

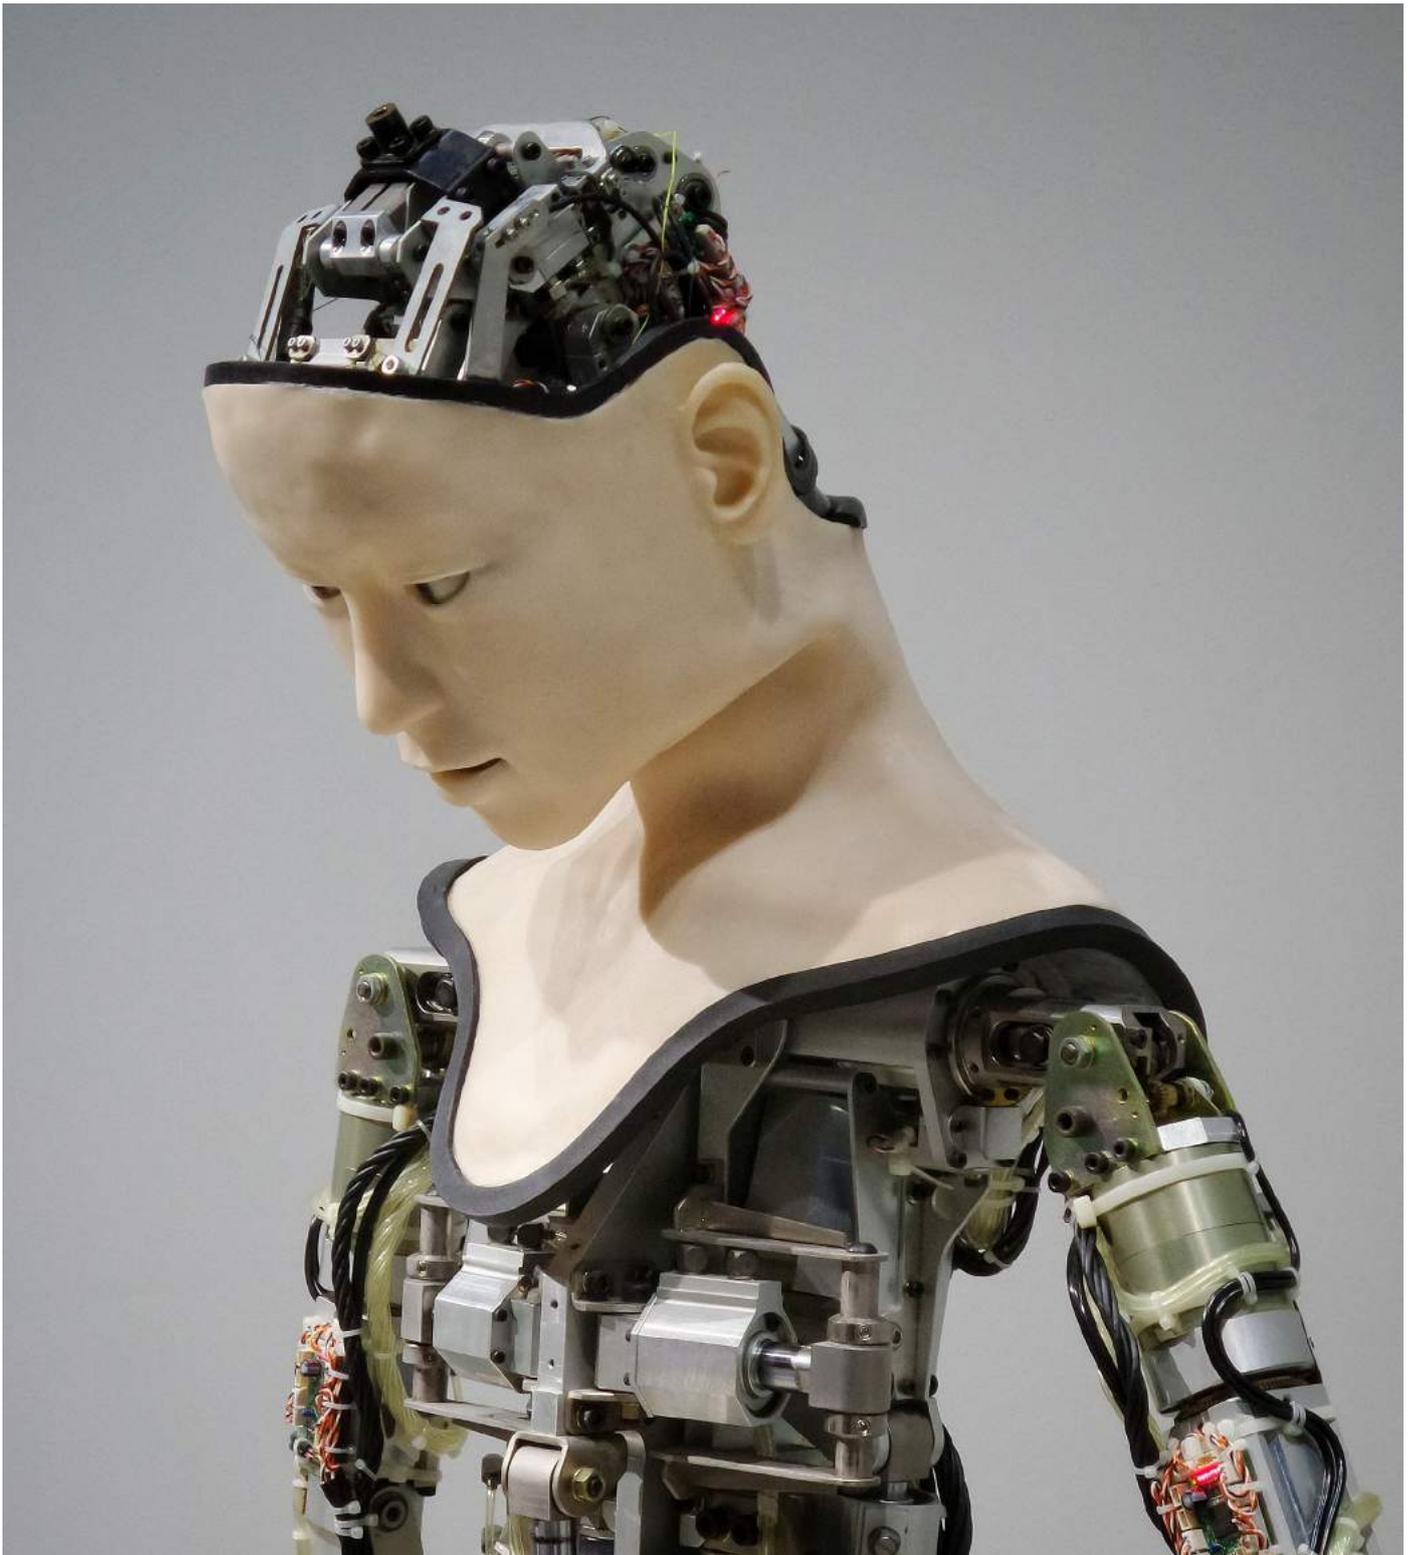


EVOLVE
ENABLING SUCCESS

ARTIFICIAL INTELLIGENCE

The past, present and future of artificial intelligence and its social and business implications

Written by Oliver Roberts

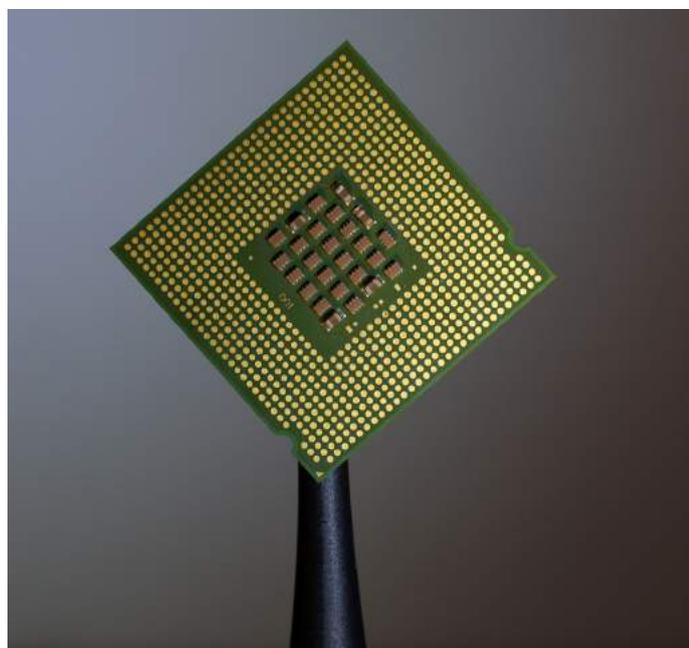


Whether it excites or unsettles you, artificial intelligence (AI) is developing and emerging faster than ever. In our highly connected and technological world, most of us use AI countless times every single day. In fact, it's highly likely you're carrying a form of AI in your pocket right now, perhaps you even spoke to Siri or Alexa this morning.

For all the speculation, excitement and reservations about AI, it is already here, it is already part of our daily lives, and its potential for positive disruption and change has already begun. From fingerprint/facial recognition and your incredibly accurate user-generated Google results, to the convenience of self-checkout counters at your local grocery store, AI makes your daily life easier, quicker and safer. And of course, the implications for your business are almost limitless.

For the past few decades, there have been countless starts, stalls and speculations regarding the escalation of AI; now it seems to be steadily flourishing and the social and professional ramifications are far-reaching and fascinating. By 2030 AI will contribute \$13-trillion to the global economy by 2030 and at least 70 percent of companies will have adopted some sort of AI¹.

The British AI industry is strong too. Our network of AI startups is the largest on the continent. Just on its own, London has three times more than any other city in Europe². And in 2014, Google bought Deep Mind Technologies—a London-based AI firm established in 2010—for £400-million. Later on, we'll look at a few of the ground-breaking AI startups operating in the UK.



There's a lot of noise about AI at the moment—Will it make humans redundant? How will it affect my industry? Will computers become conscious and take over the world?—so, from a personal and business standpoint, it's important to understand what AI is and what it isn't, where it's headed and what it can and can't do (for now).

This paper will touch on or expand on the following:

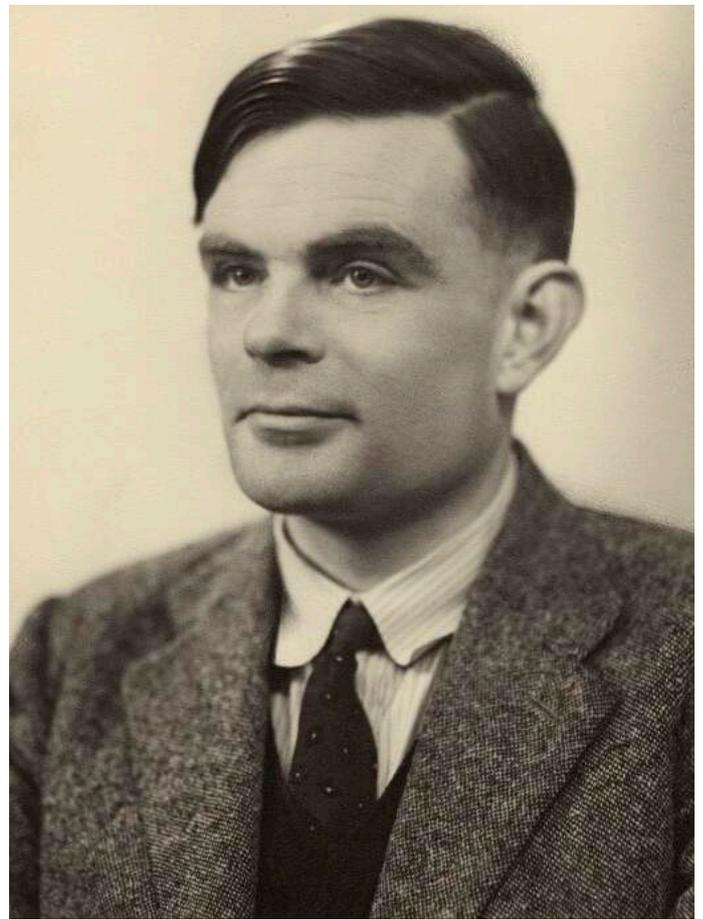
- The origins of AI and what it actually is
- Why AI is better than ever before
- The different branches and systems born out of AI
- AI systems and approaches at Google, Amazon, Microsoft, Apple, Netflix and Baidu
- Thoughts from people working at some of the UK's brightest AI companies and startups
- The future

The origins of artificial intelligence

American computer scientist John McCarthy is widely credited as the father of AI. He came up with the term in 1955, and a year later organised The Dartmouth Summer Research Project on Artificial Intelligence, at Dartmouth College in New Hampshire. Lasting six weeks, the conference was basically an extended brainstorming session among eleven mathematicians and scientists.

In his funding proposal for the conference, McCarthy stated the group study would: “...proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find out how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves... we will concentrate on a problem of devising a way of programming a calculator to form concepts and to form generalisations.”³

Many of the conference’s attendees would become leaders of AI research but the process was a rocky and uncertain one. Not long after some Dartmouth researchers predicted it would take no more than a generation for a computer to become as intelligent as a human, it became increasingly plain this was extremely far from being realised, if not impossible. The result was a withdrawal of funding from both the US and British governments in 1973 and a seven-year period known as an ‘AI winter’, during which little significant progress was made in the field.



“I believe that at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted.” – Alan Turing⁴

In 1980, thanks to funding from the Japanese government, there was a minor resurgence into AI but the relatively weak processing powers of even the most advanced computers thwarted progress. It wasn’t until around 2011, following the rapid development of faster computers and machine learning techniques that AI really began to show the promise that John McCarthy and his small cadre of engineers and scientists had been so excited about nearly fifty years before.

Though McCarthy and the Dartmouth conference were the apparent spark for AI’s realisation and development, AI history would be incomplete without mentioning Alan Turing. Apart from his code-cracking Enigma machine, his Turing machine is one of the earliest models of what we know as a computer.

Turing published a paper in 1950 in which he discussed creating a machine that ‘thinks’. From this, he devised his famous Turing Test—essentially a way of gauging a machine’s capacity for reason and logic that is on the same level as, or indistinguishable from, a human’s. The Turing Test is seen as the first serious proposal into the ideas and meaning of AI.

AI as an idea is older than you might think...

It's interesting to note that even though AI as a 'thing' first entered public consciousness 60-odd years ago, the notion of mechanical beings and the transference of mind into matter dates as far back as ancient Greece. The god Hephaestus—blacksmith of the gods—built robots or automatons to work for him and to serve and protect other gods. His self-animated Golden Tripods carried feasts to and from the gods, a bronze giant named 'Talos' patrolled Crete to protect it from invasion and Hephaestus made two golden maidens that followed him around and performed chores.

An automata is mentioned in the circa 3rd century BC Chinese Daoist text 'Liezi':

"It walked with rapid strides, moving its head up and down, so that anyone would have taken it for a live human being... it turned out to be only a construction of leather, wood, adhesive and lacquer, variously coloured white, black, red and blue. Examining it closely, the king found all the internal organs complete—liver, gall, heart, lungs, spleen, kidneys, stomach and intestines; and over these again, muscles, bones and limbs with their joints, skin, teeth and hair, all of them artificial⁵."

"It is unworthy of excellent men to lose hours like slaves in the labour of calculation which could safely be relegated to anyone else if machines were used." – Gottfried Leibniz⁶

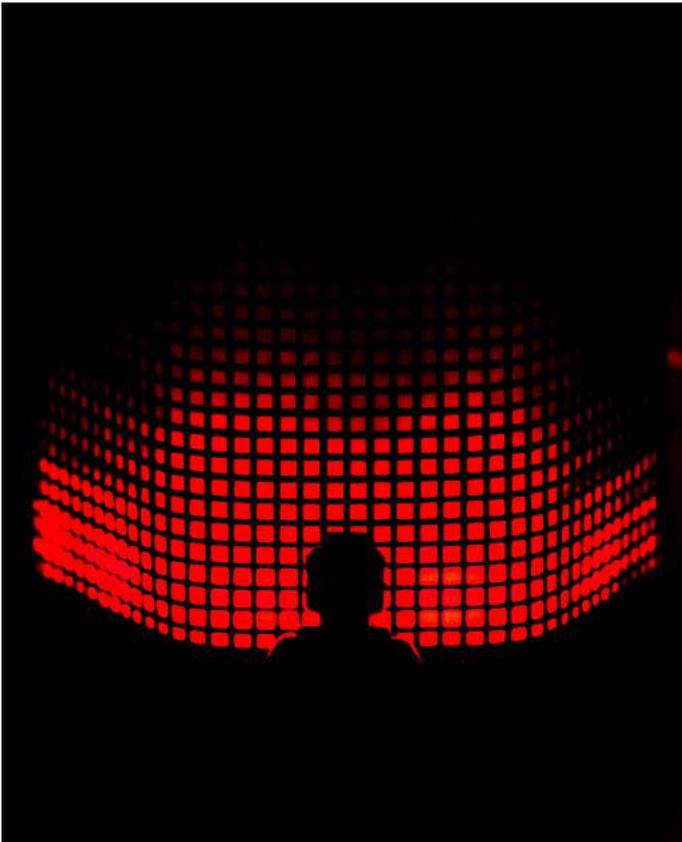


In 1206, Muslim polymath and inventor al-Jazari published 'The Book of Knowledge of Ingenious Mechanical Devices'. Apart from mere mechanical devices like pumps and camshafts, the book also featured a waitress robot endowed with human features.

And from the 17th century, philosophers like Gottfried Leibniz, Thomas Hobbes and René Descartes were exploring systematic reasoning, which sought to prove that we could reduce all rational, human thinking to numbers and symbols. This gave birth to the physical symbol system wherein symbols and expressions are processed to produce new expressions and induce intelligent action.

The basic premise of the physical symbol system is that human thinking is based on the manipulation of symbols (the expression of that manipulation being 'thought') and therefore by feeding the corresponding symbols into a computer processor and a similar manipulation occurring, the computer itself is said to 'think'. The physical symbol system gave birth to computer science, which laid the foundations for AI.

The present – Enter big data and machine learning



As we've seen, there have been surges of interest and promise in AI in the past, only for another 'AI winter' to descend. So what's different this time? What factors have led to the surge of enthusiasm and development we're witnessing? Out of all the positive drivers, one stands out: Big data.

Big data refers to ways of processing data that are too massive and/or complex to be dealt with by more traditional data-processing systems and software. Following the rapid development of the internet, the enormous data that was being collected needed new processing models. This gave birth to machine learning, an application of AI. It's no surprise then that Google was an early adopter and developer of machine learning and that Amazon, Microsoft and Baidu have joined Google as frontrunners in machine learning application and are all significant investors in AI. Giants like IBM, Oracle, SAP, HP and Dell have spent close to \$20-billion employing specialist firms to manage, analyse and apply their data through machine learning.



In terms of its ability to drive the economy, and its immense value as a resource, big data is being compared to oil. Its many applications include governance, healthcare, education, media, retail and sport. Huge amounts of data in these fields give rise to desirable outcomes from greater efficiency, personalised service and precise marketing, to advanced systems management, targeted training and simulation.

Big data is key for advancing machine learning because in order for anything to learn it needs information, lots and lots of it. And the more it knows, the more it's able to understand concepts, make connections and evolve. Feeding big data into an advanced AI is very similar to a person making their way through school to university and mastering a specific subject or skill—their brain receives almost infinitesimal information, sorts through it, keeps what it needs, discards what it doesn't, and the result is higher intelligence, the ability to analyse and interpret a variety of facts, and, ultimately, autonomy.

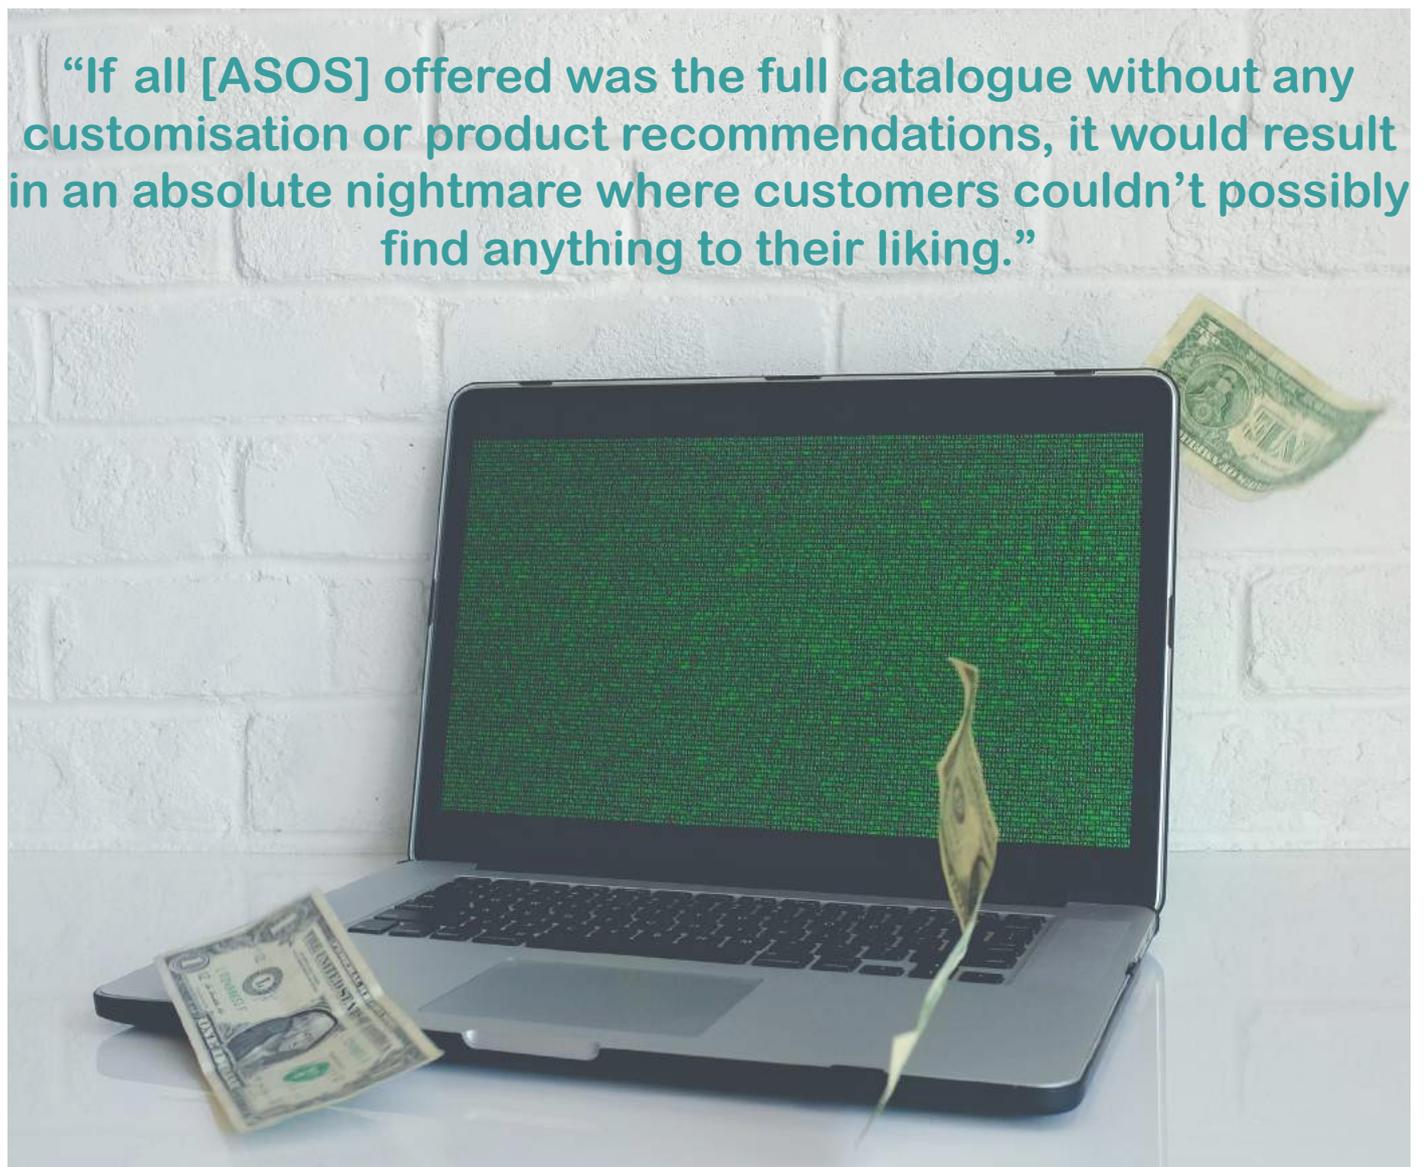
In a 2019 article on Forbes.com⁷, Roy Yammer, co-founder and CEO of GlassesUSA.com—the fastest growing online glasses retailer—writes about how consumers increasingly expect a personalised and intuitive shopping experience attainable only through data collection and analysis. If you can't produce this personalised experience, you are very likely to lose customers to a competitor. It is therefore imperative for brands not only to use big data, but to fully understand it so they can convert it into sales.

Yammer says:

“Each user who visits your website provides individual, yet anonymous information that allows you to group them based on similarities, such as demographics, geographic location, etc. Before coming up with any selling techniques or marketing strategies, laying the groundwork begins with knowing and categorizing your customer.

“Recommending products based on [the consumer's] search history and purchases brings you another step closer to sealing the (next) deal. A classic example is ASOS, which stocks tens of thousands of items on its platform. If all it offered was the full catalogue without any customisation or product recommendations, it would result in an absolute nightmare where customers couldn't possibly find anything to their liking. Compare it to shopping in a department store the size of Macy's in New York (or even bigger), but without clear sections, signs or any sales clerks to guide you.”

A small business owner doesn't really need to use big data on the same scale as one running a larger business, if only because of the cost involved in collecting and utilising it. However, all companies, big or small, will benefit from diligent data collection, such as through simple CRM software.



“If all [ASOS] offered was the full catalogue without any customisation or product recommendations, it would result in an absolute nightmare where customers couldn't possibly find anything to their liking.”

*Thoughts from... Darragh Jones –
CTO of Deputi.ai, a London-based
company using AI and blockchain
to help run and grow businesses*

What does Deputi.ai do?

We're building an AI that targets great businesses—SME scale-ups that are typically tech savvy—and growing scale by using AI to take care of the administrative burden of growing a company.

So if you think about all the things involved in running a company, for example managing your government returns—VAT or tax returns or Companies House documentation—we make sure you're not missing any deadlines and can help you prepare those documents. If you've got critical bills or clients coming up, the AI helps ensure you're not neglecting them over time. It's like automating your COO or admin manager. And we link your core business processes such as your finance and HR systems, through their APIs so it can check whether you've got an overdue bill or PAYE, etc.

What kind of AI systems are you using?

We use Slack as our mechanism that imbeds within the workflow of the businesses that we are working alongside. Then we use various bits of technology to build up a picture of your business. We hold your critical business documents in what's called your 'corporate memory' and that maintains your source of truth, all your key business documents. We use natural language processing to pull key information from the documents, such as dates and clauses, and we store those in our database. Textual classification and machine learning is at the core of this.

Where do you see Deputi.ai in five years' time?

We think it's still early days in terms of using AI in the SME growth sector, but critical to that is the fundamentals of getting your documents in line and your business processes boarded onto a platform which you can start to automate various areas of your business. In five years' time we want to be central to



that ecosystem. We're effectively building a platform that you can plug different components into that will drive automation across your business. At the moment, there's nobody out there that's really targeting the SME growth businesses. Almost all the case studies you see are done with the likes of Walmart and Asda and banks, big firms that have big budgets, and large data tech. Small businesses struggle because very often it's data that's in their filing cabinets or strewn across emails of three or four people within the organisation; we're trying to pull all that data together, unify it and then build an AI platform on top of it.

How does the current level of AI development and application compare to what's happening in the US?

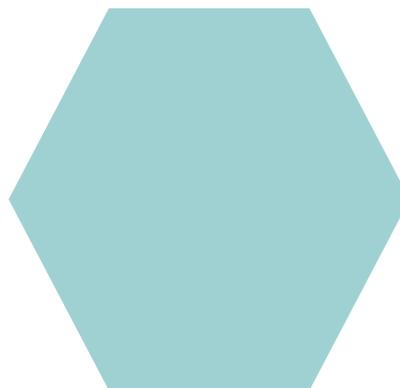
From a UK government perspective there's a lot of interest and a lot of work going on in terms of developing AI, machine learning and all these progressive technologies within the business environment. We're working with the Department for Digital Culture, Media and Sports, which has an AI department with-

in it that focuses particularly on how to encourage AI within business in the UK. So the government is putting a lot of thought behind it and we're starting to see that come through in terms of grants, innovation funds and other bits and pieces. In terms of funding, the US is always going to be a little bit more favourable, although that is changing. We were in an accelerator for blockchain-based companies called State Zero, and they were offering grant money for startups. So there's encouraging signs that things are changing in the UK around private money, and a couple of bigger UK firms are leading the way in terms of business process automation. It's a young industry and it's a healthy environment, part of it is an education and process.

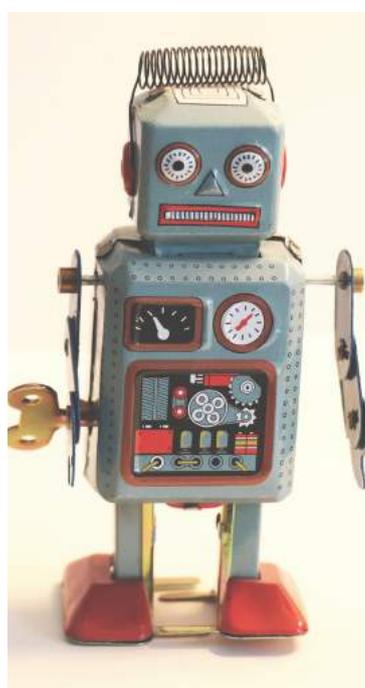
What about the social implications of automation—does AI pose a threat to jobs?

AI isn't here to take away people's jobs, far from it. If you look at the stats, AI is projected by 2035 to add over £630-billion to the UK economy, so it's a massive growth area and that comes in the form of either productivity gains or cost-savings; this area is growing and it's huge and it's only going to become more and more important. When we're pitching we look at it in terms of freeing up businesses from the tedium of business admin, allowing those business owners or CEOs to focus on being their most creative, strategic, innovative selves. If they're focusing on their business we're going to see much more growth in the UK economy.

“From a UK government perspective there's a lot of interest and a lot of work going on in terms of developing AI, machine learning and all these progressive technologies within the business environment.”



“AI is projected by 2035 to add over £630-billion to the UK economy, so it's a massive growth area and that comes in the form of either productivity gains or cost-savings.”



What is machine learning?

Because anything to do with AI at the moment is a hot topic, there's a tendency to throw a lot of terms about with little thought given to their inherent differences from AI. Machine learning is one such term.

As mentioned above, machine learning is an application of AI rather than a form. The simplest way to explain the difference is this:

- Artificial intelligence is evidenced by a machine designed to perform a given task or tasks. The machine, in turn, thinks and makes decisions in a way that mimics human cognisance, but only within its programed parameters.
- Machine learning is evidenced by a computer system that can perform specific tasks without explicit human instruction. So while the machine will have been designed and made by a human, it can learn and evolve independently according to its input. It is because of big data and the incredible processing powers of modern day computers that machine learning is rapidly advancing.

IBM's Watson and deep learning

IBM's question answering (QA) computer system, Watson, is probably the most advanced example of the machine learning offshoot, deep learning. Conceived and developed in the mid-2000s the system was named after former IBM CEO Thomas J. Watson. In 2011 Watson competed in a three-round match of the US game show Jeopardy! Playing against two of the show's most successful contestants, Watson triumphed quite easily. Since then, the system's abilities have progressed, through programming updates, to go beyond being a mere QA system into software and hardware that can not only talk, read, learn and interpret but also 'see' and 'taste'.

To build Watson's knowledge, programmers feed millions of documents, including encyclopaedias, dictionaries, thesauri, news articles and novels into its software. In its current guise, it can process 500 gigabytes per second, the equivalent of one million books⁸. Isn't this simultaneously amazing and unnerving?

Watson uses what's called a 'distributed cognitive growth system', meaning that its 'mind' is constantly connected to the cloud, gathering data everywhere and every time it's used. The more people plugged into the system, downloading and uploading information, the more Watson learns and gets better at its job.

Watson's applications are vast and ever-advancing. As well as collecting and implementing data in healthcare, education, weather forecasting and advertising, businesses are using Watson's AI tools to boost growth, automate tasks and create new products and services. Locally, Nationwide and RBS use virtual agents powered by Watson to deal with customer queries⁹. Lufthansa recently used it to create an interactive ad campaign¹⁰. Clothing brand The North Face uses Watson's 'cognitive retail' know-how to ask customers a series of questions so they can find exactly what they're looking for online¹¹. In 2017, a Japanese life insurance firm utilised Watson to calculate policy payouts. The result was much more timely payouts and increases in employee productivity. However—and this is where AI sceptics' concerns come in—the implementation of the system led to the redundancy of 34 brokers at the firm...¹²



“Computing will never rob a man of his initiative or replace the need for creative thinking. By freeing man from more menial or repetitive forms of thinking, computers will actually increase the opportunities for the full use of human reason.” – Thomas J. Watson, Jnr¹³

Thoughts from... Patrick Davis – Software developer at Greenwood Campbell, a digital transformation company based in Bournemouth

“

AI has already dramatically changed the technical landscape in [several areas of finance]. For example Chatbots automate a large majority of customer service and banks use AI to combat fraud. The insurance behemoths are next up for more AI integration. Currently, masses of paperwork and contracts, historical records and analysis of trends lead to an incredible amount of stats. But someday soon, our belongings and home will be so totally intertwined with the Internet of Things that these finer details will just become second knowledge to companies.

For example, customers may one day not even need to be quoted for prices on their car insurance—the pricing system will already know what they have paid previously from just a few personal details (or even automatically gathered beforehand) due to pre-analysed data held by AI, exposing any good or bad habits. The future’s connectivity will be so advanced and accessible that AI can pretty much infinitely analyse data from drivers. Their real time actions could even effect their risk category and their driving record.

Furthermore, what of the smart home market? A compelling forward-looking insurance policy from London-based company Neos protects your home from pretty much any issue a smart home sensor could detect. You’re on holiday and a gas leak develops... no problem—a sensor notifies you and the insurance company via a smartphone whilst simultaneously arranging for a contractor to fix it as soon as possible. No one is hurt and nothing is lost, all because of the AI behind these sensors analysing data in real time.

By adopting AI and automation in this respect, insurance companies should expect to overcome automation in other areas. The idea of industries now moving away from what is termed categorical data towards source data means that companies know their customers better than ever before, and this is just the beginning.

”



Mimicking the human brain

One of the more exciting things to emerge from deep learning are artificial neural networks (ANN), which are connection systems loosely based on the neural networks and functions of our own brains. We say “loosely based” because while an advanced ANN might have thousands of ‘neurons’ (processor units), the animal brain has billions. And though most working with ANN are not trying to get networks to behave entirely like a human brain or perform its biological functions, researchers have, for instance, used ANNs to accurately simulate the workings of the human eye.

The model is based on a collection of artificial neurons. And like the animal brain, each neuron transmits information from one neuron to another using synapses or ‘layers’. To process an ANN, data is repeatedly passed from an input layer, through a hidden layer and then onto an output layer. It is in the hidden layer where the learning and decision-making takes place. An ANN is therefore not an algorithm but a structure for machine learning to process vast and intricate sets of data, allowing for more thorough and intuitive problem-solving. Deep learning is implied when there are multiple hidden layers.

Applications within deep learning and ANN are complex and profound and have led to significant advancements and disruption in finance and marketing in particular. The 4th annual Deep Learning in Finance Summit, held in London in March, sought

to explore the impact and corollaries of these disruptions. The post-event report listed five key takeaways and talking points from the summit:¹⁴

- Trust and responsibility are key to building a customer base who are loyal and confident in the product.
- Unbiased data ensures all customers are treated fairly.
- The importance of being able to explain to customers why a certain decision was made, particularly when dealing with a customer’s personal finances.
- Using deep learning to improve customer experience must not be underestimated or overlooked
- Collaboration between industry, academia and government bodies are really important in progressing research in all areas and bringing research to market.

Through predictive analytics, ANNs are enhancing marketing tools and strategies like never before. One example would be using an ANN to sift through and identify key trends from a previous marketing campaign—How many people did it reach? Who was the main demographic? Did it go viral?—and use the data to plan a new campaign and make accurate predictions about its outcome and provide unprecedented guidelines for content creation, social media strategy and product distribution.

“The main problem we’re trying to solve with machine learning is helping customers find the right answers to their queries and helping agents to diagnose and respond to customer queries swiftly. We’ve automated this by validation testing. We are continuously making predictions in production and we test to ensure that they make the expected predictions. If this fails, we get alerted.” – Neal Lathia of Monzo, speaking at the Deep Learning in Finance Summit





Case study: Netflix's thumbnail trick

One very interesting example of a company using AI and neural networks is the extensive research that Netflix did into the thumbnails for its programs. Did you know that the thumbnails you see for your favourite series and movies are not the same as someone else's? Just like the rest of the user-generated content on Netflix, the thumbnail artwork is tailored specifically to your tastes and watching trends.

You might wonder why, if Netflix already has a very effective recommendations system, they bother with the thumbnails. After all, given you might look at the thumbnail for two seconds before clicking on it to play the program, what difference does it make? As it turns out, quite a lot.

In 2014, Netflix conducted studies¹⁵ on the importance of a thumbnail and found that thumbnail artwork was not only the biggest influence on a

user's decision about what to watch, it also accounted for 82% of their focus while browsing. So, whether you're aware of it, that little square icon with the name of the program of it significantly determines your watching habits.

Netflix's concern is that, statistically, if a user's attention isn't grabbed within 90 seconds of logging on, there's a good chance he or she will go looking for content on another platform, such as Amazon Prime. Generating a compelling thumbnail is key to prevent this from happening. But how customisable can a thumbnail really be, and how is it created? Let's find out...

The average hour-long program is made up of around 86,000 frames, and to select the best, Netflix uses an application of AI called ‘aesthetic visual analysis’. The process can be broken down to two basic steps:

- 1) A program analyses each static frame of the film, and image recognition algorithms create meta-data according to the information gathered from each shot. This information includes things such as light, contrast, faces, skin tone, camera angle, tight and long shot, and instances of nudity. The best frames are then placed into a database where they’re grouped into three categories—visual, contextual and compositional.
- 2) The AI uses the database’s metadata to pick out the best and most suitable images. The basic requirements of these images—e.g. sharp, symmetrical and non-offensive—will already have been decided by Netflix and programmed into the AI. Then, a creative team (of humans!) add text and other visual elements to complete the thumbnail.

From here, relentless A/B testing determines which thumbnail to assign to each user, which is why a thumbnail for a program might change from day to day or even hour to hour, depending on what you watched. Your general preferences also come into play in very interesting ways. For instance, if your collected data shows you watch a lot of comedies, you may find a comedic actor showing up on a thumbnail for a drama he or she only has a bit part in. Or, if you’re a fan of romance, an image of two actors

kissing or holding hands could show on a thumbnail, even if the romantic element is a side-story to the main plot of the film. There are regional differences too, wherein a thumbnail’s artwork might reflect the general aesthetic of the country you live in.

However, when data mining and algorithms go as far as this, AI can either prove to be too smart for its own good, or show itself up for its lack of subjective understanding. This happened when certain demographics noticed their thumbnails featured actors of the same ethnicity, even if they only had a few minutes of screentime. The AI was only doing what it was supposed to, yet some users felt Netflix was being disingenuous or even patronising by matching thumbnails to a viewer’s ethnicity, especially if they’re considered a minority or marginalised group. This is a minefield that could easily spill over rather messily into gender politics and feminism.

This error aside—and one which will undoubtedly be smoothed out as the model improves—Netflix’s vast research and development into something as seemingly insignificant as a thumbnail is a great example of how you can use AI to not only identify and improve every minutiae of your business but also enhance your customer’s experiences and keep them coming back for more.

Here are the three main things the Netflix learnt about what makes a thumbnail a good thumbnail: Faces should be emotionally expressive and close-up, villains should be shown instead heroes and don’t have more than three characters in the shot.



“We have been working hard the past few years to create a framework that allows us to effectively intersect big data with creative, ultimately helping members discover stories they will enjoy even faster. As a result of that work, we now have the unique ability to understand how to most effectively tell our members why a story is right for them—all through a single image.” – Nick Nelson, former global manager of creative services at Netflix

Thoughts from... Tim Groot, co-founder and CEO of London-based Grip, a B2B event matchmaking app

If you look at the global trends, what we're moving towards, AI is already used a lot where there is huge amounts of data. I think what we're seeing now with B2B is that the data is becoming more structured and it's becoming more possible to start doing interesting things with that, slowly but surely.

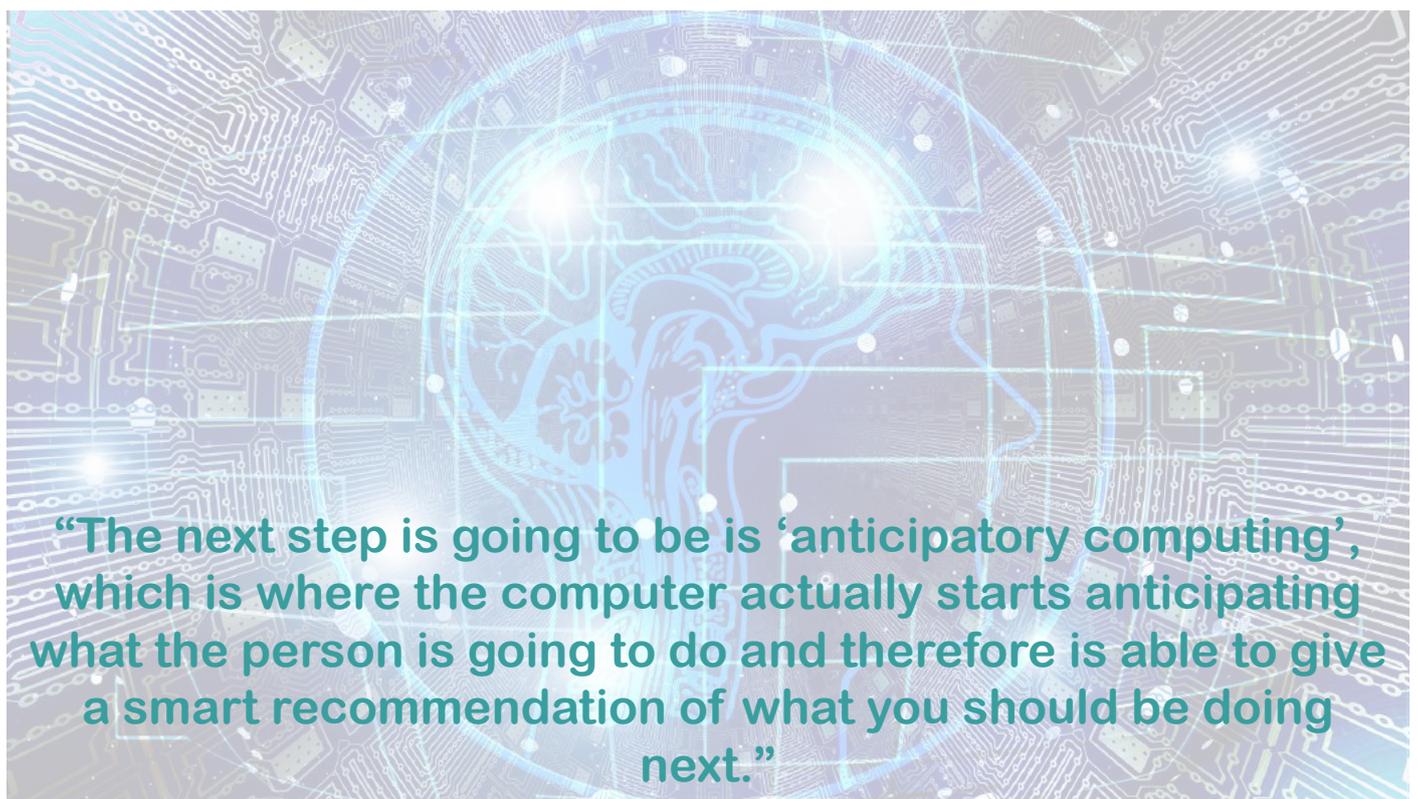
A lot of what we've seen up until now has been based on people pooling information from a system and then a system giving smarter recommendations. For example, Google search results are getting better thanks to machine learning. But I think what the next step is going to be is 'anticipatory computing', which is where the computer actually starts anticipating what the person is going to do and therefore is able to give a smart recommendation of what you should be doing next. So it could be a simple thing like an AI assistant predicting traffic conditions and taking this into account for when you should leave for your next meeting. Machine learning allows you to build smart predictions into that.

Using Grip for events, it becomes possible to really recommend—based on, meetings, connections, sessions, registration data, all these different data points—who is most relevant for you to meet with at an event. I think that's a big part of where AI will

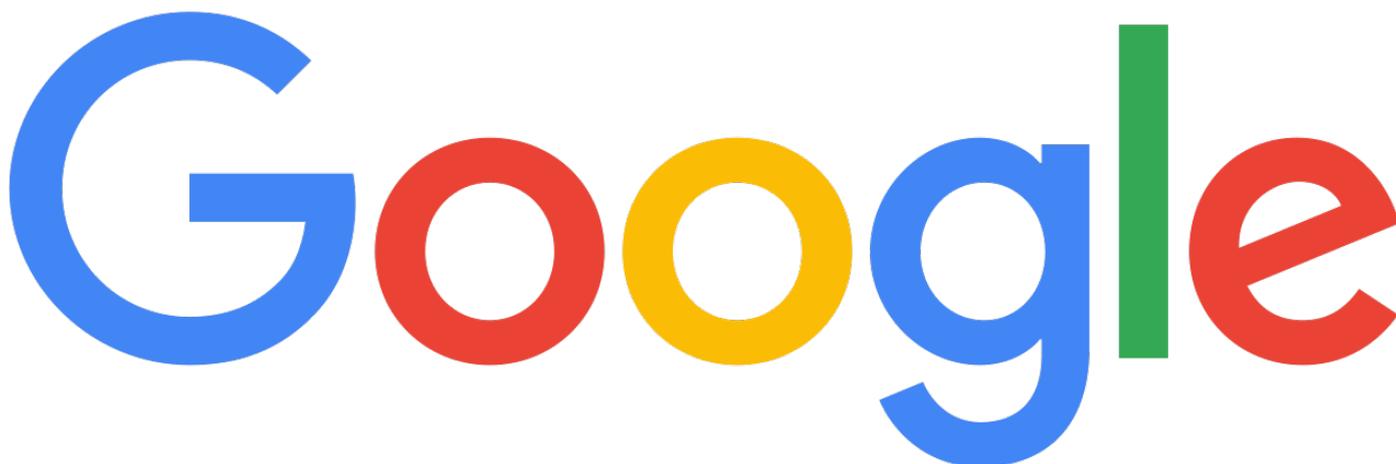
sit—giving recommendations almost before you've realised that you actually need them.

We've seen a real impact in terms of how many recommendations convert into meetings and how much higher that percentage is compared to what they've done in the past. Grip does another thing, called Must-Meets, where we recommend a small number of people that we think are very relevant for you at an event, and what we've seen is people saying, "One of these five people you recommended is someone that I know but I didn't know he was here so it's amazing that you were able to pick out that he is relevant for me to meet with." Face-to-face interactions are still such a crucial part of doing business. What we're doing is face-to-face matchmaking.

In the UK, data cleanliness is an issue. A lot of companies have a silos with a lot of different data that is not properly formatted and people don't understand the best practices when it comes to data collection and how they can set themselves up for future use. That's a very important thing to understand for businesses in general—you need to start being smart with your data right now so that, if not now, in the future you can take maximum advantage of it to improve your business processes.



What are Google, Amazon and Microsoft up to?



Google AI

Google AI is the laboratory from which Google conducts research to advance the technology. It has made clear its intention to make AI available to everyone in the world and wants to use its collected big data to solve problems, not only in everyday scenarios but also in the realms of healthcare and scientific discovery. By publishing its research and open sourcing its systems, Google AI is encouraging knowledge-sharing to address societal challenges. It has issued an open call for businesses and organisations around the world to submit ideas for using AI for good. Those selected receive Google's expertise and funding from a \$25-million pool.

Some of Google AI's business focus areas include:

- Algorithms and optimisation—Research into graph mining, large-scale SEO optimisation and market algorithms, all under the helm of machine learning. Interests include addressing online ad allocation problems, design for effective advertising exchanges and the development of real-time pricing for ad auctions.
- Cloud AI—Using the latest machine learning programs for data analysis and improving the interactions between humans and machines. Research and development of natural language processors that learn from their interactions with humans, thereby producing highly accurate insights to determine business processes.
- Language—Combining machine learning and natural language processing to research and refine areas such as questions answering, document understanding and summarising, and writing/content generation.

“We want to use AI to augment the abilities of people, to enable us to accomplish more and to allow us to spend more time on our creative endeavors.” – Jeff Dean, Google Senior Fellow

DeepMind

As mentioned in the introduction to this paper, DeepMind was founded in London in 2010 and purchased by Google in 2014 for a cool £400-million. Now part of Google's Alphabet group, DeepMind is, as its name suggests, heavily focused on the development of deep learning and deep neural networks.

DeepMind still operates out of its original London office, but has also expanded to include research centres in Canada and Mountain View in California. DeepMind's stated goal is to 'solve intelligence', not only so it can advance machine learning but also to gain greater knowledge about the human mind and consciousness. To date, some of DeepMind's more significant research has seen systems programmed

to learn to play video games from the 1970s and 80s, such as Pong and Space Invaders, using only raw pixels as its data input.

Through trial and error, and working with a rewards system, these systems have eventually mastered these games without human input.

In October 2015, DeepMind's AlphaGo—a program designed to learn and play the ancient Chinese board game Go—defeated a professional Go player five games to zero¹⁶. Go is considered a more difficult game for computers to master than, say, chess, because it features more possibilities than there are atoms in the universe¹⁷. AlphaGo defeated several other professional players before effectively defeating itself in 2017 with the improved AlphaGo Zero, whose self-taught strategies triumphed over the original system 100-0. To give you an idea of the speed of AlphaGo's AI development, AlphaGo took several months to learn its strategies, while AlphaGo Zero took just three days while using less processing power...



Amazon Web Services

Launched in 2002, Amazon Web Services (AWS) started out as a platform that offered some basic tools and services. Thanks to advancements in AI and cloud computing, by the end of 2018, AWS—a subsidiary of Amazon since 2006—reported an annual revenue of \$25.65-billion. To put the speed and enormity of this growth into perspective, consider that AWS's 2017 revenue was \$17.46-billion—that's an \$8.19-billion leap in a single year!¹⁸ No wonder Amazon CEO Jeff Bezos is excited.

He says:

“We are now solving problems with machine learning and artificial intelligence that were in the realm of science fiction for the last several decades. And

natural language understanding, machine vision problems, it really is an amazing renaissance¹⁹.”

Today, AWS provides a low-cost IT infrastructure platform that powers more than a million customers in 190 countries. Among its business applications is Alexa for Business, which is essentially an intelligent office assistant. Other products include:

- Amazon Comprehend, a natural language processor which uses machine learning to seek insights and connections in text, particularly unstructured data.
- Amazon Forecast, a business forecaster.
- Chatbot builder Amazon Lex.
- Amazon Personalize for customer engagement.
- Amazon Polly, which turns text into lifelike speech.





In January 2017, the first Amazon Go convenience store opened in Seattle, following a year of development. The concept behind Go is straightforward but its potential long-term effects are massively disruptive.

Under the watch and supervision of computer vision, deep learning and loads of sensors, a customer signs in via the QR code on the Amazon Go app as they walk into the store, buys whatever they need, and leaves without ever joining a checkout line of any kind. It works similarly to the self-checkouts found in most of the big brand convenience stores today, i.e. using sensors to detect when you've placed or removed an item onto/from the parcel tray, except with Go it's all in real-time (there are weight sensors in the shelves) and your Amazon account is charged as soon as you pass through the turnstiles to exit the store.

Seven months after the first store, a second one opened, also in Seattle. There are now several stores in Chicago and San Francisco, with one due for New York soon. According to a Bloomberg report, Amazon is planning to open as many as 3,000 stores in the US by 2021. Also in the pipeline are assemble-on-site Go compact stores for office lobbies and airports.

The only factor in the way of more expansion is the fact that each store costs about \$1-million to build, due to all the tech. This means it would take about

two years for Amazon to break-even. However, you can bet that Bezos and his team of geniuses at Amazon will soon come up with ways of overcoming this issue.

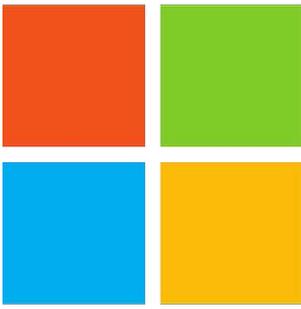
- In April, Sainsbury's in London's Holborn Circus opened a similar store, allowing shoppers to use an app to pay for products without using a tillpoint.

Robotic buildings

As of late-2018, 26 of Amazon's more than 175 fulfilment centres worldwide are considered 'robotic buildings'. Here, employees are assisted by not only by Amazon's AI and computer vision systems, they are joined by a hundred or so of the more than 100,000 robots Amazon deploys in its centres to move shelves of items around with no human input.

Kiva Systems originally designed the robotic units—which look similar to a robot vacuum cleaner—and Amazon bought Kiva in 2012 for \$775-million, its second-largest acquisition to date²⁰. The robots work by following a series of barcoded stickers on the floor and use sensors to avoid collisions.

Inside the Amazon fulfilment centres, each robotic unit calculates its pace in correspondence to each employee's historical packing and processing speed, thus avoiding any delays. They also self-charge and, if an item falls out of a pod or is misplaced, AI and computer vision will alert the robot.



Microsoft

In a 2019 Forbes report on AI in conjunction with Microsoft²¹, Mitra Azizirad, corporate vice president of AI marketing at Microsoft says: “The full potential of AI is realised when organisations foster a culture where AI experiences are democratised and scaled across the organisation. By providing every employee in an organisation with AI capabilities and self-service AI, employees are empowered with the insights most relevant to their roles at a scale and speed never seen before.”

In simplified terms, Microsoft AI is pushing the idea it is not only imperative for businesses to adopt AI right now, it’s essential—both for the future success of AI and the workforce—that the leaders of these businesses instil an AI culture within the organisation.

Fears about AI taking jobs needs to be replaced by the understanding that working together with AI will not only bolster the business but also emphasise and speed up the human talent that already exists within the business. A harmonious system that sees humans providing the input and AI processing the output could advance learning, insight and profits in

unprecedented ways.

“You have to understand the institutional fears around AI and design comprehensive, long-reaching programs to prepare your workforce. You can’t let it happen accidentally—you need to deliberately plan for what your company will look like in the next generation,” says Simon Thomas, Global Head of Data and AI, Avanade, in the same report. “The dynamic of the human-machine relationship is not just about how AI will influence workers; it’s how workers will influence AI.”

Dr. Athina Kanioura, Chief Analytics Officer and Global Lead, Accenture Applied Intelligence, remarks: “There will be jobs where humans complement machines. For example, where workers train machines to perform tasks, explain the machine outcomes and sustain the machines in a responsible manner.” The report states that establishing this accord between man and machine, instituting data supply chains and finding the right partnerships to help you keep up with the rapidly growing advances in the field are key to building and maintaining AI success within a business.

“We’re right at the point on the growth curve before it goes vertical. Organisations across the board will make significant investments in these technologies as they look to scale across their organisations and markets. They know there’s a lot at stake. The time for companies to act on AI is now—or risk quickly falling behind and losing competitive advantage.” – Dr. Athina Kanioura, Chief Analytics Officer and Global Lead, Accenture Applied Intelligence



AI for Good

Microsoft's AI strategy very much aligns with Bill Gates' philanthropic attitudes. Whereas a lot of the other big AI research hubs are almost completely focused on the economic impact of AI, i.e. how AI can drive profits, Microsoft seems genuinely concerned about the societal impact AI will have on jobs, privacy, safety, inclusiveness and fairness.

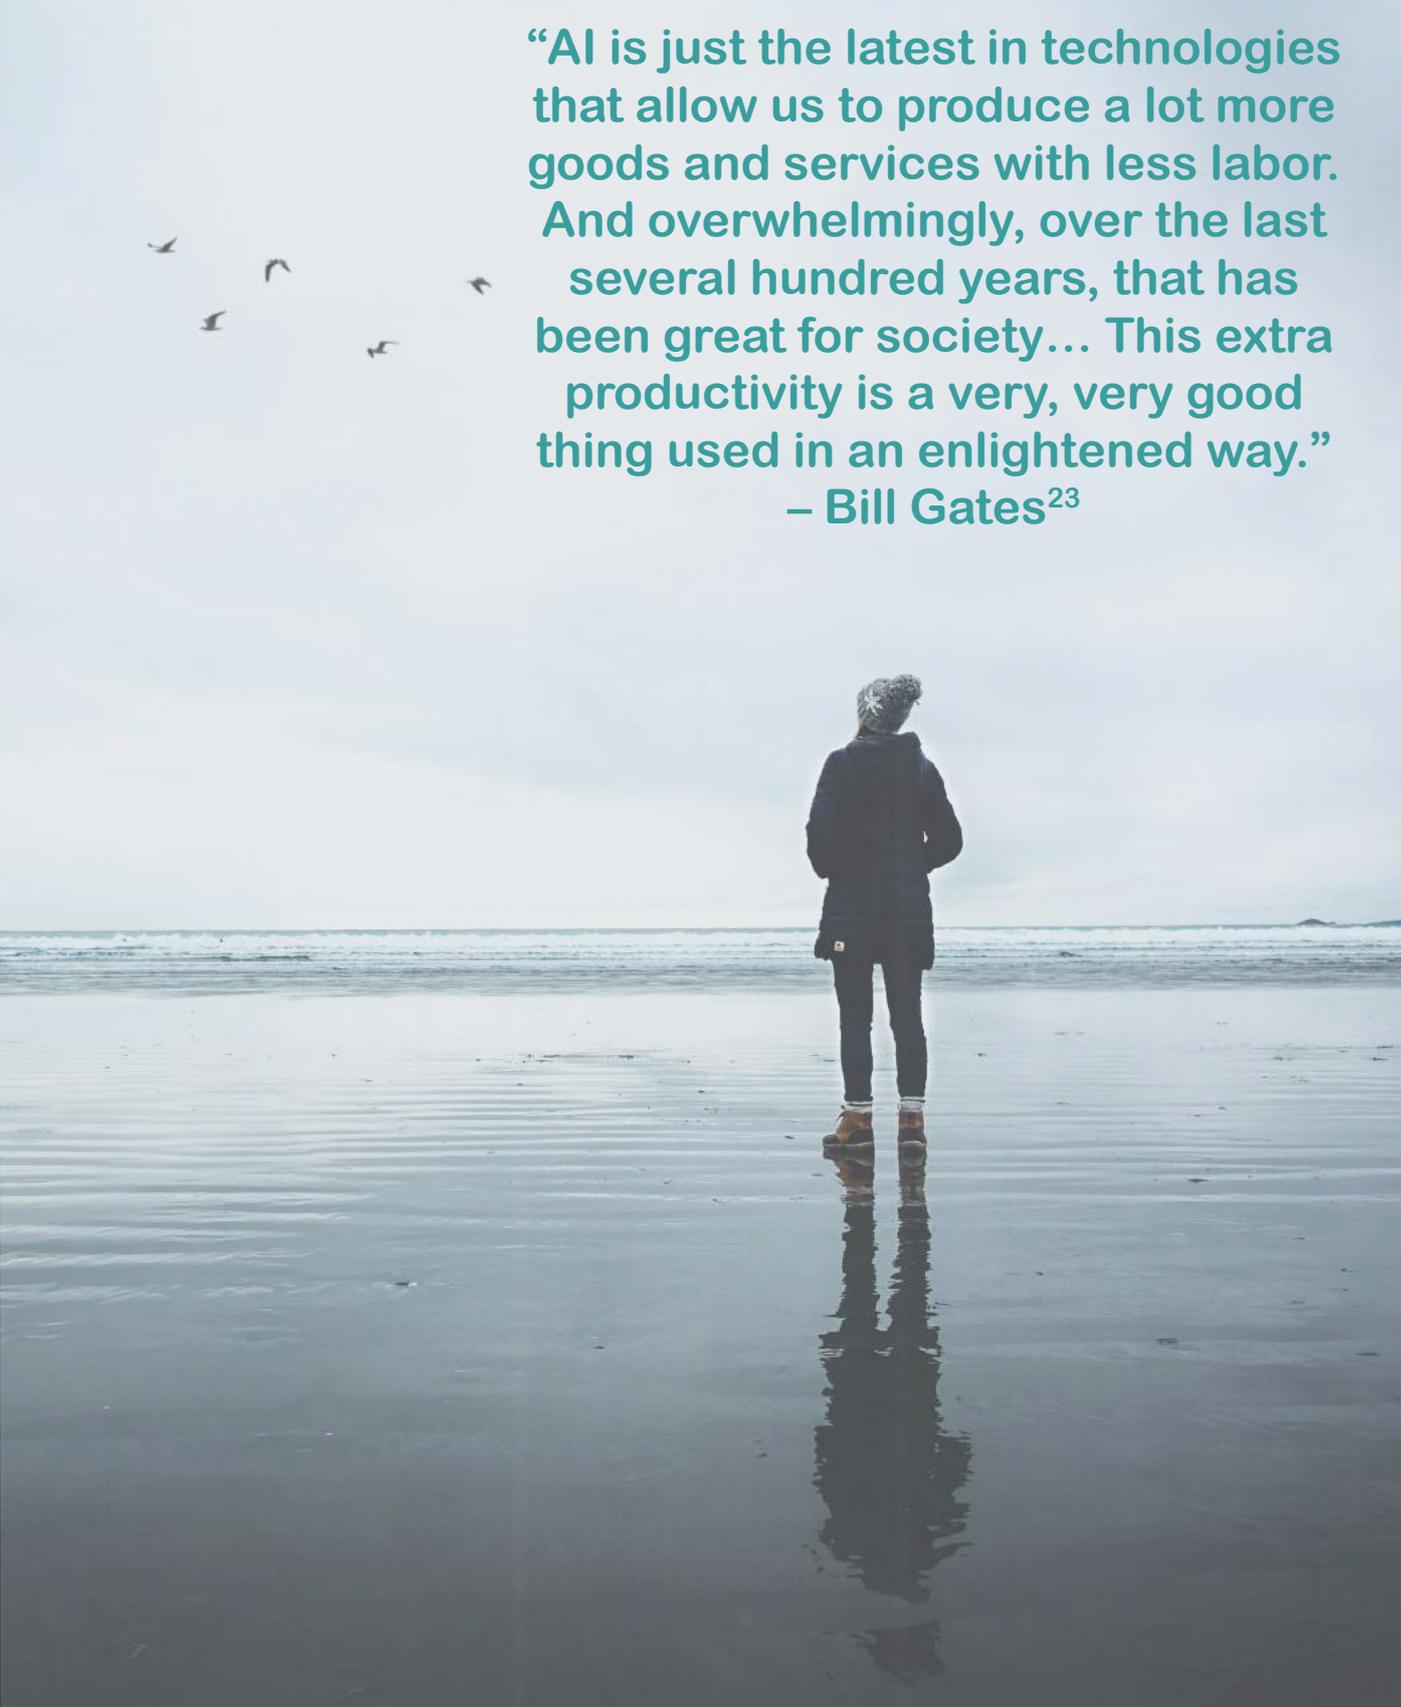
In a recent executive summary²², the company asked: When AI augments human decision-making, how can we ensure that it treats everyone fairly, and is safe and reliable? "Ultimately, for AI to be trustworthy, we believe that it must be 'human-centered'—designed in a way that augments human ingenuity and capabilities—and that its development and deployment must be guided by ethical principles that are deeply rooted in timeless values."

The summary states six principles by which the above can be accomplished:

- 1) **Fairness:** When AI systems make decisions about medical treatment or employment, for example, they should make the same recommendations for everyone with similar symptoms or qualifications. To ensure fairness, we must understand how bias can affect AI systems.
- 2) **Reliability:** AI systems must be designed to operate within clear parameters and undergo rigorous testing to ensure that they respond safely to unanticipated situations and do not evolve in ways inconsistent with original expectations. People should play a critical role in making decisions about how and when AI systems are deployed.
- 3) **Privacy and security:** Like other cloud technologies, AI systems must comply with privacy laws that regulate data collection, use and storage, and ensure that personal information is used in accordance with privacy standards and protected from theft.
- 4) **Inclusiveness:** AI solutions must address a broad range of human needs and experiences through inclusive design practices that anticipate potential barriers in products or environments that can unintentionally exclude people.
- 5) **Transparency:** As AI increasingly impacts people's lives, we must provide contextual information about how AI systems operate so that people understand how decisions are made and can more easily identify potential bias, errors and unintended outcomes.
- 6) **Accountability:** People who design and deploy AI systems must be accountable for how their systems operate. Accountability norms for AI should draw on the experience and practices of other areas, such as healthcare and privacy, and be observed both during system design and in an ongoing manner as systems operate in the world.

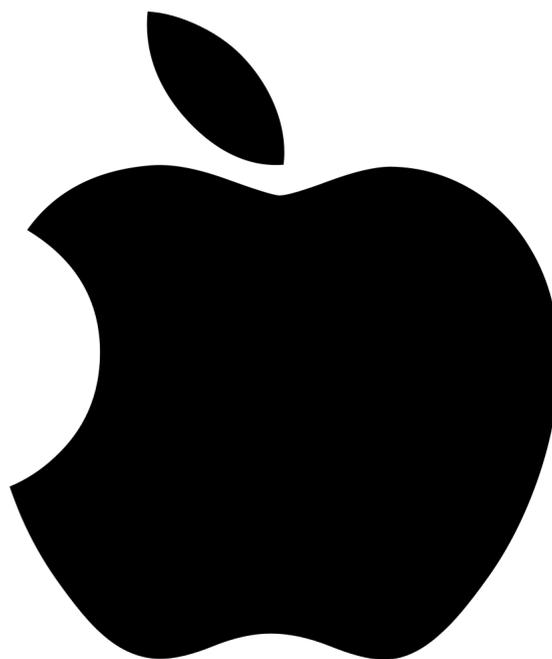
In September last year, Microsoft announced its AI for Humanitarian Action program. It forms the third aspect of the AI for Good action plan, the other two being AI for Earth, which addresses climate issues and AI for Accessibility, which is enhancing the lives of people with disabilities.

With a budget of \$40-million a year, the computing giant's aim is to use AI to address the fallout from natural disasters, ensure the welfare of children, refugees and displaced people, and encourage greater respect for human rights. The program is part of Microsoft's reported \$115-million, five-year plan to solve some of society's biggest challenges using AI.

A person wearing a dark jacket and a hat stands on a wet beach, looking out at the ocean. The person's reflection is visible in the shallow water on the sand. Several birds are flying in the sky above the person. The text is overlaid on the right side of the image.

“AI is just the latest in technologies that allow us to produce a lot more goods and services with less labor. And overwhelmingly, over the last several hundred years, that has been great for society... This extra productivity is a very, very good thing used in an enlightened way.”
– Bill Gates²³

What's happening at Apple?



For a brand so connected to originality and innovation, Apple has been surprisingly lax in developing its AI. For example, its CreateML app, which app creators can use to train AI models on Mac, arrived three years after Google's TensorFlow—its own cloud-based, open-source AI framework—was released. The result is that the latter is miles ahead in terms of development.

“TensorFlow is leagues and leagues ahead of [Create ML], which currently looks like a toy compared to it,” Reza Zadeh, CEO of startup Matroid. “No serious developer or researcher is even considering [using the Create ML technology]²⁴.”

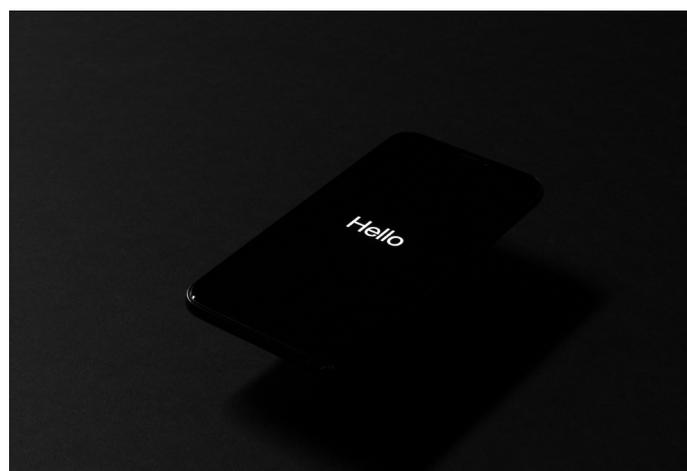
And the ground-breaking Siri—for many of us our first taste of an AI assistant when it launched on the iPhone 4S in 2011—and its offshoot, HomePod, have been eclipsed by similar offerings from Google (Google Home) and Amazon (Alexa). It's perhaps no surprise, then, that in 2018, Apple poached John Giannadrea—former Senior Vice President of Engineering at Google—to head up Apple's machine learning and AI strategy, focusing specifically on Siri and CreateML.

To further bolster its AI strategy, Apple acquired the startup Silk Labs late last year. The company was reportedly impressed by Silk Labs' AI platform's ability to modify its behaviour over time, according to what it learnt about its users²⁵.

In another move, in January 2019 Apple dismissed

over 200 employees from its autonomous vehicle group, Project Titan, and moved many of them over departments focused on AI and machine learning. Then, in April, Apple appointed another Google employee, Ian Goodfellow, and made him Director of Machine Learning. Goodfellow is known for inventing a generative adversarial network—a form of machine learning training that basically places two neural networks in competition with one another to use raw data to produce photos and videos that look real.

One reason Apple appears to be lagging behind the likes of Google and Amazon is its stance on privacy. Unlike Google, which trains its AI using the big data it collects from its users, Apple has long taken a vow to keep its users' information as private as possible. The challenge for Apple, then, is how to develop its AI in meaningful ways while still keeping to those principles.



Baidu 百度

We need to talk about Baidu

Launched in 2000, Baidu—the Chinese equivalent of Google—could become the world’s leader in AI development. The reason? Masses and masses of data and a government that not only allows but encourages data mining of its population.

To put into perspective just how massive we’re talking, the US has around 275-million internet users. China? Over 800-million.

Amongst other projects, Baidu are sharply focused on three areas of AI—autonomous driving, cloud computing for businesses and natural language and speech recognition. DuerOS—Baidu’s voice and image recognition program—easily outperforms Alexa and Siri. Baidu has developed an AI that can mimic a person’s voice after just 50 seconds of sample audio. Apollo—Baidu’s autonomous driving AI platform—is the only open-sourced autonomous driving simulation in the world. Its goal is to have fully autonomous cars on the roads by as early as 2021.

The power China could wield with its data tech-

nology has become a recent cause for concern, as evidenced by the UK government’s recent claims that Huawei and its 5G technology pose a security risk. Recently, Huawei launched a lawsuit against the US government following a ban on its products there. Huawei believes these restrictions—which have also been enforced in Germany, Italy, Australia and New Zealand—are part of a political move to suppress China’s development.

Perhaps to allay these fears—and avoid future similar bans or restrictions—last October Baidu became the first Chinese firm to join the Partnership on AI, a US group of companies that collaborate on ethical AI practices and study the effect of AI on society. Among its members are Amazon, Apple, Google and Facebook.

In a press release following the announcement, Baidu’s president Ya-Qin Zhang said: “Ensuring AI’s safety, fairness and transparency should not be an afterthought but rather highly considered at the onset of every project or system we build.”



Thoughts from... Kirsten Dent –

Finance manager at Greenwood Campbell

Companies who offer credit or loans will soon be using AI to predict customers' worthiness based on their past behaviour. The current AI technologies don't often provide reasoning behind how they reach their decision, but it won't be long before it improves and we'll see a lot more of it than perhaps we'd desire.

Fintech drifting into the AI world doesn't just keep the tech-savvy, convenience-seeking millennials happy, it's easing the mind of the fraud conscious too. Facial Recognition will help us keep our pennies safe by combating fraud. Users will soon be authenticated by who they are rather than what they know. The app will be using 'liveness detection' avoiding photos being used to unlock the barrier.

The future of savings has been served to us on a very easy, convenient and appreciated platter from the

app world. The years of mums and dads keeping on to get their children to save rather than spend has come to an end. Companies like Plum and Chip have made trying to save simple, quick and inattentively.

The idea is to connect to your bank account where the setup process is easy and a lot more appealing than the forms of setting up a savings account. Small amounts are automatically withdrawn to save without us consciously working out this month's budget. By analysing your everyday transactions, the AI technology learns about your income and spending pattern and uses this knowledge to withdraw the right amount for you. Goodbye to the days of trying to save and finding it necessary to 'dip in to your savings' for tonight's dinner.

A creditor's dream and a spendthrift's nightmare? It only remains to be seen...



The Future

It seems there's some evolutionary twinge inside humans that makes us afraid of anything new that we don't understand. We're wary of its potential to destabilise what we've become used to and see it as a threat to our survival. This is especially true when that 'something new' involves a machine or device.

When Johannes Gutenberg invented the printing press in 1439, sparking a revolution in the spread of knowledge, there was much opposition. Members of the church descended into a panic. Suddenly, the scripture that had only been for the 'privileged' few was available to all. Imagine what would happen when all those peasants could learn to think on their own... Uprising! Chaos! The breakdown of society! Gutenberg clashed with monks, who accused him of spreading the word of God about 'like dung'. In his 2000 book, 'The Justification of Johan Gutenberg', Blake Morrison recounts a confrontation between Gutenberg and the head monk of a monastery where Gutenberg had set up one of his presses:

Gutenberg: "I do not wish to despoil the Word."

Head monk: "But it will happen. To hand it about to all and sundry is languorous. Would you have ploughmen and weavers debating the Gospel in taverns?"

Gutenberg: "If that is what they want to do."

Head monk: "But what of the dangers? It would be like giving a candle to infants."

Gutenberg: "Such copies we make of the Bible would first be for monasteries and churches."

Head monk: "The Bible? You plan to make the Bible as well?"

Gutenberg: "I have considered it."

Head monk: "The Bible, to have authority, must be written by monks, not by some heretic machine."

Gutenberg: "With my press, it will look as though a monk has written it."

Head monk: "But it will be counterfeit, the work of an engine. And God does not inhabit an engine."

Our long held misgivings about the ‘heretic machine’ has since seeped into many inventions. When trains first came out, some people believed humans shouldn’t travel at such speeds (20mph) and that they would be obliterated in their seats. The telephone could electrocute you if touched, was a conduit for evil spirits and, according to the New York Times, telephones would ruin our privacy and personal relationships and turn us into “nothing but transparent heaps of jelly.”

Fast forward to more recent years and you’ll witness similar suspicions and dread over television, personal computers, fax machines, computer games, mobile phones, Wi-Fi and the internet. All of them technologies that’ll ruin society and reduce us to unproductive idiots.

Granted, just about any object and technology can be used for harm—another of Man’s unfortunate instincts—but if we’re not moving forward as a species, we’re at risk of redundancy. Embracing artificial intelligence is a way of ensuring that doesn’t happen. Instead of seeing it as the thing that’s going to kill us off, we should look at all the ways it will make our lives more streamlined, how it will improve healthcare, transform the lives of the disabled, and take care of the myriad mundane tasks we hate doing and thus free up more of our creative potential.

Perhaps the reason we’re so especially suspicious of AI is because it has the possibility to be like us. For the first time in history, here is something that can mimic us and do things that until now we believed only we could do. It’s easy to see how this is a scary idea. However, if we shift our perceptions, if we look back and remember all the good that the human mind has created, why can’t a network of artificial neurons that, remember, is still vastly unmatched to our own, do the same?

It’s time to believe.



Notes

- 1 McKinsey & Company (2018). Notes from the AI Frontier: Modeling the impact of AI on the World Economy. Available at: <https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-modeling-the-impact-of-ai-on-the-world-economy>
- 2 Asgard (2017). The European Artificial Intelligence Landscape. Available at: <https://www.linkedin.com/pulse/european-artificial-intelligence-landscape-more-than-400-westerheide>
- 3 https://en.wikipedia.org/wiki/Dartmouth_workshop
- 4 Alan Turing (1950). Computing Machinery and Intelligence.
- 5 <https://en.wikipedia.org/wiki/Automaton>
- 6 <http://www.history.mcs.st-andrews.ac.uk/Quotations/Liebniz.html>
- 7 <https://www.forbes.com/sites/forbestechcouncil/2019/03/07/four-ways-big-data-converts-into-sales/#42fdc6632c10>
- 8 [https://en.wikipedia.org/wiki/Watson_\(computer\)](https://en.wikipedia.org/wiki/Watson_(computer))
- 9 Computer World UK (2019). How UK companies are using the IBM Watson AI system. Available at: <https://www.computerworlduk.com/galleries/it-vendors/innovative-ways-companies-are-using-ibm-watson-3585847>
- 10 Fast Company (2018). Lufthansa and IBM Watson think AI can convince you to use your vacation days. Available at: <https://www.fastcompany.com/90253969/lufthansa-and-ibm-watson-think-ai-can-convince-you-to-use-your-vacation-days>
- 11 Medium (2017). Watson: Your cognitive fashionista. Available at: <https://medium.com/cognitivebusiness/watson-your-cognitive-fashionista-d786880798ff>
- 12 The Guardian (2017). Japanese company replaces office workers with artificial intelligence. Available at: <https://www.theguardian.com/technology/2017/jan/05/japanese-company-replaces-office-workers-artificial-intelligence-ai-fukoku-mutual-life-insurance>
- 13 Medium (2018). Artificial intelligence and the law: Friend or foe? Available at: <https://medium.com/@hpap2/artificial-intelligence-and-the-law-friend-or-foe-1de525daf558>
- 14 <https://www.re-work.co/events/deep-learning-in-finance-summit-london-2019>
- 15 Becoming Human (2019). How Netflix uses AI, data science, and machine learning—from a product perspective. Available at: <https://becominghuman.ai/how-netflix-uses-ai-and-machine-learning-a087614630fe>
- 16 <https://en.wikipedia.org/wiki/AlphaGo>
- 17 Business Insider (2016). Why Go is so much harder for AI to beat than chess. Available at: <https://www.businessinsider.com/why-google-ai-game-go-is-harder-than-chess-2016-3?r=US&IR=T>
- 18 https://en.wikipedia.org/wiki/Amazon_Web_Services

19 CNBC (2017). AI is in a 'golden age' and solving problems that were once in the realm of sci-fi, Jeff Bezos says. Available at: <https://www.cnbc.com/2017/05/08/amazon-jeff-bezos-artificial-intelligence-ai-golden-age.html>

20 https://en.wikipedia.org/wiki/Amazon_Robotics

21 Forbes Insights (2019). Human AI is here. Available at: <https://www.forbes.com/forbes-insights/our-work/human-ai/>

22 <https://news.microsoft.com/futurecomputed/>

23 CNBC (2018). Bill gates: 'AI can be our friend'. Available at: <https://www.cnbc.com/2018/02/16/bill-gates-artificial-intelligence-is-good-for-society.html>

24 CNBC (2018). Apple's AI strategy stands apart from the rest of big tech, for better or worse. Available at: <https://www.cnbc.com/2018/06/13/apples-ai-strategy-devices-not-cloud.html>

25 Tech Crunch (2018). Apple puts its next generation of AI into sharper focus as it picks up Silk Labs. Available at: <https://techcrunch.com/2018/11/21/apple-silk-labs/>