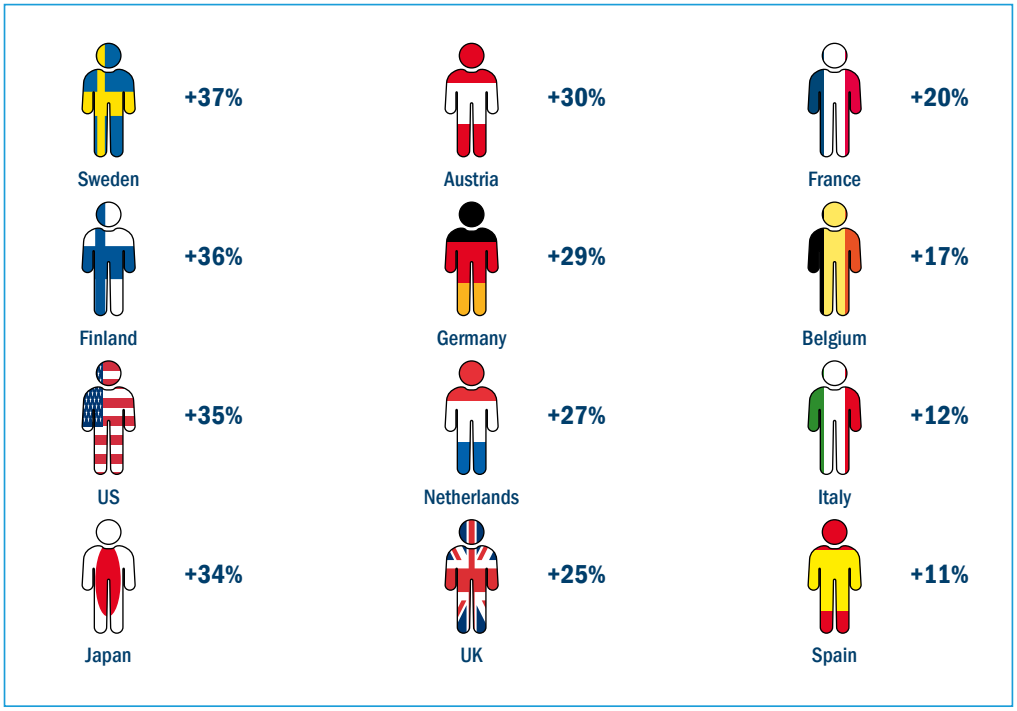

Source: Accenture and Frontier Economics Sepember, 2016.
Percentage increase in labour productivity with AI, compared to expected baseline productivity levels by 2035.

Unlocking revenue growth

Robson suggests the second major source of gains from AI will come from the ability to use it to accelerate revenue growth, by developing new products and services based on data that companies own, generate and, in some cases, buy in to augment their proprietary resources. The results of this process lie further in the future and, at this stage, it is hard to predict the outcomes of AI-enhanced corporate R&D and product development. Nevertheless, the potential opportunities here are very large, he argues.

The revenue gains that successful new products and services deliver could be significant. But Robson also points out that in order to apply AI effectively to their data and unlock new sources of value, companies will have no choice but to move their data resources from corporate silos into central 'data lakes,' which are most likely to be held in the public cloud.

The benefits of this transition are two-fold. Not only does it make all of the company's data AI-accessible in one place, facilitating the product development process, but it also significantly reduces the business's technology infrastructure costs. "As you go towards cloud solutions, you drop your (IT) cost structure 20%-plus and you have the ability to innovate and develop new products because all of your data is in one place," he says. "That might make your data even more valuable, which means you might be able to charge more." Successful adoption of AI therefore holds out the prospect of developing revenue-enhancing products and services while simultaneously enjoying the benefits of a leaner technology cost base. This points to another route through which AI could

> "The competitive gains AI delivers will reinforce the skew of returns on equity towards the top decile of companies, further concentrating gains among the leading players, where sustained RoE of 30%-40% will be achieved."

Paying dividends: Proven premium value

In 2019, Accenture surveyed 1,500 C-suite executives at companies with a turnover of more than US$1 billion from across 12 countries and 16 industries. Overall they reported a positive return on their investments. Accenture dug deeper.

Was there any relationship between successfully scaling AI across the enterprise and key market valuation metrics? What was the 'premium' for being a leader?

Using survey data combined with publicly available financial data, Accenture's team of data scientists created a model to identify the premium for companies in its sample that successfully scale AI, controlling for various characteristics of the companies.

Accenture discovered a positive correlation between successfully scaling AI and three key measures of financial valuation: Enterprise Value/Revenue Ratio, Price/Earnings Ratio, Price/Sales Ratio.

Companies that were identified as Strategic Scalers realise a success rate of 70% or more in their AI scaling initiatives and a return on their AI investment of 70% or higher.

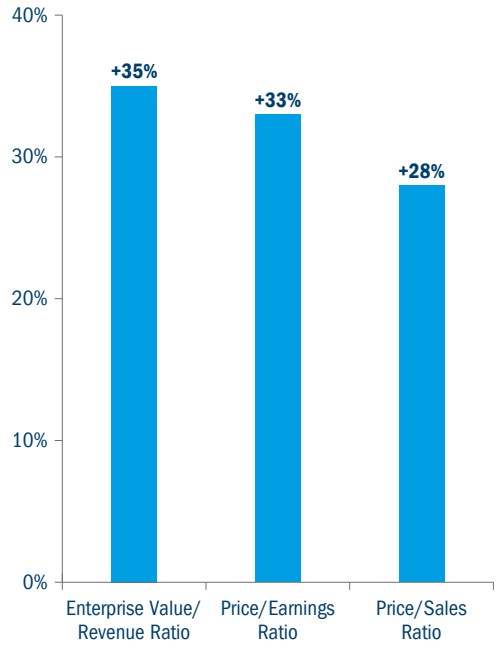

Source: Accenture 2019.
https://www.accenture.com/_acnmedia/Thought-Leadership-Assets/PDF-2/Accenture-Built-to-Scale-PDF-Report.pdf#zoom=50

enhance returns on invested capital for successful adopters.

Taken together, these predicted gains from applying AI to multiple business processes are likely to secure major gains in operational efficiency and revenue growth for the most successful companies. This dynamic should be expected to reinforce a trend that has been apparent for some time, in which a subset of what the consultancy McKinsey has called 'superstar companies' open a steadily growing lead over the rest in terms of returns on equity. Robson suggests that the competitive gains AI delivers will reinforce the skew of returns on equity towards the top decile of companies, further concentrating gains among the leading players.

Who are the winners?

The possible applications of AI are diverse, to say the least: from an AI-enabled oven that can decide how to cook whatever is placed inside it to the use of facial recognition technology to combat child trafficking. "When you read the use-cases you realise this is everywhere. Every single stock in your portfolio is going to be impacted in some way," says Robson. Equally, as other technologies such as 5G wireless networks mature, opportunities will proliferate to apply an 'AI layer' to huge new data flows – for example, from the 5G-enabled networks of sensors that will form the 'internet of things.'

From transport, energy and communications networks to virtually every manufacturing process, the scope for AI-based, real-time decision-making will vastly increase. As this happens, new products will emerge and consumer behaviours will change. Robson compares what will happen to the changes that followed the switch from 3G networks to data-enabled 4G. “Who would’ve guessed that making a phone call wouldn’t even be in the top 10 things we do on our telephones any more?”

He views the next three years as a period dominated by the build-out of AI infrastructure, including wider adoption of cloud computing, the introduction of 5G networks and increased efforts by companies to aggregate and structure their data. Potential winners during this phase are relatively easy to spot. Major data owners are well placed, particularly if they can use AI to enhance their offering in data analysis, he suggests, as are leading providers of essential hardware such as NVIDIA, maker of the leading programmable chip sets deployed for machine learning. Similarly, he believes existing leaders in fields such as industrial automation (eg, Keyence of Japan) or gene sequencing (eg, Illumina of the US) can use AI to enhance their existing competitive advantages.

However, the clearest winners are the oligopoly that dominate cloud computing: Amazon and Microsoft in front, followed by Google. “Total enterprise spend on computing is over a trillion dollars a year and Amazon Web Services in 2020 had annual revenues of around $45 billion, and Microsoft grew by 50% last year. So maybe they’re at $80 billion-plus between them. That’s going to go to $500 billion-plus and I don’t see any new entrants in this business. It’s almost impossible because you would just bleed cash.”

Robson suggests that although Google trails the top two in cloud services, AI adoption is likely to make it a stronger player as it leverages its heavy investment in AI talent to provide cloud services augmented by AI tools – the ‘AI as a Service’ model. “What we’re hearing back from companies now is that there’s a renewed look at Google because of their AI tool sets. As you move towards AI, Google’s cloud business should improve.” As AI adoption spreads, the leading cloud providers are likely to see their returns on invested capital settle in the high-30% range, further entrenching their long-term outperformance.

How much of a threat do these tech giants pose to leading players in other industries?

“Autonomous vehicles present a global market opportunity for the tech giants that parallels the opportunity Microsoft seized in operating systems for personal computers.”

There are clear risks, Robson suggests, for example in Google’s ambitions in autonomous vehicles. If Google became the dominant supplier of control systems for autonomous vehicles, the effects could be dramatic. At a licensing cost of $3,000 per vehicle, Google would consume the entire EBITDA margin of the typical auto industry original equipment manufacturer, which stands at around 12%, he says. “There’s a shift in the value chain. It probably won’t be as dramatic as that but it’s a shift in the value chain that is absolutely vital for investors to get right.”

Autonomous vehicles present a global market opportunity for the tech giants that parallels the opportunity Microsoft seized in operating systems for personal computers. In other industries, however, Robson believes they are much less likely to displace sector specialists because they do not have the necessary domain knowledge or data to compete effectively. Instead, they will provide the enabling infrastructure and tools, rather than attempting to dominate the entire value chain.

The challenge for investors, therefore, will remain the same: working out which companies have the strongest competitive advantages and are best positioned to benefit from the dynamics at work in their industries. “It’s easy to identify the data providers, the tool providers and the tech giants who are clearly going to win in AI,” says Robson. “But when you apply it to different industries it gets much harder. Who’s going to do best in banking? Will it be the existing banks or someone from outside the industry? That’s where it comes down to investors having conversations with companies to understand what they’re actually doing.”

Neil Robson biography

Neil Robson is Head of Global Equities at Columbia Threadneedle Investments. He took up this role in July 2017. He joined the company in 2011 as a portfolio manager within the Global Equities team.

Robson is the manager of several global equity funds and mandates for institutional clients. He is also the co-manager of the Threadneedle Global Extended Alpha.

Before joining the company, Robson worked as a fund manager at companies including Martin Currie, Barings and Citibank. In addition, he was Head of Global Equity at Pioneer Investments from 2003 to 2009.

He has an Economics degree from the University of Bristol.

An investment perspective:

AUGMENTING, NOT REPLACING, HUMAN INVESTMENT INTELLIGENCE

While AI is a powerful data analysis tool, it lacks the 'general intelligence' to be anything more. Interview with James Waters, Senior Data Scientist.

QUICK READ

- **Huge increases in the volume of data and computing power make it possible to gain an investment edge through swifter and better interpretation of information.**
- **Columbia Threadneedle Investments puts computing power into the hands of portfolio managers, enabling them to analyse unstructured data to gain a better understanding of a business's fortunes.**
- **AI is just one tool in the data analysis armoury, although a powerful one.**
- **In the field of investment analysis, the possibility of machine intelligence replacing rather than augmenting human intelligence remains remote.**

James Waters
Senior Data Scientist

When managing money, knowledge is everything. In an age of instant information and regulated company disclosures, it's rare to gain an investment edge from discovering facts about a business or the economy ahead of the market. Yet today's huge increase in volumes of data and computing power have made it possible to do so through swifter and better interpretation of information.

With human activity generating vast amounts of data, mining this for nuggets and patterns of information can provide insight into trends such as changing sentiment towards a stock or consumer behaviour. This is achieved through different forms of data analysis, the most advanced of which are the cognitive learning and problem-solving abilities of machine or artificial intelligence.

Columbia Threadneedle Investments takes a specific approach to this in line with its DNA as an active asset manager. It puts computing power in the hands of its portfolio managers, enabling them to analyse unstructured data and gain a better understanding of a business's fortunes. Almost three years ago, it set up a data science team dedicated not only to gathering data more efficiently, but also to finding fresh insights. AI is just one tool in its armoury, but a powerful one for spotting patterns and anomalies.

> "The information we gather, including through the use of AI, provides investment insights that help the team produce investment ideas. Our approach to investing remains the same. However, data analysis enables our investment teams to gain more insight in a more efficient manner."

> "We have used natural language processing (NLP) for sentiment analysis, looking at how we can infer consumer sentiment changes towards brands from social media data. This provides our analysts with more information on the industries and companies they cover. This, in turn, may or may not provide insights that affect their recommendations. It is another source of information that analysts can use."

How an asset management company chooses to harness AI depends on its investment style. At Columbia Threadneedle Investments, AI is dubbed "augmented intelligence", depicting its use as an extra source of information that helps portfolio managers to make investment decisions. While AI – especially machine learning – has developed quickly in fields such as natural language processing and image recognition, it lacks a human's "general intelligence" which is needed to understand the broad nature of a company.

Explaining the approach, James Waters, Senior Data Scientist at Columbia Threadneedle Investments, says: "We have a data science team in our investments group that helps to extract information from data sources, particularly those that analysts themselves have difficultly manipulating, for example, unstructured data sets or large data sets. We also develop processes to extract information in a far more efficient manner, freeing up time for analysts to focus on analysis rather than data gathering.

"The information we gather, including through the use of AI, provides investment insights that help the team produce investment ideas. Our approach to investing remains the same. However, data analysis enables our investment teams to gain more insight in a more efficient manner."

This way of using AI – as part of a broader data analysis tool kit – reflects Columbia Threadneedle Investments' nature as an active asset management company, where portfolio managers use fundamental analysis to identify high-quality growth companies. In an asset management company specialising more in quantitative analysis, it's likely that AI would be more central to investment decision making.

Waters gives several examples of how he has used AI to assist portfolio managers. For instance, natural language processing (NLP), a branch of machine learning, has been deployed to help identify changes in sentiment towards companies from social media feeds or financial statement filings.

"We have used NLP for sentiment analysis, looking at how we can infer consumer sentiment changes towards brands from social media data," he says. "This provides our analysts with more information on the industries and companies they cover. This, in turn, may or may not provide insights that affect their recommendations. It is another source of information that analysts can use."

"This way of using AI, as part of a broader data analysis tool kit, reflects Columbia Threadneedle Investments' nature as an active asset management company."

It is important to keep in mind that no formula, model or tool can in and of itself be used to determine which securities to buy or sell, or when to buy or sell them. There is no guarantee that any investment objective will be achieved or that return expectations will be met.

> “As analysts become more and more comfortable with some of these tools, techniques and their efficiencies, it will become more important to take them into consideration.”

Beyond investment analysis, Columbia Threadneedle Investments harnesses AI to enhance portfolio construction, using optimisation algorithms.

But in the field of investment analysis, the possibility of machine intelligence replacing rather than augmenting human intelligence remains remote. AI remains just one tool in the data science armoury.

Indeed, at Columbia Threadneedle Investments, it is portfolio managers and investment analysts who drive the use of AI, using it to help them judge a company’s competitive advantage, as well as any switches in sentiment towards the stock or its products.

As computers get substantially more powerful and the amount of data grows exponentially, so the importance of data science in investment seems only likely to grow. Does that mean portfolio management teams will increasingly turn to AI’s talent for analysing patterns?

“As analysts become more and more comfortable with some of these tools, techniques and their efficiencies, it will become more important to take them into consideration,” says Waters. “But it’s not just the techniques, it’s the volume of data that is more awesome. Everyone gets carried away with machine learning and AI, but the quantity and quality of data makes the greatest difference.”

Annual growth of the global datasphere

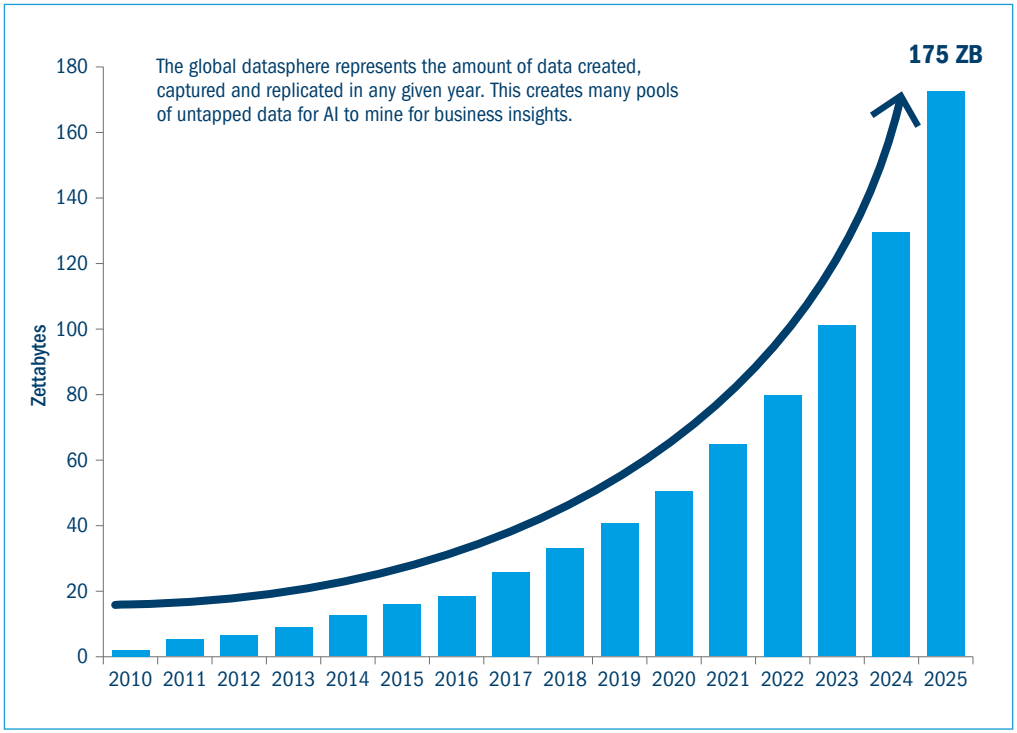

Source: Data Age 2025, sponsored by Seagate with data from IDC Global DataSphere, Nov 2018.

James Waters biography

James Waters joined Columbia Threadneedle Investments in October 2012 as a member of the fixed income team.

In 2018, he formally moved to the newly established data science team as a senior data scientist.

Before joining the company, Waters worked at Goldman Sachs Asset Management, initially in the technology team and then within the fixed income group.

He has both a master's in mathematics and a master's in computing, along with a degree in actuarial science.

A business perspective:

WHY AI WILL CREATE NEW BUSINESS LEADERS

Artificial intelligence is set to improve economic productivity and endow some companies with powerful competitive advantages. Interview with Anand Rao, PwC's Global AI Lead.

QUICK READ

- **AI is set to boost economic productivity within three to five years from now.**
- **Businesses are deploying AI to cut costs, grow revenues and enable disruption.**
- **While businesses are already using AI, it remains in its early stages of development, reminiscent of personal computers in the mid-1980s.**
- **Those businesses that harness AI successfully will achieve powerful competitive advantages.**

Anand Rao
PwC's Global AI Lead

Whether by powering digital virtual assistants, enhancing medical image screening or allowing chatbots to simulate human conversation, artificial intelligence (AI) is entering the business world. These are just the early days of a technological revolution that is accelerating and will reach across all industrial sectors in many different forms. AI is what's called a "general purpose technology" – like electricity, the personal computer or the internet – meaning that it has the potential to affect the entire economy.

In contrast to the natural intelligence of humans and animals, AI describes machines that mimic human intelligence. From the university lab to the business R&D department, machines are rapidly learning human-like cognitive skills. Technology, financial and medical businesses have quickly started to deploy them, while other sectors are beginning to implement them to cut costs, boost top-line growth and disrupt established business models.

In time, businesses that pioneer AI are likely to gain such great competitive advantages that it will prove hard for others to catch up. Anand Rao, a Boston-based PwC partner who is the firm's Global AI Lead, predicts that a wave of AI adoption will shortly have significant economic effects. Noting that it is too early to see AI boosting broad economic productivity, he thinks its effect will be apparent in macroeconomic data within three to five years. He also predicts that some companies will build "economic moats" that are hard to breach, disrupting entire sectors as they do so.

"Economic studies show that there is a virtuous circle at play with AI technology, where some companies earn an advantage through their data or people expertise," Rao explains. "Let's say you have lots of data and you are building an AI or machine learning (ML) algorithm that's somewhat better than others. As a result of better personalisation, more customers come to you, your data gets still better and your profits increase. With more profit you hire better people. Your AI gets better still, you get more data, more customers and so you have this virtuous circle."

Seeking competitive advantage

So far, the race to harness AI is at an early stage, comparable to the personal computer in the early 1980s or the internet later that decade. Yet already

"Let's say you have lots of data and you are building an AI or machine learning algorithm that's somewhat better than others. As a result of better personalisation, more customers come to you, your data gets still better and your profits increase. With more profit you hire better people. Your AI gets better still, you get more data, more customers and so you have this virtuous circle."

the tech and social media giants have huge amounts of data gained from their day-to-day businesses that they can use to develop AI through ML. This means creating algorithms that can learn from data, or to put it another way, computer programmes that can programme themselves by looking at information. "Deep learning" is the type of machine learning that is enabling today's progress. It uses many layered artificial neural networks – software that roughly copies the way that neurons work in the brain.

Today's powerful "massively parallel" computer processors help AI neural networks to learn faster than ever. Having huge amounts of data helps to train neural networks because they continually recalibrate their settings, gradually becoming more accurate. For example, through a technique called "supervised learning" a machine can be taught to recognise a dog by being fed thousands, or even millions, of images labelled "dog". Other ways to train machines are "reinforcement learning", which means learning through trial and error, and "unsupervised learning", which means teaching machines to learn from data coming from their environment. The former can be used to train self-driving cars in simulators, but the latter is still at an early stage.

So, where does the competitive advantage for companies lie? According to Rao, the answer is twofold. For the consumer tech companies and social media companies like Amazon, Google and Facebook, it lies in the data. Those that have more data have better machine learning algorithms. But in other sectors there is less data for training algorithms. In these cases, human expertise or "cognitive capital" becomes equally important, as the expertise of people can be combined with data to create AI algorithms.

"Take medical diagnostics," explains Rao. "Various clinical experts have done the diagnostics, there's a huge volume of image data. Now take the image data and combine it with the expertise of the people. That is critical, because otherwise you don't know what is cancerous or what is not cancerous, so human expertise is involved in labelling the data or saying what the things are that you really want a system to learn. So, human expertise, combined with the specific data in that particular domain, can build something that's tangible. Therefore, that cognitive capital is where competitive advantage will lie over the next decade, if not more."

Where will the value gains come from with AI?

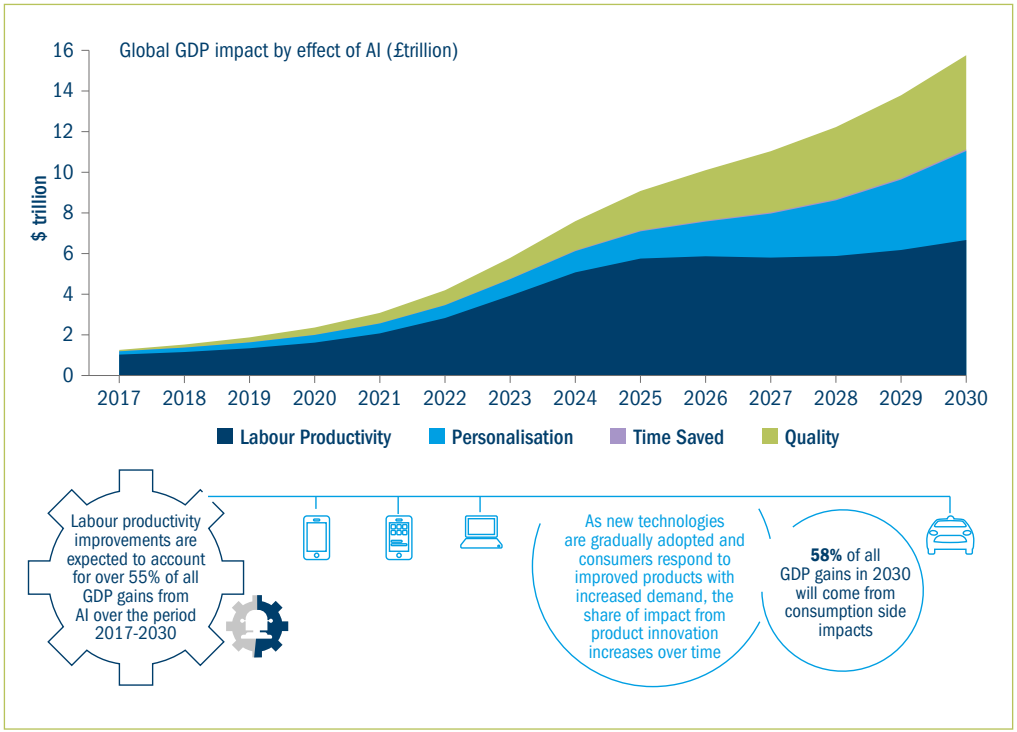

Source: PwC analysis 2019.

How industries are adopting AI

Tech companies are at the forefront of developing AI, followed by media companies. The former are also using it to collect data about consumer behaviour – what people are buying – whether it is goods, services or entertainment.

The second wave is broader, including financial services, healthcare and retail. Financial services companies, from hedge funds and capital markets banks, to retail banks and insurers, are harnessing AI for a broad range of purposes from investment analysis through to customer care. Healthcare companies are developing a variety of ways that AI can help them provide better care at lower cost and retailers are using AI to gain a better understanding of their customers.

The third wave is heavy industry, where the industrial internet of things (IoT) is the catalyst. Aerospace manufacturing, oil, gas and utilities are all sectors deploying sensors and IoT to connect various types of equipment. These sensors are collecting huge amounts of data. AI can make sense of this data, thereby predicting where maintenance will be required and preventing expensive breakdowns.

From a functional perspective, AI is being used in three areas. The first of these is the front end of a business, including strategy development, customer analysis, customer experience and distribution. The second is product development, operations, pricing and customer service. The third spans the back office, including finance, HR and, in the case of financial services, the risk function.

> "If you don't automate your business processes you are going to be irrelevant because your costs will be out of whack with the competition."

From automation to disruption

As businesses are digitised, this inevitably generates lots of data, which leads to standardisation and then automation. As back office paperwork is digitised, so processes are automated and, at some point, AI is used to optimise processes. Similarly, operations and customer service are being automated, smoothing the flow of goods and services. Everything becomes more efficient and more effective.

At the front end of a business, AI is improving the customer experience. For example, chatbots are making it simpler for people to interact with businesses, and AI techniques are being used to turn audio and video into text.

But organisations can overestimate the benefits of AI, according to Rao. There is a tendency to think that cutting the amount of time taken to perform an administrative task should lead to an equivalent reduction in headcount. Often employees use some of the time saved to benefit the customer in other ways.

While it's still early to see the benefits of efficiency gains feeding through to productivity numbers, Rao predicts this is not far off. "In specific sectors, people are looking at how to use some of these technologies. For example, how do we use them to reduce the load on call centres? I think you'll see very specific back office functions – customer service and support functions – where there is a huge volume of this happening. There you will see the productivity gains.

"In fact, if you don't automate your business processes you are going to be irrelevant because your costs will be out of whack with the competition."

When it comes to boosting sales, AI is being used to give customers a better experience. The way Amazon and Netflix use AI to personalise customer experiences by making recommendations is a good example of this. Helping customers to reach decisions in this way should make them more loyal.

But the third way that the technology can boost sales, "disruptive AI", is more fundamental. It seeks to answer the question: how can I disrupt my industry or a neighbouring sector? For example, PwC has used AI to help a major car manufacturer develop a strategy for car-sharing. This led to the establishment of a multi-billion-dollar business unit devoted

to car-sharing, electric vehicles and autonomous vehicles.

Another example is oil and gas exploration, where some exploration companies are using AI to interpret the data in seismic surveys. The exploration companies with the most seismic survey data are in the best position to exploit the technology. This allows them to judge precisely where to drill.

Taking AI out of the lab

Looking forward, AI is still a relatively new concept in the business world. Excepting the tech companies' uses of AI in social media and consumer electronics devices such as mobile phones and virtual assistants, there is a yawning gap between what the science of AI has made possible and the uses that businesses are putting it to.

It's not just that businesses may not follow AI's progress in the university labs closely enough. Even more importantly, AI must be adapted for business. Companies need to have executives who understand AI, who can cleanse the data and label it, as well as understanding the risks. Academics, for their part, are not interested in writing papers about methodologies for adapting different strands of AI to the business world.

"So, AI has to move from being an academic discipline of better algorithms to essentially software engineering, with appropriate methodologies, processes, controls and governance," notes Rao. "I think businesses will catch up with what the academics are providing but they'll also be charting a new course forward." He predicts that start-ups will spring up to build the tools needed to commercialise AI.

Additionally, AI is just one part of the broader business technology ecosystem. While AI will make decisions, there will need to be processes and protocols concerning collecting data, organising data, presenting it to people and integrating the decisions made by the AI machine with those still made by humans. There is also computer hardware and software to be developed.

That said, AI is rapidly making a difference in some sectors, especially considering it was hardly used in businesses outside the consumer tech and social media industries 18 months ago. So, how can one judge whether AI is about to turn the competitive dynamics of an industry upside down, with a start-up usurping the leadership of a long-dominant giant?

> The interesting question is how does regulation, and some of the things that are happening economically and politically, play into that? Many politicians and nations are thinking about this

"I think you should look at how AI is being used, both from a top-line and a bottom-line perspective and how it's affecting performance," Rao says. "Then there is the notion of whether AI is going to disrupt an industry, with an organisation coming in from outside the sector and creating new competitive advantages, new cognitive capital."

As companies start to build competitive advantages, many of them could become unassailable, having accumulated the best data, algorithms and people. "Left purely to the market, I think that's what's likely to happen," observes Rao. "The interesting question is how does regulation, and some of the things that are happening economically and politically, play into that? Many politicians and nations are thinking about this. How do I protect my citizens? About 30 countries are developing national AI strategies and we have been talking to quite a few of them. This is at the top of their mind: if foreign companies dominate our industries through the use of AI, what will it mean for the economy of our country, for our security? The biggest question is whether there will be laws to prevent this from happening. My guess is there will."

To what extent AI disrupts sectors remains to be seen, but it is likely to transform the global economy and the fortunes of businesses. Like the general-purpose technologies that preceded it – from electricity to the personal computer and the internet – AI will prove a powerful tool for creating leading businesses.

Dr. Anand S. Rao biography

Dr. Anand S. Rao is a Partner in PwC's Advisory practice. He is the global artificial intelligence lead, cross-vertical analytics champion, and the co-sponsor for the AI Centre of Enablement within PwC. With more than 33 years of industry and consulting experience, Rao leads a team of practitioners who work with C-level executives at some of the world's largest organisations, advising them on a range of topics including global growth strategies, marketing, sales, distribution and digital strategies, behavioural economics and customer experience, and statistical and computational analytics. As the global lead for AI, Rao is responsible for research and commercial relationships with academic institutions and start-ups, as well as research, development and commercialisation of innovative AI, big data and analytic techniques. With his PhD and research career in AI, and his subsequent experience in management consulting, he brings business domain knowledge and statistical and computational analytics to generate unique insights into the practice of "data science".

Prior to joining management consulting, Rao was the chief research scientist at the Australian Artificial Intelligence Institute. He has held board positions at start-ups and not-for-profit companies. He has received widespread recognition for his extraordinary contributions in the field of consulting and artificial intelligence research. He received the Most Influential Paper Award for the Decade in 2007 from the Autonomous Agents & Multi-Agent Systems organisation for his contribution on the Belief-Desire-Intention Architecture; MBA Award of Distinction from Melbourne Business School, 1997 and University Postgraduate Research Award from University of Sydney, 1985; and the Distinguished Alumnus Award from Birla Institute of Technology and Science, Pilani, India. He was recognised as one of the top 50 data and analytics professionals in the United States and Canada by Corinium; one of the top 50 professionals in InsureTech; and his recent paper on "A Strategist's Guide to Artificial Intelligence" has won the National Gold Award by ASBPE for the best technical article in 2017 and the FOLIO editorial award.

He has co-edited four books and published more than 50 papers in refereed journals and conferences. He is a frequent speaker on AI, behavioural economics, autonomous cars and their impact, analytics, and technology topics in academic and trade forums.

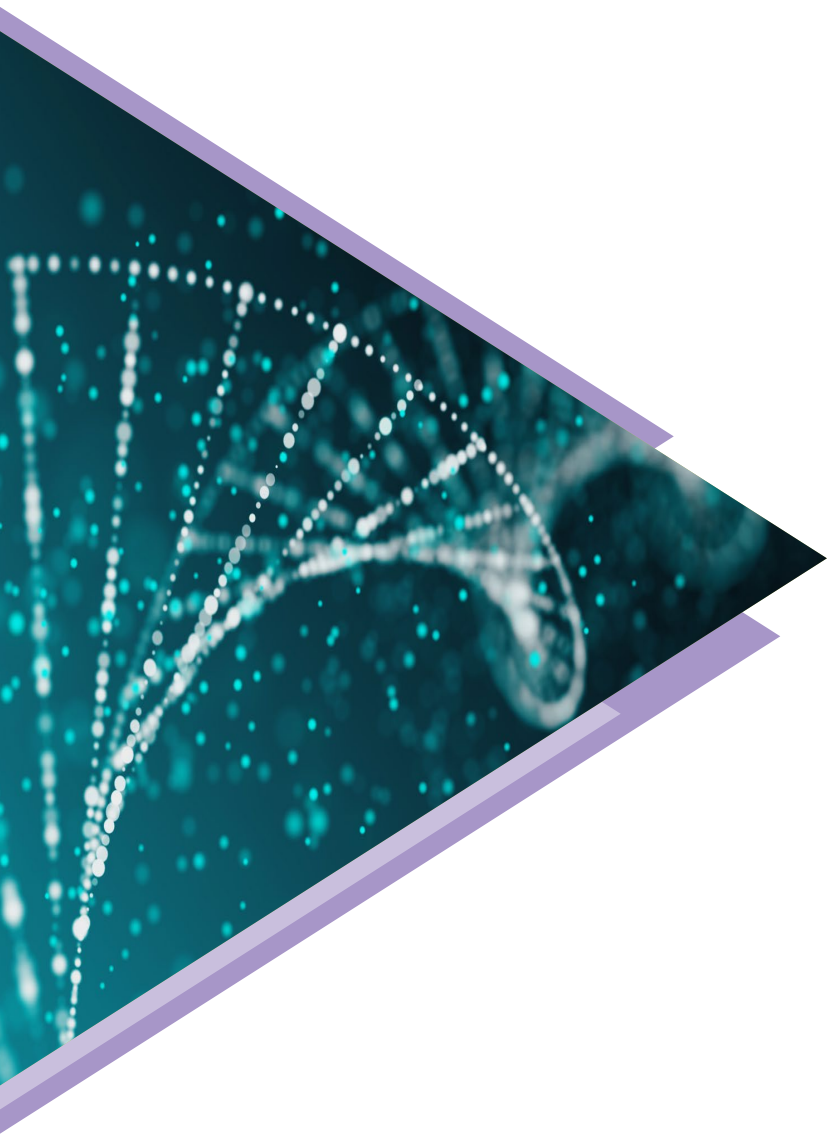

An academic perspective:

AI'S LEAP FORWARD, HUGE POTENTIAL AND UNDOUBTED LIMITATIONS

Advances in computing power and data availability have accelerated AI's evolution and it is now entering our daily lives. Even so, it is still in the foothills of the Himalayan task of developing systems with deep understanding. Interview with Professor David Barber, Director of the UCL Centre for Artificial Intelligence and fellow of the Turing Institute.

QUICK READ

- **AI suffers from a blurring of the boundaries – it's about machines being able to mimic the way humans work, rather than simply the analysis of large data sets.**
- **Computing power and data availability have combined to enable a period of accelerated development in machine learning, a major data-driven sub-field of AI.**
- **AI works best when carrying out limited and well-defined tasks where large amounts of data are available to train the algorithms effectively.**
- **While the prospect of "artificial general intelligence" remains distant, the real-world applications of AI are going to be hugely economically significant.**

Professor David Barber
Director of the UCL Centre for Artificial Intelligence

Our everyday lives are increasingly spent interacting with technology that replicates abilities. Advances in AI have resulted in voice recognition software that allows us to give instructions to Siri or Alexa. They have created sophisticated predictive text functions in email programs, online customer service chatbots and the telephone-based systems now being deployed in call centres, all of which depend on natural language processing. Translation tools and digital assistants that are capable of turning speech into text work in the same way. Image recognition software, as deployed in facial or number plate recognition systems and autonomous vehicles, also represent everyday examples of AI in action.

Innovations in robotics, a closely related field that offers exciting possibilities in areas such as driverless cars, warehouse automation and personal care for the elderly or infirm, also depend heavily on advances in AI. These machines all use AI to replicate humans' ability to interpret and interact with the physical environment, as well as drawing on insights from neuroscience into how humans function.

The emergence of AI as an increasingly common feature of modern life suggests we are on the cusp of a transformation that will produce vast changes in the way humans live and work. However, to interpret and navigate the effects that AI is likely to have on society and the commercial world, we must explore how and why it appears to have undergone such a major leap forward in recent years, and appreciate its current limitations as well as its undoubted potential.

Why are we talking about AI now?

Dr David Barber, Professor of Machine Learning at UCL and Director of the UCL Centre for Artificial Intelligence, points out that attempts to create human-like abilities in man-made systems date back centuries. Indeed, Prof Barber is a fellow of the Turing Institute, which recognises the pioneering role of Alan Turing, who died in 1954, in the development of the discipline. Turing and fellow mathematician and economist David Champernowne wrote their groundbreaking chess program, Turochamp, in 1948 during their research into AI. But the algorithm that powered Turochamp was too complex to run on the computers of the time and Turing was only ever able to execute the program manually, using paper calculations.

Computing power's unstoppable rise

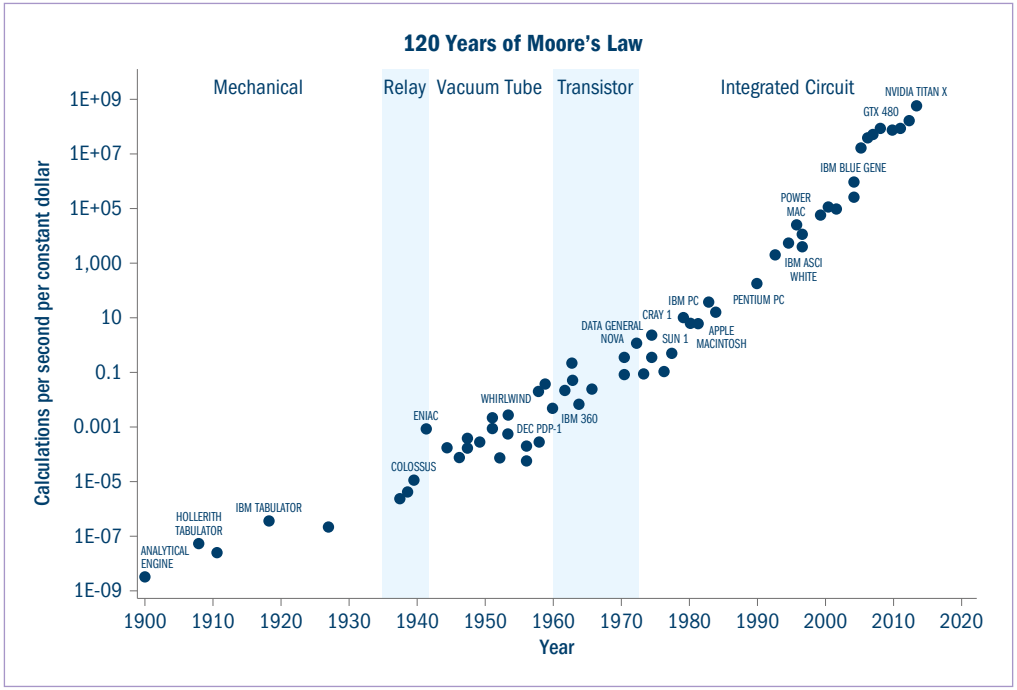

Source: As of 2018. https://www.britannica.com/technology/Moores-law.

This anecdote illustrates an important point. The basis of many of the algorithms in use today are not new. What has unlocked their potential and, therefore, that of AI, is the abundance of computing power that has become available in recent years as processing speeds have increased. An advanced image recognition system takes about a week to train using today's single NVIDIA GPU-based computers. Performing the same number of calculations using the best workstations available in the early 1990s would have taken hundreds of thousands of years. Gains in computing speed, compounded over decades, have delivered hardware capable of allowing AI to operate in real-time.

The second vital factor in the emergence of AI has been the increasing availability of data. As the quantities of digital data created and stored have multiplied rapidly over recent years, so data sets large enough to train algorithms to high levels of accuracy and proficiency – for example, images for use in teaching object recognition – have been created.

Machine learning emerges as the dominant approach

Taken together, these two factors – computing power and data availability – have combined to enable a period of accelerated development in machine learning (ML), a major, data-driven sub-field of AI. As a result, over the past 15 years or so, ML has become the dominant paradigm within AI and is largely responsible for the advances that underpin the applications we are most familiar with today.

Within ML, a key avenue of development dating back decades involved neural networks – systems loosely based on the structure of the human brain. Having been largely out of favour for many years, neural networks became the key focus of ML research once again in 2006, when a small group of researchers demonstrated that with access to enough computing power, the technique offered significant improvements in the results obtained.[1] Other major advances followed quickly. Soon afterwards, researchers were able to adapt graphics processing units (GPUs) developed for computer gaming to speed up the process of training ML algorithms 100-fold.

Armed with rapidly improving technology tools, a research group led by Geoffrey Hinton made a breakthrough in image recognition using ML in 2012 and was immediately acquired by Google. Hinton's team then went on to quickly create a speech recognition system far ahead of any previous system. ML and its variants, such as deep learning, had become the key technique within AI.

Deep Blue, AlphaGo and the limits of gameplaying

The best-known landmarks in AI tend to be moments such as the victory of IBM's Deep Blue computer over world chess champion Gary Kasparov in 1997, or of Google-owned Deepmind's AlphaGo over Korean Go champions Lee Se-dol in 2016 and Ke Jie in 2017. These highly symbolic events naturally appeal to our fascination with the idea of machines that can outwit humans. But how significant are they, in fact?

Prof Barber observes: "A story that ML researchers like to tell is that we had a machine that could beat the best human chess player in 1997, but we still don't really have a robot that can smoothly and reliably pick up a chess piece and move it." High-profile achievements like these are important to generate attention, but in pure research terms they have turned out to be far less significant than many assume, he argues. "What's ultimately important is not the ability to play chess or Go, but the delivery of systems that will be useful in our daily lives. The rest is largely entertainment."

The most important challenge for AI, he argues, is moving it beyond the enclosed, rules-based world of games and making it good enough to operate alongside humans in the much more complicated environment of our daily lives.

> "What's ultimately important is not the ability to play chess or Go, but the delivery of systems that will be useful in our daily lives. The rest is largely entertainment"

[1] One particular milestone is reported here https://science.sciencemag.org/content/313/5786/504 in which neural networks were shown to vastly outperform traditional methods for image compression.

> "Speech recognition is a good example where we have quite good performance now. But it's still very superficial – the machine doesn't actually understand in a deep sense what you're saying"

Where today's AI works best

If the research significance of AI systems that can beat chess or Go champions may be overstated, however, in one respect perhaps these achievements do carry some wider significance. Games such as Go or chess are extremely complex, rules-based problems for which very large quantities of game-play data are available that the ML algorithms can learn from. A newer version of AlphaGo, AlphaZero, taught itself to play Go, chess and shogi to a higher level than its predecessors simply by applying the rules it had been given to play practice games, thus removing the need for data from completed human games.

It is no coincidence that the real-world fields in which ML-based AI systems have been most successfully applied have tended to be those that share certain characteristics with games. The scope of the task the AI is asked to carry out is limited and well defined, and large amounts of data are available to train the algorithms effectively. Applications ranging from facial and number plate recognition, to the machine's ability to recognise and decode the phonemes that make up human speech, or even the visual characteristics that define everyday objects all, to a greater or lesser extent, exhibit the same combination of characteristics.

However, in areas such as object recognition by autonomous vehicles, which must operate to extremely high levels of accuracy to meet safety concerns, the performance of ML systems still remains below what will be required. The task of correctly interpreting every object the system encounters in our highly complex everyday environment is neither limited nor well defined. Consequently, even the most advanced image recognition systems fall short [as we explore in the interview that follows with Dr Ali Shafti].

A chatbot can successfully handle simpler banking or insurance enquiries, because the range of tasks it needs to perform is limited by the nature of the conversation and the data it requires to do so is readily available from the customer or the bank's records.

But if a lonely bank customer were to phone the call centre in search of someone to talk to, a skilled and empathetic human operator could meet their need. Today's chatbots would stand no chance – the task is way beyond their ken. The best they could do would be to hand over the call to a human.

The Himalayan tasks to come

"Speech recognition is a good example where we have quite good performance now," says Prof Barber. "But it's still very superficial – the machine doesn't actually understand in a deep sense what you're saying. Similarly, translation appears largely successful. You can now get pretty good translations automatically from one language to another, particularly if the languages are quite close. But does it really understand what you're saying?

"There's nothing wrong with where we are right now. The progress we've made is extremely impressive but we're just at the foothills of this Himalayan task. We still don't know how to make that leap to systems that are much better at deep understanding. It's a leap that the tech giants are well aware of and are investing heavily to crack, because if you can crack that, then the utility of things like digital assistants will vastly increase."

Besides AI's inability to understand context, intuit or unspoken meaning, other challenges loom large. The data-hungry nature of ML-based systems necessarily limits the contexts in which they can be used most effectively. In data-poor environments, they struggle. This is one of the chief problems with so-called reinforcement learning, where AI-based systems learn from their environment rather than by crunching through huge volumes of training data learning to associate decisions with their long-term consequences. A man-made system that could learn from environmental stimuli, in the way humans do, would have to rely on a far smaller volume of information than current AI systems require. Making that leap to a more data-efficient style of learning is a key current research goal.

"Machine learners are intellectually fascinated by the idea of reinforcement learning because it is in some sense the 'mother problem' for AI: how to train systems with only very limited feedback on the eventual success or failure of a current decision," says Prof Barber.

The future of AI

Even though the prospect of "artificial general intelligence" remains distant, Prof Barber argues that real-world applications of AI, for example, in fully autonomous vehicles, are going to be hugely economically significant. Similarly, the development of robots that can accurately pack goods for despatch from warehouses, which remain largely manual environments, will have a major impact.

> "The progress we've made is extremely impressive but we're just at the foothills of this Himalayan task. We still don't know how to make that leap to systems that are much better at deep understanding. It's a leap that the tech giants are well aware of and are investing heavily to crack, because if you can crack that, then the utility of things like digital assistants will vastly increase."

Proponents of AI argue that widespread application of robotics and AI in business will free humans from many of the dull, repetitive, physically demanding tasks that we currently perform simply because human labour remains cheaper than robots. "Ever since the industrial revolution and even before that, we've been using humans as if they were machines," says Prof Barber. "Is it a good use of a human being to be stuck behind the wheel of a truck eight hours a day for 30 or 40 years? Humans are much better than that. We have amazing abilities – empathy, compassion, creativity. These are things that machines are rubbish at and probably will be rubbish at for a long time. I'm passionate about freeing humans to do the kinds of things that we are uniquely good at. So in that sense, AI is a very positive pursuit."

The transition he sees unfolding will involve humans working increasingly alongside AI-based machines, performing the tasks requiring greater experience and skill, while leaving the mundane, repetitive elements to the machines.

In areas that still involve largely manual processing of standardised tasks, such as the back-office operations of large banks and financial institutions, there is obvious potential over time to replace tens of thousands of human roles with so-called robotic process automation, a repeat of the automation of factory floors through the 20th century.

The prospect of an AI-induced transformation of the workplace provokes understandable fears, Prof Barber acknowledges. "I don't know if revolution is necessarily the right phrase. These things are always over-hyped. I think it's going to be an evolution to some extent and I think humans are always remarkably resilient in creating meaning in their work and personal lives, despite the transformations that happen in society. Not surprisingly, he warns against getting "too fearful" about the number of today's jobs that could be replaced by AI and robotics.

However, these issues will not disappear. The big questions for investors, companies and their employees will be less concerned with the potential effects of AI and robotics on economic activity, which in time will prove profound, but much more concerned with how far and how quickly governments and regulators choose to respond to the searching questions AI will pose for our societies.

What is 'true AI'?

As the buzz around this branch of computer science has intensified over the past few years, terms such as AI, along with related techniques such as ML, have been bandied around freely. AI is now often applied to projects that purists would not regard as AI, even though they use some of the same tools, such as the statistical analysis of very large data sets. Given the levels of excitement among businesses, governments and the public about the potential of AI, attempts to cash in on this wave of interest and excitement are not surprising. Indeed, some might argue on this basis that algorithmic trading does not constitute "true AI", but is simply data analysis on a huge scale.

In the academic community there is a clearer consensus about what constitutes "true AI", says Prof Barber. "AI for me is about the ability to replicate human perception and reasoning and our abilities to interact with each other and the physical world."

Does this blurring of boundaries matter? Prof Barber argues that for academics it is important to be clear about the scope of the discipline, not least so that funders can understand what their money is paying for. But beyond that, other issues weigh more heavily: "What's more important [than rigid definitions] is that we actually make progress in making systems that are of practical utility for mankind, things that people find interesting and really make a positive difference to our lives."

Professor David Barber biography

David Barber is director of the UCL Centre for Artificial Intelligence, which aims to develop next generation AI techniques.

He has broad research interests related to the application of probabilistic modelling and reasoning.

He is also chief scientific officer for re:infer, which is a natural language processing startup that "turns unstructured communications into structured data to drive action."

He received a BA in Mathematics from the University of Cambridge and subsequently a PhD in Theoretical Physics (Statistical Mechanics) from the University of Edinburgh.

An academic perspective: Case study

AI POWERS THE SPREAD OF INTELLIGENT ROBOTICS – AND DEFINES ITS LIMITS

Robotics brings AI into the physical world, with developing prospects including autonomous vehicles, carebots, surgical robots and cobots. Interview with Dr Ali Shafti, Senior Research Associate in Robotics and AI, at the Brain & Behaviour Lab, Imperial College London.

QUICK READ

- **Robotics transmits the data-processing and decision-making capabilities of software into the physical world.**
- **Autonomous vehicles are robots and therefore stand out as arguably the most economically significant field of robotics research.**
- **Care robots and surgical robots are other key areas of development.**
- **So-called “co-learning” is a potential way for robots to learn more about the context in which they are operating and move closer to human intelligence.**

Dr Ali Shafti
Senior Research Associate in Robotics and AI, at the Brain & Behaviour Lab, Imperial College London

Robotics combines insights primarily from computer science, mechanical and electronic engineering and neuroscience. It aims to produce 'intelligent machines' capable of replicating humans' ability to sense the physical environment, interpret and make decisions based on those stimuli in real time, then translate those decisions into actions.

Looked at from the perspective of AI, robotics therefore transmits the data-processing and decision-making capabilities of software into the physical world. In robotics, as in other areas where AI is being applied, the dominant forms of AI are based on the data-driven techniques of machine learning (ML), which have experienced a period of greatly accelerated development over the past 15 years, as discussed in our interview with Prof David Barber.

"It's this ability to act back on the physical environment by moving something or making something happen that's important. That defines a robot," says Dr Ali Shafti, Senior Research Associate in Robotics and AI, at the Brain & Behaviour Lab, Imperial College London. "Up until that point the machine's no different to a computer or a smartphone."

Based on this definition, autonomous vehicles are robots and therefore stand out as arguably the most economically significant field of robotics research. However, the challenges of developing AI sufficiently powerful enough to allow robots to operate alongside humans in the highly complex environments that humans inhabit are significant. The dream of vehicles that can pilot themselves through a city rush hour to carry us home is more distant than some of its advocates care to admit.

Autonomous vehicles push the limits of AI

Rapid recent progress in ML has fuelled huge interest and investment in the development of autonomous vehicles, an area both carmakers and the world's biggest tech companies are aggressively pursuing. But producing fully autonomous vehicles represents one of the greatest challenges for robotics researchers, notably because of the difficulty of developing the AI required to control them.

Dr Shafti suggests there will be big advances in autonomous vehicles over the next decade but cautions that the vision of robots driving independently alongside normal traffic and without human safety drivers on board is still decades away. The key problem, he says, is that deep learning, the variety of ML

> "It's this ability to act back on the physical environment by moving something or making something happen that's important. That defines a robot. Up until that point the machine's no different to a computer or a smartphone."

“Some of the best outcomes of deep learning have been in computer vision, particularly in real-time object recognition and labelling, which is an essential module for many autonomous systems. But these systems can make mistakes and are easily fooled.”

that dominates the visual recognition systems used in autonomous vehicles, is reaching its limits. This issue is made more difficult because so much driving takes place in extremely complex, densely populated environments that were not originally designed to accommodate cars.

“Deep learning has taken us a long way forward, but the progress is slowing down now. It’s plateauing. Some of the best outcomes of deep learning have been in computer vision, particularly in real-time object recognition and labelling, which, of course, is an essential module for many autonomous systems – one example being self-driving cars. But these systems can make mistakes and are easily fooled.

“There’s a famous example where, if you put a few small stickers or a small graffiti drawing on a stop sign, it will be mistaken for other signs such as speed limits. That wouldn’t happen with a human because we understand the context. The system doesn’t – it’s just looking at pixels. It’s not intelligent beyond the very specific task it is trained for, so it sees a slightly disfigured sign and is easily fooled into thinking it is not a stop sign.

“Having an autonomous car operating in the same environment as a non-autonomous car is a very difficult problem. There’s a lot of talk but there are no real examples of an autonomous car operating extensively without a safety driver in a mixed environment. That alone, to me, shows that we’re way behind on this.”

The most likely intermediate stage of development is to designate certain traffic lanes or zones of cities for autonomous cars in order to avoid the challenges of letting robots drive alongside humans, he suggests. In the long term, however, he believes that the shift to autonomous cars will greatly reduce injuries and deaths on the roads and will create major benefits by enabling autonomous vehicles to communicate with each other. This will allow for optimal traffic management, as vehicles are all connected as a network and in constant communication, leading to the possibility of higher vehicle density while retaining efficiency and speed.

“Imagine arriving at a multistorey car park, getting out of your car and leaving the car to go and park itself. That’s a lot of time saved and a lot of optimisation gained. Cars can park much tighter to each other because, when you want to summon your car, the others will all move out of the way so it can come out.”

The life robotic: global industrial robots

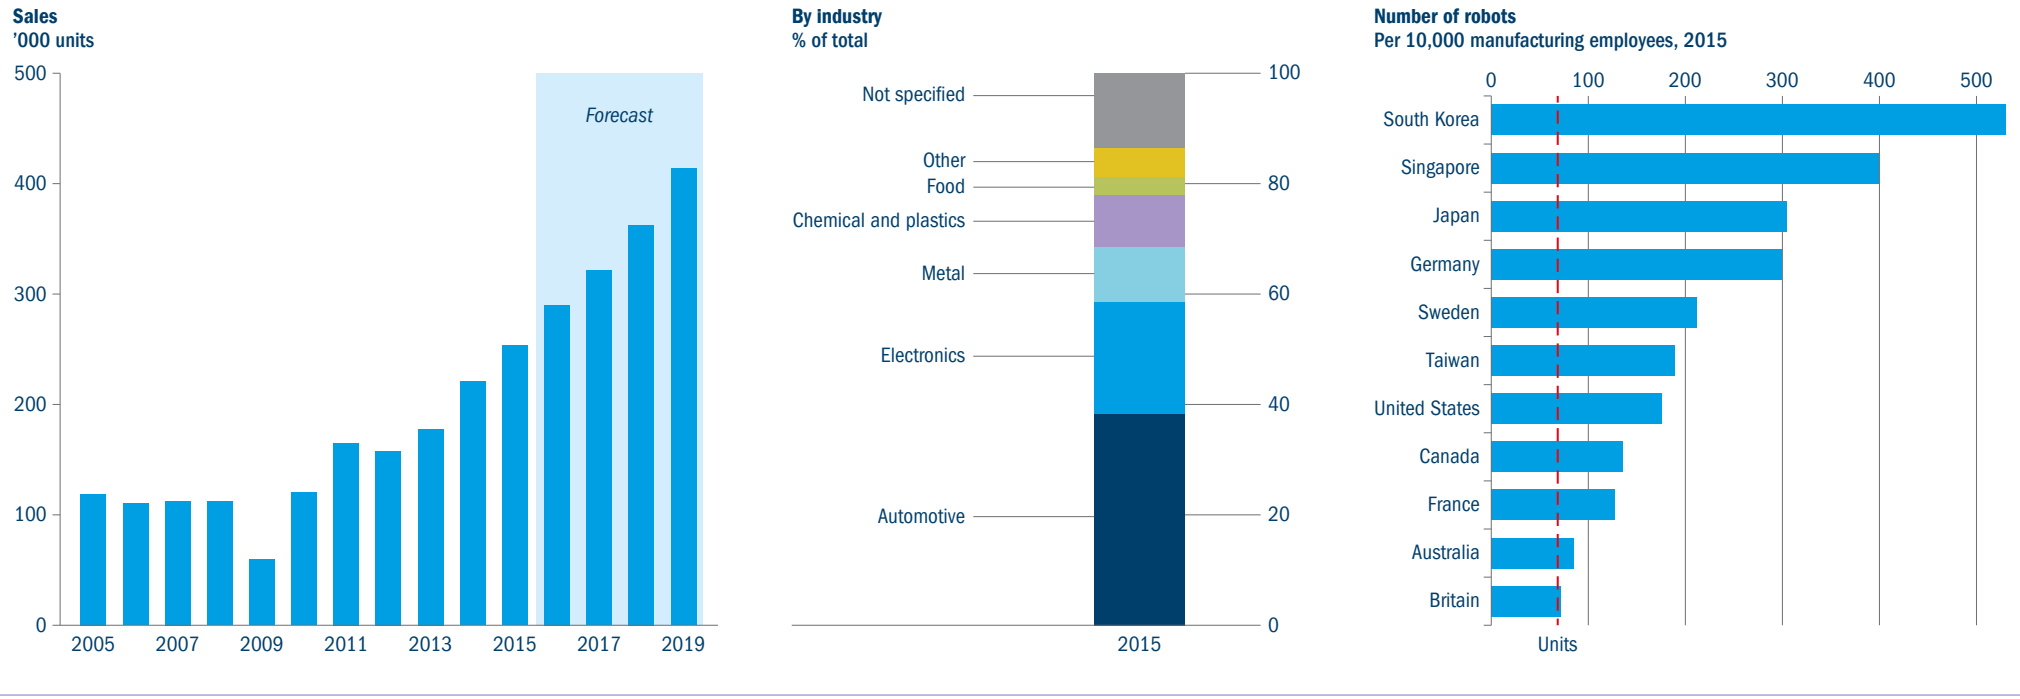

Source: International Federation of Robotics.

> “There’s an accelerated push to develop intelligent social robots and I think in the next decade or so we’re going to start seeing a lot of older people having these types of robotic systems in their home.”

Other key areas of development in robotics

Robots have been used in industry for decades, but more recently they have started to advance into other real-world settings. Two areas of development stand out.

Social and care robots

Dr Shafti believes that robots designed to interact with and monitor people who are lonely or suffering from conditions such as dementia will start to spread in the near future. This is one of the few areas of robotics where it will be relevant to create anthropomorphic whole-body robots, which many wrongly assume is a central goal of researchers in the general field of robotics.

The ability to use natural language processing to enable conversation between humans and robots is central to this area of robotics and recent advances have brought the widespread use of social robots significantly close, he says. “There’s an accelerated push to develop intelligent social robots and I think in the next decade or so we’re going to start seeing people suffering from loneliness or social anxieties, as well as those with conditions such as dementia, having these types of robotic systems in their home.” As well as providing company, these systems will be able to monitor human behaviour and aid people with declining cognitive abilities, for example, by reminding them to take essential medication.

Surgical robots

In laparoscopic surgery, robots are becoming established, the leading producer being Intuitive of the US, whose Da Vinci machines are the most advanced on the market.

These master-slave systems allow a surgeon sitting at a terminal to make extremely precise movements, translating a hand movement of several centimetres into a far smaller movement of the surgical instrument inside the patient. They also provide multi-tool instruments, allowing surgeons to operate through a single keyhole incision, rather than the three required in manual surgery.

Dr Shafti says current computer vision research in surgical robotics has been focused on issues such as enabling

3D vision and automated recognition of organs and features or defects from camera images. This allows surgeons to gain a lifelike view inside the patient, instead of the 2D, hard to recognise screen images they currently rely on. Researchers are also adding haptic feedback to these systems, so that surgeons can sense how hard or soft organs and tissue inside the patient's body are, possibly indicating the presence of a tumour, for example.

However, the major barrier to the take-up of surgical robots is likely to be the professional conservatism of senior surgeons, he says, who have spent decades operating manually and prefer familiar methods. Ultimately, surgeons are likely to be trained to operate both manually and using robots, at which point adoption would spread more rapidly.

Collaboration vs replacement

Robotic arms have been used in industry for years, but only in the past decade have they become safe enough to leave the segregated areas they had previously been confined to and operate in the same space as humans. Advances in sensing and mechanical engineering to prevent potentially fatal collisions have made this development possible.

The result has been rapid growth in the adoption of these 'collaborative robots,' or 'cobots,' in industry over the past few years. Key producers include Universal Robots of Denmark, Munich-based Franka Emika and Kuka, a Chinese-owned company also based in Germany. The latter is an established maker of old-style industrial robots, but has also moved into cobots. The arrival of robots that can safely work next to humans

Surgical robots market

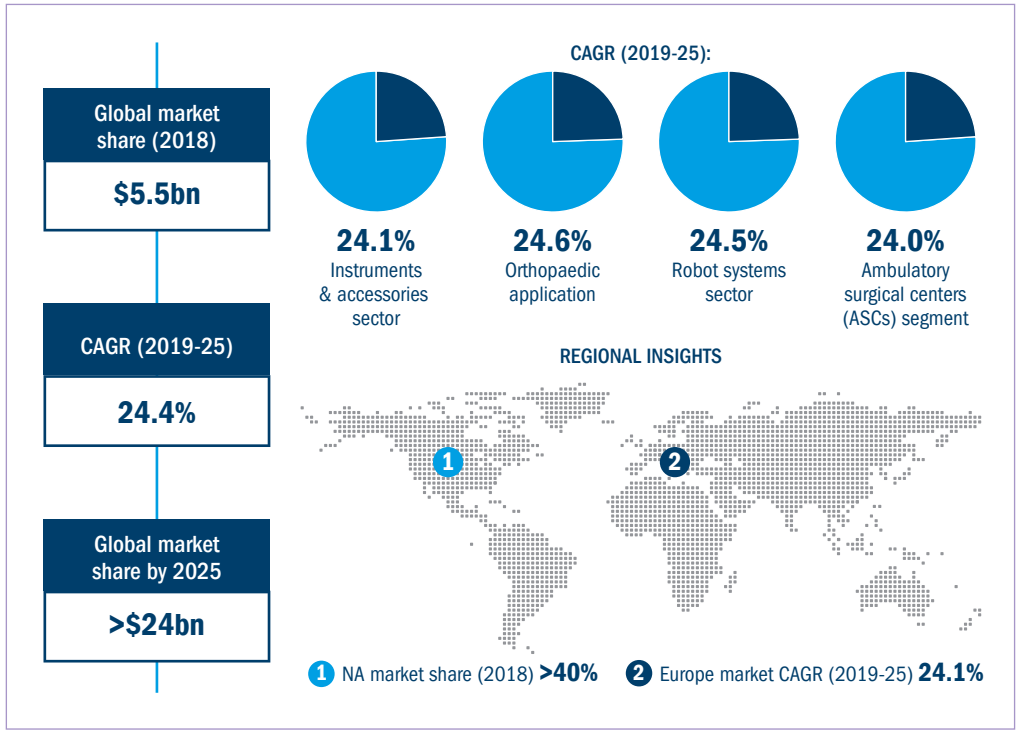

Source: https://www.gminsights.com/industry-analysis/surgical-robots-market. March 2019.

“I believe the next big thing is having humans more involved. In human-in-the-loop methods, you don’t put humans aside and have everything run autonomously end-to-end. The system runs autonomously but there’s a human within the loop being asked to monitor and intervene, which allows you to optimise for the human.”

represents a major advance in the technology and opens up large new markets for robotics manufacturers in many more industrial settings.

However, the appearance of collaborative robots also opens the way for robotics to develop using the ‘human-in-the-loop’ methods that Dr Shafti specialises in. He argues that human-in-the-loop approaches provide better outcomes, both for society and in terms of the development of robotics.

“Deep learning is reaching its limits and people will have to come up with the next big thing,” he says. “I believe the next big thing is having humans more involved. In human-in-the-loop methods, you don’t put humans aside and have everything run autonomously end-to-end. The system runs autonomously but there’s a human within the loop being asked to monitor and intervene, or collaborate in real-time, which allows you to optimise for the human. This is the way to make progress happen faster and with fewer negative effects on human lives.”

Combining human and robot intelligence, Dr Shafti believes, will create an intermediate stage in the development of robotics, where tasks that humans are less suited to such as repetitive actions, heavy lifting and accurate, precise movements that are physically difficult or tiring can be performed by robots, while humans use their knowledge and experience to direct the activity. This means that less general robotic intelligence is required, which in turn means existing intelligent algorithms can already be adapted and deployed to work with humans towards better work environments, while reducing the need for high computer power and its resulting carbon footprint. “Trying to have robots learn and generalise how to work in factories is not yet feasible with end-to-end deep learning approaches, and would also be very computationally heavy, consuming a lot of power. We should look into other methods.” he says.

This vision of the medium-term future of robotics envisages a two-way process in which robots augment human capabilities without fully replacing them, while over time the robots learn through the process of working alongside humans and become able to perform more complex tasks. So-called ‘co-learning’ is an area of keen interest among robotics researchers, who see it as a potential way for robots to learn more about the context in which they are operating and, in doing so, move closer to what we would think of as human intelligence.

Dr. Shafti biography

Dr. Shafti is a Senior Research Associate in Robotics and Artificial Intelligence with the Brain and Behaviour Lab at the Department of Computing & Department of Bioengineering, Imperial College London.

He studies physical collaboration and interaction between humans and intelligent robots – or embodied AI. He looks into making these interactions intuitive and natural for increased synergy, and augmented capabilities on both sides, leading to explainable, trustworthy and productive human-robot interaction. To this end, he implements machine intelligence in the context of robotics, while conserving the role of human intelligence as an essential part of the action/perception loop and the interaction.

He researches methods in robotics, machine intelligence and human behaviour analytics, as well as ways to integrate the outcomes through human in-the-loop methods. He has applied these findings in different scenarios, including collaborative robots, assistive robots, and autonomous vehicles.

Shafti has a PhD in Robotics from King's College London, where he focused on human in-the-loop physical human-robot interaction.

To find out more visit

columbiathreadneedle.com

COLUMBIA THREADNEEDLE INVESTMENTS

Important Information: For use by Professional and/or Qualified Investors only (not to be used with or passed on to retail clients). This is an advertising document. This document is intended for informational purposes only and should not be considered representative of any particular investment. This should not be considered an offer or solicitation to buy or sell any securities or other financial instruments, or to provide investment advice or services. **Investing involves risk including the risk of loss of principal. Your capital is at risk.** Market risk may affect a single issuer, sector of the economy, industry or the market as a whole. The value of investments is not guaranteed, and therefore an investor may not get back the amount invested. **International investing** involves certain risks and volatility due to potential political, economic or currency fluctuations and different financial and accounting standards. **The securities included herein are for illustrative purposes only, subject to change and should not be construed as a recommendation to buy or sell. Securities discussed may or may not prove profitable.** The views expressed are as of the date given, may change as market or other conditions change and may differ from views expressed by other Columbia Threadneedle Investments (Columbia Threadneedle) associates or affiliates. Actual investments or investment decisions made by Columbia Threadneedle and its affiliates, whether for its own account or on behalf of clients, may not necessarily reflect the views expressed. This information is not intended to provide investment advice and does not take into consideration individual investor circumstances. Investment decisions should always be made based on an investor's specific financial needs, objectives, goals, time horizon and risk tolerance. Asset classes described may not be appropriate for all investors. **Past performance does not guarantee future results, and no forecast should be considered a guarantee either.** Information and opinions provided by third parties have been obtained from sources believed to be reliable, but accuracy and completeness cannot be guaranteed. This document and its contents have not been reviewed by any regulatory authority.

In Australia: Issued by Threadneedle Investments Singapore (Pte.) Limited ["TIS"], ARBN 600 027 414. TIS is exempt from the requirement to hold an Australian financial services licence under the Corporations Act and relies on Class Order 03/1102 in marketing and providing financial services to Australian wholesale clients as defined in Section 761G of the Corporations Act 2001. TIS is regulated in Singapore (Registration number: 201101559W) by the Monetary Authority of Singapore under the Securities and Futures Act (Chapter 289), which differ from Australian laws.

In Singapore: Issued by Threadneedle Investments Singapore (Pte.) Limited, 3 Killiney Road, #07-07, Winsland House 1, Singapore 239519, which is regulated in Singapore by the Monetary Authority of Singapore under the Securities and Futures Act (Chapter 289). Registration number: 201101559W. This document has not been reviewed by the Monetary Authority of Singapore.

In Hong Kong: Issued by Threadneedle Portfolio Services Hong Kong Limited 天利投資管理香港有限公司. Unit 3004, Two Exchange Square, 8 Connaught Place, Hong Kong, which is licensed by the Securities and Futures Commission ("SFC") to conduct Type 1 regulated activities (CE:AQA779). Registered in Hong Kong under the Companies Ordinance (Chapter 622), No. 1173058.

In the USA: Investment products offered through Columbia Management Investment Distributors, Inc., member FINRA. Advisory services provided by Columbia Management Investment Advisers, LLC. Collectively, these entities are known as Columbia Management.

In Switzerland: Issued by Threadneedle Portfolio Services AG, Registered address: Claridenstrasse 41, 8002 Zurich, Switzerland.

In EMEA: Issued by Threadneedle Asset Management Limited. Registered in England and Wales, Registered No. 573204, Cannon Place, 78 Cannon Street, London EC4N 6AG, United Kingdom. Authorised and regulated in the UK by the Financial Conduct Authority. This document is distributed by Columbia Threadneedle Investments (ME) Limited, which is regulated by the Dubai Financial Services Authority (DFSA). For Distributors: This document is intended to provide distributors' with information about Group products and services and is not for further distribution. For Institutional Clients: The information in this document is not intended as financial advice and is only intended for persons with appropriate investment knowledge and who meet the regulatory criteria to be classified as a Professional Client or Market Counterparties and no other Person should act upon it. **Columbia Threadneedle Investments is the global brand name of the Columbia and Threadneedle group of companies.**