


Are you ready for the AI revolution?



JARMANY

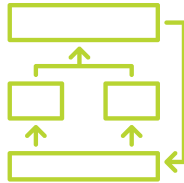


It seems that Artificial Intelligence is on the brink of impacting nearly all businesses that deliver their service or product using a digital channel. The question is how, what and when?

In this report we have pulled together the thinking of Microsoft and the Government and added our take on what lies ahead. If you are at the stage where you want to know more about AI; want to know how to get started, how you can combine AI and your people, how to implement an AI solution and what the benefits could be then read on.

We hope you enjoy the report and it piques your interest in AI, if it does then please get in touch. We are a specialist data analytics and insights agency and count ourselves fortunate that some of the world's largest organisations trust us to manage their data and help them predict needs, deliver efficiencies, drive growth and connect people.

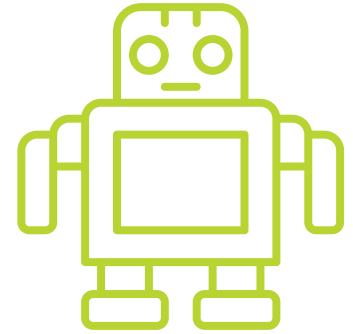
Keith Varty
Director



Contents

Introduction	4
1. AI in action	5
2. What AI is and isn't	7
3. Coming clean	10
4. Addressing prejudice	11
5. What skills will you need	12
6. How to get started	13
7. Trusting the machines	14
8. Will AI make us more productive?	15

Introduction



Aware of it or not, Artificial Intelligence is already deeply embedded in our lives and it's access to large quantities of data that is fuelling this boom. AI is by most people's assessment both an opportunity and a threat, so what should you be doing to allow your business to take advantage of AI whilst ensuring it complements the people you value so highly?

Firstly, let's give it a definition that at least for this report we can use as our guide:

'A set of technologies that enable computers to perceive, learn, reason and assist in decision-making to solve problems that are similar to that of humans'.

The last 10 years have seen rapid advances in AI solutions thanks to the confluence of three factors:

And it's machine learning that's the real gateway to AI, enabling computers to learn without being explicitly programmed. Instead they use algorithms – lines of code to spot patterns in data and then behave in a predictive way.

1. Ubiquitous cloud computing
2. Vast amounts of data
3. Breakthroughs in machine learning algorithms



1 AI in action

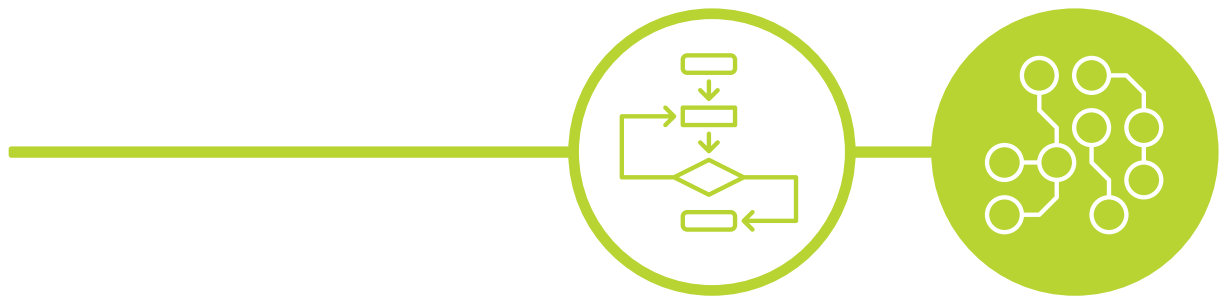
Speech recognition, natural language recognition, search recommendations are all now very much in the public domain. More advanced AI applications use advances in the field of deep learning, this is a type of machine learning that is inspired by how neural networks in the human brain process information. So, in seeing a picture of an object, the machine will first detect a shape from a matrix of pixels, then it might identify the edges of that shape, then contours, then the object itself and so on until it identifies the image.

Medical start-ups are using this technology to read large numbers of X-rays, MRI's and CT scans more rapidly and accurately than a radiologist – meaning they can get to the patients who need emergency treatment quicker and put those who don't minds at rest in a similar timeframe.

What is becoming quite apparent is that the immediate potential of AI is in it's ability to improve the speed and accuracy of vital processes and tasks.

Consider the following...

1. You wake up, refreshed, as your phone alarm goes off at 7:06am, having analysed your previous night's sleep to work out the best point to interrupt your sleep cycle.
2. You ask your voice assistant for an overview of the news, and it reads out a curated selection based on your interests.
3. Your local MP is defending herself—a video has emerged which seems to show her privately attacking her party leader. The MP claims her face has been copied into the footage, and experts argue over the authenticity of the footage.
4. As you leave, your daughter is practising for an upcoming exam with the help of an AI education app on her smartphone, which provides her with personalised content based on her strengths and weaknesses in previous lessons.
5. On your way to work, your car dashboard displays the latest traffic information, and estimates the length of your journey to the office, based on current traffic conditions and data from previous journeys.
6. On arrival, you check your emails, which have been automatically sifted into relevant categories for you.



7. A colleague has sent you several dense legal documents and software automatically highlights and summarises the points most relevant to a meeting you have later.

8. You read another email, sent by your partner, asking if he can borrow your bank login details to quickly check something. On closer inspection you decide it is probably a fake, but still, you hesitate before deleting it, wondering briefly how the spammers captured his writing style so unerringly.

9. You have other things to worry about though, as you head to a hospital appointment. However, after a chest x-ray, you are surprised when the doctor sits you down immediately afterwards, explaining that you look to have a mild lung infection— you had expected it to take weeks before the results came back.

10. Your relief is short lived—a notification on your phone warns you of suspicious activity detected on your bank account, which has been automatically stopped as a result.

11. You call the bank, and someone called Sarah picks up, and helps you order a replacement card. Except, you

soon realise, Sarah is not human at all, just a piece of software which sounds just like a real person

12. You are a little unnerved you did not realise more quickly, but still, it got the job done, so you do not particularly mind. After a quick detour to the local supermarket, where the products on the shelves have all been selected automatically based on previous customer demand, current shopping trends and the likely weather that day, you drive home.

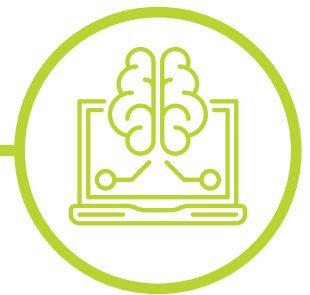
13. On your way back, your car detects signs that you are feeling slightly agitated, and chooses some music you have previously found relaxing

14. After dinner, you and your partner watch a film suggested by your TV, which somehow strikes just the right note for both of your normally divergent tastes

15. After dozing off, your house, predicting you are asleep by now, turns off the bathroom light and turns on the washing machine, ready for another day.

Now that's quite a day right, but it works to demonstrate what AI driven processes could look like.

2 What AI is and isn't



Firstly, it's worth just providing some clarity on the most common used terms in AI:

Algorithm

A series of instructions for performing a calculation or solving a problem, especially with a computer. They form the basis for everything a computer can do and are therefore a fundamental aspect of all AI systems.

Expert system

A computer system that mimics the decision-making ability of a human expert by following pre-programmed rules, such as 'if this occurs, then do that'. These systems fuelled much of the earlier excitement surrounding AI in the 1980s, but have since become less fashionable, particularly with the rise of neural networks.

Machine learning

One particular form of AI, which gives computers the ability to learn from and improve with experience, without being explicitly programmed. When provided with sufficient data, a machine learning algorithm can learn to make predictions or solve problems, such as identifying objects in pictures or winning at particular games, for example.

Neural network

Also known as an artificial neural network, this is a type of machine learning loosely inspired by the structure of the human brain. A neural network is composed of simple processing nodes, or 'artificial neurons', which are connected to one another in layers. Each node will receive data from several nodes 'above' it and give data to several nodes 'below' it. Nodes attach a 'weight' to the data they receive and attribute a value to that data. If the data does not pass a certain threshold, it is not passed on to another node. The weights and thresholds of the nodes are adjusted when the algorithm is trained until similar data input results in consistent outputs.

Deep learning

A more recent variation of neural networks, which uses many layers of artificial neurons to solve more difficult problems. It's popularity as a technique increased significantly from the mid-2000s onwards, as it is behind much of the wider interest in AI today. It is often used to classify information from images, text or sound.

Now there are genuine concerns about AI being a replacement for everyday jobs and that being the case what will the people that do those jobs do next? The positive take is that whilst some of these jobs will disappear, others will spring up on the back of AI – that



said we are to a degree entering the unknown, so nobody really knows.

As of today, it is worth noting that AI is only capable of carrying out a single and specific task – having first been designed by a human. AI isn't capable of intuition, empathy or emotional intelligence.

You could therefore reach the conclusion that AI will only in fact make people better, smarter and freer to do the things that matter to them. Take the example of the burgeoning use of AI in back office systems to scan documents, autofill forms and perform a host of other low-level tasks. This frees up employees to focus on the more 'human' requirements of their job like critical thinking, empathy and creativity. So, AI may well grow to become a pervasive technology which underpins our daily existence.

PwC concluded in a 2017 report that up to 30% of existing UK jobs are at 'high risk' of automation by the 2030's. These risks are highest in sectors such as transportation and storage (56%), manufacturing (46%) and wholesale and retail (44%). However, economists believe that due to the additional jobs which are likely to be created through economic growth over this period, the net effect on employment is likely to be neutral.

But there's a flip side which we can see happening today and it's directly linked to the vast data churning systems of Facebook and Google and the social changes that are being driven by the data footprints they are exposing. Many say that without these two organisations, both the Trump election in the US and the Brexit vote in the UK would never have happened. The principle argument is this; that by revealing so much about ourselves (without really knowing) we have exposed so much personal data, that we can be targeted with only news and views that reinforce our own emotions and prejudices and often nudge us along to potentially more extreme positions.

And more opinionated media attracts more engagement which in turn attracts more media dollars!

'I don't believe either Facebook or Google consciously tries to manipulate your behaviour. I do believe they've created platforms that allow advertisers to do that' –
Roger McNamee, illustrious Silicon Valley Investor and writer of the book 'Waking up to the Facebook catastrophe'



This does threaten the long-term viability of such digital systems and as a result, privacy and security will continue to dominate business, governments and consumers alike. And large companies with this level of control over vast quantities of data must be prevented from becoming overly powerful.

Add to this, the research carried out by **Imperial College** London that states that children have reduced attention spans, shallower cognitive capabilities and experience a loss of identity because of time online and using social media. It warns that the idealised world represented on social media 'leads to many illnesses including eating disorders ... and serious mental illnesses. The increasing use of AI would they say only add to this problem.

Microsoft has recently been putting it's weight behind, 'responsible AI' – the principle is that when designing AI solutions, we should be very mindful of whether we actually want machines to act and speak on our behalf, their blueprint for responsible AI includes the following:

- **Fairness**
- **Reliability & Safety**
- **Privacy & Security**
- **Inclusiveness**
- **Transparency**
- **Accountability**



3 Coming clean

A number of experts believe that AI provides added impetus to the need to better educate the public on the use of their data and the implications for their privacy. Professor Peter McOwan, Vice Principal, Queen Mary University of London, thinks that AI systems have become better at automatically combining separate datasets, and can piece together much more information about us than we might realise we have provided. He cited as an example 'pleaserobme.com', a short-lived demonstration website, which showed how publicly accessible social media

data could be combined to automatically highlight home addresses of people who were on holiday.

As a result, there are those that believe that far more emphasis should be placed on individuals owning or controlling their own personal data. For example, Sir Tim Berners-Lee is currently working on 'Solid' - a proposed set of standards and tools, based on the idea of linked data, which would allow individuals on the internet to choose where to keep their personal data and how it should be used.



4 Addressing prejudice

There is a genuine concern that the prejudices of the past must not be unwittingly built into automated systems, and such systems must be carefully designed from the beginning. Access to data needs to change, so that innovative companies, big and small, as well as academia, have fair and reasonable access to data, while citizens and consumers can protect their privacy and personal agency in this rapidly evolving world.

To do this means not only using established concepts, such as open data and data protection legislation, but also the development of new frameworks and mechanisms, such as data portability and data trusts. Large companies which have control over vast quantities of data must be prevented from becoming overly powerful within this landscape.

The current generation of AI systems, which have machine learning at their core, need to

be taught how to spot patterns in data, and this is normally done by feeding them large bodies of data, commonly known as training datasets. These systems are designed to spot patterns, and if the data is unrepresentative, or the patterns reflect historical patterns of prejudice, then the decisions which they make may be unrepresentative or discriminatory as well. This can present problems when these systems are relied upon to make real world decisions. Within the AI community, this is commonly known as 'bias'.

AI can help us fix some of the bias as well. Humans are biased; machines are not, unless we train them to be. AI can do a good job at detecting unconscious bias as well. For example, if feedback is given in performance reviews where different categories of people are treated differently, the machine will say, 'That looks weird. Would you like to reconsider that?'



5 What skills will you need?

The World Economic Forum's 'Future of Jobs Report' demonstrates the change that is coming in the top 10 skills by 2022 and what's on the decline:

Trending, 2022	Declining, 2022
<ul style="list-style-type: none">• Analytical thinking and innovation• Active learning and learning strategies• Creativity, originality and initiative• Technology design and programming• Critical thinking and analysis• Complex problem solving• Leadership and social influence• Emotional intelligence• Reasoning, problem solving and ideation• Systems analysis and evaluation	<ul style="list-style-type: none">• Manual dexterity, endurance and precision• Memory, verbal, auditory and spatial abilities• Management of financial, material resources• Technology installation and maintenance• Reading, writing, math and active listening• Management of personnel• Quality control and safety awareness• Coordination and time management• Visual, auditory and speech abilities• Technology use, monitoring and control

6 How to get started



As with any new technology driven paradigm there is significant amount of inertia within all industries. Figuring out how to get started, identifying a problem or process to address, getting to grips with the costs and understanding any impact on people makes it easy for AI to be parked in the too difficult to do pile.

Whatever the answer, the key is to make the problem the starting point – not the AI. And before embarking on any serious process of AI based transformation it is imperative for an organisation to get it's data in order. A number of organisations are already on this journey and the advent of cloud computing has made capturing, storing and accessing it easier than ever before. Data integration, automation and predictive analytics are relatively well known and to some degree used.

But the process of managing data must be on the basis the data is clean; accurate, unbiased and legitimately sourced.

'Great AI needs great data. But if you start by asking what you have, you are predisposed to the solution. Much better to start with the business problem, then see if you have the right data to tackle it' – **Norm Judah, Microsoft Chief Technology Officer, Enterprise**

And that quote provides a great basis on which to get started with AI. Clare Barclay, Microsoft UK's Chief Operating Officer believes

that if you approach AI with 'excitement' you are more likely to reach a positive outcome. So, there you have it:

1. Establish the business problem
2. Have integrity in your data
3. Do the right thing
4. Approach with excitement

That's a great methodology and mindset to approach your AI project with!

Next up – figure out what it is you want AI to do for you.

Our founder Matt Rhodes firmly believes that you must first identify those 'decision led tasks' inside your organisation. What it means by this is; what are the decisions that are being made (as in the radiology example of earlier) that could have a simple algorithm applied to them to deliver that decision for you. You'll then need to ensure you have the data to support such a solution and if not, get access to it. So that's the machine doing the work for you, the AI piece is the action that delivers this – there's often confusion between machine learning and AI – think of it as the manifestation of the machine! That action could be as simple as an alert to do something or more detailed such as the size of the opportunity and what to do to achieve it.

You can see from this example how AI really can super charge you and your teams to achieve more.

7 Trusting the machines



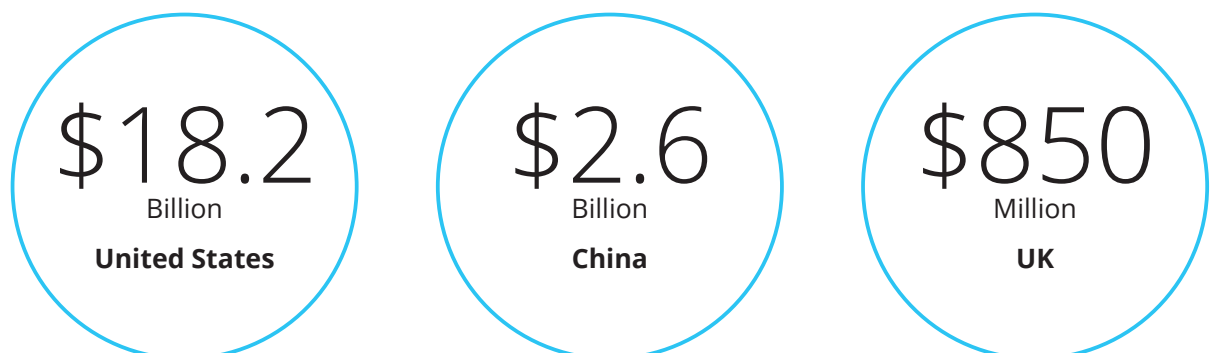
Investment in AI development

A number of countries are currently investing significant sums in AI research and development, with varying degrees of state support and co-ordination. Estimating the size of these investments across the public and private sectors is difficult, but according to Goldman Sachs, between the first quarter of 2012 and the second quarter of 2016, the United States invested approximately \$18.2 billion in AI, compared with \$2.6 billion by China, and \$850 million in the UK, the third highest investment by a country in this period. While the current United States administration appears to have no national level AI strategy, individual government departments are continuing to invest in AI, with the Department

of Defence, for example, spending approximately \$2.5 billion on AI in 2017. Meanwhile, China has explicitly committed itself to becoming a world leader in AI by 2030 and aims to have grown it's AI ecosystem to \$150 billion by then.

And in terms of the approach to AI; Germany for example is strongly influenced by it's flagship Industry 4.0 strategy for smart manufacturing. This strategy seeks to use AI to improve manufacturing processes, and to produce 'smart goods' with integrated AI, such as fridges and cars. As Professor Wolfgang Wahlster, CEO and Scientific Director of the German Research Centre for AI (DFKI), states, 'this is quite different from the US approach, which is based more on internet services'.

Approximate investment between the first quarter of 2012 and the second quarter of 2016:





8 Will AI make us more productive?

Hopes that AI will improve productivity must be set against a backdrop of low productivity growth across the developed world, and almost non-existent productivity growth in the UK since the 2008 financial crisis. The Royal Society for the encouragement of Arts, Manufacturers and Commerce (RSA) told us of 'lacklustre productivity levels, with UK workers on average 35% less productive than their counterparts in Germany and 30% less productive than workers in the US'. Sarah O'Connor from the FT said this was 'a puzzle that is taxing the best minds in technology and economics right now'. She also states that 'if you do not have productivity growing at a decent clip then you cannot have sustainable increases in living standards' and AI 'could mean a step change in productivity'.

The Institute of Chartered Accountants in England and Wales (ICAEW) said: 'we must not lose sight of the reality of most businesses, who are a long way behind in their adoption of many technology trends, including AI'.

Sage's research demonstrates 'that companies currently spend an average of 120 working-

days per year on administrative tasks. This accounts for around 5% of the total manpower for the average small and medium-sized businesses. The research suggested that the amount of time spent on such tasks is because digital tools have not been adopted, and that if UK businesses could be 5% more productive, GDP could increase by £33.9 billion per year.

In 2017, 5.7 million businesses in the UK were classified as SMEs (99% of all businesses), with 5.4 million of those employing fewer than 10 staff. Such micro-businesses accounted for 33% of employment and 22% of turnover. The CBI, which represents over 190,000 businesses in the UK, said: 'Digital innovations are at the heart of economic, social and cultural development across the UK. They drive productivity, help to raise living standards and lay the foundations for tomorrow's world'. If businesses are not using existing technology, SME's, we are concerned that the potential benefits to productivity offered by artificial intelligence could bypass significant portions of the business community in the UK.



There's also the subject of how you judge the success of an AI implementation. **Professor Neil Sebire who is a Pathologist and Chief Information Office at Great Ormond Street Hospital** provides a compelling example of how fears around AI are not always shaped by fact.

'Currently it is estimated that around 70% of medical conditions are diagnosed correctly first time in standard medical practice. This means about a third of first diagnoses are initially wrong. Patients seem to accept that. However, if we were to bring out an AI system and say one out of every five of the diagnoses were totally wrong, I am sure that it would be rejected as a terrible system, even though it has a better success rate than with a human doctor.'

Inherently and this may be a generational thing, there is an over-riding sense that people can be trusted more than machines.

And finally, there is a potential 'AI arms race' as various countries seek to develop more sophisticated AI, and potentially disregard concerns around safety and ethics in the process.

President Vladimir Putin's speech on AI in 2017 attracted attention worldwide, when he observed that alongside 'the colossal opportunities' it also brought 'threats that are difficult to predict', and that 'whoever becomes the leader in this sphere will become the ruler of the world'.

Sources:

HOUSE OF LORDS, *Select Committee on Artificial Intelligence – 'AI in the UK: ready, willing and able?'*

MICROSOFT UK – *'Maximising the AI Opportunity'*

