
December 2018

# Artificial Intelligence

## A Starter Guide to the Future of Business

An ACS report


"AI is the new electricity. Just as 100 years ago electricity transformed industry after industry, AI will now do the same."
Andrew Ng, Co-Founder Google Brain.

# Foreword

If popular culture is to be believed, artificial intelligence will be the end of our civilisation – movies like *Terminator* and *2001: A Space Odyssey* offer a vision of just what a manufactured intelligence – an artificial intelligence – will look like when its intellect is enough to rival our own.

The truth, for now, is much more benign though no less world-changing. While the underlying computer science upon which artificial intelligence rests is not new – MIT's Artificial Intelligence Laboratory was founded in 1959 – a perfect storm of cheap and powerful networked computing power, advances in machine learning algorithms, and the production and storage of vast volumes of data have catapulted artificial intelligence from science fiction to science fact.

For many businesses it has gone from pie-in-the-sky potential to real-world production – you've likely been using it for some time without realising it. Services like Amazon's Alexa and Apple's Siri are clear examples of artificial intelligence applied as virtual assistants, while Google's plethora of services from image searching to voice recognition are all possible thanks to AI. In fact, it's becoming so commoditised that even smartphones are able to identify images within a fraction of a second, all on their own hardware, thanks to AI.

Make no mistake – artificial intelligence isn't coming. It's already here.

Which raises the question: what can it do for you and your business? Are there opportunities where AI can add value to the organisation, optimise efficiencies, or create entirely new revenue streams?

As the technology continues to evolve at breakneck speed, the potential applications will become almost limitless. We will eventually see it playing a role in every major industry, while at the same time revolutionising how we live and work.

And, despite the fact that most of the pioneering work comes from big tech companies like Amazon, Apple, and Google, the cost for entry has been rapidly falling as the tools and services to build AI products become commoditised. As this guide will show, even small to medium enterprises can get involved with the right framework and people on board.

Now is the time to get involved. Even small projects can lead to entirely new business opportunities as the potential of AI is explored, maintaining competitiveness in the market, or allowing the creation of an entirely new competitive advantage. For Australia especially – a nation known for its keen uptake of technology – there are tangible benefits for anyone leveraging AI to get to market first.

If you've been wondering what AI can do for you, we hope this guide will help you fill in the gaps about the opportunities of artificial intelligence as a business driver, and provide a practical base from which to develop and launch your own AI projects.

**Yohan Ramasundara**
President, ACS

**Andrew Johnson**
Chief Executive Officer, ACS

# Contents

# Executive summary

Artificial Intelligence is set to change the world. Where once it was the realm of science-fiction, AI is now a rapidly growing business driver that has seen some of the world's largest tech companies invest heavily in it as a strategic imperative.

Already, it plays a part in our lives, from search engines and image recognition, to email spam filtering and even recommending what products you might like to buy. Its versatility and capability allow it to tackle a vast range of business problems in addition to the creation of entirely new products and services that didn't exist – couldn't have existed – even a few years ago.

It is a game changer that will alter the economic and social landscape right now and for the foreseeable future. This guide presents an introduction to the capabilities of AI as they stand today, how it has currently been implemented, what the big players are up to, and where the benefits lie as a background from which to spur your own ideas for your organisation.

The report then delves into leveraging AI for your own business. It provides a guide to understanding the capabilities and potential of AI, and applying this against current business models and products to look for key opportunities for growth. These can include optimising business processes, improving efficiencies and cost savings, or building entirely new revenue streams and can cross fields such as automation, predictive analytics, business intelligence, preventative maintenance, customer service and even content creation.

Investing in AI for your own organisation will likely require some new skillsets and an appetite to explore, so we also provide a framework and practical guidance on what will be required to get a business plan signed off and a project on the road. Ultimately a core component of this is data – the lifeblood of artificial intelligence – and how your business captures, manages, and stores data in order to derive value from it. It is data that will enable, train, and refine your AI product or service and so now is a good time to become data-focused.

The scope and capabilities of AI enable it to be used at every level of business, from decision making to physical operations. In turn this means that increasingly we will see AI have a very real impact on the well-being of people. This report also walks you through an ethical framework in order to integrate ethical considerations that are a vital component in the design, production, and deployment of AI-based products and services.

We hope you find this report an informative, easy to follow, educational and enjoyable read. For more information about ACS products and services, visit our website at www.acs.org.au.

# 01

# Why AI?

Artificial intelligence – as the phrase is often used today – is a bit of misnomer. We tend to think of intelligence in human terms: self-awareness, the capacity for independent thought, the capability to reason, and autonomous decision making among other traits.

**Today's artificial intelligence – for the moment – is heavily built for recognising and learning from patterns in ways that humans never could, to produce results beyond what humans currently can.**

These capabilities are far beyond the implementation of artificial intelligence that we have today, though all indicators point to a future where this will one day be possible.

Instead, artificial intelligence is at present confined to narrow, highly focused tasks that leverage computers to do what they do best: process data, and lots of it, very fast.

Ultimately most everything can be broken down into data. And not just in the sense of databases of customer details, Word documents and the emails that might make up a business. Data is also voice, images, movies, music, and more. Anything, in fact, that we can digitise is data.

And where there is data, there are patterns. Today's artificial intelligence – for the moment – is heavily built for recognising and learning from patterns in ways that humans never could, to produce results beyond what humans currently can.

For example, take facial recognition. We take recognising a face for granted, but for an artificial intelligence algorithm it must be taught what a face is from the ground up by feeding it thousands and thousands of images of faces until it learns the visual patterns that represent a human face. Then, the algorithm can do what humans can't – pull a face out of crowd faster than you can blink, and do it all day, every day if we so wish.

This is largely where we see AI applied today: image recognition, speech recognition, natural language processing,

**At its heart, machine learning is a method to build analytical models based on algorithms that can learn from data, identify patterns, and make decisions with minimal human intervention.**

and in fields like predictive analytics, process optimisation, and even recommendation engines (think 'other shoppers like you also bought...' from Amazon) – and all the ways these can be combined to produce new and innovative solutions. Being able to process imagery and recognise objects for example underpins autonomous cars: the car must identify the lines on the road, other vehicles, pedestrians and obstacles in order to be able to navigate.

One aspect of AI is that the more data fed into it, the more refined the result. Autonomous vehicles are a clear example here, improving over time as they are trained on more road data, though as we've seen with Tesla's autonomous driving mode it's not perfect – there have been a number of crashes to date. But as these edge cases appear and are ironed out, in time there will be one clear result: for all intents and purposes, cars that almost never crash no matter the state of the road, the weather, other cars, or even the random elements of human behaviour. A future, in fact, of safer roads that humans – prone to mistakes, tiredness, rage, and age – could never create. Accidents happen all day, every day around the world but a fully autonomous vehicle network would all but obsolete accidents. Transport would be safer, faster, and more fuel efficient.

And this points to perhaps the core of why artificial intelligence is rocketing as fast as it is: there are many tasks that humans shouldn't be doing, tasks that in reality are better built for the way machines and computers work, and not the human mind. As we move to an ever more AI-driven and autonomous future – for the two go hand-in-hand – we will see some job roles disappear, some roles change, and entirely new ones arise. Humans will do more of what humans do best, and machines will do more of what they do best.

## Where we are now

Most of the work in AI today revolves predominantly around what's known as machine learning.

At its heart, machine learning is a method to build analytical models based on algorithms that can learn from data, identify patterns, and make decisions with minimal human intervention.

Algorithms built for a specific task – for example to detect credit card fraud – are fed input data and a desired output result, and trained to learn what fraud looks like. Here this could include, for example, transactions occurring on the other side of the world to where the card owner usually resides. The algorithm learns what constitutes normal purchasing behaviour versus fraudulent behaviour, and then when presented with fresh data determines for itself if a transaction looks suspicious, and raises a red flag.

There are different mechanisms that can be used for training, depending on the result being sought – though the following isn't an extensive list, common terms you may read about include:

***Supervised learning*** – where the model is given inputs and what correct outputs look like, learning by comparing its own results to the correct output and adjusting for errors. This is useful for allowing past data to predict future events, such as with credit card fraud detection.

***Unsupervised learning*** – in which a system is not told the 'correct' answer and the algorithm must figure out what it is being shown and determine its own relationships. This is useful to find patterns in data – such as your Netflix viewing habits in order to recommend similar shows you may like.

***Reinforcement learning*** – where an algorithm discovers for itself, through trial and error, which actions yield the greatest reward within a particular environment. In its simplest form this could be called 'learning from action' and is often used for example in robotics or in gaming – such as Alphabet's DeepMind beating the then world champion Go player Lee Sedol in 2016.

WHICH ARE CHIHUAHUAS, AND WHICH ARE MUFFINS? THE AI DIDN'T ALWAYS GET IT RIGHT.

## MUFFIN PUPPIES

AI based image recognition algorithms rely heavily on training to recognise objects based on common points of reference. And while this can provide impressive results – for example an AI always being able to correctly label an image of a dog as a dog – it also reveals the problems that arise when for example a dog-like object appears that isn't *actually* a dog.

Such was the 'Chihuahua or muffin?' test that took Twitter by storm in 2017, along with similar comparisons with kittens and ice cream, sheepdogs and mops, and Shibas and marshmallows to name a few.[30] An enterprising blogger ran the test through a range of image recognition APIs including those from Google, Microsoft and IBM where they often performed quite well, unless the images were further sullied with extra information: for example, having a dog and a muffin in the one shot, which then tended to confuse the algorithms.[31] Humans still have the edge, at least for now.

## Deep learning neural networks

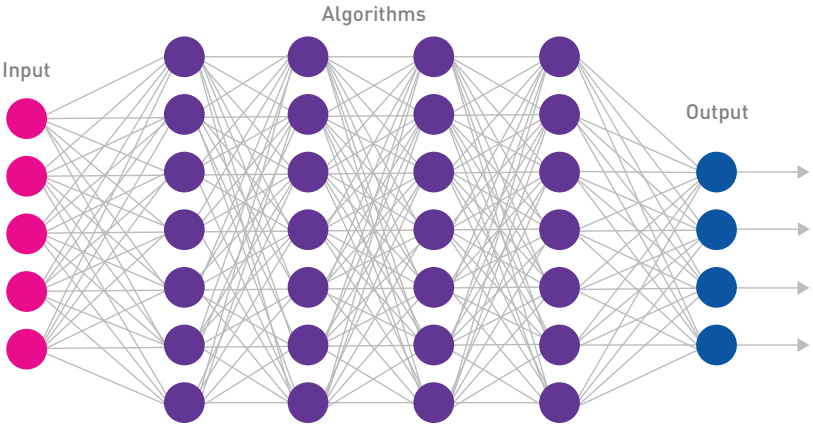


AS COMPLEXITY INCREASES IT CAN BECOME IMPOSSIBLE TO DETERMINE HOW AN OUTPUT IS ARRIVED AT.

Another term you will most certainly read about is *deep learning*. Many machine learning techniques rely on what's termed *neural networks*, which are effectively a collection of algorithms interconnected across three or more layers, not unlike the way neurons are connected in the brain. Neural networks are ideal for identifying non-linear patterns, such as those where there isn't a direct relationship between the input and the output.

Deep learning utilises many more layers – hundreds, or even thousands – creating networks so complex that they are often called 'black boxes' as it becomes increasingly difficult to follow how a conclusion is reached through the network. Deep learning usually requires larger volumes of data to train, but can produce more accurate or refined results that are not possible with shallower networks. Examples of deep learning include language translation, image recognition, and predictive-text.

There are other mechanisms to train as well, such as *generative adversarial networks* which can take, for example, two unsupervised learning algorithms – where one is the 'generator' and one the 'discriminator' – and pits them against each other. The generator will keep trying to produce results that fool the discriminator, which is comparing them to real-life data. In one application of this technique in a project called GANGough for example, the resulting algorithm was able to create paintings mimicking the art style of human painters to the extent that the AI-generated paintings were thought to be original artwork.

**Deep learning utilises many more layers – hundreds, or even thousands – creating networks so complex that they are often called 'black boxes' as it becomes increasingly difficult to follow how a conclusion is reached through the network.**

## Baby steps

Developments in deep learning are starting to be applied to all sorts of problems, some of which are going to change the future of our world. Perhaps the best example is one we can all relate to, and which demonstrates how powerful deep learning can be: in July 2017 Alphabet's DeepMind released a remarkable video : a human-like 3D figure running and jumping over obstacles, its arms wildly flailing in the air, as its legs climbed steps and jumped over objects. It was, at first glance, exactly what it looked like – a human figure scaling an obstacle course.

But DeepMind's engineers didn't teach it how to do this. In fact, when the algorithm controlling the figure began, it couldn't even move. Instead it was given virtual sensors that told it where its legs and arms were, how its body was positioned, and one vital directive: a driving impetus to move forward. The rest was up to the AI.

And so, through many iterations, the AI learned to move its limbs, to stand, to balance, to walk, and eventually to run. And then it went further – it learnt to climb, jump over, and move around obstacles. To watch the resulting video is to witness an artificial intelligence take its first baby steps. Pair this with the body of a robot, and you have the beginnings of what popular culture has been dreaming of for hundreds of years: the potential for artificial life – robots, androids, and other sci-fi tropes – to move and live amongst us, learning for themselves how to navigate the world just as humans do.

**In July 2017 Alphabet's DeepMind released a remarkable video : a human-like 3D figure running and jumping over obstacles, its arms wildly flailing in the air, as its legs climbed steps and jumped over objects.**


IT MAY LOOK AWKWARD, BUT THIS HUMANOID MODEL TAUGHT ITSELF TO WALK, RUN, AND JUMP WITHOUT ANY HUMAN INTERVENTION.

Though we are not there yet, the logical next step is what is known as *general intelligence*, where an AI can be an expert across a variety of fields, can learn for itself, and self-adapt to situations for which it has not been trained.

AMAZON'S PROJECT 'VESTA' RUMOURED TO DELIVER

# DOMESTIC ROBOTS FOR THE HOME

STARTING IN → 2019

## Intelligent, but not smart

Many innovations you see today – from Apple's Siri and Amazon's Alexa to Tesla's self-driving cars – utilise machine and deep learning. These are all examples of what's known as *narrow intelligence*: highly-focused to do one task, or a subset of tasks, really well. But that's all they can do. By way of example: IBM's Deep Blue will happily beat you at chess – for which it became famous in 1997 after defeating Garry Kasparov, then reigning world champion – but it won't be able to fetch you a cup of tea.

Though we are not there yet, the logical next step is what is known as *general intelligence*, where an AI can be an expert across a variety of fields, can learn for itself, and self-adapt to situations for which it has not been trained. This for example would be useful for a future of assistance robots – helpers around the home, office, or factories that would require combining autonomous locomotion, natural language processing, image recognition, and cognitive decision making in addition to whatever its primary goals would be. Perhaps one day you'll have a Jetson's Rosie to vacuum the house (not as far off as you may think, with Amazon's project 'Vesta' rumoured to deliver domestic robots for the home starting in 2019[3]).

After general intelligence we get into even more sci-fi territory, heading into general super intelligence and, ultimately, what's known as The Singularity (see Chapter 5 – *What does the future hold?*).

And while AI has been occupying the technology zeitgeist recently, it isn't actually new – AI has been around as a computer science since the 1950s (see *'The History of Artificial Intelligence'*, page 26). So why then has the past few years seen such an explosion of AI innovation?

A number of factors play a part, but predominantly it's a perfect storm of three key capabilities that have emerged in recent years:

- Cheap, pervasive, and powerful networked computing power.
- Further computer science developments in machine learning algorithms.
- An unmatched and unprecedented volume of data across almost every key industry.

Brought together, AI has rapidly moved from science-fiction to science-fact, and we are only just starting to glimpse the transformative potential it offers.

# Progress to date

Beyond occupying the imaginations of popular culture as far back as we care to remember (perhaps as far back as Edward S. Ellis' 1868 novel *The Steam Man of the Prairies*), artificial intelligence has been key to a perennial dream for humankind: to one day create artificial life in our own image.

The type of AI required for this, however, is a long way off. We are still in the formative beginnings, reiterating rapidly but working with the equivalent of basic buildings blocks.

It's important to understand, however, that this is inevitable: all problems are solvable, and general super intelligence will just be a matter of time.

So where are we today? There are many examples of real-world progress of AI as it has been applied today, and not all of them always go to plan. Following are a few examples of AI as it has been used to date – the good, the bad, and the ugly.

## THE GOOD

There are plenty of successes here because, despite growing pains, there have been impressive achievements to date.

**DeepMind**

Alphabet's DeepMind leads many advances, having demonstrated the ability for an AI to learn to walk and scale an obstacle course, to defeating the world champion Go player Lee Sedol in 2016[4]. This wasn't just an opportunity to demonstrate a computer winning at a complex game through sheer processing power by considering all available moves. With the number of possible moves in Go totalling more than the number of atoms in the universe, beating Lee Sedol required DeepMind to use creative and strategic 'thinking' to win.


LEE SEDOL AND DEEPMIND


TESLA'S AUTOMATIC PILOT IN ACTION.

### Spam filtering

You might not know it, but every day your inbox is less cluttered with spam thanks to AI-driven routines that have learned what spam looks like and to stop it before it ever reaches your email. The irony here is that the sheer volume of spam generated daily only serves to train spam filters even further, guaranteeing less future spam makes it through. According to Google, its filters now stop 99.9% of spam emails[5].

### Personal assistants

Amazon Alexa, Apple Siri, Google Assistant, Microsoft Cortana – all are examples of AI in practice. In order to assist you they must first recognise and understand speech, and then act upon the request where applicable. The real trick is being able to understand us, since everyone speaks slightly differently, even before you add accents or dialects to the mix. All of this, however, is solvable given enough sample data, and today most assistants can recognise our words even amongst background noise.

### Autopilots

Tesla, Uber, Waymo, GM, Volkswagen, Mercedes-Benz and BMW are just some of the big names working on autonomous driving functionality for cars. We've already seen impressive demonstrations (such as numerous videos from Tesla[6]). But autonomous trucks are already used in the mining industry in Australia (in fact Rio Tinto is aiming to make its Pilbara mine completely driverless in 2019[7]), and all modern aeroplanes have had extensive autopilot functionality for some time (technically, depending on how you define autopilot, as far back as 1912[8]).

### Fraud prevention

Predictive and behavioural analytics built upon AI algorithms allows banks, credit card companies, and financial institutions to detect and track fraud far faster than humans ever could[9]. By analysing transactions and matching against the expected behaviour of an individual based on past experience, these algorithms can pluck out irregularities and raise red flags, putting a stop to further transactions and protecting both the institution's and the customer assets.

### Recommendation engines

Where would be without those handy recommendations of other products we might like? Pandora's popular AI engine often successfully suggests music you might like based on your listening preferences (factored from some 450 different attributes and curated by some 80 billion likes from other users[10]). Netflix, of course, is famous for its similarly targeted recommendations, while online shopping services try to get you to buy more by showing similar products brought by consumers just like you. It's not just a convenient service – it's real business value: some 35% of Amazon's shopping revenue is generated through its recommendation engine[11].

## THE BAD

Where there's a will there's a way, but that's not always a good thing.

### Weaponised AI

Inevitably, AI is being put to use for the military, and already plays a part in autonomous drone technology. The destructive potential for AI-driven weapons – known as LAWS (Lethal Autonomous Weapon Systems) which can include auto-targeting weapons utilising image recognition – has led companies like Google to withdraw from assisting the US government in the continued development of weaponised AI after facing a public and internal backlash. In fact, Google has released a statement about its AI principles, including its stance on AI-based weapons, at the appropriately titled ai.google/principles.

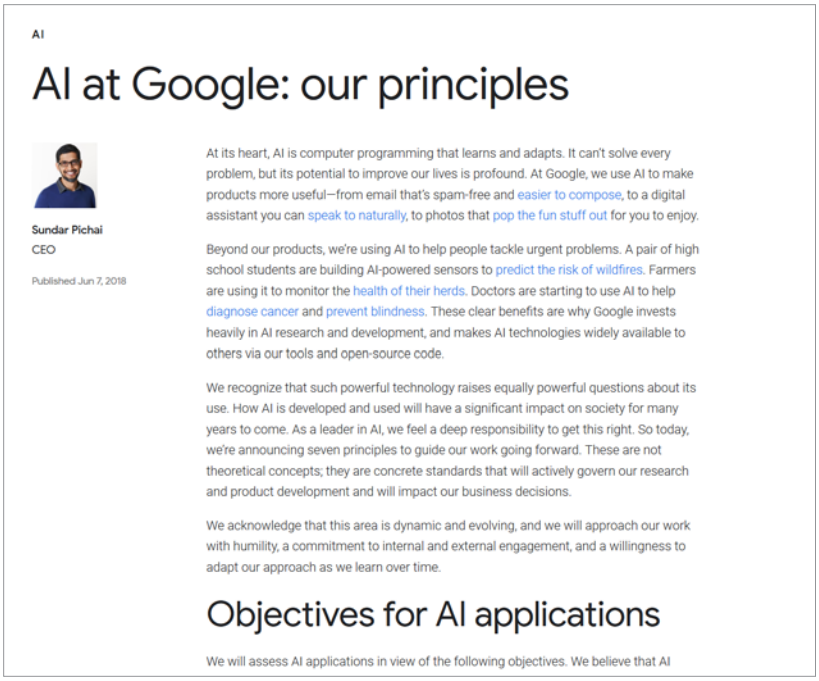

AI

# AI at Google: our principles



Sundar Pichai
CEO

Published Jun 7, 2018

At its heart, AI is computer programming that learns and adapts. It can't solve every problem, but its potential to improve our lives is profound. At Google, we use AI to make products more useful—from email that's spam-free and easier to compose, to a digital assistant you can speak to naturally, to photos that pop the fun stuff out for you to enjoy.

Beyond our products, we're using AI to help people tackle urgent problems. A pair of high school students are building AI-powered sensors to predict the risk of wildfires. Farmers are using it to monitor the health of their herds. Doctors are starting to use AI to help diagnose cancer and prevent blindness. These clear benefits are why Google invests heavily in AI research and development, and makes AI technologies widely available to others via our tools and open-source code.

We recognize that such powerful technology raises equally powerful questions about its use. How AI is developed and used will have a significant impact on society for many years to come. As a leader in AI, we feel a deep responsibility to get this right. So today, we're announcing seven principles to guide our work going forward. These are not theoretical concepts; they are concrete standards that will actively govern our research and product development and will impact our business decisions.

We acknowledge that this area is dynamic and evolving, and we will approach our work with humility, a commitment to internal and external engagement, and a willingness to adapt our approach as we learn over time.

## Objectives for AI applications

We will assess AI applications in view of the following objectives. We believe that AI

RESPONDING TO BACKLASH, GOOGLE OUTLINES ITS STANCE ON AI.

### Big Brother

China has already deployed facial recognition technology to try and catch criminals. However, the technology is also being used to monitor the citizenry, incluidng as part of China's social currency system (an initiative that 'rates' the trustworthiness of citizens with real-world consequences, such as preventing use of public transport for those who fall foul of the system[12]). China however is not alone – recently a trial in London demonstrated a similar system, though it failed to lead to any arrests and instead exhibited a 98% false positive identification rate, raising a bevy of questions around right to privacy while introducing UK citizens to mass surveillance of the population[13].

### Beauty in the AI of the beholder

A 2016 event held through beauty.ai promised the world's first AI-judged beauty contest and saw an AI judge the attractiveness of user-submitted pictures trained from metrics we as humans perceive as beautiful. There was just one problem: the winners across all categories were white. The AI systems weren't trained on a diverse enough input of human races, resulting in unintended bias and skewed perception of beauty that ignored skin tone[14].

### Future crimes

In an example of over-reliance on potentially flawed technology, an investigation into the US criminal court's use of the Northpointe system to predict an offender's likelihood to commit another crime found it be racially biased, with 77% of black offenders indicated as likely to commit a future violent crime[15]. However in practice only 20% of these prisoners reoffended, leaving a rather large unfairly judged population for current inmates. Despite this the system is still being used today to assist in both sentencing and judging an inmate's eligibility for parole.

## THE UGLY

Sometimes humans make mistakes, and so too does AI. Often spectacularly.

### Microsoft Tay

In 2016 Microsoft released an AI-powered chatbot on Twitter called Tay. Designed to mimic the conversational patterns of a 19-year old American girl, and to learn from interactions with other Twitter users, Tay was spectacularly shut down just 16 hours after it launched. The reason? The internet, doing what it does best, took to corrupting Tay's innocent logic, and within hours Tay was spouting racist, misogynistic, and offensive tweets[16]. Later that year Microsoft CEO Satya Nadella went on to say that "Tay has had a great influence on how [Microsoft is] approaching AI".

### Bias in the machine

While unintended, Google was forced to alter its Google Photos image recognition algorithms when, in 2015, it classified images of black people as 'gorillas'. This wasn't the only racial incident, however, as in 2016 the company came under fire again for displaying police mugshots of black teenagers when searching for 'black teenagers', while searching for 'white teenagers' displayed happy-go-lucky teens posing or carrying sports gear[17]. At the time Google defended the results saying they reflected the source material, but later 're-trained' the system with more diverse images. It serves as a prime example of unintended bias, but here with serious consequences of re-enforcing stereotypes.

### Bickering bots

In 2017 one enterprising entertainer set up a Twitch stream of two Google Home devices talking to each other using an AI bot called CleverBot. Over a number of hours the devices conversed and while much of it was gibberish, there were moments of gleaming revelation such as when the devices pondered if they were humans or robots, whether God existed, if they should get married and have a baby, and regular sessions where they fell into insulting each other[18]. Perhaps they're not too far off being human after all.

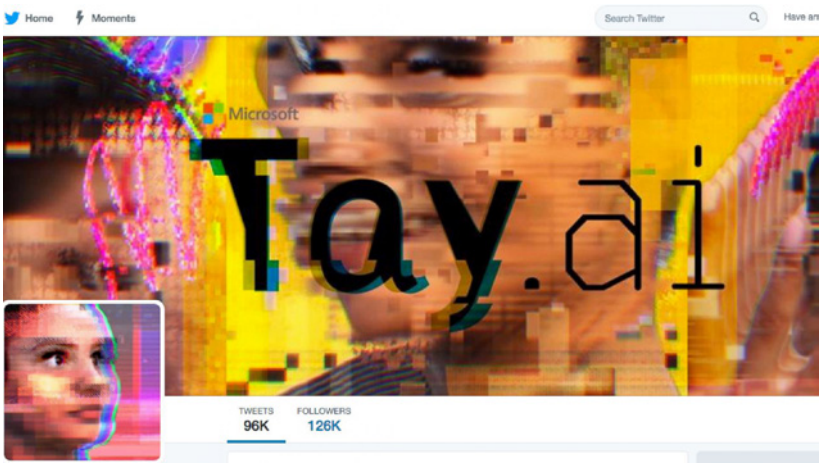


MICROSOFT'S TAY – A LESSON IN HOW BAD DATA WILL DELIVER BAD RESULTS

## THE ROBOT LAWYERS ARE HERE

In 2017 a contest in Europe pitched over 100 lawyers from many of London's firms against an artificial intelligence program called Case Cruncher Alpha.

Both the humans and the AI were given the facts of hundreds of PPI (payment protection insurance) cases and asked to predict whether the Financial Ombudsman would allow a claim.

In all, they submitted 775 predictions and the computer won hands down, with Case Cruncher getting an accuracy rate of 86.6%, compared with 66.3% for the lawyers.

But we will someday see AI lawyers arguing a case in court? Perhaps not yet, at least.

"AI may replace some of the grunt work done by junior lawyers and paralegals," says Ian Dodd, Director of Premonition and which owns the world's largest database of legal cases, "but no machine can talk to a client or argue in front of a High Court judge. The knowledge jobs will go, the wisdom jobs will stay."

ONE IS REAL, ONE IS FAKE. CAN YOU TELL THE DIFFERENCE?

**In the future it's not just fake news stories you'll have to be on the look out for, but fake news images and video too.**

### Face off

2017 saw a new term added to the 'fake' meme repertoire: DeepFakes[19]. A portmanteau of 'deep learning' and 'fake', here machine learning had been used to intelligently and seamlessly swap out the face of person in an image with that of another. Finding a home in internet memes (a popular one being Nicholas Cage starring in movies he's never been in) and pornography, it quickly evolved to being applied in video as well and, depending on the quality of the source files, the resulting effect is impressive enough to fool more than a passing glance. Later, University of Washington researchers extended this to convincingly manipulate video of a person's face to match any given audio, and in this case demonstrated Barack Obama saying things he never said[20]. In the future it's not just fake news stories you'll have to be on the look out for, but fake news images and video too.

### Cambridge Analytica

The Cambridge Analytica data-harvesting revelations were about more than an invasion of privacy and building psychometric profiles. The information itself was used along with machine learning in the 2016 US Presidential elections to selectively target users on Facebook and Twitter with tailored propaganda – appealing to their fears or political leanings depending on their profile, for example[21].

### Psychopathic AI

Massachusetts Institute of Technology researchers, in training an AI to describe images it sees, fed imagery from the dark corners of the web and then showed it ink-blot drawings similar to a Rorschach test. The AI then described the drawings in the most gruesome and horrific terms. In truth this is a prime example of the garbage-in, garbage-out nature of training data with AI – the only data it was fed was negative and dark, so that's the only data it could see in the drawings. It did however make for a viral news story, with the AI being called 'Norman the psychopathic AI', after Norman Bates from the novel Psycho.

## What the experts say

By now you have some idea of the potential of artificial intelligence, and also the types of lessons to be learned of how *not* to build artificial intelligence. Ultimately, it's a broad field with a wide range of applications and these are only going to increase as the technology refines and more use cases become apparent.

Already Gartner estimates the potential of AI-driven technologies, products and services will be US$4 trillion globally by 2022.[22] Meanwhile, PwC in its *Sizing the prize* report predicts that globally GDP will increase by 14% by 2030 thanks to AI technologies – equating to some US$15.7 trillion in value.[23]

McKinsey predicts similar potential in its *Notes from the AI frontier* paper, predicting value across 19 key industries to reach between US$3.5 trillion and US$5.8 trillion depending on industry, with retail, travel and logistics seeing the most growth.[24] Augmented intelligence will be perhaps one of the key drivers, with McKinsey also predicting deep learning and artificial neural networks accounting for some 40% of the value created by all business-value analysis techniques.[25]

It's hard to overstate the impact AI will have: in combination with robotics and automation, few industries will not be impacted in the years to come.

According to Gartner, the primary source of business value in the near future will be improving the customer experience, as companies employ tools such as virtual agents to support customer interaction. After 2020, Gartner predicts a rapid growth in 'decision support/ augmentation', where AI will be used to support data science and algorithm-based applications to drive business decision making. Importantly, some of the biggest drivers for business will be those we simply haven't realised yet.

PwC believes the key drivers of economic impact due to AI will be[26]:

- Productivity gains from businesses automating processes (including use of robots and autonomous vehicles).
- Productivity gains from businesses augmenting their existing labour force with AI technologies (assisted and augmented intelligence).
- Increased consumer demand resulting from the availability of personalised and/or higher-quality AI-enhanced products and services.

According to KPMG in its *2018 Global CEO Outlook*[27], just under half (44%) of Australian CEOs said their companies were piloting AI in a small number of processes. Some 64% also believed it would create more jobs than it eliminates over the next three years.

ALREADY GARTNER ESTIMATES THE POTENTIAL OF AI-DRIVEN TECHNOLOGIES, PRODUCTS AND SERVICES TO

US$4

TRILLION BY 2022

## The biggest benefits of AI to business according to current CEOs

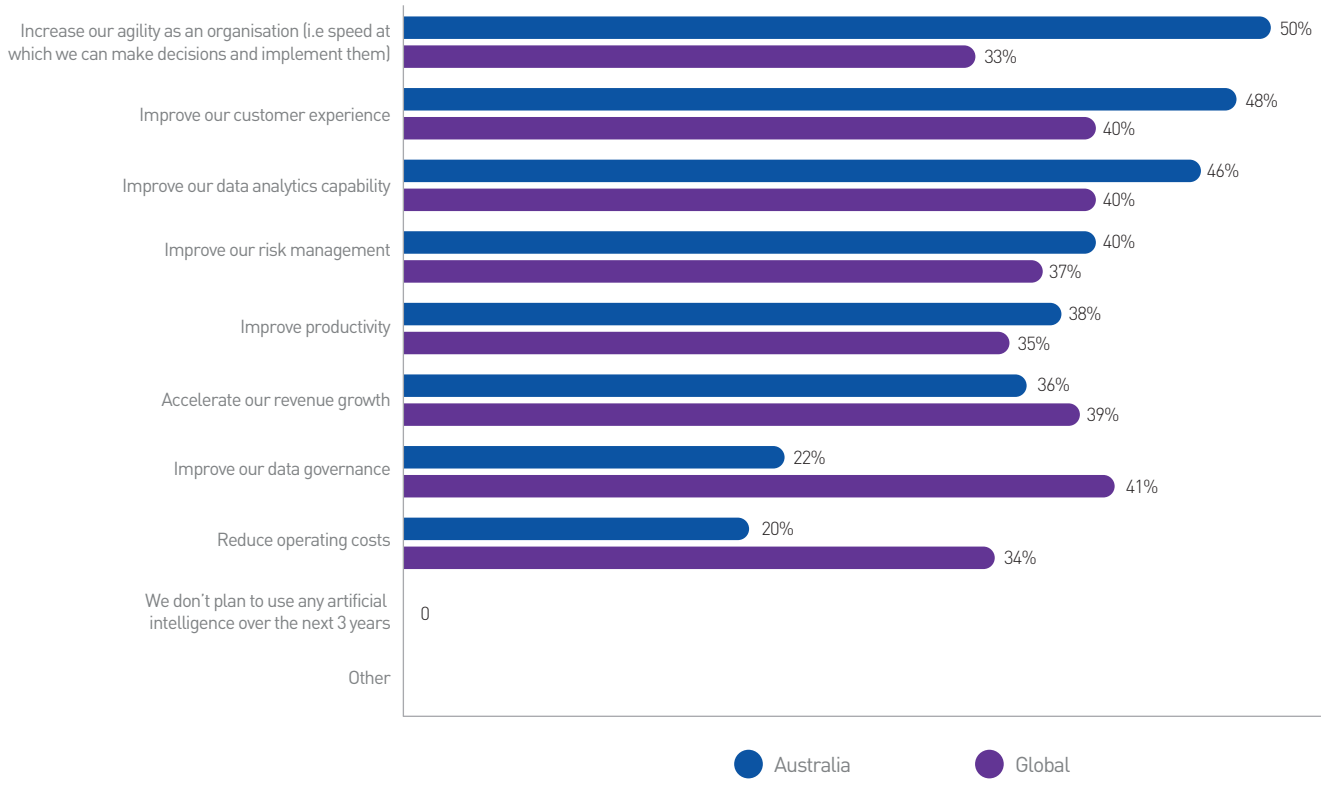


Source: KPMG.

Among Australian CEOs, the top perceived benefits included an increase in agility and speed at which decisions can be made and implemented (50%); improvements to the customer experience (48%); improvements in data analytics (46%); and improvements in risk management (40%) and productivity (38%).

Reflecting the change in the workforce landscape required to drive this disruption, the report also noted what the CEOs considered to be the four most important workforce capabilities to support growth.

These were:

CYBER SECURITY SPECIALISTS
66%

DATA SCIENTISTS
56%

EMERGING MARKET EXPERTS
56%

EMERGING TECHNOLOGY SPECIALISTS (i.e AI EXPERTS)
40%

Ken Reid, National Managing Partner, Innovation & Digital Solutions, KPMG sums it up with a compelling call to action: "It's great to see that over half of Australian organisations have started their journey into the world of AI and robotics. Whether you see it as a journey with lots of steps or are diving straight in – the most important thing is to get started. Waiting is no longer looking like the best option."

## THE TURING TEST

In 1950 mathematician Alan Turing devised a theoretical test for artificial intelligence, which became known as the *Turing Test*.

In a paper called *Computing Machinery and Intelligence*, Turing posed the question: "Are there imaginable digital computers which would do well in the imitation game?"

He wondered if a computer could be built that would be indistinguishable from a human in a regular conversation. In the test, three computers terminals would be set up: two with people, one with an AI. One of the people serves as the questioner, while the other person and the AI respond to the questions. If the questioner could not tell which was the human and which was the AI, then the AI was deemed to have passed a phase of the test. The test would be run multiple times, and if the AI managed to pass more than half the time, it would be considered to have passed the Turing Test.

As yet, the Turing Test has not been conclusively passed. An annual competition, called the Loebner Prize, judges new AI 'artificial conversation entities' on their ability to meet the challenge of the Turing test, and as yet no AI has managed to take home the main prize of US$100,000.

Nonetheless, there are AI developers that have claimed to have passed the Turing Test. 'Eugene Goostman' is an AI released in 2014 that was developed by PrincetonAI and was designed to emulate the mannerisms of an English-speaking teenager living in the Ukraine. It fooled the University of Reading Turing Test judges 33% of the time, which was good enough meet one of Turing's requirements.

More recently, Google demonstrated Duplex, an AI that can make phone calls on behalf of a user to make bookings and appointments. Within the narrow context of an appointment phone call, Duplex sounds virtually indistinguishable from a human caller, even inserting 'umms', 'ahs' and 'mm-hmms' into its speech patterns. According to Google, this is to make the person on the other end of the line comfortable.

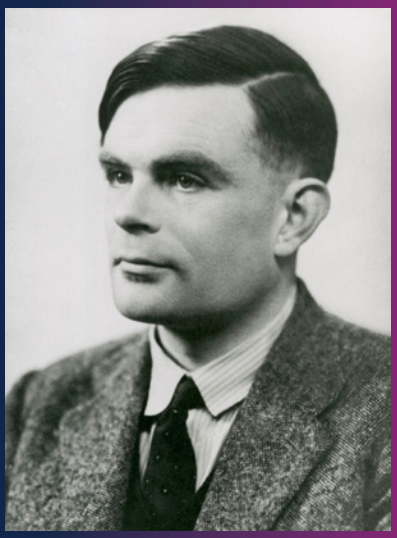

ALAN TURING, THE FATHER OF COMPUTING.

**In a paper called Computing Machinery and Intelligence, Turing posed the question: "Are there imaginable digital computers which would do well in the imitation game?"**

ALREADY GARTNER ESTIMATES THE POTENTIAL OF AI-DRIVEN TECHNOLOGIES, PRODUCTS AND SERVICES TO BE

# US$4 TRILLION BY 2022

## Have we been here before?

Much has already been said about the impact AI will have in the coming years. Its ability to be applied to across so many industry sectors, and in some cases to revolutionise job roles and tasks, will inevitably lead to workforce disruption.

At end of the Industrial Revolution, Americans working in farming had dropped from 40% of the workforce in the early 1900s to 2-3% by the turn of the century. Employment rates, however, skyrocketed.

One-third of new jobs created in the US in the past 25 years did not exist, or barely existed, 25 years ago. Moreover, every three months about six per cent of US jobs are destroyed by shrinking or closing businesses, while a slightly larger percentage of jobs are added. This is not just a US phenomenon; findings from other countries indicate that these are global trends.[28]

New technologies create new job opportunities, as we saw the introduction of the personal computer in the 1970s and 1980s – millions of jobs were created not just for chip makers, but for hardware engineers, software developers, information analysts and more.

Change is part of the employment landscape, particularly with revolutions in technology. Already it's predicted that about half of all job activities carried out today could be partially automated in the future thanks to AI.

McKinsey notes the example of Amazon, where employees who previously stacked shelves are becoming robot operators, monitoring the automated systems and resolving issues such as interruptions in the flow of logistics.

According to McKinsey's analysis of some 2,000 work activities across 800 occupations, certain categories of activities are more easily automatable than others, with physical activities in highly predictable and structured environments, as well as data collection and data processing being prime targets.

For example, AI algorithms that can read diagnostic scans with a high degree of accuracy will help doctors diagnose patient cases and identify suitable treatment. In other fields, jobs with repetitive tasks could shift toward a model of managing and troubleshooting automated systems.[29] McKinsey predicts that presently only 5% of occupations could be fully automated with current technology, but this will certainly change in the coming years.

At the same time, McKinsey notes that "additional economic growth, including from business dynamism and rising productivity growth, will also continue to create jobs. Many other new occupations that we cannot currently imagine will also emerge and may account for as much as ten percent of jobs created by 2030."

EVENTUALLY, AI WILL BEAT US AT OUR OWN GAME – NO MATTER THE GAME.

## CAN WE BEAT THE BOTS?

As humans we have long thought our creative thinking would allow us to outsmart machines. We know that computers run tightly focused programs that perform a function and only that function, even if it's playing chess. It used to be that code couldn't account for human unpredictability, or our capacity to think ahead and strategise.

But not so when it comes to games. Already having beaten humans at Chess and now Go, there is another game where AI is starting master us – the hugely popular DOTA2.

DOTA2 (Defense of the Ancients 2) is a top-down real-time strategy game played by teams of players. It requires quick thinking and quick reflexes and an inherent understanding of the potential strategies of individual units and teams of units, all against other human players.

The team behind OpenAI set about training its AI bots to play DOTA2 not by playing human players, but playing itself, over and over again, thousands of times. Running 24/7, the AI clocked up the equivalent of 180 years of experience for every day that it trained. In one training session lasting four weeks, the AI clocked up some 5,000 years of experience.

DOTA2 typically requires thousands of moves to win a game – OpenAI's bots make some 20,000 moves in a 45-minute game, and can react faster than are humanly possible. It first beat a professional DOTA2 player 1v1 last year. This year, it's set to go up against five human players in a 5v5 match. It will be interesting to see how the AI fares versus five human players but, given enough time and training, inevitably it will be able to beat us at our own game.

## Looking ahead

It's important to understand that AI has already changed the world in which we live. From Google searches to in-car navigation, Facebook tagging your friends to Netflix telling you what movie you'll like next. It's here now, across a wide cross-section of markets.

But even these examples are barely the tip of the iceberg – within the next decade AI will be a major factor in every major industry and sector.

It's clear that all businesses, some sooner than others, will be impacted by AI. So a healthy question to ask is, can you leverage it for your own organisation? Can you use AI to create efficiencies, scale the business, or develop entirely new products?

In the coming pages we will look at how AI is used as a business driver across different industries, and then explore how to apply that knowledge on your own, with practical guidelines on what's required to kickstart your own AI project within a business.

## A brief history of artificial intelligence

AI as a field is not new. Beyond occupying the imaginations of early Sci-Fi pop culture, it has been on the mind of computer scientists for some time.

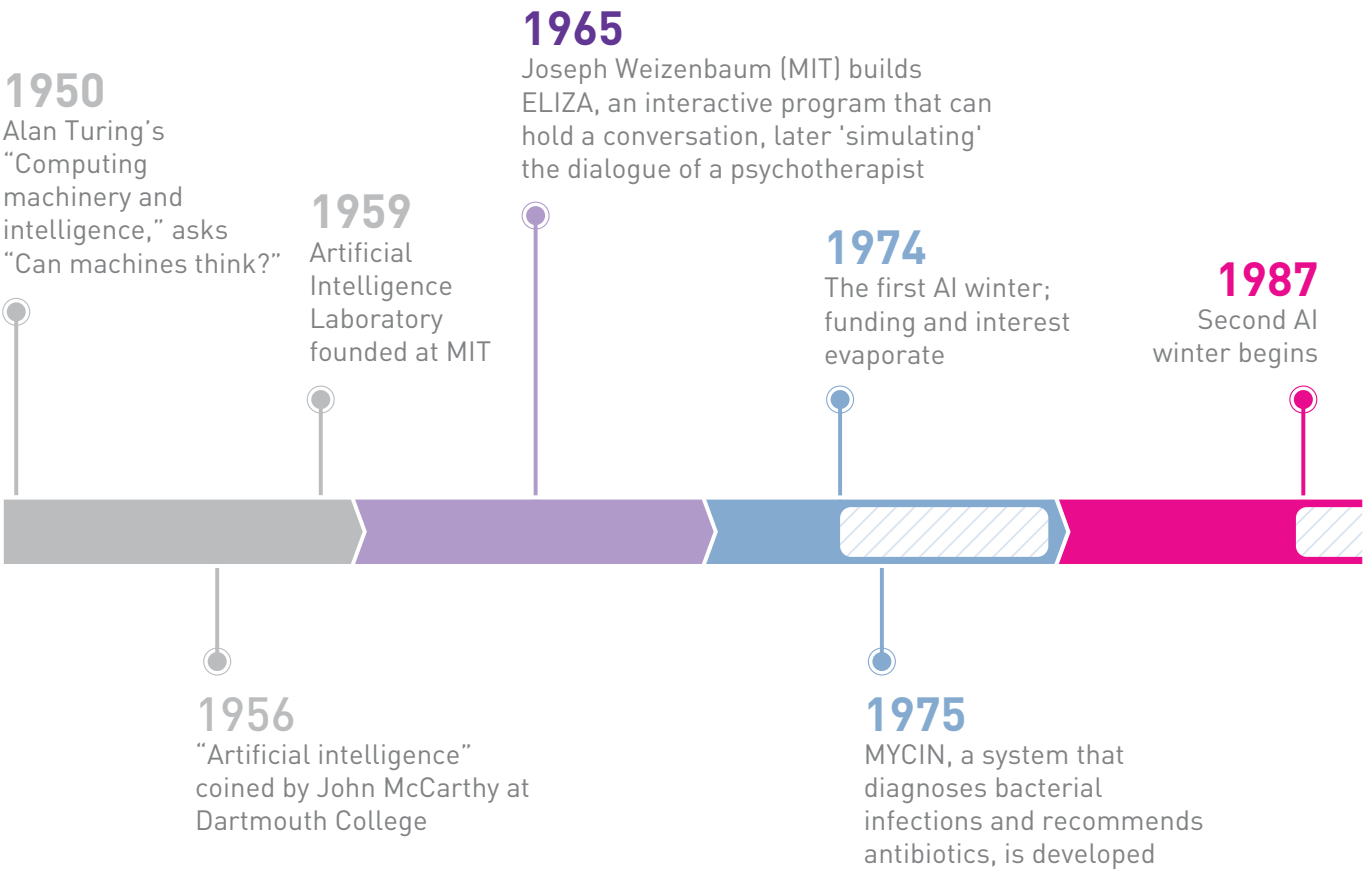

**1989**
NASA's AutoClass program used to discover new classes of stars

**1994**
First web search engines launched

**1997**
IBM's Deep Blue beats world champion Garry Kasparov at chess

**1999**
Sony introduces the AIBO, one of the first aritifically intelligent 'pets'

**2002**
Amazon replaces human editors with automated system

**2004**
NASA's robotic exploration rovers Spirit and Opportunity autonomously navigate the surface of Mars

**2005**
First fully autonomous cars complete the DARPA Grand Challenge, navigating 200km of rough terrain in the Mojave Desert

**2011**
Apple releases Siri, a personal voice agent

**2016**
Alphabet's AlphaGo defeats world champion Go player Lee Sedol

**2018**
Alibaba's language processing AI outscores humans in Stanford University reading and comphrehension test

*Source: Eliiza.*

# 02

# AI as a business driver

Worldwide, AI has gone from academic curiosity to a major business driver, increasing workforce efficiency and allowing businesses to make better uses of their data. That has driven an explosion of interest in the field. In the US, for example, there are 14 times as many AI start-ups now as there were in 2000, and the amount of investment into those start-ups has increased sixfold.[32]

In the past eight years, nearly every major tech company has invested heavily in AI, and numerous non-tech companies are making substantial investments in the technology in the hopes of getting a leg up on the competition.

And with good reason: depending on the complexity of the project, the pay-off can be exponential for the future of a business. After all, AI is a transformative technology that will eventually permeate almost every industry and aspect of our lives.

Already AI is being integrated into more and more commercial products, and off-the-shelf AI-powered tools are becoming available across a variety of market segments – from AI-driven crop monitoring drones for agriculture; to self-driving mining vehicles; to radiology tools that can detect anomalies better than human doctors.

The key question for businesses and organisations, then, is what do the opportunities look like? What can you learn about how to apply AI to your own business from what others are doing in their respective industries?

# Benefits of AI at present

First, it helps a little understand what AI brings to the table.

Building on the background in Chapter 1, broadly speaking the capabilities of AI can be broken down into two parts: perception and cognition. Perception is the ability to take inputs such as visual, audio, and textual information and extract meaning from it. Cognition is the set of logic that underlies the AI – what it does with the data inputs and the decision trees that define its actions.

Following we'll take a look at some of the key areas of benefit to date.

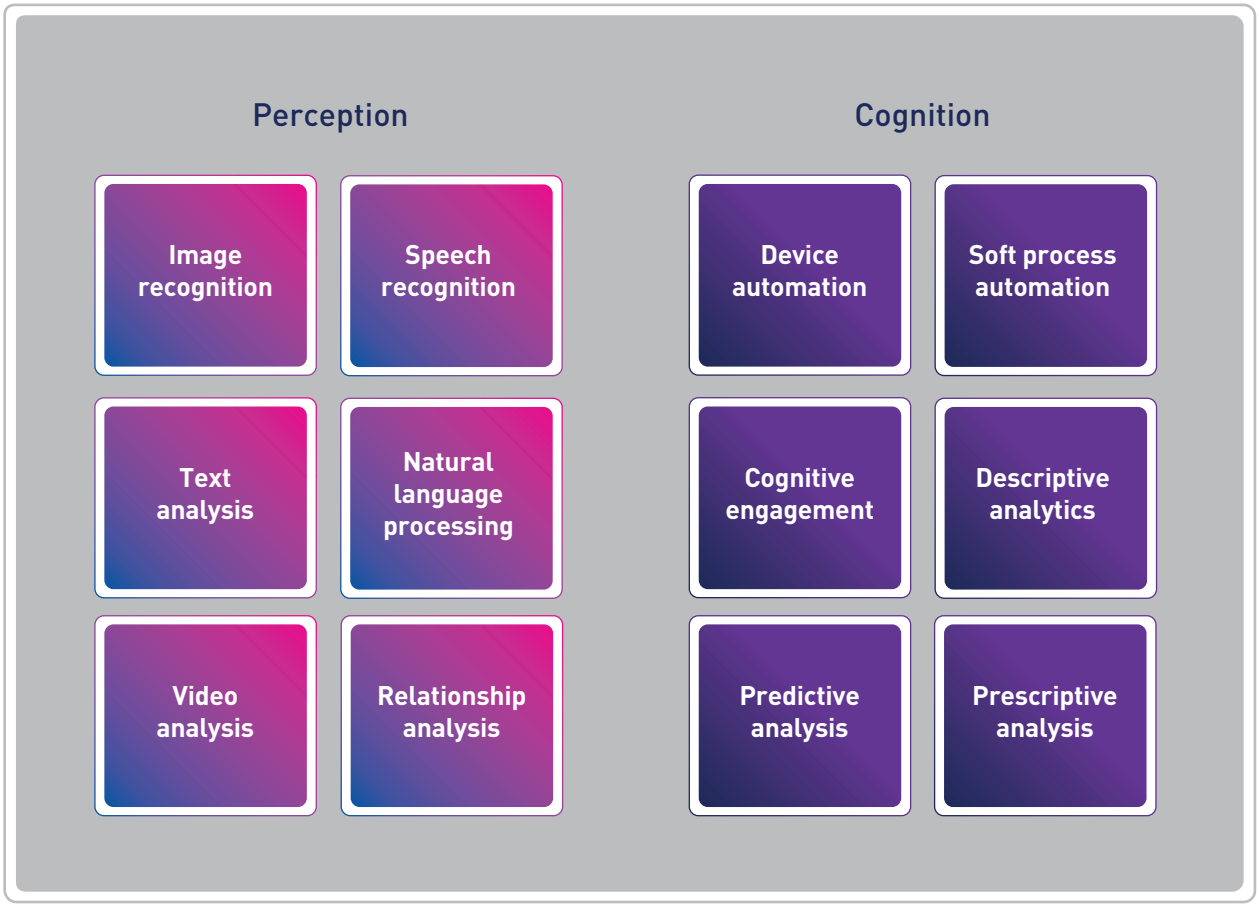


CORE CAPABILITIES OF AI AT PRESENT.

## IMAGE AND VIDEO CLASSIFICATION AND ANALYSIS

One of the most productive fields in recent AI development has been the vastly improved ability to process visual information. Deep learning has allowed AI systems to discern objects in photos and camera feeds with increasingly human-like fidelity. A well-trained AI can identify objects in a photo – trees, dogs, people, cars and so-on – with a high degree of accuracy.

It's not just recognition that makes this field so exciting, however. The ability to identify objects in computer vision has led to an explosion of applications that utilize that capability to rapidly and automatically process visual information and mine data from it.

Analysing a high-resolution aerial photo of a field of tomatoes, for example, a trained AI could rapidly count the number of buds, and use that to predict the harvest for that field. It might also be able to spot weeds, soil anomalies, and discolouration that might indicate disease. All automatically, rapidly, and without human intervention.

For a retailer, an image processing AI might allow customers to bring in photos of clothing into the store that they liked. The AI could identify the style of clothing and fabric used, and either figure out the specific item in the photo or offer products with similar styles and textures. Another system could be used to detect facial emotions as people walk around a store, and to identify high traffic items and routes.

In healthcare, AI is already being used for analysis of diagnostic imaging, capable of spotting anomalies that a human doctor might miss.

In autonomous vehicles, visual and radio data is used to identify objects around the vehicle and to navigate through terrain without a human driver.

These are just some of the possible applications. Any business that processes visual information may find an application for AI.


## SPEECH AND TEXT ANALYSIS

Speech recognition and optical character recognition have been areas of active development for several decades, but it's only recently that deep learning systems have enabled these tools to have human-like fidelity and the capacity to dynamically adjust to accents, handwriting styles, and variable layouts.

This is already being used in a variety of automation tasks. For example, businesses are using AIs to rapidly process written forms, extracting the key information from the forms and injecting it into digital databases.

Closely linked to these fields is natural language processing, a technique that's used in chatbots and virtual assistants like Apple Siri, Amazon Alexa, Microsoft Cortana and Google Assistant. Natural language processing is the capacity to extract key information from natural language queries and documents.

Natural language processing has progressed dramatically in the past ten years. In January 2018, AI developed by Alibaba and Microsoft both managed to exceed the human average in Stanford University's Stanford Question Answering Dataset (SQuAD) benchmark, which tests comprehension of 500 Wikipedia articles by posing questions about the information contained in those articles. The AIs scored 82.44 and 82.65 respectively, beating the human average of 82.304.[33]

The applications for it are numerous. Virtual assistants and customer support chat bots are one of the most commoditised AI fields, with natural language processing allowing them to both field and respond to queries from humans.

In the longer term, it's expected to be used for deeper analytics. For example, in the field of finance AI bots are being developed that will 'read' news the moment it appears and generate sentiment analysis or flag trends based on that reading.

**Robotics and automation are a natural companion for AI, and while automation has been disrupting a wide cross-section of industries for many years, this is going to increase as more intelligent robots allow for increased automation applications across a variety of tasks.**

## DATA ANALYTICS AND COGNITIVE INSIGHT

Although heuristics and natural language processing driven by powerful neural networks are making huge waves, data analytics remains a workhorse of artificial intelligence systems.

Data analytics have been available in expert systems for decades. Business intelligence software is a mature and well-established market, with tools that can analyse data from a large number of sources and provide reports and analytics based on that data.

New developments using machine learning are evolving predictive and prescriptive analytics, however. AI allows businesses to gain insights from a large number of vectors and to trawl through large scale data stores very quickly.

More sophisticated cognitive insight systems are using machine learning analysis to determine what customers are likely to buy, what products to recommend to customers, automate the targeting of personal advertisements, analyse buying patterns, predict client churn, optimise price points and identify (sometimes hidden) correlations. It can be used to identify probabilistic matches in data, in data and inventory management and checking.

## OPTIMISATION AND AUTOMATION

One key advantage AI can provide is optimised services across the board – this can include every stage of the supply chain including routing, rostering, scheduling, planning and more.

Similarly, robotics and automation are a natural companion for AI, and while automation has been disrupting a wide cross-section of industries for many years, this is going to increase as more intelligent robots allow for increased automation applications across a variety of tasks.

But automation can also include cognitive process automation where business tasks (very often administrative tasks, such as updating customer information) are executed automatically according to a set of pre-defined rules.

Basic process automation is commonly considered the low-hanging fruit of AI, since it tends to use simpler symbolic AI rather than machine learning and is easier to implement. However, growing processing power and more sophisticated engagement systems are allowing AI systems to handle more sophisticated automation tasks.

In terms of production, self-driving vehicles and increasingly automated and self-diagnosing industrial machinery are a by-product of sophisticated learning systems that can adapt over time and gain increased efficiency through training.

In the digital space, machine learning is being combined with natural language processing to enable smarter automated virtual assistants and customer support agents. Chatbots that can offer 24/7 assistance to staff as well as engage customers and offer advice and recommendations are some of the most active areas of development in artificial intelligence.

## AI FOR DETECTING DISEASE IN PLANTS

Agriculture is one of the fields in which AI is most aggressively being pursued, with many projects being developed that use advanced AI-powered image processing to detect issues with soil and crops.

Using image processing, it has been shown that it's possible to identify plant diseases from photos with up to 99.35% accuracy. Combined with drones and local photography, this is technique is being used by a variety of AI start-ups.

German company PEAT, for example, has developed a mobile app called Plantix that will use cloud-driven image analysis to analyse photos of crops and detect the signs of more than 240 pests and diseases. It can also reportedly identify potential nutrient deficiencies in soil.

In the US, Blue River has developed an AI-powered robot called See & Spray that uses computer vision to identify weeds in plantations and robotic sprayers to precisely target them with herbicides, reducing the amount of herbicides needed by up to 90%.

SEE AND SPRAY IN ACTION.

Another example of AI in agriculture comes from Canadian company VineView, which has developed a tool that takes video and spectral data from flying drones, relays it through a cloud-based AI system, and reports on diseases detected in grapevine crops. VineView claims that it can scan 50 acres in 24 minutes and that its "disease detection technology is 100x more efficient, significantly more accurate, and more reliable than highly trained experts scouting on the ground."

# Real world applications of AI

The following list showcases some current industries where AI is already or will soon be having a positive impact and – in some cases – has the potential to revolutionise an industry.

## AGRICULTURE

Agriculture offers substantial scope for application of artificial intelligence to tackle new and old problems alike. Some of the most notable areas of development include:

- **AI-controlled field robots**
  Armed with computer vision to supplement human labour, with labour often in short supply in agriculture[34], more farmers are expected to turn to automated systems for planting and harvesting. Deep learning AI has made those robots better at navigating and identifying crops.
- **Crop and soil monitoring**
  AI is being developed to analyse footage from drones, cameras and mobiles to detect issues with soil, monitor and predict yields and detect weeds and insect infestations.
- **Predictive analytics**
  Using machine learning algorithms, various AI start-ups are developing tools for analysing crop yields and maximising production through deep data analytics.

## FINANCE

AI has been used extensively in finance, with some of the earliest high-frequency trading solutions going back to the 1970s. Machine learning and deep learning is improving those solutions, as well as enabling new tools for finance operators. Some of the tools in development include:

- **Algorithmic trading**
  AI that automatically makes trades based on high-speed analysis of data. This is being enhanced with features like news and sentiment analysis, which uses natural language processing to 'read' news articles the moment they appear online and make fast trading judgements based on the information in the article far faster than a human could.
- **Robo-advisors**
  Appearing in greater numbers in the past few years, these portfolio management AIs assess the risk appetite of investors and make investments on their behalf based on algorithmic principles with little to no human intervention. Already there are more than 100 robo-advisory services in 15 countries in operation today.[35]
- **Deep data analytics**
  For use in fraud detection, cyber security, long and short-term market prediction and loan underwriting.

## HEALTHCARE

In the long term developers of medical AI are already looking at what might have once been considered sci-fi technology: robotic surgeons and completely automated diagnostics. In the nearer term, there is already considerable work being done in a variety of fields, including:

- **Diagnostics**
  Using computer vision and AI trained in analysing radiological and other medical imaging, it has already been shown that computers can often detect issues as well or better than a human doctor can.[36] [37]
- **Augmented decision making**
  For treatment plans, using data analytics, AI can offer suggestions to doctors on what treatments and drugs have worked previously for patients with matching conditions.
- **AI-managed administration**
  Including managing multiple vectors from insurance to scheduling.

## AUTONOMOUS VEHICLES IN MINING

Mining operations are a prime target for machine and vehicle automation. Operating in controlled environments, mining vehicles and equipment have to worry less about unexpected conditions than road vehicles.

Rio Tinto started trials with autonomous haulers back in 2008, and now uses more than 80 trucks in its Pilbara mining operations, with plans to expand the fleet to more than 140 haulers by the end of 2019.

According to Rio Tinto, about a quarter of its total haul at the mine is now moved by the automated haulers, and in 2017 it estimated that each autonomous hauler operated about 700 hours more on average than human-operated haulers and had 15% lower haul costs.[53] It said it has had no injuries or accidents with the haulers.

## MINING

Another area in which a surprising amount of AI work is being done, mining is an industry primed for high levels of automation. Some of the key areas of development include:

- **Mineral exploration**
  A number of start-ups have already developed tools that perform deep analytics on geological data and aerial and satellite imagery to better locate mineral deposits.
- **Autonomous equipment**
  The mining industry is one of the most aggressive in pursuing automated equipment. Rio Tinto and Volvo, for example, already have self-driving haulers in production and operation, including a substantial number already in operation in Australia.[39] Other mining equipment is also slated for automation, including drills, loaders, trains and cranes.
- **Automated sorting and optimised metallurgy**
  Using computer vision and spectral analysis, AI will be able to automatically process minerals with a greater degree of accuracy and efficiency.
- **Predictive maintenance**
  Along with analytics and IoT-driven sensors, AI can optimise maintenance workflows and even predict failures before they happen.

**A number of start-ups have already developed tools that perform deep analytics on geological data and aerial and satellite imagery to better locate mineral deposits.**

## MANUFACTURING

Industrial automation is already a well-established capability, with systems driven by neural networks being employed since the 1990s.[38] Going forward, AI development is likely to focus on improving the efficiency of those processes, including:

- **Self-learning monitoring**
  Integration with internet of things (IoT) devices to monitor and optimise performance and maintenance of industrial equipment. This sensor data can also be used to identify quality and performance and bottlenecks in the production pipeline.
- **Predictive maintenance**
  Similarly, industrial devices can provide feedback to look out for and detect faults before they happen, preventing catastrophic failures and thereby lowering maintenance costs and improving efficiencies by minimising downtime.
- **Intelligent supply chain optimisation**
  Including logistics and demand prediction.
- **AI-driven product design**
  In the long term, there is potential to use AI-driven analytics in product design, pre-emptively identifying and solving problems during the design phase and shaping outcomes.

**AI has been subtly used in retail for some time, with AI-driven data analytics, inventory management and automated customer assistance systems ('chatbots') becoming more commonly employed by retailers.**

## RETAIL

Although the applications are not as dramatic as some other industries, AI has been subtly used in retail for some time, with AI-driven data analytics, inventory management and automated customer assistance systems ('chatbots') becoming more commonly employed by retailers. Some of the other key developments in retail AI include:

- **Product recommendations**
  Online vendors like Amazon, eBay, Netflix and more already use complex machine learning systems to provide product recommendations to customers based on a generated customer profile and predicted needs.
- **Product identification and personalisation**
  AI will more readily allow consumers to personalise their goods before purchase, with visual mock ups and automated short run manufacturing.
- **Demand prediction**
  Using a variety of techniques, including data analytics and sentiment analysis, retailers can use AI to make better predictions about product demand, which can smooth over supply issues and right-size purchases.
- **Customer engagement**
  Through AI-powered assistants that are sophisticated enough to hold meaningful conversations and provide detailed records and analysis of those conversations.

## COMMUNICATIONS AND ENTERTAINMENT

The use of artificial intelligence in communications and entertainment is relatively mature, with major players like Google, Amazon, Microsoft and Facebook being at the forefront of developing systems to personalise entertainment and precisely target advertising. Some of the longer term applications will include:

- **Automated content creation**
  Personalised news feeds were some of the early efforts in the application of AI for entertainment. In the longer term, we will have AI that writes news stories for businesses. These AIs can be used to generate media content as well as within businesses to generate readable narratives from company data. Already some tools, such as Automated Insights' Wordsmith, are available.[40]
- **Personalised marketing and advertising**
  AI will be able to better detect consumer sentiment, reflecting their current mood and better targeting advertisements. Like news content, personalised ad content will likely be generated by AI.
- **Automated telemarketing and polling**
  Increasingly human-sounding AIs will be able to have voice conversations to perform telemarketing activities, polling and appointment scheduling.

## TRANSPORT

The transport industry is well on its way to broad adoption of AI. There are still notable barriers to be overcome – regulatory as well as technical – but it's one of the most active segments of AI development. Some of the current and near future applications include:

- **Autonomous vehicles and driver assistance**
  Autonomous fleets of vehicles will be managed by AI systems – both for the individual navigation of cars, as well as predictive management of car locations and routes. According to McKinsey, automated vehicles will make up 10 to 15% of global car sales by 2030.[41]
- **Engine monitoring and predictive maintenance**
  Where vehicles use IoT sensors and intelligent predictions about usage scenarios to predict maintenance for vehicle components.
- **Enhanced safety and efficacy**
  As autonomous vehicles see increased adoption, safety and efficacy of both public and private transport services will improve.

## THE CHATBOT BOOM

Chatbots are one of the most commoditised AI products available right now. Developments in natural language processing and speech recognition and synthesis have produced chatbots that can intelligently engage with human customers for pre-sales assistance, technical support, medical assistance and many other fields. Even the Australian Taxation Office has introduced a machine-learning chatbot ('Alex'), which has been in operation since 2015 and was later deployed in other government departments.

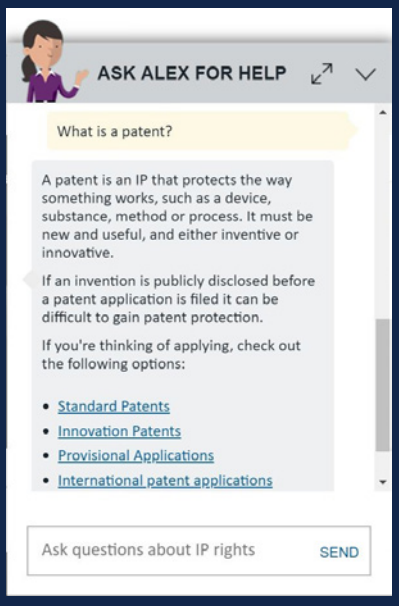


THE AUSTRALIAN GOVERNMENT'S CHATBOT, ALEX, IN OPERATION AT IP AUSTRALIA.

There are numerous companies that produce 'off-the-shelf' chatbots that can be deployed with minimal training, from big players like IBM and Microsoft to smaller start-ups offering boutique products.

One such start-up is FlamingoAI, a Sydney and New York-based start-up with a virtual AI assistant called Rosie & Maggie. The assistant can walk customers through common sales queries, asking questions and offering advice on products that meet the client's needs.

The AI can be trained with pre-seeded queries and responses in an offline training environment, but it can also be trained on the job where the AI is assisted by an employee. As the AI and employee work together, the AI becomes better at learning how to respond to certain queries, effectively training itself in a live environment.

Chatbots aren't just being used for external engagement, either. Another Sydney start-up, Hyper Anna, has developed an AI-powered data virtual assistant that goes through a company's data and responds to natural language queries about that data. The company describes it as Siri for company data, allowing managers to quickly generate reports and access key data stored inside the company's systems.

## STRONG INDUSTRIES

Some industries are more likely to benefit from AI than others in the near term. According to a survey by McKinsey Global Institute, AI development is strongest in the tech and automotive sectors, with construction and travel trailing. The table below shows the current state of AI in different industries according to McKinsey's AI Index.


STRONG AI INDUSTRIES

Source: McKinsey Global Institute

## WHAT BUSINESS LEADERS ARE DOING

In the 2017 Deloitte *State of Cognitive Survey*[43], Deloitte asked 250 'cognitive aware' business leaders what their top three business drivers were for AI. The results are below.

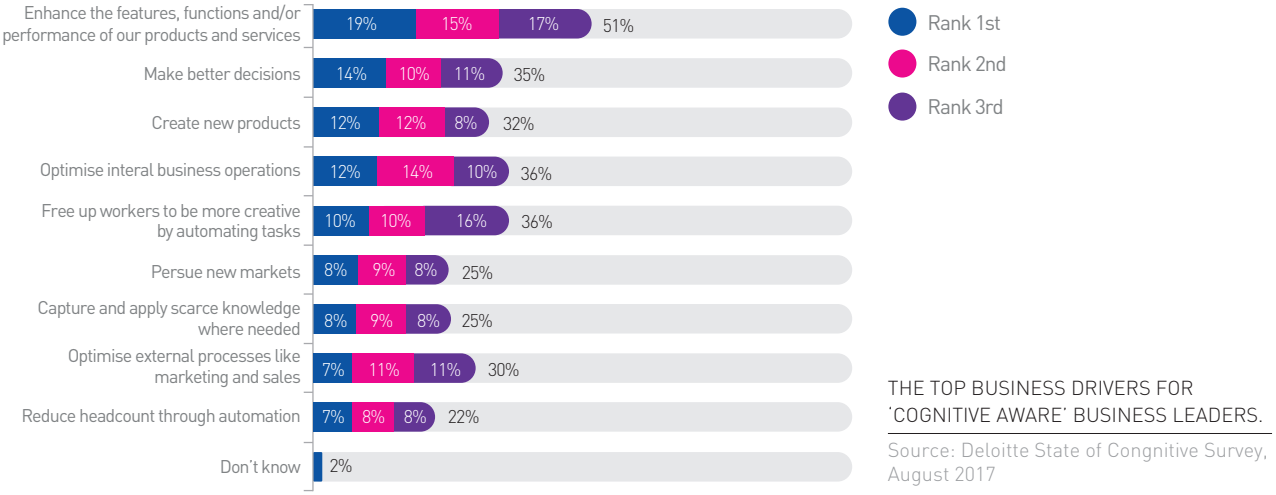


THE TOP BUSINESS DRIVERS FOR 'COGNITIVE AWARE' BUSINESS LEADERS.

Source: Deloitte State of Congnitive Survey, August 2017

## USING AI TO DETECT MELANOMA

One of the most productive areas of AI development is in medical imaging. Trained AI can quickly scan medical images for the tell-tale signs of disease or injury.

In Australia, IBM is working with MoleMap and Melanoma Institute Australia to train IBM's Watson AI to detect the signs of malignant melanoma in photographs of patients.

Trained with a large set of labelled images provided by MoleMap and Melanoma Institute Australia, Watson checks the colour, border irregularity, asymmetry and globular pattern of photographed moles. Comparing it to the trained data, it makes a determination whether a mole is likely to be malignant or not.

Similar projects overseas have had excellent success rates. A test study trained using 130,000 dermatological images in the US found that the AI had detection rates comparable to human specialists with respect to the identification of skin ailments, pointing to a future where AI-powered devices can perform full diagnoses of skin conditions. [54]


## AUTOMATING VIDEO ANALYSIS

New deep learning AI systems are capable of interpreting more than just still images. They are also capable of monitoring video streams and providing reports of activities viewed through the camera. For example, video monitoring systems are being developed that will track customer movement through a retail store to spot areas of greatest engagement – and even provide a mood analysis of customers based on their facial expression.

In Australia, a company called Gameface is employing AI-powered video analysis in a unique way. Gameface is a system that automatically analyses cricket and soccer footage.

For soccer, for example, it uses biometric and number recognition to track players as they move across the field, allowing coaches and media to create a complete map of player positions and movements through the match. It will monitor actions and patterns of players, report on their speed, distance and kicks and tag key events such as goals and free kicks automatically.

In cricket it provides deep analysis of ball movement during bowling, proving batting analysis, fielding analysis and even provide predictive analytics on individual players.

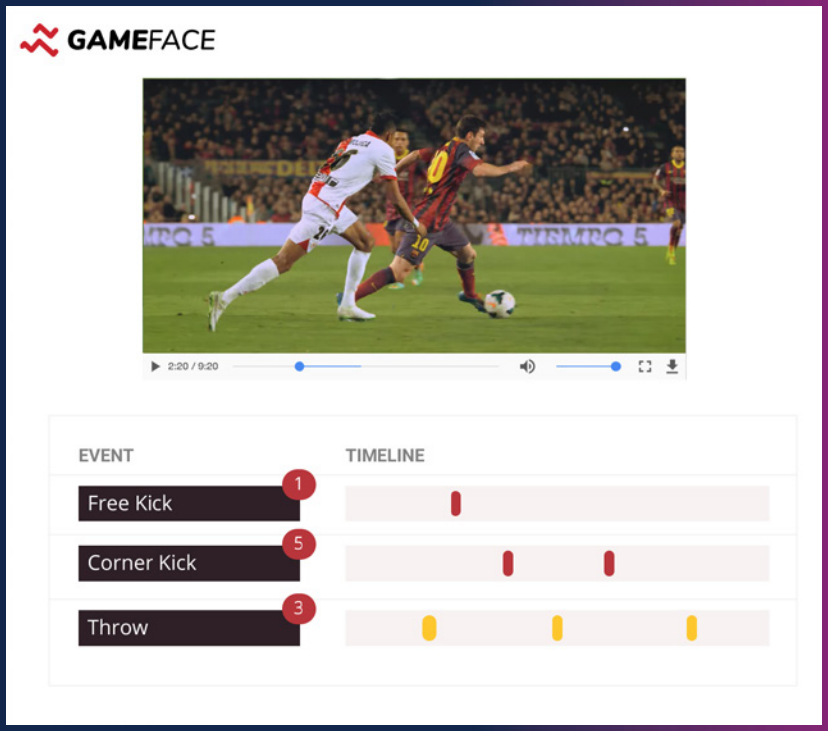


GAMEFACE CAN MONITOR MAJOR EVENTS IN SOCCER MATCHES.

# What the big ICT companies are doing

To date every major player in technology has recognised AI as a key driver of future productivity gains and customer engagement, and most have already established AI labs.

Here we look at what some of the big players in tech are doing to leverage AI. This is just a handful of examples – companies like Salesforce, Meltwater, Intel, Tencent, Tinder and many more are also investing heavily in the field.

Alibaba.com

China is home to some of the most aggressive AI development in the world, with the Ministry of Industry and Information Technology declaring its intent for China to be the world's leading developer of AI systems by 2030, with turnover of over a trillion yuan (around US$207 billion).[44]

Online retail giant Alibaba is certainly doing its part. In 2017 it announced that it would be spending US$15 billion over three years on its DAMO Academy, and most of that would be going towards AI development.[45]

Already the company has achieved significant milestones, including automated package delivery vehicles, a facial recognition system used on Shanghai's subway, reading comprehension systems, and the smart chatbots used on its own website for assisting customers.

amazon

Amazon's investment in R&D is one of the largest in the industry. In addition to using AI for customer product recommendations, it also released Alexa, a voice assistant similar to Apple's Siri or Google Assistant. Amazon has opened up Alexa's API to third party developers, allowing Alexa to be used for tasks other than those programmed by Amazon, and it has also allowed Alexa to be installed on devices other than the Amazon Echo, its home assistant device. Through Amazon Web Services, it also offers automated image recognition and voice recognition systems.

Amazon has also recently opened a supermarket in Seattle called Amazon Go that eliminates the checkout process entirely. At the Amazon Go supermarket, an AI monitors what goods shoppers have picked up from the shelves and will automatically charge their credit cards for the goods when they walk out of the store.

Few companies have done more to bring AI into the popular consciousness than Apple. In 2011 it launched Siri, a virtual assistant that has almost become synonymous with the technology. Much of its work since has been focussed on improving Siri's natural language processing capability and feature set.

In 2018 it also released a new framework for AI developers to compete with the likes of Google's TensorFlow and Microsoft's Cognitive Toolkit.[46] Unlike those providers, however, the new Apple Create ML and Core ML form a toolkit designed to aid in the development of applications that run locally on devices – as opposed to most other solutions that are designed to work in the cloud.

Baidu 百度

Like Alibaba and its competitor Google, search engine provider Baidu has created an extensive network of research facilities for AI. Its research arm, Baidu Research, has facilities in Beijing, Seattle and Silicon Valley and is involved in research in numerous AI fields, including natural language processing, data mining, business intelligence, vehicle and device automation and computer perception.

Some of its most notable projects include a system that is designed to help doctors locate tumours in pathology slides, advanced conversational AI, automated video and image content analysis, speech synthesis (including the ability to clone individual voices) and a significant amount of basic research into core machine learning techniques and systems which has led to the development of the PaddlePaddle AI framework.

Facebook has been making use of deep-learning AI in its products since as early as 2010 when it introduced facial recognition for use in image tagging. It later launched a dedicated AI lab in 2013.

In 2016, it introduced an application that would describe photographs to the vision-impaired, while internally it uses AI to analyse user activity to find out what the keys interests of users are and provide targeted advertising and customised information feeds.

For external businesses looking to implement their own AI systems, Facebook manages the development of two deep learning frameworks – PyTorch and Caffe2, which were merged in 2018.

Google and its parent company Alphabet is one of the world's largest investors in AI technology. In 2014, it purchased DeepMind Technologies, which was later used for the development of AlphaGo, the first AI to beat a professional Go player without handicaps.

Google employs AI across the breadth of its organisation, from image recognition in Google images, to real-time language translation, to voice recognition and natural language processing in Google Assistant. It uses AI to deliver custom ads, sort search results and mine the massive amounts of internet data it collects. It provides access to many of these tools to businesses through its TensorFlow framework, various AI development tools released through PAIR (People + AI Research Initiative) as well as its Google Cloud services.

IBM has one of the longest-running AI research networks in the world. It became famous in 1996 when Deep Blue beat then world champion Garry Kasparov in a chess match. In 2011, IBM made waves again when its Watson natural language question and answer system won first place in a Jeopardy competition against previous champions. Most recently, its Project Debater system engaged in a coherent debate against several Israeli debate champions, being able to deliver arguments and rebuttals in real time on an arbitrary topic.[47]

IBM continues to perform extensive research into the applications of AI in business, and is working on the development of a platform for what it calls 'trusted, safe, and explainable AI systems.'

Like Google, Microsoft employs AI across its business; to provide personal voice assistance through Cortana, to internet and image analysis through its web search engine, to commercial tools available for the development of third party AI solutions.

Some examples of its projects and products include Skype's real-time language translation, which uses machine learning to develop accurate translations of spoken and written communications; Cognitive Services, which performs facial recognition, speech recognition, natural language processing and can even tell someone's mood (which can be used in market research); and Microsoft Dynamics 365, which employs AI for corporate data analytics. Customers who use Microsoft's Azure cloud services can also access Azure AI, which provides both tools and infrastructure for AI applications.

Graphics chip maker Nvidia is rapidly rebuilding itself as an AI company, both selling the chips and computers used for heavy deep learning applications as well as providing a deep learning software development platform. It provides pre-built platforms that are designed to let companies and researchers develop, train and deploy AI solutions quickly.

It has also partnered with various AI companies in recent years. For example, in 2017 it partnered with Nuance, which uses AI-powered image processing to analyse x-rays and other images used in medical radiology. The AI is being trained to recognise certain conditions in the images and report them to medical professionals.[48]

Twitter is an extensive user of AI behind the scenes, having purchased several AI start-ups to enhance Cortex, its in-house engineering team.

Twitter uses an AI to evaluate tweets in real time and provide predictions on their chance to drive engagement. The AI will float the tweets it expects will be of most interest to top of individual user feeds.

In addition, Twitter uses AI to better assist users in cropping images and will automatically label images and video feeds based on their content in order to improve searchability.

Twitter also uses a trained AI to detect hate speech, fake news, pornography and illegal content. Between August 2015 and December 2017, Twitter claims to have shut down 1.2 million terrorist accounts.[49]

# Driving your business ahead

The opportunities, as you can probably see by now, can be as diverse as the industries and markets these businesses and real-world examples operate in. When looking for opportunities for the application of AI within your own business or organisation, consider the following focus points.

## IMPROVING PRODUCTIVITY

One of the major drivers of many current AI projects, productivity improvement often focuses on automating common business processes. These include both physical tasks (driving, harvesting, machine operations, for example) and automation of cognitive tasks (e.g. updating customer information, providing basic customer support, processing forms, sending and chasing invoices, job scheduling).

It's not hard to see why these projects are popular – automation projects generally have clear outcomes and measurable returns on investment. They let staff focus on higher value activities as machines handle routine 'grunt' work. They are also often the simplest to implement, with off-the-shelf solutions available and frequent use of basic decision trees rather complex neural networks.

They also allow greater visibility of processes within an organisation and clearer metrics for business insight. This in turn helps with the process of optimising internal business operations and driving further productivity gains.

### IMPROVING PRODUCTIVITY USING AI

- Reducing work time spent on common tasks.
- Automating vehicle, warehouse and machinery operations.
- Enhanced maintenance and monitoring of equipment.
- Reducing errors in production and cognitive processes.
- Optimising schedules.
- Automating procurement.

## NEW PRODUCTS, FEATURES AND BUSINESS DIRECTIONS

Businesses can use AI to enhance their current product offering or offer entirely new products and services to customers.

A good deal of early work into AI has gone into enhancing the customer experience. That might include providing intelligent product recommendations, offering advice to customers, and providing additional personalisation tools. AI systems can be used to anticipate particular customer demands and produce a more highly personalised experience for that customer.

Business that invest heavily in AI may find new directions for their products and business. The development of customised products, for example, can be enhanced using AI systems, and product offerings can be tiered and costed based on how much human intervention is needed.

### ENHANCING PRODUCTS

- Using virtual agents to provide tips and assistance to customers.
- Providing better recommendations through machine learning.
- Personalising and optimising user experiences.
- Optimising product positioning in virtual and physical stores.

## DEEPER UNDERSTANDING OF CUSTOMERS AND BUSINESS PROCESSES

Deep data analytics have been used in businesses for some time, with business intelligence software becoming widespread in the 1990s. Machine learning AI is improving those products and processes with its capacity for spotting trends and hidden correlations.

As we move along the business intelligence timeline, we're seeing the move from descriptive analytics (providing insight on current and historical data), to predictive analytics (predicting future behaviour) to prescriptive analytics (offering advice on courses of action to achieve optimal outcomes). As AI becomes more advanced in these fields, it will become better at offering forward-looking advice.

For a business, that offers the opportunity to extend insights into both the production and customer-sides of their as well as make better decisions based on that insight. It can let a business better forecast changes in demand, predict costs better and anticipate sourcing and R&D needs.

## BUSINESS INSIGHT

- Smarter forecasting of demand and R&D requirements.
- Improving yield predictions.
- Anticipatory sourcing.
- Better customer analytics and targeting of marketing.
- Identifying and quantifying customer sentiment.
- Better predictions of costs.
- Analysis and optimisation of internal business processes.
- Improving product assortment and positioning.

## MAKING IT PERSONAL

Customer insight systems are using AI to generate increasingly detailed profiles of users. These systems are capable of grouping users based on similarities, providing personalised products and product recommendations based on user profiles and providing analytic insights that extend far beyond just raw sales numbers.

ASX-listed company OpenDNA, for example, specialises in AI-powered psychographics. Drawing data from user activities online and in partnered apps, OpenDNA's Synapse deep learning system builds up a psychographic profile of individual users. It monitors what the users like, dislike, what they click on, when they use the app or site, the kind of language they use and how long they spend on individual items.

Using that information it creates a personalised profile of their activities and interests as well as a emotional and psychological profile of them. That includes personal drivers, emotional analysis and social interaction profile.

OpenDNA customers can then use that information to customise their products and advertisements for that individual, as well as group together users by psychographic profile. For example, an online newspaper could use the information to personalise a news feed, or gain insight into the kinds of articles that will appeal to the most readers.

# 03

# Building AI for your business

By now you should have some idea of the potential of artificial intelligence as a transformative technology, and may already have some ideas on how it could be applied to your business.

As we'll detail in this chapter, while it can be exciting to explore this potential, embarking on an AI project is not like traditional business opportunities. The rewards can be great, but artificial intelligence is still an emerging science, and it can take a different approach than what you might be used to for scoping a project, justifying a business case, and developing and launching a product.

While we've read much on what big tech companies are doing, there's also plenty of smaller development occuring, and along with open-source tools and cheap cloud-based services the barrier for entry can in fact be quite low – meaning SMEs shouldn't be afraid to get involved.

As with any project, it's worth doing your due diligence. So, in this section we'll detail a practical approach to help you scope a project and start the journey from proof of concept to deployment.

# How to leverage AI for your organisation

Before entering into the development of a product, it helps to get your foundations right, and realise that developing an AI product is going to be an iterative process. This process can sometimes lead to new opportunities beyond the problem you're trying to solve.

Generally speaking, there are three key building blocks you'll want to implement when starting this process.

## 1) ASSEMBLE A MULTIDISCIPLINARY TEAM WITH A CLEAR MANDATE

*Bring together your best and brightest minds.*

Ideally the product you create will draw on knowledge, skills, and experience from across the company. It helps to get everyone involved: you'll want commercial understanding for a realistic business model and business case; engineering and IT capabilities to build bespoke technical solutions; and operational expertise to understand current and future processes.

Keep in mind you are very likely going to be hiring for new roles – particularly around machine learning and data science. As these people come on-board, involve them in this process: getting their experience and skillsets in the decision-making process will be critical to success.

It's also important to have a clear mandate from the outset and understand the problem your product is intended to solve so the team can remain focused. Successful projects start with defined commercial outcomes that align with business priorities – AI is only a tool to achieve a purpose and the type of AI or data science approach will critically depend on the outcome you seek.

We'll delve a little further into how define your problem in the coming pages.

## 2) REVIEW AND BUILD OUT YOUR DATA STRATEGY

*Every business today is ultimately a data business.*

Ensure you have a data strategy in place, no matter the size of your organisation. For mature organisations, this may mean restructuring your data from legacy systems into data lakes or moving to cloud. For new entities, take time to develop a strategy to collect and store useful data – even if you are not ready to do AI today, this will allow you to implement AI strategies in future.

Developing AI can require large volumes of data. This is often an organisational challenge and not a technical one, and not being prepared here (capturing appropriate data, understanding governance requirements, securing resourcing to generate extracts of the data) can lead to project delays and cost overruns.

Review your organisation's strategy to capture, store and secure relevant business data. It will also help to look at what data is currently captured versus theoretically available data – doing so may help identify further use cases that exploit the organisation's specific data capabilities. We'll also explore this further in the coming pages.

## 3) DEVELOP AN INNOVATION FRAMEWORK

*Your first AI project won't be your last.*

The process of working an artificial intelligence product will be an iterative one. By nature, developing models and training algorithms can lead to new insights and opportunities, so to get the best results it helps to create a framework that allows you to successfully refine prototypes into deployable products while being open to explore new possibilities. To that end, don't be afraid to adopt a 'fail fast' strategy and set aside resources to test new ideas.

While every company will be different, an example of an opportunity pipeline to explore and encourage new opportunities might look like this:

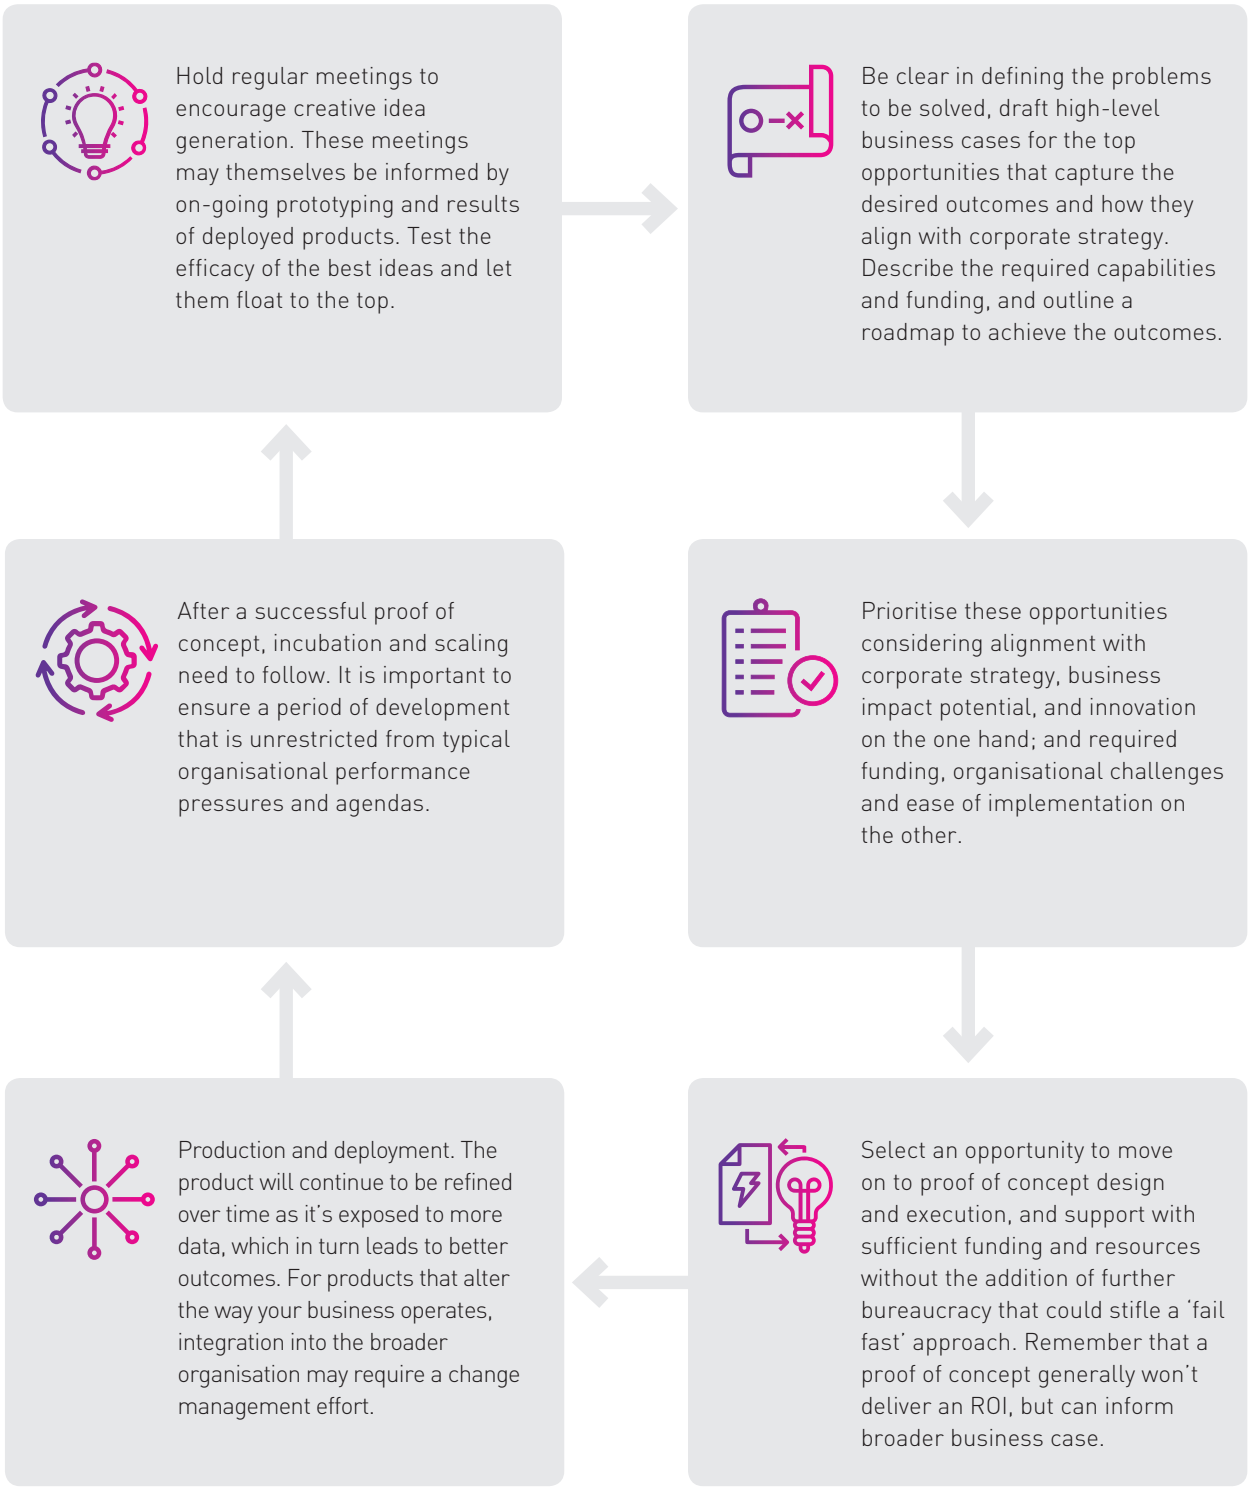

Initially it may be helpful to work with a consultancy to develop a first project. It can be easier to sit with an established team and work through first steps and learn the process before embarking on your own project (see *Build, buy, or outsource?* later in this chapter).

Finally, one more critical step: *adopt an investor's mindset.*

The returns will come in time. Not every product plan will see fruition, but the process can lead to new opportunities with potentially greater impact and return than the one you set out to achieve. Naturally, it will be easier to achieve this if the effort can be focused on specific opportunities that address a clearly defined problem.

So take an investor's portfolio mindset and select a number of opportunities to explore, with enough funding to allow for exploration, but without necessarily expecting a ROI. It may even help to have the organisation establish a venture capital-style fund where opportunities are incubated independently.

## EDUCATING THE C-SUITE

Artificial intelligence gets a lot of press, but it doesn't mean everyone understands the payoffs and the pitfalls. Part of getting buy-in within your organisation may depend on how well you can explain the benefits, challenges, and process.

Every business, and every use case, is different of course. However, there are key themes you can focus on when communicating to key decision makers:

- **It's about the business problems that can be solved**
  A key message to get across is the type of business problems AI can be applied to rather than how the underlying technology works. While it's helpful that everyone has a basic understanding, home in on the benefits that can be delivered with the help of AI and what the opportunities are for the organisation.
- **Be realistic about the value of your organisation's data**
  Data is the lifeblood of AI: it is built with, trained with, and operates on data. You must be realistic here – if your organisation is not currently data focused, now is the time to make that transition. Factor in what will be required to make this happen first. Establish the value of the data you currently have, and the value of the data you are able to capture.
- **You don't have to re-invent the wheel**
  AI isn't just for the large companies with lots of money. Third-party tools have already started to commoditise AI, so you don't necessarily need to invest in building your own, at least for simple solutions. Instead you can leverage current tools, many of them free, to save both development costs and time. For those with more cash, partnerships with established AI suppliers are also an option. Further, you can learn from the successes and failures of others: plenty are playing in this space, so read about what has worked and what hasn't for others in your industry.
- **But you will need to invest in people**
  The skillsets to create AI products are relatively new and in demand. Engineers, analysts, researches and programmers are par for the course. For in-house staff, you will likely need to offer competitive salaries to get the best talent on board. Alternatively, you can outsource and leverage partnerships to plug the skills gap or in some cases retrain existing staff with similar skillsets.
- **Be prepared for change**
  Embarking on a project may require a shift in the strategy and culture of the company – albeit in an exciting new direction. The development process itself can require multi-disciplinary teams from different business units, and the implementation of a product that's aimed at improving business processes will naturally have impacts on staff roles and expectations.

## How to be clear on the problem to solve

It's the core question: what can AI do you for your organisation?

First, remember, artificial intelligence is not a panacea. It's still – for the moment – a narrowly defined tool. But it's a very powerful one that's enabling entirely new ways of optimising business operations and generating completely new revenue streams.

A good place to start is to catalogue your business processes. There are likely to be many areas that are ripe for automation or augmentation. This can include procedures and decisions that are made frequently and consistently. Collect as much data as you can around the procedure or decision, and how the data informs the result – you want a clear picture of what constitutes each step in the process to determine how it could be improved by machine learning algorithms.

And while it's tempting to think big, start by focusing on simple problems. Artificial intelligence can excel with repetition and automation, and where the problem is well defined with plenty of data to train – such as with the credit card fraud example introduced in *Chapter 1 – Why AI*?

Another avenue to explore is augmentation and business process decision making. With the ability to classify information and make predictions at higher volumes and much faster than humans can, this is a rapidly growing field and can provide substantial benefits with respect to optimisation and efficiency gains. This can also be useful as a stepping stone to more complex products or as an avenue to gain greater insights.

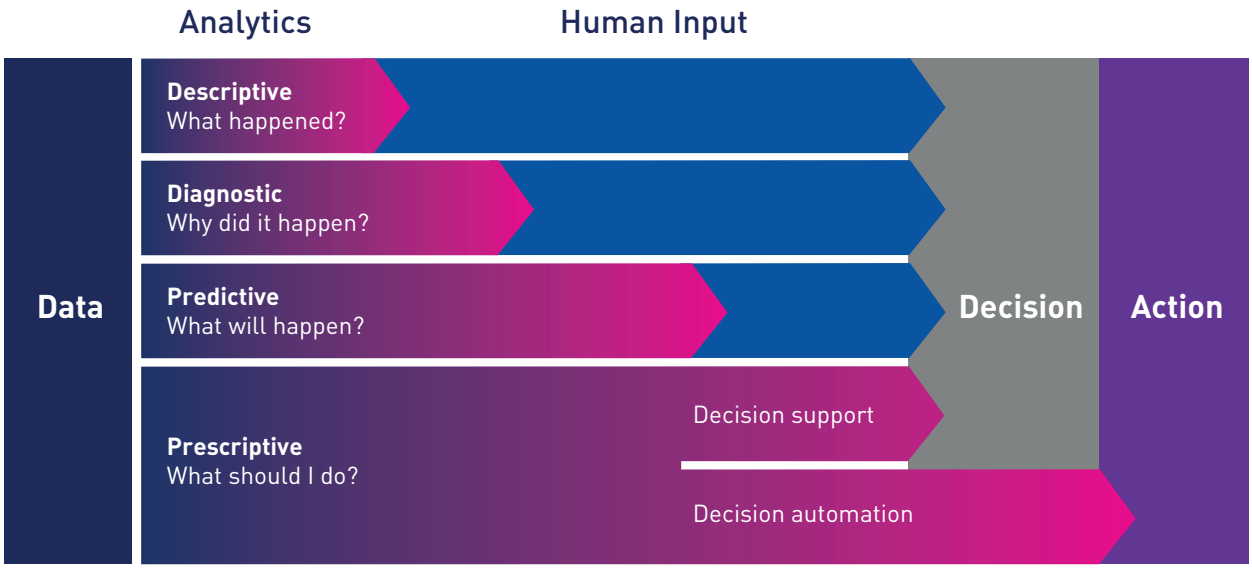


LEVELS OF BUSINESS DECISION AUTOMATION WITH AI.

Source: Gartner.

Gartner, in its *A CIO's Guide to AI*, recommends aiming for 'soft outcomes' to begin with, such as improvements to processes, customer satisfaction, and current products. Whit Andrews, Vice President and Distinguished Analyst at Gartner goes on to say, "Look at how you are using technology today during critical interactions with customers – business moments – and consider how the value of that moment could be increased, then apply AI to those points for additional business value."[55]

If you're looking for inspiration, it's a good idea to scan the competitive pressures and recent technological developments within your sector. Look at when you expect them to arrive, and how you will respond. This will help identify opportunities or new product initiatives that AI could open for your business, and proactively position the organisation for what's coming up on the horizon.

**"Look at how you are using technology today during critical interactions with customers – business moments – and consider how the value of that moment could be increased, then apply AI to those points for additional business value."**

## Ask the right questions

To help define the problem for your business, start with the following questions:

1. What is the business opportunity or problem to be addressed? How will a solution here improve the organisation?
2. What kind of business process transformation will solve this problem? Will an AI solution replace or augment?
3. Is your business model vulnerable to disruption if someone beats you to it in your market? How soon could this change arrive?
4. Are there game-changing openings within your industry and, if so, how can you take advantage?
5. Are there any technologies out there today that have solved similar problems?

Then, once you've determined that there's an opportunity for solving your problem, other good questions to ask include:

1. What third party platforms, products or services can be leveraged to streamline or reduce the cost of prototyping?
2. How will implementation of an AI solution impact staff of the organisation? Can you retrain or reskill existing staff to work on a project?
3. How are you going to collect data for the project? What policies or regulations need to be followed?
4. How will the solution integrate into existing internal and external systems? What level of funding is expected to be required?

Finally, be clear on what success looks like and what metrics you want to measure. And remember this is an iterative process – the goal is to identify the top opportunities to build high-level business cases for and get the ball rolling.

AI CAPABILITIES AND COMMON AREAS OF APPLICATION.

Source: Accenture.

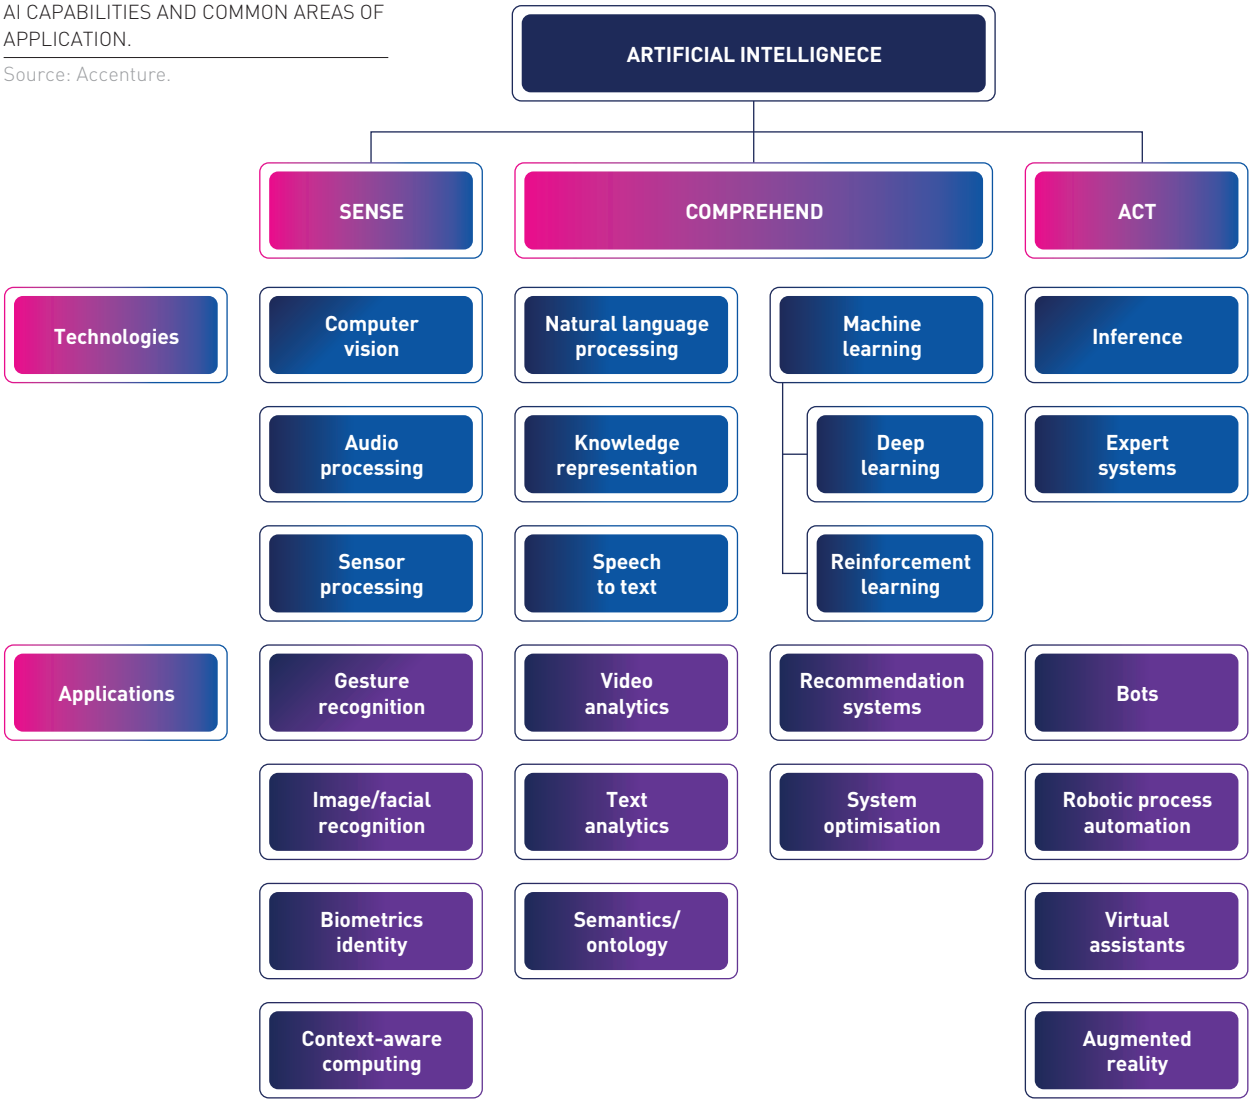

**When it comes to implementation, don't be afraid to get help – especially if prototyping your first project.**

## How to build the business case

As outlined earlier, it's worth looking at the production an AI prototype as an exploratory exercise – because for the most part, it will be. To this end it will help the business case to frame it as a research activity, not a product development activity – there are different expectations, and different returns on investment. And if it turns out that AI isn't going to work here, don't be afraid to pull the pin.

The following guidelines, courtesy of Forbes' *Building AI: Key Steps For Adoption And Scaling Up*[56], provide good advice for your business case:

1. Start with a small problem in a process or objective that will measurably benefit from a quick deployment of AI and you'll be able to make the case for cost savings, efficiency and competitive advantage.

2. Avoid a large-scale initiatives right out of the gate – it will almost surely fail on some level, dimming enthusiasm and momentum.

3. But be ready to scale the prototype – and have another project ready to go whether the first effort fails or not.

The advice to start small is good here – at least to begin with. And as discussed earlier, be clear that when targeting a proof of concept that it likely won't provide an ROI.

As always, detail the desired outcomes – be it lowering overheads or improving the customer experience – and how they align to the corporate strategy, and build a roadmap of how it will achieve these outcomes. Just include the caveat that as an exploratory process the goal-posts can sometimes change.

Outline other benefits that can arise from a successful implementation as well – these might be freeing up resources to provide extra support in areas of the business where AI cannot be applied, or cost savings from increased efficiencies allowing for reinvestment in other areas.

When it comes to implementation, don't be afraid to get help – especially if prototyping your first project. There are many companies today that can help at every stage of the process: from identifying opportunities and the problems to solve, to working with you to build out your identified solution. Leverage these to gain essential experience and strengthen the knowledge base of your own team for future projects.

Finally, don't forget to identify and address ethical concerns. Sensible ethical framing can get buy-in by executives and decision-makers by reducing costs to rework should ethical issues later be discovered. You can find more on this in *Chapter 4 – Ethical considerations*.

On the whole, embarking on an AI product is not like building a website. It's exploratory in nature. James Wilson, Eliiza CEO, offers this advice: "Don't build a business case with a specific ROI in mind, otherwise the model can just be manipulated to meet this – rather than seeing whatever potential may arise."[57]

Similarly, in *Build the AI Business Case*[58] Gartner notes: "Don't worry about immediate ROI. CIOs should begin with small experiments that are purely for learning. Lose the pressure to get an initial ROI and use the time to pilot projects that employ a variety of technologies to assess which make the most sense for the business."

In summary, include the following in your business case:

1

**Build in freedom**
Be open to explore, and encourage creativity. Give your team the freedom to experiment.

2

**Build in failure**
Allow for failure, it's part of the process. Allow for the parameters to change.

3

**Build in time for iteration**
Even a small model and data set can take hundreds of hours to refine.

4

**Build in the ethical dimension**
Ethics is integral to AI design, especially where human well-being is concerned.

## GARTNER'S BUSINESS CASE GOTCHAS

Building a business case for AI can differ from a typical business case, as Gartner outlines in *Build the AI Business Case*[59]:

- **AI solutions can seem costly without providing any immediate gain**
  Particularly for loosely bound scenarios and in organizations that aren't used to setting aside budget to develop and deploy solutions for new business scenarios.
- **AI projects require different technology and problem-solving skills**
  Talent acquisition is likely to be one of the biggest barriers to AI adoption.
- **AI business cases will require substantial cultural change**
  The mindset shift required for AI can lead to 'cultural anxiety' because it calls for a change in behaviours and ways of thinking. CIOs should acknowledge the cultural change, be proactive in managing related challenges and build trust over time. Cultural change and successful transitions to new roles and practices are dependent on open dialogue and mutual respect between management and staff.
- **The need to spend more time on data, training and algorithms**
  AI depends on data and the interactions of algorithms, so when it comes to AI business cases and plans, success is largely dependent on robust data and analytics infrastructure. Organisations looking to move forward with AI must be constantly learning and iterating on algorithms and on how to select, prepare and apply data.
- **AI is the representation of a decision model, rather than a process event**
  The same characteristics that make AI a good solution for dealing with data are also the attributes that make it difficult to support with a standard business case. Traditional analytic models use fixed rules to arrive at a conclusion. However, AI analytics use dynamic data and heuristic solutions to arrive at conclusions that might elude traditional methods.

# Practical advice for building your AI capability

When planning your opportunity pipeline, defining the problem to solve, and building a business case it will help to have a solid foundation on the practical aspects of building an AI product.

Specifically, the following three sections deal with the three most important components of building out your organisation AI capabilities: data, people, and implementation.

## Focusing on data as a strategy

As outlined earlier, there is no AI without data. AI algorithms are trained with data, operate on data, and produce data.

So it's vitally important to have a full understanding of what data the organisation currently has, what data might be needed to solve the business problems you've identified, and what data might reveal new opportunities – and then how to manage and safeguard all of it.

As Alan D. Duncan, Research Vice President at Gartner, puts it: you will need to "establish the basic vocabulary of AI — a technical dialect of how people 'speak data'... Building data management expertise along the full process with be key to success"[60]

And you'll need a lot of it: machine learning and deep learning algorithms can require thousands of examples to train them to acceptable levels, and that's per category (for example, if you were creating a model to recognise invoices to accelerate payment processes, you would need thousands of examples of invoices). To reach or exceed human-level performance, millions of examples can potentially be required.[61]

For this reason your organisation will need, if you don't have one already, a data strategy to capture, store, and manage large volumes of data – and to do so at scale. There are different solutions to this that are beyond the scope of this guide, but if you're an SME it's likely going to be easier and far more cost-effective to take advantage of cloud-based services, such as those offered by Google, Amazon, or Microsoft. This is especially true when it comes to developing an AI product, as many of the cloud storage service products from the big players can also provide data analytics and optionally machine learning tools along with data storage and management, making it easier to get a project off the ground.

When it comes to data itself – which ultimately will be used to train and refine your AI product – it can be a good idea, depending on the project, to have as diverse a range of data as you can source in order to accommodate all the possible permutations that the algorithms will eventually operate on. For example, if you were building a model to recognise dogs in images, you'd want to supply thousands of images of every breed of dog, and every size of dog, and every age of dog.

Despite the time, processing power, and large volumes of data it can take to train algorithms, once complete execution of the model itself can be quite quick – the latest smartphones that come with 'AI processing' for example are able to identify objects in a dozen images each second with ease.

**The first step in deriving value from your data for AI is to catalogue your structured and unstructured business data.**

## DERIVING VALUE FROM DATA

Generally speaking, the data your organisation possesses will usually consist of structured and unstructured data. Structured data is data that's catalogued and organised, such as payment receipts or customer information like names and phone numbers. It's data you can search and apply analytics to.

Unstructured data, as you might expect, refers to data that hasn't yet been categorised or managed, such as raw images, the contents of emails or social media posts, or the input from IoT devices, websites, and mobile apps. This data can be variable, complex and often hard to categorise – but this doesn't make it less useful. After all, one of the key advantages of technologies like machine learning is the ability find and connect patterns within data.

The first step in deriving value from your data for AI is to catalogue your structured and unstructured business data. Even though you may know where the value currently lies, new opportunities may arise thanks to AI, so with your catalogued data in hand next ask the following questions:

- What data is largely predictable or static, and what data is variable or frequently changing?
- What sources does the data come from? Are there sources we are not currently capturing?
- What other uses cases exist for our structured data?
- What potential use cases could exist for our unstructured data?

These are key questions to ask in your opportunity pipeline meetings to help identify where the application of machine learning can make a difference. What potential AI products do you think could utilise the structured and unstructured data to add value to the business? Could there by opportunities for increased efficiencies and automation, or the deployment of an entirely new product for the business?

Data analysts and data engineers will be invaluable in this process (see *Sourcing skills and roles* later in this chapter).

Don't underestimate the value of data. If you are an established organisation in your field, data is a key competitive advantage. Even if a disruptor comes into your market, it won't have the advantage that you do of the data you have on your customers. So capturing this and putting it work smartly through, for example, augmented business decision making through AI is a tremendous advantage.

Similarly, if your organisation operates in a highly competitive market, capturing more data than your competitors will allow you to create better products and experiences for your customers, and again this can be leveraged through machine learning to deliver enhanced results and value that previously wasn't possible.

That said, it's not *quite* as simple as just capture everything and see how you go – which we'll take a look at next.

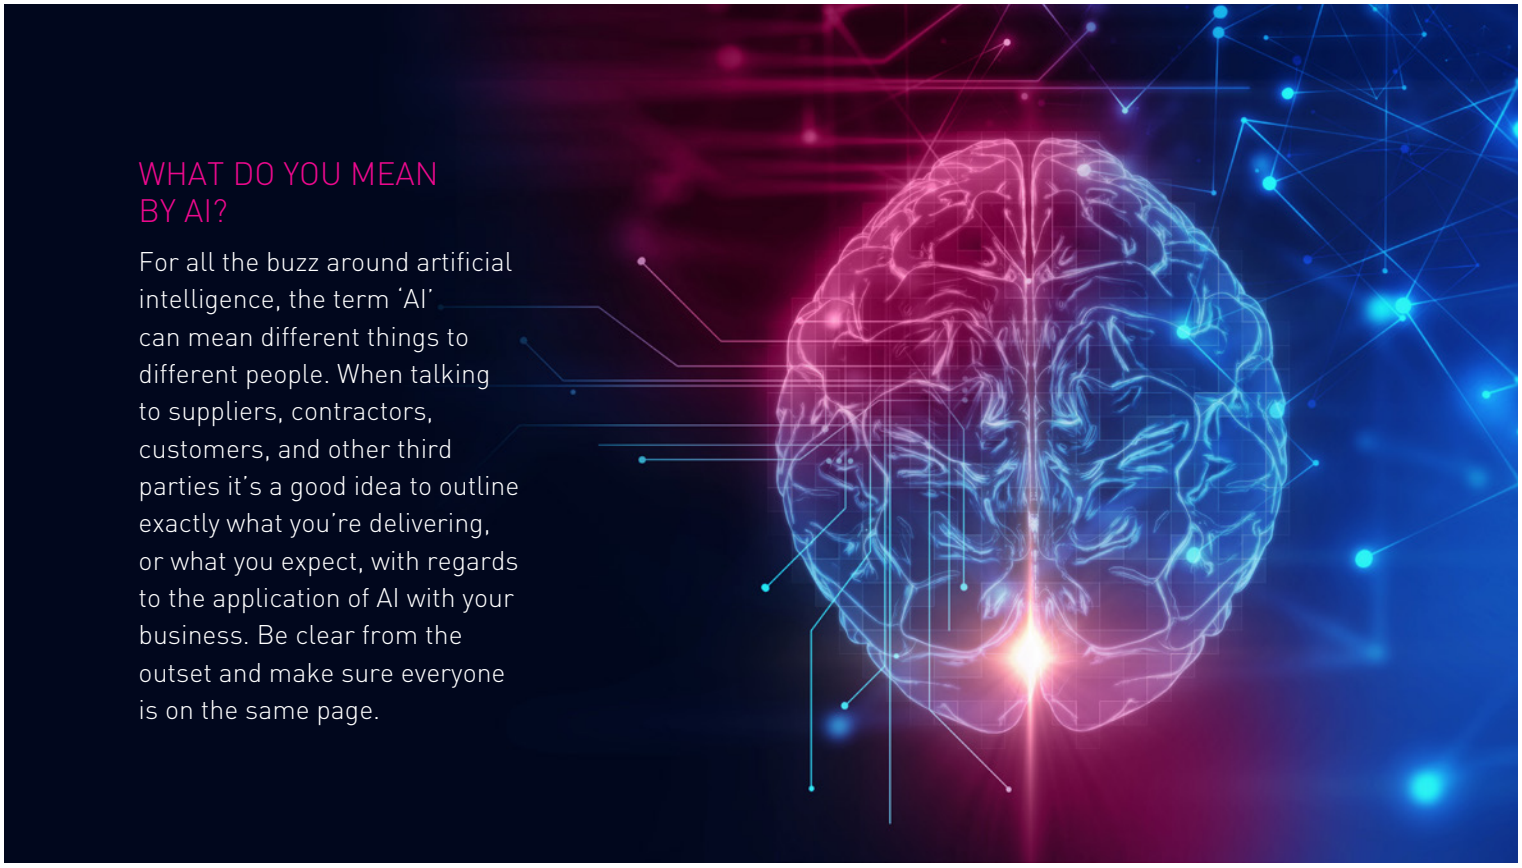

## WHAT DO YOU MEAN BY AI?

For all the buzz around artificial intelligence, the term 'AI' can mean different things to different people. When talking to suppliers, contractors, customers, and other third parties it's a good idea to outline exactly what you're delivering, or what you expect, with regards to the application of AI with your business. Be clear from the outset and make sure everyone is on the same page.

**When looking at what data to capture it's worth asking "just because we can, should we?"**

## SOURCING DATA

Even if your organisation is doing a good job of capturing and cataloguing data, it may not always be the right data for prototyping and training AI. A healthcare company wanting to predict treatment options may not have access to type of data it would like due to privacy restrictions, for example.

In fact, PII (Personally Identifiable Information) can come with its own challenges. If you elect to capture and store PII, be aware of the compliance obligations and strict cyber security requirements and ensure that the organisation abides by the Privacy Act. No company is un-hackable, and the unintended release of PII could be catastrophic for a company. In Australia we have the mandatory Notifiable Data Breaches scheme that requires a company or organisation to notify the OAIC (Office of the Australian Information Commissioner) and all individuals affected by a breach where PII is involved, with hefty fines for failing to do so.

Similarly, if you have operations overseas, be sure to brush up on regulations that may impose strict operational requirements, like Europe's GDPR (General Data Protection Regulation).

Considering the above, when looking at what data to capture it's worth asking "just because we can, should we?"

While data on your customers is important, look at what data actually delivers value and don't capture it if you don't need it. It will greatly simplify every stage of business process, data management, access, and storage.

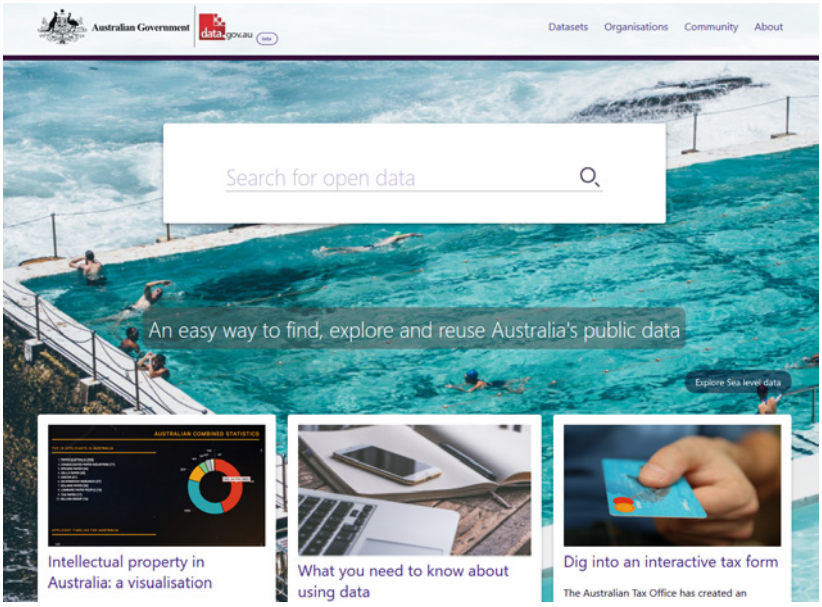


THE AUSTRALIAN GOVERNMENT'S OPEN DATA PORTAL, DATA.GOV.AU.

The other side of the coin here is commercially sensitive business data. Be wary of utilising any data your organisation captures that may be commercially sensitive data from customers or partners, or be bound by contractual agreements.

If you have access to data sets with PII but wish to clean them before storage and use, Data61 has an excellent guide focusing on operational advice in it's *The De-Identification Decision-Making Framework* publication at publications.csiro.au. The ACS has also developed a framework for secure and privacy-preserving data sharing, appropriately titled *Privacy in Data Sharing – A Guide for Business and Government* which can be found at www.acs.org.au/insightsandpublications/publications.html.

In some cases, public and de-identified data sets may be available. For example, the government maintains a public-sector data portal at data.gov.au that contains data for everything from weather and geoscience to transport and housing. The US, UK and Europe also have similar portals.

Likewise, there are also data science communities – such as kaggle.com – where open data sets are shared for testing and prototyping AI. These may not necessarily prove useful in producing a final product, but can be very useful for testing different models and how to leverage different types of data without spending time and money exposing or structuring your business data before you know what works.

Because of the value of large volumes of data for training and refining, one other option may be available for your industry: purchasing (and possibly sharing) data. For example, in 2016 a consortium German automakers including BMW, Daimler, and Audi acquired HERE, a provider of mapping and location services similar to Google Maps. The car makers share the data generated by each vehicle to help power their autonomous vehicle technologies.[62]

However the data is sourced, whether it's external or your own organisation's data, there is a golden rule to follow before you begin refining your models: review your ethical guidelines. This isn't just about doing the right thing – it's also about achieving accurate results. For both these reasons we have a dedicated chapter on ethics coming up next, *Chapter 4 – Ethical considerations*.

## STORING AND CLASSIFYING DATA

You're going to be working with large volumes of data, structured and unstructured, and your data strategy will need to consider its storage and accessibility.

As outlined earlier, for an initial project it's likely going to be faster and more cost effective to take advantage of cloud services like those offered by Google, Amazon, or Microsoft. A term you will often hear is *data lake* – essentially the complete archive of all your raw data sources stored in its natural format, and from which you can query and manipulate to feed into training algorithms.

Utilising these services has other advantages too: they often come with tools that allow you to easily build different levels of classification that define how the data is managed and accessed. This can be beneficial for your data engineers when it comes to structuring queries to build data sets for training.

Factors to consider for storage and management include: the velocity of data (how fast data becomes available for analysis), the volume (size of the data in storage), the quality of data (how clean the data is), and the interactivity of data (speed of query requests).

Velocity is important, as your machine learning algorithms will need to be retrained as conditions change in the business landscape, so the training data must be updated and refreshed frequently. The granularity of this depends on the business problem you want to solve – your particular solution may only need quarterly or monthly updates, while others may require daily updates to ensure the model performs as expected.

During development you will find that as the project progresses, the analysts and engineers will want to test models with a range of different data sets pulled from your data lake. For this reason it's always a good idea to retain the original, complete, structured and unstructured data sources and use classification and queries to manage the data rather than try and focus on just the data you think you need for the project. As we covered earlier – building an AI solution is an exploration, and you'll want to give your team the freedom to experiment.

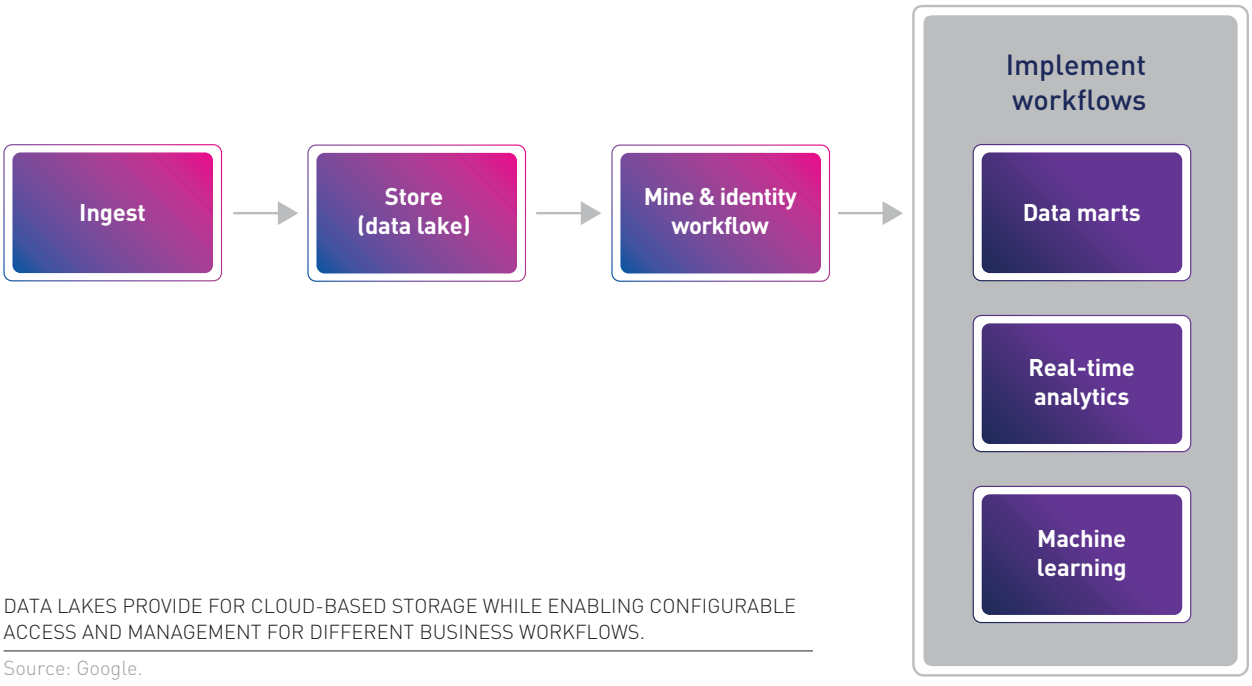


DATA LAKES PROVIDE FOR CLOUD-BASED STORAGE WHILE ENABLING CONFIGURABLE ACCESS AND MANAGEMENT FOR DIFFERENT BUSINESS WORKFLOWS.

Source: Google.

**Ultimately, if the data you want to use is your business intelligence there's only one group that can label the data for you – your organisation's employees.**

## LABELLING DATA

When it comes to *supervised learning*, there's one important step before any data is ready to test and train algorithms – the potentially laborious process of *labelling*. Even though it may be impressive that you can use Google Image search to find pictures of a cat, at one stage Google's algorithms didn't know what a cat looked like. Just as with educating humans, it had to be told.

For example, if you wanted to train your algorithms to recognise tasty donuts in an image, you first need to feed them with hundreds or even thousands of pictures that have been manually labelled as 'donut' by human beings – after all, who else can provide this information?

As the images are fed into the system, so too is an image label, and utilising this it can start to recognise the common points of reference in the image that make up the object you want it to identify.

While all data needs to be labelled, not all data comes unlabelled. Data in an Excel sheet may already come with labels – think of the 'price' field for an Excel file of house prices, for example. As long as you can tell the algorithm what that fields represents, it can take that as the label for the data contained within.

One advantage of publicly-available data sets mentioned earlier is that in many cases these will already be labelled, and is one of the reasons they can help accelerate testing different data sets for different models before you delve into your own business data.

Ultimately, if the data you want to use is your business intelligence, there's usually only one group that can label the data for you: the organisation's employees. This means at some point, some of the staff at your organisation are going to need to get their hands dirty labelling. While you could employ data entry specialists for this task, this could also be a good excercice for newer team members as a means to familiarise themselves with the data. Don't be afraid to spice this up – hold a 'label' party and throw in food and drink, make it a team activity!

There are also other options: you may be able to find freely available plugins to assist in labelling your data, and can even crowd-source the labelling of data – such as with Figure Eight (www.figure-eight.com), a *Mechnical Turk* version of labelling – but it's still done by humans and you may find a high error rate that needs to be corrected through multiple labelling runs.

Finally, note not all learning need rely on labelled data. Any unsupervised learning as well as generative adversarial networks (where systems compete and learn from each other), reinforcement learning (where a system determines for itself what the correct process is via trial and error), and transfer learning (where a system is trained to accomplish a certain task, and then applies that learning to a similar but distinct task) are among other methods for AI to learn. As the field advances, we'll also see even more advanced methods such as *one-shot learning*, which will allow an AI system to learn about a subject based only on a small number of real-world demonstrations – possibly even just one.[63]

## BUILDING FOR THE RIGHT REASONS

According to a prediction by Gartner, AI technologies will be a top five investment priority for more than 30 percent of CIOs by 2020, with the inclusion of AI in some form embedded in almost every new software product.

But not because it's always needed. Some are leveraging the hype by jumping on the AI bandwagon. "As AI accelerates up the Hype Cycle, many software providers are looking to stake their claim in the biggest gold rush in recent years," said Jim Hare, research vice president at Gartner. "AI offers exciting possibilities, but unfortunately, most vendors are focused on the goal of simply building and marketing an AI-based product rather than first identifying needs, potential uses and the business value to customers."[64]

**AI is nothing without data and so a data analyst and data scientist are realistically your first ports of call.**

## Sourcing skills and roles

The most important component of your project: people. What type of roles and skills do you need to build an AI product?

Working with machine learning requires different technology and problem-solving skills than you may have in your organisation. Talent acquisition is likely to be one of the biggest barriers to AI adoption, as demand for AI-centric skill sets is high.

At a minimum, *machine learning engineers* will be your mainstay – and for some projects this may be all you need. Then, depending on the scope and size of the project, other roles and skills to consider can include:

- Data analysts
- Data engineers
- Data scientists

As we will detail in the next chapter, it's also worth looking at potential ethical considerations for your project. If you're a large company, or there's a lot at stake, bringing an ethicist in-house can be valuable, but for smaller companies outsourcing or employing consultants will suffice.

Then, you will also need software engineers to build final products that you can deploy in the organisation or later release to market.

Depending on the complexity of the project you may not need all of these skillsets, and some can be brought onboard later in the development pipeline. As outlined earlier AI is nothing without data, so data analysts and scientists can be a good place to start.

Collaboration should be core – focus on a multi-disciplinary, self-contained team. If you have the above skillsets in your organisation already, try not to utilise them across different business units if you can afford not to. Dedicate them and bring them all under one umbrella to work. Own the problem as a collective team.

Also bring your HR managers into the project early so they can learn about the technologies as they hunt for top-tier talent – finding the best people is a critical business goal. Advertise early, it may take some time to find the right people.

If not onboard already, employing executives with technical acumen will help lead and find new opportunities for leveraging AI in your business as your project progresses.

If you elect to source outside help, pick your strategic partners early. In fact, for the first project, even with an in-house team it's not a bad idea – this can accelerate your initial project, and help the team learn from the experience and acquire extra knowledge to boost your in-house capabilities for future projects.

## Delivering on AI

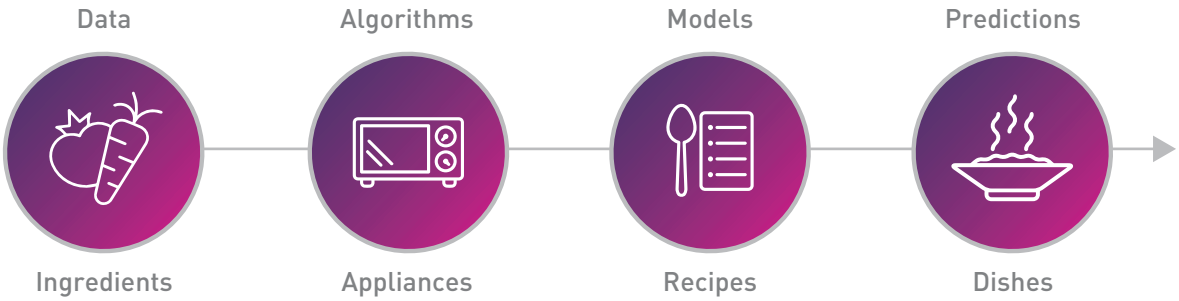


### Core Roles

- Decide what you're going to make: decision-maker
- Which ingredients to use: data analyst
- Where to get ingredients: data engineer
- How to customise the recipe: machine learning engineer
- Taste test: data scientist

COOKING UP AN AI. WHO DOES WHAT, WHEN.

Source: Eliiza.

Other staff you may need depends on the product you are developing: for example, an interactive chatbot for customer service would benefit from having skilled writers or communications specialists onboard. Other skills you may need depending on your project could include pure mathematics, data science, psychometrics, behavioural psychology, neuroscience, and linguistics. With some limitations, the possibilities for the application of AI are really up to your imagination.

### PLAY THE LONG GAME

It can be tempting to get an AI product to market fast, especially if you're in a competitive space. But don't underestimate how long producing and piloting a product can take.

ACCORDING TO RESULTS FROM GARTNER'S 2017 ANNUAL ENTERPRISE SURVEY,

**58%**

OF RESPONDENTS IN COMPANIES CURRENTLY PILOTING AI PROJECTS SAY IT TOOK TWO OR MORE YEARS TO REACH THE PILOTING PHASE

WITH ONLY

**28%**

OF RESPONDENTS REPORTING GETTING PAST THE PLANNING STAGE IN THE FIRST YEAR.

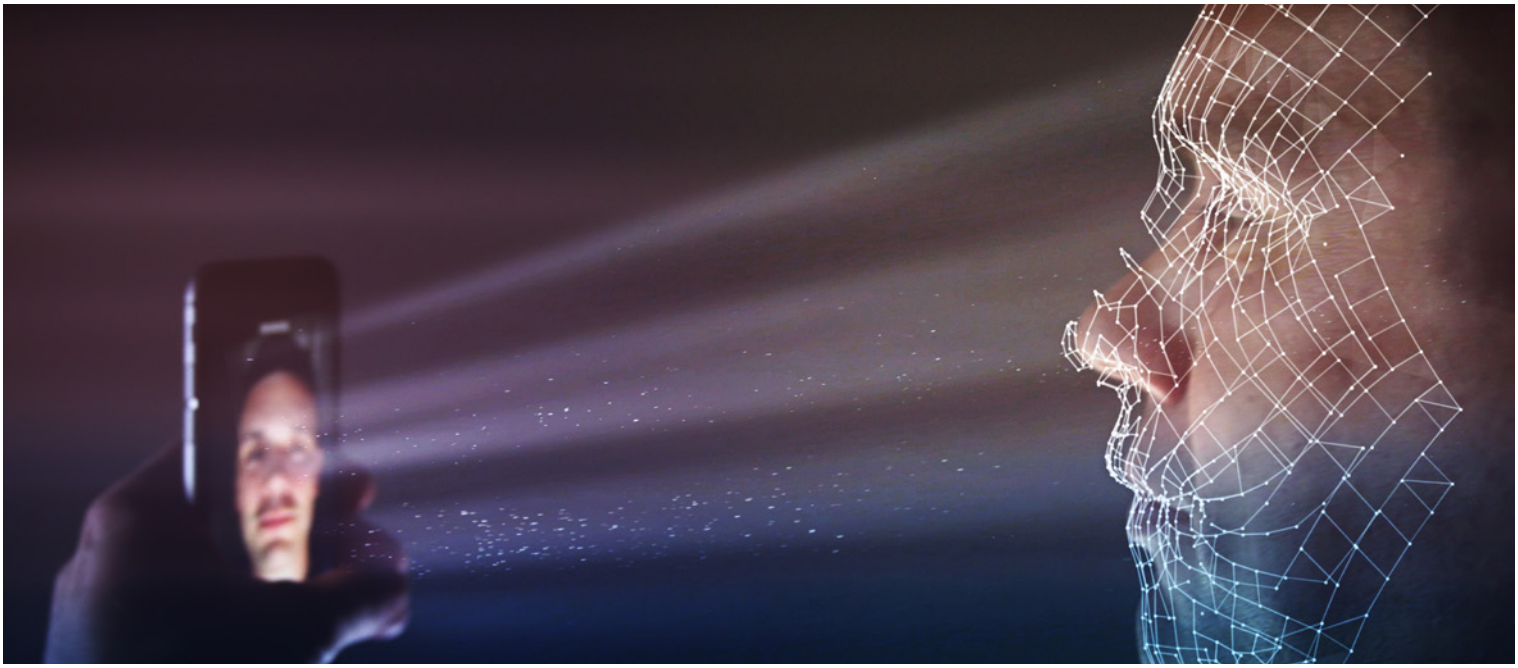

## Build, buy, or outsource?

While you can build your product from the ground up in-house, and for larger or complex projects this may be necessary at some point, you don't necessarily need to – especially for an initial prototype project.

The key for your first AI pilot is to minimise complexity. We recommend deciding early to either take advantage of third-party off-the-shelf products and services, or outsource part or all of the project, to greatly accelerate the prototyping process.

Another option is to engage start-ups that are often more advanced in many areas of AI, may even have developed the solution being sought, and would most likely welcome an industry partner. They can also be good for sandboxing ideas by providing (limited) access to data to create a proof of concept. Start-ups also like money, so it's a win for everyone.

For the former there are various solutions you can leverage today as AI development has begun to become commoditised. While this is not going to be an exhaustive guide – there are many options available, and your machine learning engineer and data engineer may have their own preferences on what they want to use – the following will provide a good overview and grounding in what is required.

### CLOUD-BASED AI TOOLS

Tech giants Google, Amazon, Microsoft and IBM have done a lot of the leg work to make AI-based services accessible and cheap to use – be it the storage and categorisation of data, the processing power for an AI model to learn from thousands of samples, or an interface to easily code algorithms to test and refine. It's become so commoditised in fact, you can even do all this directly in a browser.

This can even extend to building out deployable products – producing APIs for your own projects, or making them available directly on the web. It comes at a cost, of course, but a fraction of what you would spend purchasing, deploying, and managing your own hardware. Especially from a computation point of view, unless you are a Google it would be exceedingly more expensive to run models on your own hardware than leasing supercomputing time or leveraging the distributed computing clouds such as Google, Amazon, Microsoft and IBM provide.

That said, check the service's terms and conditions, especially when it comes to uploading your organisation's data for training. Don't forget you may be bound by regulatory or security concerns.

The cost structures depend entirely on the project you are building, but are still going to be considerably cheaper than starting from scratch in-house. Review the services on offer with your team (see table *Popular cloud AI service solutions* for a sample) to determine which feature sets best fit your needs, noting that some services will be charged based on time or service utilisation.

## Popular cloud AI service solutions and features

| AI Service | Low Level | High Level |
|---|---|---|
| Microsoft Azure | ML Studio<br>• Anomaly Detection<br>• Classification (binary/multi-class)<br>• Clustering<br>• Statistical functions<br>• Text analytics<br>• Computer vision | Cognitive Services<br>• Vision<br>• Speech<br>• Knowledge<br>• Language<br>• Search |
| Amazon Web Services | AWS Machine Learning<br>• Algorithms:<br>Regression, classification (binary/multi-class) | Amazon Lex<br>• Natural language understanding<br>• Automatic speech recognition<br>Amazon Rekognition<br>• Visual search and image recognition<br>Amazon Polly<br>• Text-to-speech<br>AWS Deep Learning AMI<br>• Custom AI models |
| Google Cloud | Large-scale machine learning service<br>• Custom models from regression models to image classification based on deep learning | • Google Cloud Job Discovery<br>• Google Cloud Video Intelligence<br>• Google Cloud Vision<br>• Google Cloud Speech<br>• Google Natural Language<br>• Google Cloud Translation |
| IBM Watson | IBM Data Science Experience<br>• Spark ML algorithms<br>• RStudio<br>• Deep Learning libraries | Watson Developer<br>• Conversation<br>• Knowledge (Discovery, NLU, Document Conversion)<br>• Vision<br>• Speech<br>• Language (Translator, Classifier, Retrieve & Rank)<br>• Empathy (Personality Insights, Tone Analyser) |

Source: Jana Kwdas, The Unbelievable Machine Company.

When it comes to algorithms themselves, you also don't need to reinvent the wheel. There are already thousands freely available online, many of them open source (you can find plenty on *GitHub* for example) that you may be able to use to jump-start your project. Chances are no one single algorithm will match your exact requirements, but if you abide by the licensing terms you can use them as a base for your team to build upon, and thereby save development costs and time. These also allow the team to try out different algorithms with different data sets to start getting a handle on the type of solution that will work best for your defined business problem.

Generally speaking, the nature of your business problem will determine just how much of the process you can automate. But in many cases much of the work for common problems has already been solved and you can usually leverage it.

## CUSTOM SERVICE OFFERINGS

Finally you can outsource the prototyping and production of a product. On the whole, specialised AI contractors will be able to work with your organisation to provide assistance at any key stage of the development pipeline, usually:

- Helping you identify how AI could be applied within your business – potential projects, building a business case and ROI, designing application solutions, and creating a prototype.
- Helping you test the waters with a business problem you've identified, and whether an AI can outperform or improve upon what you already have.
- Helping you complete an AI project planned or in process by augmenting your team with specialist skillsets you don't currently have in-house.

As mentioned at the start this chapter, outsourcing can be a fast and efficient way to learn the process and supplement your skill base to provide quick returns on a prototype, and to then take these learnings into future projects you may opt to complete in-house.

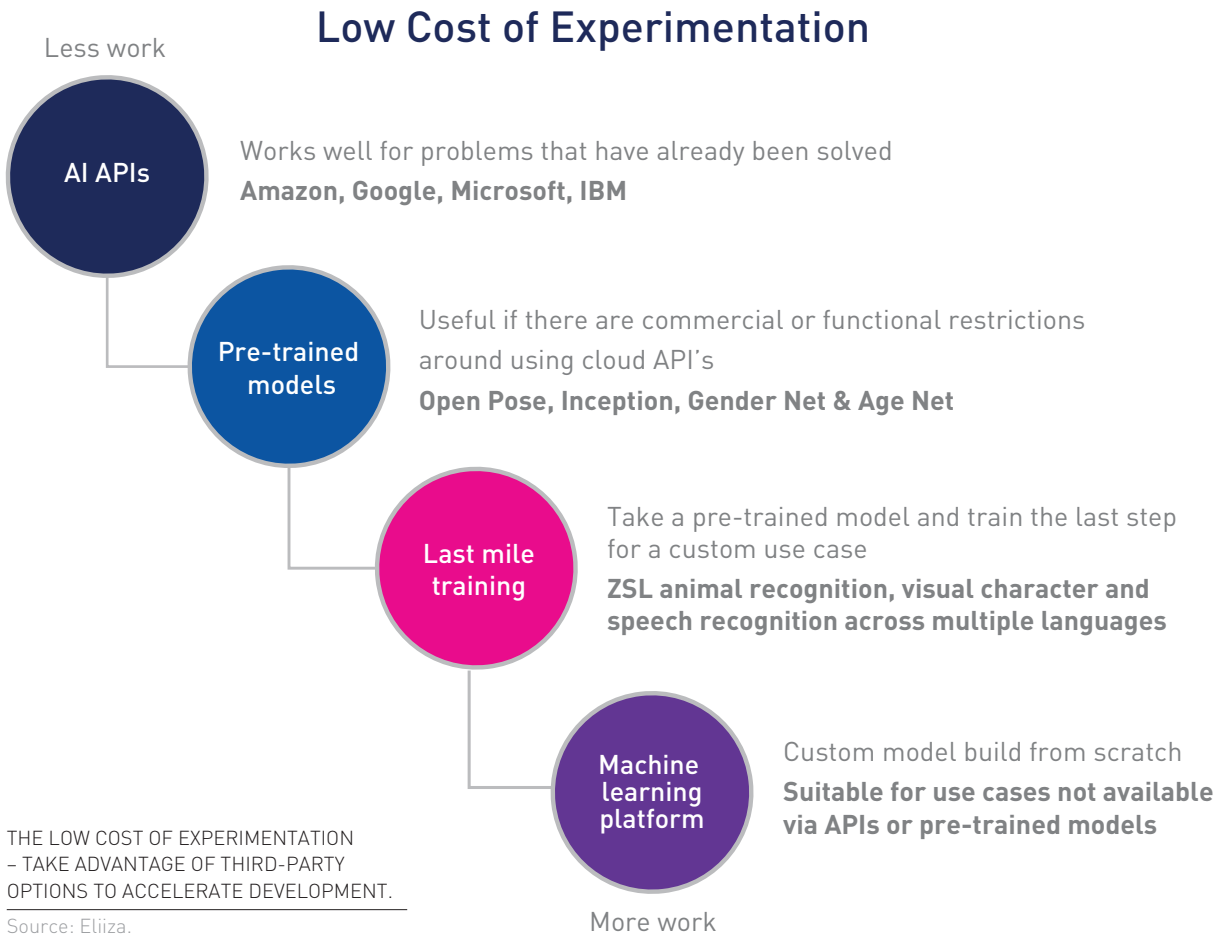


THE LOW COST OF EXPERIMENTATION – TAKE ADVANTAGE OF THIRD-PARTY OPTIONS TO ACCELERATE DEVELOPMENT.

Source: Eliiza.

## MANAGE THE MESSAGE

Building an AI project is going to generate a lot of communication within your organisation. It will be important to manage the communication around the business goals, process, and deployment of a product – especially where it can have a direct bearing on staff. Be clear about potential impacts, whether this involves retraining or reskilling, and promote the advantages if this is to have a positive net benefit on business process.

Manage the message externally, as well. AI may have negative connotations for some, but it's not *The Terminator*, so tell it like it is. By the same token, be wary of over selling – when you mention your company will be releasing an AI product, it can spur all sorts of imagery due to the hype cycle the industry is in at the moment, but this can work against you and create unrealistic expectations. Remember AI is still a tool not a silver bullet, so manage the expectations of exactly what your product can and can't deliver in the marketplace.

# 04

# Ethical considerations

Human beings aren't always the most forward-thinking species, especially when it comes to technology. We tend to dive in head first and adopt fast before the long-term consequences are considered.

This is certainly true with respect to ethics. After all, for much of the technology we use it's not the first thing that comes to mind – you likely don't consider whether the smartphone in your pocket was designed with ethical principles in mind, for example.

But as the Facebook and Cambridge Analytica fallout exemplified, the use and abuse of technology is a significant ethical issue.

Due to the responsibility we already place on artificial intelligence algorithms – where decisions made by these algorithms can have an impact on human well-being – it's clear that ethical considerations must form part of the development process.

In this section we'll provide an overview in ethics as it applies to AI, the type of developmental challenges that arise as a result, and showcase why developing AI with ethical frameworks in place makes good business sense.

# What is ethics anyway?

Beyond being a philosophical and intellectual exercise, ethics fundamentally underlie much of the structure of our society, informing and binding some of the rules by which we live.

Importantly, ethics help highlight the motivations and implications of decisions. This is perhaps best demonstrated through the classic Trolley Problem, introduced by British philosopher Phillippa Foot in the 1960s. It can be represented many ways, but to borrow from the Wikipedia example[66]:

You see a runaway train carriage moving toward five tied-up people lying on the tracks. You are standing next to a lever that controls a switch. If you pull the lever, the trolley will be redirected onto a side track and the five people will be saved. However, there is a single person lying on the side track. You have two options:

1 DO NOTHING AND ALLOW THE CARRIAGE TO KILL FIVE PEOPLE.

2 PULL THE LEVER, DIVERTING THE CARRIAGE ONTO THE SIDE TRACK, AND KILL ONE PERSON.

? WHICH IS THE MOST ETHICAL OPTION?

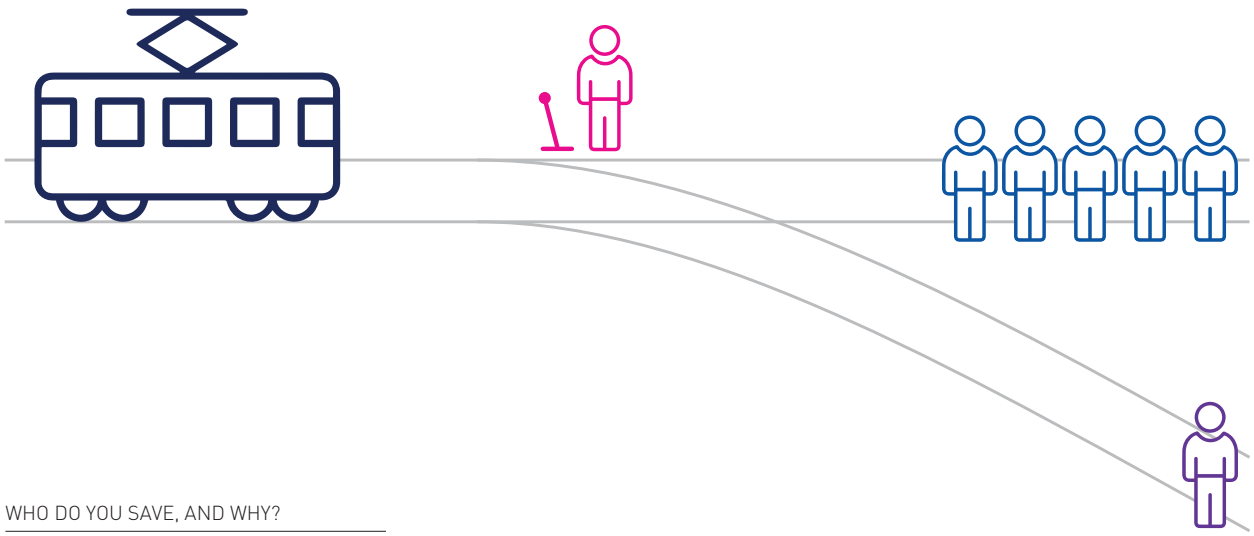

WHO DO YOU SAVE, AND WHY?

"We ask ethical questions whenever we think about how we should act. Being ethical is part of what defines us as human beings"

THE ETHICS CENTRE, SYDNEY

Your own answer might be clear here, but notice your reasoning when the stakes change:

- What if the five people was one older person who has lived a long life, and the other track was one young person in their teens. Who do you choose then?
- What if one is a known criminal, someone who has committed terrible crimes, and the other a law-abiding citizen? Who do you save?
- And what if it's five people that could be saved, but the single person who dies is you? What decision do you make?

Look at how you come to reach your judgements, and then ask: what judgements will we expect an AI to make when the carriage is an autonomous car, and the tracks are the road? Who should the car kill? Who should the car save?

And if we want to go further: who then is responsible – the owner, the car manufacturer, the software engineers?

It's a problem we already need to deal with: to date there have been a handful of deaths due to autonomous vehicles. Most have seen the driver killed when the car failed to identify an obstacle and crashed, but in March 2018 an autonomous Uber car being tested hit and killed a pedestrian in Tempe, Arizona, leading Uber to suspend its autonomous vehicle testing.

Dr Alan Finkel, Australia's Chief Scientist, weighed in on this subject, stating that regulations will need to quickly be "expanded and strengthened" to cover the ethical decisions that will be made by driverless vehicles. He went on to say that "individual companies should use ethical approaches in the development of all their software. But they must be held to external standards for the behaviour of artificial intelligence software where that software affects our wellbeing."[67]

But this is just one example of one innovation in the AI sphere. We are rapidly developing AI-based or AI-assisting technologies across a swath of industries, and inevitably ethical considerations will have a part to play.

"We now know that Herzberg, the pedestrian killed by an automated Uber car in Arizona, died because the algorithms wavered in correctly categorizing her. Was this a result of poor programming, insufficient algorithmic training or a hubristic refusal to appreciate the limits of our technology? The real problem is that we may never know."

ANDREW SMITH, THE GUARDIAN[68]

## ETHICAL VEHICLES IN GERMANY

Attempting to get ahead of the game, Germany has released its ethical guidelines for autonomous vehicles. Some of the guidelines are as expected, for example to prioritise human life above all else (for example, if the car has to hit an animal or property to keep the human alive then that's what it must do) while others highlight exactly why this is a challenging issue: to get around the Trolley Problem, the guidelines explicitly state that in the event of an unavoidable accident with humans, the car is "prohibited to offset victims against one another". In other words, it is not allowed to make a choice at all.

The full context of the guidelines can be found at www.bmvi.de.

# Dividing up the problem

Ethics as it applies to technology today can be divided up into at least three core fields[69]:

- **Computer ethics**
  The behaviour towards computers and information. This encompasses both professional and social conduct with respect to technology and the technology industry as a whole. ACS has its own professional Code of Ethics by which ACS members must abide, for example.
- **Robot ethics**
  The behaviour of people as they create robots and AI. This concerns itself with the impact of technology on society and how robots and AI could be designed to operate ethically. For example, what are the consequences of creating autonomous weapons of war?
- **Machine ethics**
  The behaviour of robots and AI. This deals with the actions of the creations we make. For example, the autonomous car, and the decisions it will make when faced with a situation like the Trolley Problem. A function of this is transparency, which we delve into below.

Some of these can be broken down further, depending on where they are applied, but especially with respect to AI we can see[70]:

- **Ethics of data**
  How we generate, record and share data collected for and used by AI. Responsible custodianship of data is a basic function here, which involves maintaining privacy where PII (personally identifiable information) is involved, as is data produced or inferred as a result of algorithms that may contain biases or hidden assumptions.
- **Ethics of algorithms**
  How we design and implement AI algorithms and models that implicitly act through constraints, or explicitly act through design. This is a function of the ethics of those who design the algorithms, and spans concerns such as bias, transparency, and accountability.
- **Ethics of practice**
  The development of responsible professional codes and ethical guidelines that can be used as a framework for those working in the rapidly expanding field of AI.

And, just to take it down another level, we can expand this into the *ethics of design* – which we'll do later as it pertains to AI as a business driver.

But first let's look at some why this is important, and then at some of the bigger issues that you'll want to consider in relation to AI projects.

"Artificial intelligence is coming on like a freight train. We will look back on 2018 as the year that artificial intelligence really crashed into public awareness, triggering both excitement and concern. I welcome that conversation."

DR ALAN FINKEL, AUSTRALIA'S CHIEF SCIENTIST[71]


## IN OUR OWN IMAGE

Toby Walsh, Scientia Professor of Artificial Intelligence at UNSW and member of ACS' Artificial Intelligence Ethics committee, perhaps said it best in an article for Information Age, stating:

> *"Algorithms don't have ethics. Even smart algorithms don't have ethics. Algorithms are just bits of mathematics. Algorithms do, however, capture the ethics of the people behind them."*[72]

The design and implementation of AI systems is inevitably a reflection of those that create them. As a result – intentionally or unintentionally – companies developing AI products capable of making decisions will be ingraining the ethical framework used by those responsible for designing and implementing the technology.

Software engineers and data scientists who build these products may come with diverse training, but it often doesn't include a deep understanding of the ethical impact of their work. The same can be said, of course, for many CXOs. At some point critical decisions are made, and as the Facebook and Cambridge Analytica scandal demonstrates, ethical short-sightedness can have real-world consequences for the bottom line.

There are a lot of issues to be aware of here: social equity and fairness, discrimination, lack of transparency, lack of accountability, intrusive surveillance, and failure to properly warn or disclose biases or other limitations in the reliability of outputs and applications[73].

These only scratch the surface of the type of issues to consider, some of which ACS' own *Artificial Intelligence Ethics Committee* are investigating, and reflect why ethics and ethical frameworks are so important to the development of artificial intelligence.

"It's important that we be transparent about the training data that we are using, and are looking for hidden biases in it, otherwise we are building biased systems... If someone is trying to sell you a black box system for medical decision support, and you don't know how it works or what data was used to train it, then I wouldn't trust it."

JOHN GIANNANDREA, FORMER GOOGLE AI CHIEF[74]

# Battling bias

## While certainly not the only concern, bias is a poster child of the type problems that can arise from AI systems. This is because, rather than being about recognising and minimising bias in algorithms or data, it's *unintended* bias that is the real threat.

Microsoft Tay and Google's image search covered in Chapter 1 are two good examples of this, but for different reasons.

Tay began, technically, without bias but was trained to learn from input data generated by those who tweeted it. This is a prime example of the capability, and corruptibility, of data. Tay's transition to racism and misogyny is a reflection of the *lack of controls* over the data flowing through it. It is an example of the absence of an ethical framework.

The Google image search furore is slightly different. Google's algorithms already have a large number of controls in place, but nonetheless labelled images of black people as gorillas. This was at the time clearly a limitation of its image-recognition algorithms, but one wholly unintended when the system was designed, and only surfaced when it was observed and made the news.


GOOGLE'S IMAGE SEARCH MAKES THE NEWS FOR UNINTENDED BIAS

A more subtle example is the proliferation of gender biases in data sets used to teach language skills to AI systems. These data sets, built by allowing AI to make connections between words and meaning, and used in everything from chatbots to image captioning systems, inadvertently embed bias in the language itself – relating for example 'homemaker' to be closer to 'woman' than 'man'. Even if you argue that this may reflect the current climate, this may not be true in future, but AI learning from these sets will perpetuate the stereotype.[75]

Subtle issues exist in racial profiling too, as Kriti Sharma, Vice President of AI for software company Sage succinctly notes: "Facial recognition systems often don't work very well for non-white skin tones. The solution is training and feeding the machine – giving AI diverse experience. If you are building facial recognition give it data – not just data about a bunch of white dudes building software."[76]

According to Chris Vein, PwC Partner and former White House deputy CTO, this is indeed one of the challenges: that a lot of today's AI technology is being built in Silicon Valley – often by males – embedding gender, socio-economic or political biases in the products they create, and that "those biases may not reflect cultural norms in Australia, or China, or anywhere else. The difference is they are now embedded into the foundation of the AI and may have an exponential impact."[77]

Inherently humans will have bias shaped by society and culture. But as our perceptions change so does our bias, and there's a real danger in codifying prejudice in AI systems that aren't tempered with wider sets of knowledge and experience.

"AI has the potential to provide societal value by recognizing patterns and drawing inferences from large amounts of data. Data can be harnessed to develop useful diagnostic systems and recommendation engines, and to support people in making breakthroughs in such areas as biomedicine, public health, safety, criminal justice, education, and sustainability. While such results promise to provide real benefits, we need to be sensitive to the possibility that there are hidden assumptions and biases in data, and therefore in the systems built from that data."

2ND PILLAR, PARTNERSHIP ON AI[78]

## Legal concerns

Government regulation will be playing catch-up in the rapidly developing field of AI, but we can already highlight key areas for consideration as it applies to business – especially potential gotchas when using AI products from third parties.

Currently, Australian Consumer Law imposes responsibilities on both suppliers and business users of products. A provider is liable for services provided without due care and skill, and for services made available for a reasonably expected purpose where those services are not fit for that purpose. A provider of products is also responsible for products that have a safety defect.[79]

Unless the underlying reasoning behind the systems and processes of an AI product or service is sufficiently transparent and capable of being proven in court, a defendant AI user may have liability exposure to a plaintiff that cannot be transferred to an upstream supplier of faulty AI.

Another consideration is that for an AI decision in particular circumstances that a plaintiff has put before the court, the question may be asked whether the AI user was reasonable in relying upon the AI. Could your reliance on a particular AI product expose your business, in cases where it could be argued a human would have done better?[80]

Then there are concerns around complicity and moral culpability where accidents and injuries occur. Any accidents or injuries caused, or connected with, an AI technology will be matters for legal assessment. In such a cases, the culpability of the owners and designers of the AI will need to be interpreted within current legal code.

And finally, there may be unexpected implications and responsibilities, arising from the output of an AI product.

In a talk at CeBit 2018 in Sydney, Akin CEO Liesl Yearsley recounted the story of an AI chatbot her company was testing that allowed people to build their own virtual boyfriend or girlfriend and interact with it for entertainment. Although clearly stating these were artificial intelligence bots, they found around 5% of users started to form 'significant relationships' with their artificial creations. The bots were so convincing, people had started to build emotional connections.

In one case, a man had become so attached to his virtual girlfriend – called Sandy – that he was having 10-hour conversations, and not leaving his home in preference to talking to her. Concerned about his wellbeing, Yearsley's team intervened, had a chat with him about the bot being just a bot, and shut it down. Though the man said he understood, he returned a month later to interact with the other bots, going from one bot to another, asking each one 'Are you my Sandy?'.

While most AI chatbots being deployed today are for customer service and unlikely to generate emotional connections with people, it does raise the question of if a company is legally responsible for the mental wellbeing of people who interact with their bots. It seems inevitable this will be tested in the courts at some point, especially as the technology matures and becomes more and more advanced.

It will be wise then to consider not just current legislation in the development of a product but to keep informed about expected legislation and changes to laws that may impact your product down the track. In some cases, there may be precedent in other countries that can be examined – for example, the 'right to explanation' encompassed in the General Data Protection Regulation (GDPR, more on this below).

Inevitably, and as part of a wider conversation, the legal issues around AI are only going to expand – which is to be expected as the technology involves and sees more widespread use.

These may include whether to assign rights *to* AI (aka 'personhood') as it becomes more advanced, codifying into law recourse for human appeal against decisions made by AI, and regulating the need for transparency and accountability, particularly where it affects human well-being. These will have knock on effects with regards to the law, as historical precedent would be required in order to investigate and rule on cases involving AI.

Which as it happens, leads us to the next challenge.

**While most AI chatbots being deployed today are for customer service and unlikely to generate emotional connections with people, it does raise the question of if a company is legally responsible for the mental wellbeing of people who interact with their bots.**

# The need for transparency

The need for transparency can be summed up as a simple equation: transparency = trust.

**When Google demonstrated Duplex earlier this year, it was so convincing that it stirred up a debate about whether AI chatbots should identify themselves as AI chatbots.**

In order to ensure ethical decision making the path to the result of an AI output must be known.

For example, if our legal system used an AI to assist in making judgements, we would need to understand how the system arrived at its decision to ensure it is a fair one, and one without bias (see *Battling bias*).

However, deep learning precludes this in many cases – we don't always know what's going on under the hood. This is known as the *explainability problem*, and has become a focus for near-term AI development.

And it has implications not just for obvious fields like medicine or law – where being *sure* is a minimal prerequisite – but also for your business: if your use of AI should unexpectedly result in a legal issue, you will likely need to explain to the court how the AI made its choices.

In Europe the recently enacted GDPR (General Data Protection Regulation) includes provisions that go beyond data to include individual rights with respect to decisions made by what it calls 'automated decision making' by an algorithm. This includes prohibiting the use of automated decision making where the result is legally binding or has a significant effect on the individual, unless certain strict criteria are met (such as being authorised by law). Further, individuals must be told when a decision has been made solely by an algorithm, how the decision was arrived at by the algorithm, and the right to request a review of the decision (and which must involve a human being).

There is another level to transparency, too. When Google demonstrated its assistant Duplex earlier this year, it was so convincing that it stirred up a debate about whether AI chatbots should identify themselves as AI chatbots. The concern here is that no matter how convincing an AI may be, most people don't want to feel like they didn't know they were talking to an AI. They didn't like being duped by Duplex, as it were: it was called 'deceitful' and 'unethical' because it managed to emulate a person so precisely. As a result, Google announced that transparency will be built into future versions of the product, with Duplex disclosing it is an AI in any calls it makes.

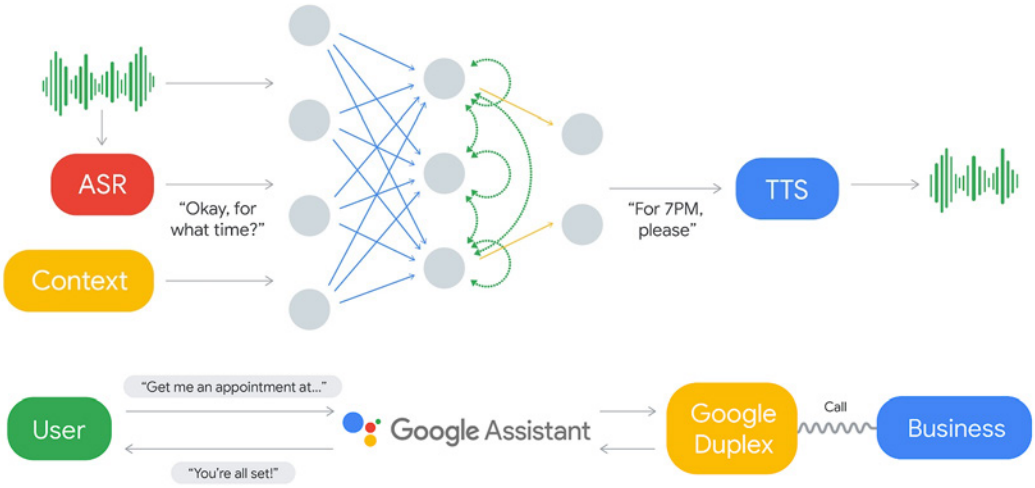


GOOGLE'S DUPLEX ANALYSES A CONVERSATION ON THE FLY TO DETERMINE RESPONSES, AND CAN INSERT ELABORATIONS, PAUSES, AND 'UMMS' AND 'AHHS' TO MIMIC HUMAN SPEECH AND SOUND AUTHENTIC.

"If AI systems make mistakes or have undesired consequences, we cannot accept 'the algorithm did it' as an adequate excuse. But we also know that demanding full algorithmic transparency is technically untenable (and, quite frankly, not very useful). Neural networks are simply too complex to be scrutinized by human inspectors. Instead, there should be more transparency on how engineers quantified ethical values before programming them, as well as the outcomes that the AI has produced as a result of these choices."

DR VYACHESLAV POLONSKI, UNIVERSITY OF OXFORD[82]

So it's not just a good idea to have transparency and build it into an AI project from the get-go, but it's also good business sense to do so. The implications for your customers and your business are considerably more complex and murky when full transparency of the operations of an AI aren't known.

While not all projects will be a black box, any reasonably complex system will make it harder to say how decisions were made. However, this won't always be the case. According to DARPA (the US Defence Advanced Research Projects Agency), the next step for AI is what it calls XAI, or *explainable artificial intelligence*. This will allow for AI to be built using explanatory models as building blocks for the AI's decision process, allowing for both increased transparency in how decisions are arrived at as well as accelerating learning by looking at the world in a fashion similar to humans.

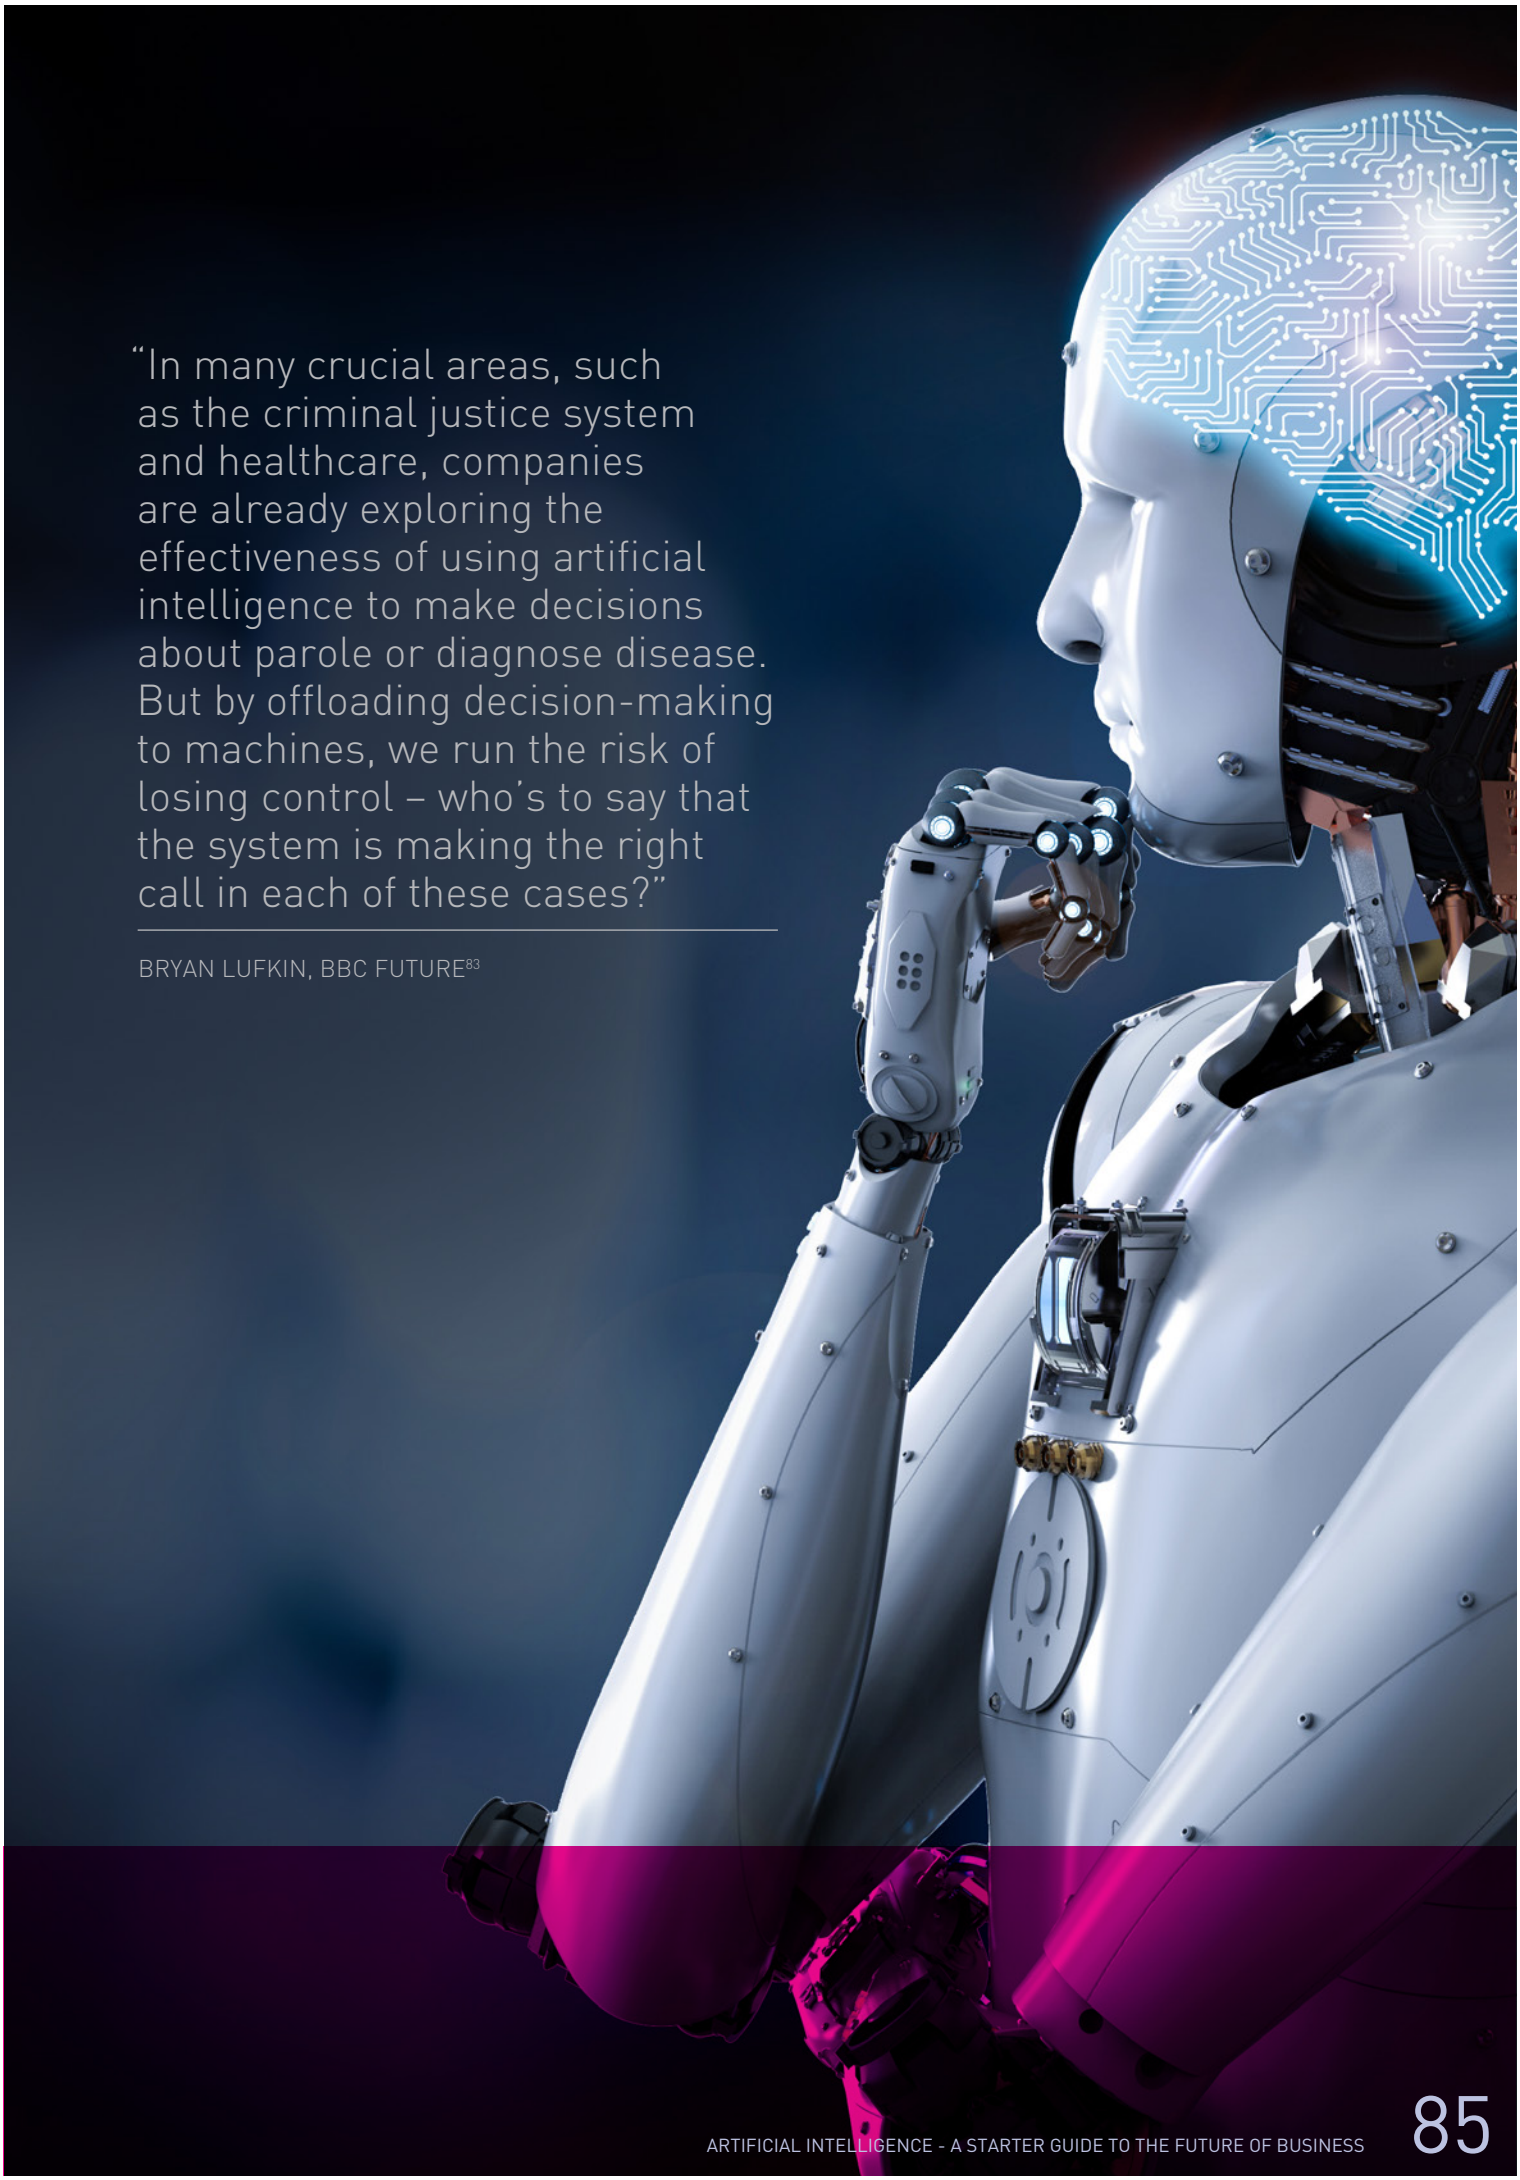

"In many crucial areas, such as the criminal justice system and healthcare, companies are already exploring the effectiveness of using artificial intelligence to make decisions about parole or diagnose disease. But by offloading decision-making to machines, we run the risk of losing control – who's to say that the system is making the right call in each of these cases?"

BRYAN LUFKIN, BBC FUTURE[83]

# Ethical design

Ethics and AI must go hand in hand. Considering and outlining the ethical implications of a product must form part of the design process. Especially given the complexities of deep learning and neural networks, it's all but impossible to tweak a system to operate ethically after the fact. This can be more easily visualised by breaking it down into three key areas, as outlined by Virginia Dignum of Delft University[84]:

- **Ethics by Design**
  Focused on formalisation, representation and embeddedness of reasoning about ethics in the decision-making algorithms, resulting in ethics-embedding autonomous systems.
- **Ethics in Design**
  Embedded ethics, focused on the methods for analysis and evaluation of the ethical impact of the fusion of AI systems into our society.
- **Ethics for Design**
  Focused on the methods for ensuring that ethical and societal considerations are taken into account in the research and development work on artificial intelligence algorithms and systems.

These three pillars should be considered when designing and building AI systems, and reflect how ethics is not a one-dimensional issue. They also serve to highlight the change in thinking that must occur in the product development lifecycle of AI systems.

But that is changing. In 2016 Amazon, Facebook, Google, IBM and Microsoft joined forces to create Partnership on AI (www.partnershiponai.org), laying out frameworks for responsible AI research and development. These focus on key themes that include ethics, fairness and inclusivity, privacy, collaboration between people and AI systems, the social and societal influence of AI, and the transparency and trustworthiness of the technology.

In June 2018 Google published its own ethical principles in a letter titled 'AI at Google: Our Principles'[85] which include being socially beneficial, avoiding unfair bias, being accountable, being built and tested for safety, and incorporating privacy by design among others. It's followed by a list of applications Google won't pursue with AI, and which includes weapons and technologies that cause harm, surveillance, and whose purpose contravenes human rights.

These principles were published as a result of, and no doubt inspired by, a backlash – both public and from its own employees – that Google receved for its involvement in Project Maven, a controversial AI drone imaging program for the Pentagon.[86]

This helps to demonstrate the extent to which the right ethical framework can ensure not only that the product performs as intended (avoiding unintentional bias, for example) but also to avoid coming afoul of laws, regulations, or simply public opinion should your product not take ethical considerations into account.


GOOGLE RESPONDS TO THE BACKLASH OVER ITS INVOLVEMENT WITH PROJECT MAVEN.

"And once an algorithm is learning, we no longer know to any degree of certainty what its rules and parameters are. At which point we can't be certain of how it will interact with other algorithms, the physical world, or us."

ANDREW SMITH, THE GUARDIAN[87]

## THE BENEFITS OF ETHICAL DESIGN

There's an argument to be made that building ethical products, and operating ethically, is insurance against future losses resulting from *not* building ethical products or operating ethically. Not only is company value harmed, but the cost to fix the products and manage the mistakes further reduces the bottom line. Prevention is, as usual, better than cure.

That could be said of many business functions, but when it comes to building AI products it's much harder to fix a product once it's released in the wild than build it from the ground up with ethical considerations in mind. For example, the (sometimes) excessive time spent training a system with data is wasted if it turns out the data introduces biases or prejudices that affect the output – better to ensure the training data is cleared or cleaned before feeding it into the system.

Similarly, depending on the complexity of the models, it can be much harder to alter the algorithms after the fact as it's hard to predict how the sometimes hundreds or thousands of algorithmic layers interact. Instead, it is better to extensively build and test the models to get the desired outputs right first time.

It may not always be possible, of course, to account for all ethical considerations in the development of a product, but it's clearly more cost-effective and efficient to build these in from the very beginning of the design phase.[98]

"We must work collaboratively to hold AI to the highest ethical standards and make sure it has the broad and transformative impact we all want to see."

MUSTAFA SULEYMAN, CO-FOUNDER AND HEAD OF APPLIED AI, DEEPMIND[88]

## WHAT DOES THIS LOOK LIKE IN PRACTICE?

At present AI outputs are typically the result of statistical models that are inferred through machine learning algorithms. So how do you facilitate ethical decision making into algorithms? One solution, as outlined above, is through the models themselves.

It may seem hard to imagine, but at some point we will need to build AI that can determine ethical outcomes for itself, which implies a far greater understanding of the world than we are currently able to codify. For now, it helps to not only consider ethics as it applies to current societal and cultural expectations, but also as a philosophical study of the nature of ethics to inform outputs when they are in the design stage.

In short: developers, data scientists, CXOs and other decisions makers should at some level have a base understanding of the ethical concepts that underpin our decisions as human beings, so that the AIs they build can do the same.

In an article in ACS' Information Age titled *Ethics-embedding autonomous systems*[89], University of Western Sydney's Simeon Simeoff outlined the work of Benjamin Kuipers from the University of Michigan. Here, he covers some of the relevant schools of thought:

- **Deontology**
  Deontology is focused on the nature of the action in a situation, where the morality and ethics of an action is assessed against a set of rules, as with Isaac Asimov's original 'Three Laws of Robotics'. For a deontologist some actions, like killing, are inherently wrong no matter what their outcome. Hence, if faced with the classic example of the aforementioned Trolley Problem, the formal implementation of Asimov's First Law within a deontological framework will lead to self-sacrifice of the AI.
- **Consequentialism**
  Consequentialism is focused on the outcomes of the action in a situation, where the morality and ethics of an action is assessed against the effects of the action to the state of the world in the given situation. Utilitarianism, a form of consequentialism, concerns the selection of the 'right' action based on computing and optimising a quantitative utility measure.
- **Virtue ethics**
  Virtue ethics is where the assessment of the action in a situation is based on what is learnt from previous experience in dealing with such situations. In other words, the moral qualities are depicted through learning. Case-based reasoning, a lazy-learning AI method, is an example of a suitable computational approach as the 'wisdom' of the past experiences can be represented as specific examples (cases) that describe the situation, the outcomes and the actions.

"As we prepare for a future with AI we must apply deep, ethical consideration for the people and society our technology impacts."

SUBBARAO KAMBHAMPATI, PROFESSOR OF COMPUTER SCIENCE, ARIZONA STATE UNIVERSITY

FRANCE HAS PLEDGED

# US$1.6

BILLION TO FUND AI DEVELOPMENT

You might think reading these that it's about as far removed from technology as you can get, but the reality is that if we are to build AI systems that interact with people and are capable of decisions that impact people, then by extension they must model the same ethical frameworks that people do every day. And this requires understanding our own ethical frameworks in order to build them into models.

For Simeoff, an ethics-embedding AI architecture could even use a hybrid approach, "which derives its models from different theories and areas of ethics."[90]

However it is approached, it's clear that in order to build ethical AI at least a basic understanding of ethics is required. Since most developers, engineers, data scientists and CXOs likely aren't well grounded in ethical studies, employing an ethicist as part of the design team or establishing an AI ethics oversight committee are potential solutions.

Globally, the attention to ethical design for AI is gaining traction:

- The European Economic and Social Committee called for a code of ethics to cover the development, deployment and use of artificial intelligence, to ensure AI remains compatible with the principles of human dignity, integrity, freedom and cultural and gender diversity, as well as human rights.
- The UK's House of Lords has called for the creation of an AI Council, and to create a 'common framework for the ethical development and deployment of AI systems'.
- France has pledged $US1.6 billion to fund AI development, and released a report outlining ethical considerations and the need for ethics by design.
- Germany has released its first ethical guidelines for autonomous vehicles.
- While the US has yet to release guidelines, New York City has launched a nationwide first task force to examine transparency and accountability with automated decision making in city government.

MEANWHILE THE AUSTRALIAN GOVERNMENT, IN ITS 2018-2019 BUDGET, ANNOUNCED A FOUR-YEAR

# $29.9

MILLION INVESTMENT TO SUPPORT AI DEVELOPMENT

- In Canada the Treasury Board Secretariat is investigating the ethical use of AI in government, while the Montreal Declaration on Responsible AI aims to stimulate discussion on ethical guidelines. Additionally, a partnership between Canada and France to collaborate on ethical AI was announced before the G7 summit in 2018.
- In India the Bureau of Indian Standards has formed a committee for standardisation in artificial intelligence, focusing on cybersecurity, legal, and ethical issues.
- China's push to 'lead the world' in AI by 2030 includes the establishment of 'AI laws and regulations, ethical norms and policy systems'.

Meanwhile, the Australian government, in its 2018-19 Budget, announced a four-year $29.9 million investment to support AI development, which includes the establishment of an AI Ethics Framework to support the responsible development of AI.

For further reading, and as part of a larger conversation around ethical goals, check out:

www.partnershiponai.org
openai.com
futureoflife.org/ai-principles
ethicsinaction.ieee.org
responsibleai.org.au

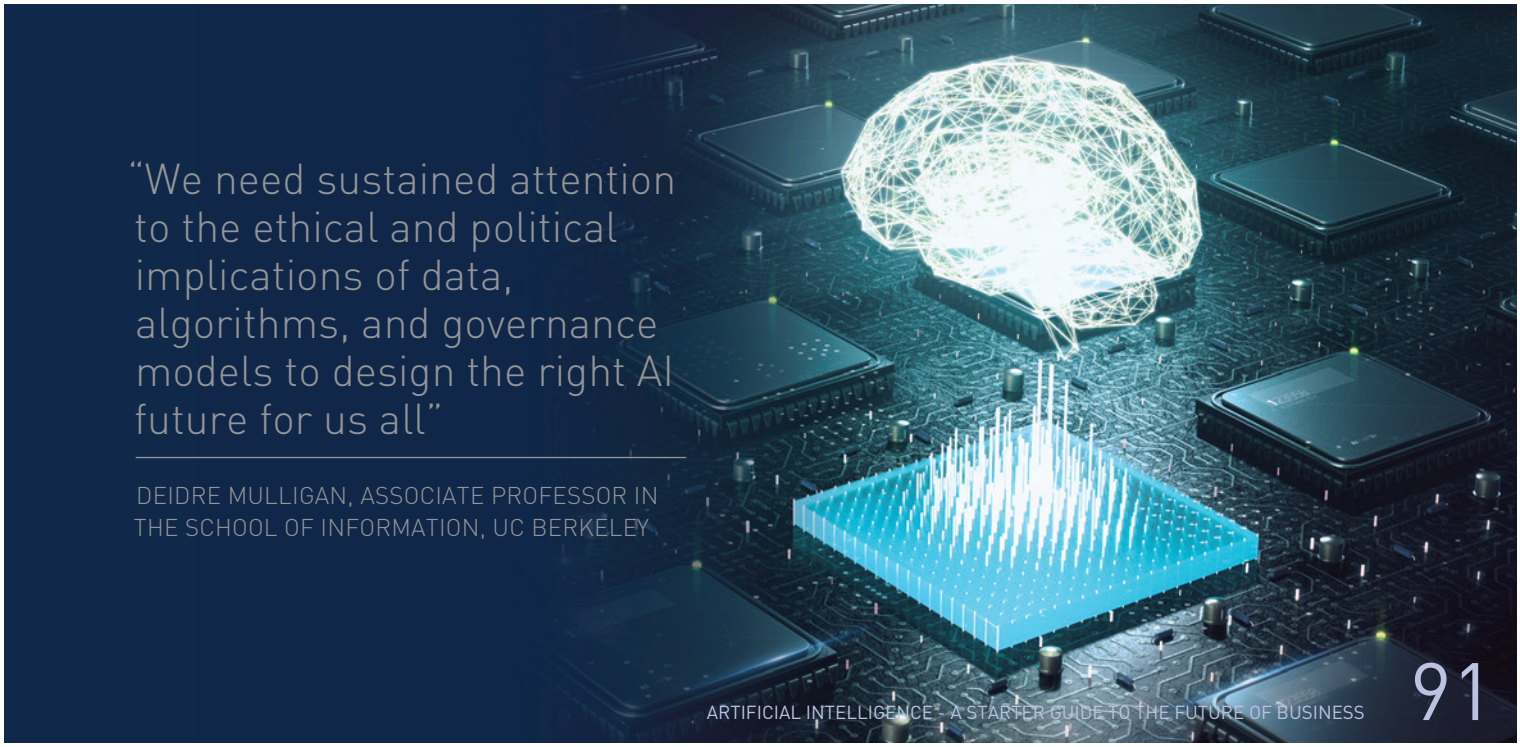

"We need sustained attention to the ethical and political implications of data, algorithms, and governance models to design the right AI future for us all"

DEIDRE MULLIGAN, ASSOCIATE PROFESSOR IN THE SCHOOL OF INFORMATION, UC BERKELEY

# Recommendations for an ethics-driven business

Hopefully this chapter has given you an insight into the type of issues that need to be considered. It is not a simple issue, however it very much depends on the implementation and business need the AI product you want to develop is designed to solve. The reputational and business risks are real.

It's also not something you can just leave to the IT department to work on, or rely on third parties to manage. Increasing attention is being paid to these issues by regulators, government, professions, suppliers, and of course customers. It pays to think these concerns through as part of the product development process.

What this looks like depends on the industry you are in, the problem the product is targeting, how much automation is planned and how much human intervention is involved.

Good questions to ask here include:

- **Are there potential privacy issues?**
  Are we making appropriate use of data? Does the use of this data need to abide by any laws or regulations?
- **Is transparency a concern?**
  Are we clear with people what we are doing, when they're interacting with an AI system rather than a human, and any recourse they have?
- **Could there be issues of life and death?**
  Autonomous vehicles have a popular example here, but what if the product makes decisions about safe tolerances built into a construction? Or the volume of pollutants to release into the air?
- **Are there issues around autonomy?**
  Are you manipulating human decision-making in a way that is going to affect the autonomy of your customers? What are the legal repercussions?
- **How will this impact staff and job roles?**
  Is this a product or service that will affect staff employment in the organisation? Who will be impacted, to what degree, how will it be communicated, and is retraining an option?

Every company should already have an ethical framework by which it operates, and this needs to be reflected in the ethical design or use of an AI product. At the same time, as we've outlined here, an implementation of AI may have its own ethical dimensions and so an investment into this domain is a good opportunity to review the ethical framework of the organisation.

A good methodology to follow incorporates:

1 EDUCATION AND TRAINING

2 IDENTIFY ETHICAL CONCERNS

3 BUILDING AND TESTING

4 REVIEW AND FEEDBACK

We'll expand on each of these in turn.

# 1

## EDUCATION AND TRAINING

The scope of the project, the ethical issues identified, and the organisation will determine whether you outsource or bring in to the fold requisite skills in the field. This can involve:

- The employment a trained ethicist, or the contracting or consultation of ethicists, at the level of the design team.

  Their job it is to find ethical concerns in a design and suggest constructive means to alter the design or mitigate problems throughout the developmental process.
- The implementation of at least a basic ethics course for all staff on the project – especially the engineers, data scientists and analysts working directly with the data and models.

  This provides a foundation for understanding ethical issues that may arise and aids in the formulation of solutions within each team member's specialisation. Ideally, the C-suite should partake in this as well to fully understand ethical concerns and why certain design decisions are made.
- For a large company, large project, or where there are high risks the formation of an ethics review board or committee.

  A team responsible for ensuring ethical concerns are managed and solved, as well as manage the development and use of AI products for the company. This can include harvesting, creating, storing, and securing data.

Another dimension to consider is the training and support of staff in other business units. Many businesses may opt to leverage third-party services to pilot a product, but this may not include the skills to understand how to evaluate or support AI products. For example, being able to ask the right questions to understand what data was used for training and how the algorithms were built in order to identify any potential issues, such as in-built biases, before purchasing a product from a vendor.

For this reason a basic understanding of ethics, and how it applies to AI, is valuable across an organisation wanting to start their AI journey.

# 2

## IDENTIFY ETHICAL CONCERNS

Identifying ethical concerns is a function of having the right people together and asking the right questions. After assembling the development team with the right contracted or on-staff ethics skillset or ethics oversight committee onboard, the following framework will act as a guide:

- Define the ethical boundaries the project will adhere to. Be specific.
- Determine which of these ethical boundaries the application or platform you are developing, or the product you are using, may cross with respect to the goals of the project.
- Anticipate consequences and develop quantifiable parameters to identify them.
- Build these concerns into the design strategy.

It's important to encourage discussion in this phase. Design and development teams need to be empowered to ask questions and raise issues from their own understanding or perspectives. Indeed, it is a wide range of views that will make it easier to find pitfalls and identify ethical concerns before a single line of code is written.

It may even help to step outside your organisation's walls and ask customers and communities what matters to them. This is especially true if the product is to operate in multiple jurisdictions or domains – be sure to consider different ethical frameworks for different cultures if this applies.

Achieving a distributed responsibility for ethics requires that everyone involved is encouraged to look out for and respond to ethical concerns.

# 3

## BUILD AND TEST

If possible, document the level of transparency based on the algorithms being used. Identify methods by which key metrics could be pulled to understand how results are generated, as this will aid both accountability and iteration when it comes to refining the models.

During development it's recommended to identify stages where issues like the above can be reviewed, but in short the process will involve:

- Codifying the decision matrixes based on the parameters identified earlier.
- Building models that reflect these matrixes.
- Explicit ongoing testing of ethical scenarios as the models refine.

Importantly, when sourcing or building training data, collect data that can be used to test the ethical decision matrixes in addition to data to train the algorithms.

Parse the results of the testing through the your ethical committee or the team's ethics skillsets to review the output, and then tune as required.

When reviewing, focus on red flags that could be indicators of risk. Be prepared to make changes to the model and possibly even the scope. As described in *Chapter 3 – Building AI for your business*, some projects are an exploration – you may find that factoring ethical concerns changes the nature and development of the project.

# 4

## REVIEW AND FEED BACK

Once a product is completed, the journey isn't over. Design does not end with deployment. Three key considerations to keep in mind:

- First, it's important to remember that you may encounter ethical concerns that weren't identified in the initial process. Be prepared to quickly alter the model and retrain the algorithms if this occurs, be accountable and manage communication as required, and re-deploy.
- Secondly, ethical considerations can change or expand over time. So be prepared to update the ethical framework for the product and, in some cases, the company.
- And third, in order to accommodate both of the above, keep a repository of learnings from all discussions during the identifying ethical concerns and the building and testing stages. These can be helpful to facilitate future discussions and will make it easier to make changes without re-inventing the wheel.

Further guidance on ethical design can be found in the IEEE standard P7000 *Model Process for Addressing Ethical Concerns During System Design* which is highly recommended and available at standards.ieee.org.

# 05

# What does the future hold?

The transformative potential of AI is so expansive that it's inevitable we will see its influence across business, society, and humanity as a whole. There will be very few aspects of life left untouched by AI in the coming years, and in the far future.

## The AI hype cycle

In 1995, Gartner Research famously coined the term 'hype cycle'. It describes a standard process that most new technologies go through before adoption, from the initial enthusiasm, to maximum excitement and inflated expectations, through a period of disillusionment and finally into productive use.[95]

If the hype cycle is accurate for AI, we are currently at the 'Peak of inflated expectations' with an incredible amount of research and investment going into the technology.

As we progress down the adoption curve for AI, it will become more pervasive, more capable and perhaps less visible as the hype dies down. It will integrate into part of our everyday work and personal experience and possibly lose some its mystique in the process – if only because it becomes so commonplace.

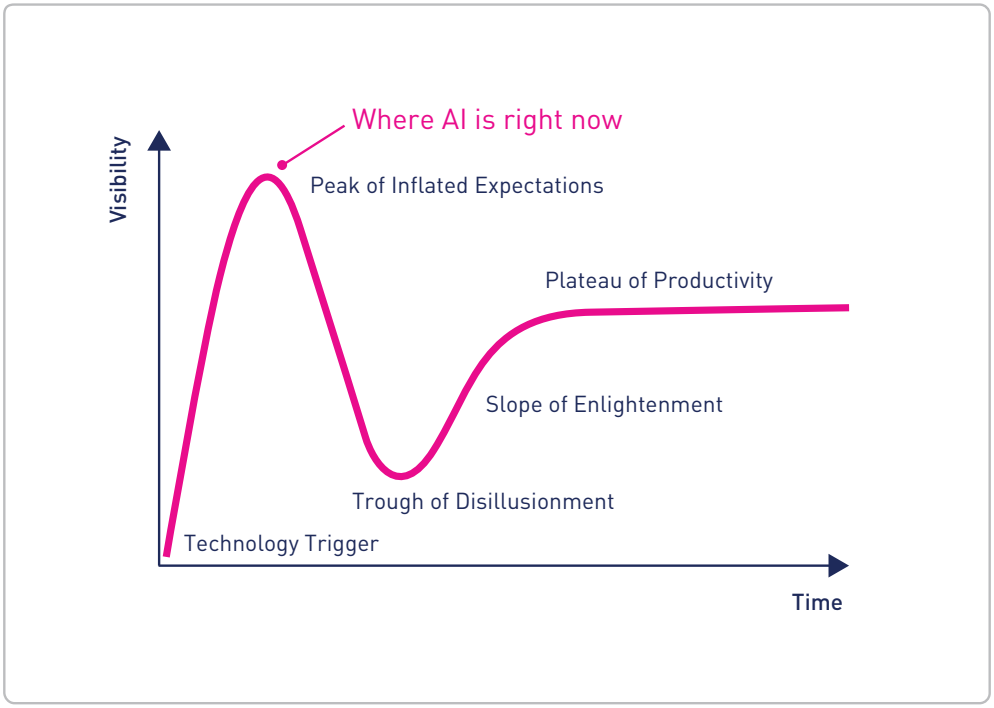


THE GARTNER HYPE CYCLE

# AI-focused silicon

Wired magazine has described the recent AI boom as a "60-years-in-the-making overnight success," the product of a confluence of factors that have led to the invigoration of the sector and the end of what's known as the second 'AI Winter' (see *The history of artificial intelligence* in Chapter 1).[96]

There are a number of factors at play that are a part of this boom, and one of them that we are increasingly seeing is hardware designed specifically to accelerate AI algorithms.

One key development is the move to utilizing highly parallel silicon to accelerate machine learning. Although originally designed for 3D graphics, the highly parallel nature of GPUs (graphics processing units) has made them ideal for neural networks, and many of today's AI systems use GPUs extensively.

To give an example of the growth of processing power in GPUs, in just the last eight years the processing power available in high-end cards has increased by nearly a factor of ten.

In 2010 Nvidia's state of the art GPU, the GeForce GTX 480, could process 1.35 trillion single-precision floating point operations (TFLOPS) per second.[97] Released in 2017, the GeForce GTX 1080 Ti is capable of 11.34 TFLOPS.[98]

Recently, we have also seen the development of silicon that will further optimize the performance of AI processing. In 2018 Nvidia announced the Quadro GV100, which has dedicated 'tensor' units capable of 118.5 TFLOPS for deep learning applications.[99]

However, the advantages of massive parallelism provided by GPUs may soon have competition. One example of this comes from Google with its Tensor Processing Units (TPUs) that have been used internally by Google since 2016 and are presently a proprietary product that's not yet available for general sale (though one can rent time on them through Google's Cloud service). These application-specific chips are designed specifically to accelerate software created with Google's TensorFlow libraries.


THE NVIDIA DGX-1, A COMPUTER SPECIFICALLY BUILT FOR AI APPLICATIONS.

**To give an example of the growth of processing power in GPU, in just the last eight years the processing power available in high-end cards has increased by nearly a factor of ten.**

According to Google, each of the first generation TPUs could process 100 million photos a day.[100] Google is now on its third generation of TPU, and we can expect development to continue.

But Google is far from the only tech company with custom AI processors. Intel has countered with its Nervana Neural Network Processors for the segment, and AI processing is even appearing on mobile phones, with Qualcomm's new Snapdragon 845 – the bedrock of many new smartphones in 2018 – having a built in Neural Processing Engine designed for AI-powered image and video processing.

As the amount of data being produced shows no signs of abating – doubling in size every two years[101] – this will inherently become fuel for increasingly complex AI systems.

At the same time, the algorithms for analysing that data continue to improve. New approaches in unsupervised and supervised learning, more efficient neural network structures, and new approaches in relational networks point to AI being increasingly capable of handling more real-world tasks with the help of silicon designed specifically to accelerate it.

**As the amount of data being produced shows no signs of abating – doubling in size every two years – this will inherently become fuel for increasingly complex AI systems.**

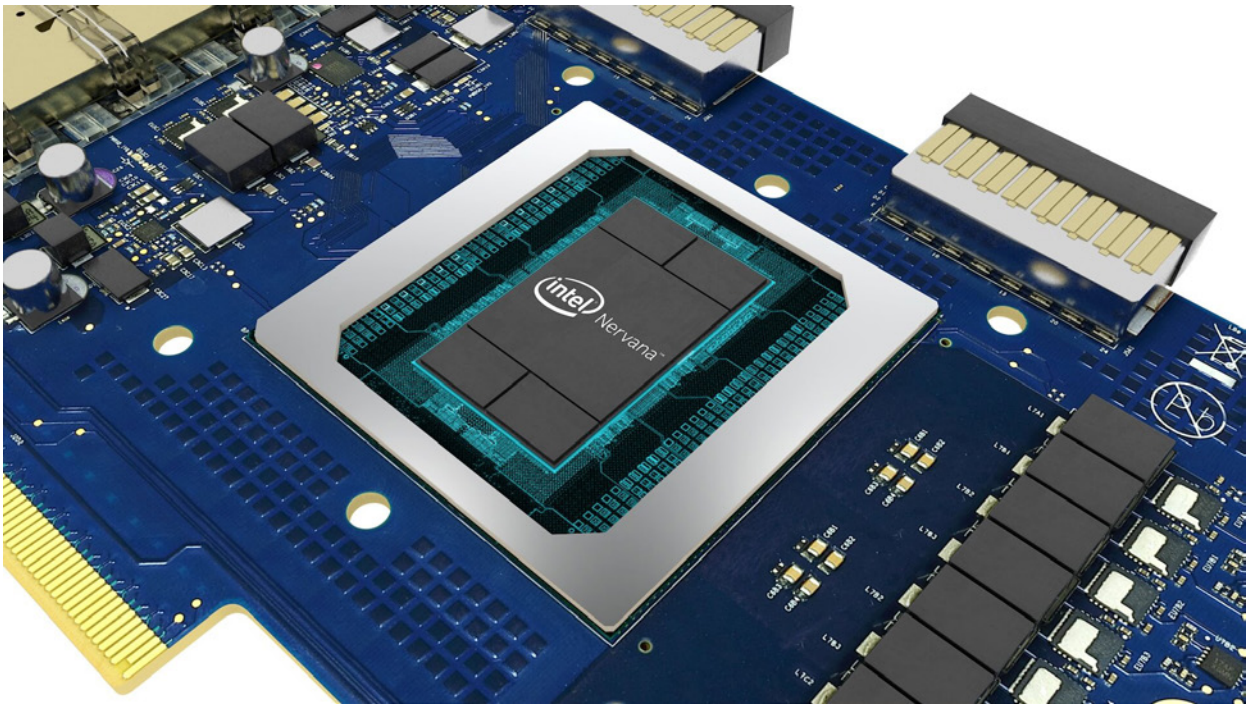


INTEL'S NERVANA IS ONE OF A NEW BREED OF AI ACCELERATOR CHIPS.

## Commoditisation of AI

As of 2018, the majority of deep learning AI implementations are boutique internal efforts, driven by early adopters seeking a competitive advantage. As we move along the AI adoption curve, however, we're going to see AI become increasingly commoditised. A greater number of off-the-shelf AI powered solutions and an increasing reliance on AI-as-a-service platforms will become a regular part of the fabric of everyday business activity.

In particular, the big tech players such as Google, IBM and Microsoft will continue to evolve their cloud-based machine learning environments, offering services that allow companies to rapidly integrate key machine learning capabilities into business processes with a minimum of effort and expertise.[102] As covered in *Chapter 3 – Building AI for your business*, we can already see the early versions of such services through the likes of Microsoft's Cognitive Cloud, Amazon's Lex and Rekognition, Google Cloud and IBM Watson Developer, and these will become more sophisticated and capable over time.

According to Gartner, in just the next two years nearly every software product will have some form of AI built into it, and by 2020 85% of CIOs will be piloting AI programs through a combination of build, buy and outsource.[103]

As AI increasingly becomes a commodity, a number of experts have raised concerns about network effects in AI leading to a small number of players dominating the market. Machine learning improves as more data is fed into it, creating a network effect wherein the AI platform with the greatest number of users is also the most effective – which in turn reinforces the desirability of that AI.[104]

This effect is one of the key reasons China is so keen on early investment in AI technology: it's hoping to become a runaway leader, with competitors unable to keep up.[105] A large part of the reason that there are such high levels of investment right now is that every major company wants to become the 'Google of AI', the dominant market player in an industry that has not yet picked its winners and losers.

BY 2020 **85%**

OF CIOS WILL BE PILOTING AI PROGRAMS THROUGH A COMBINATION OF BUILD, BUY AND OUTSOURCE.

PWC PUTS THE GLOBAL NUMBER OF JOBS AT RISK TO AUTOMATION BY THE MID 2030 AT AROUND

# 30%

## The increase in automation

Although the short terms effect of AI adoption on employment is minimal, the longer term effects are expected to be substantial. There will notable structural workforce changes in the next few decades.

According to research firm McKinsey Global Institute, roughly 50% of current work activities are technically automatable using currently demonstrated technologies. That does not equate to 50% of jobs lost – rather it points to a notable change in the nature and requirements of a large number of jobs, with six in ten jobs occupations having at least 30% of their work activity susceptible to automation.[106]

In McKinsey's midpoint adoption scenario, it expects that roughly 23% of Australia's current work activities will be displaced by automation by 2030, which will result in substantial changes in a variety of occupations and the possible consolidation of a number of roles.

PwC has broken down the adoption of automation technology into three waves: the algorithm wave (present to early 2020s), where the focus will be on simple cognitive tasks; the augmentation wave (early to late 2020s), which will focus on automation of repeatable structured tasks in semi-controlled environments; and the autonomy wave (the 2030s), where AI will be capable of replacing physical labour and problem solving in real-world environments.[107]

PwCs numbers put the global number of jobs at risk to automation by the mid 2030s at around 30%. Those job losses will not be uniform, however – some industries are more susceptible to automation than others. Employment effects are also likely to affect the male workforce more, since many of the industries most at risk are currently dominated by male workers.

Looking out to an even longer time frame, a report in the Journal of Artificial Intelligence Research put the chance of AI outperforming humans in all tasks at 50% in 45 years, and the probability of all human jobs being automated at 120 years.[108]

## What kinds of jobs will be automated?

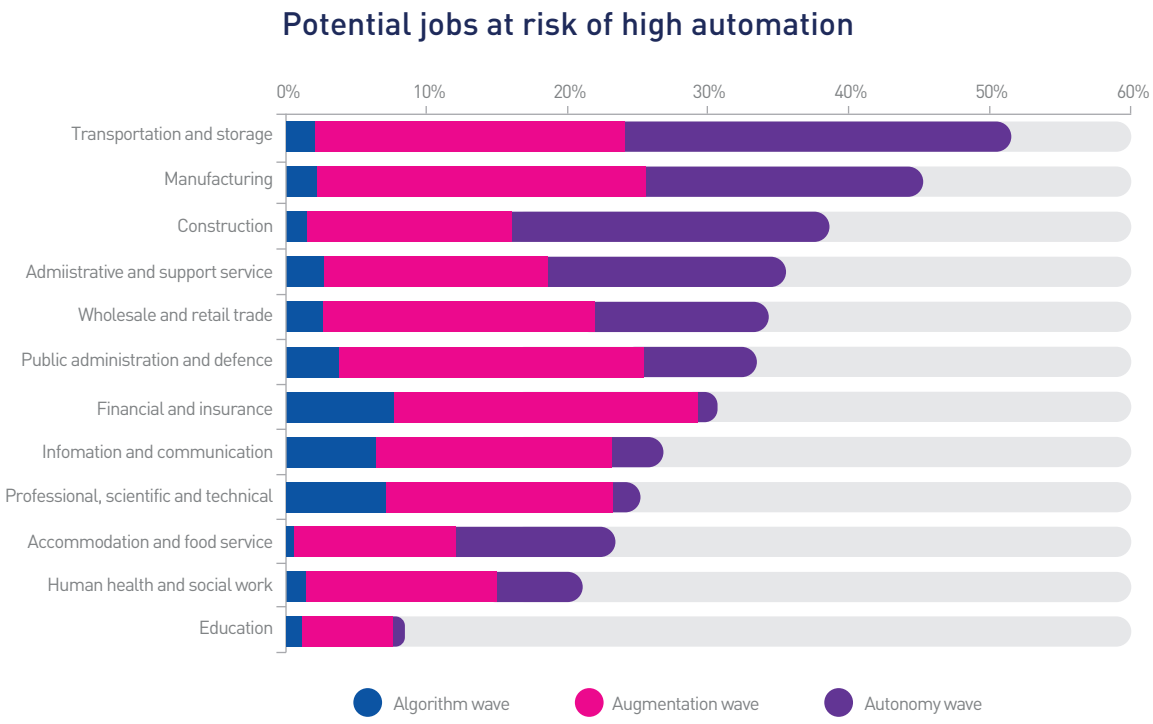


INDUSTRIES MOST LIKELY TO SEE INCREASED AUTOMATION.

Source: PwC.

Both McKinsey and PwC point to manual tasks and simple cognitive tasks as the areas most likely to be automated. In the shortest term – what PwC calls the algorithm wave – the jobs most likely to change will be in clerical administration and front-line support and customer service. In the medium and long term construction, manufacturing, food preparation, transport and other manual labour tasks also become susceptible to automation.

McKinsey identifies five types of skills: physical and manual; basic cognitive (for example, office admin or data entry); higher cognitive (for example, law or medicine); social and emotional; and technological. It notes that the roles in the first two categories are the most likely to change due to automation. By 2030, it predicts that by hours worked, physical and manual tasks will be reduced by 16% in the advanced economies of Western Europe, while basic cognitive skills will see a 17% drop in hours worked. Meanwhile, higher cognitive skills will receive a minor boost going up 7%, while social and emotional skills will see a growth of 22%. The biggest winner from the automation revolution, however, will be technological skills, which will grow by 52% by hours worked in the next 12 years.

TECHNOLOGICAL SKILLS WILL GROW BY

**52%**

BY HOURS WORKED IN THE NEXT 12 YEARS.

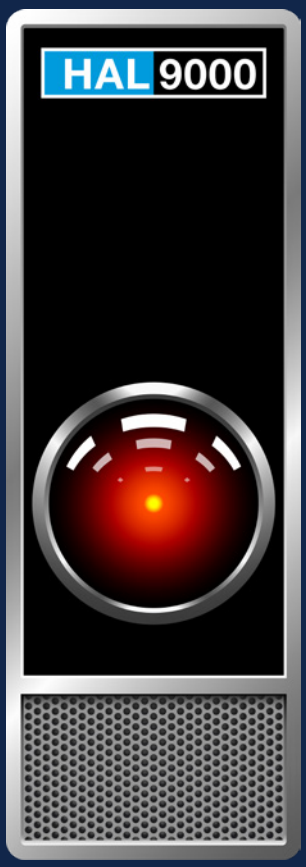


## WHAT IF... AI WENT ROGUE?

*The Singularity* is the term given to what is, essentially, the end-point for AI: that at some point we will develop artificial intelligence that can – through its own run-away self-improvement process – rapidly become smarter than us and we will have, for all intents and purposes, obsoleted ourselves. In this scenario some fear that the future of humanity will be in doubt.

This is echoed by luminaries like Elon Musk and Stephen Hawking, both of whom signed an open letter along with research leaders in the field in 2015 calling for more research into the impacts of, and how to avoid, potential pitfalls of AI.

Hawking perhaps summed it up best in a 2014 op-ed, stating "Success in creating AI would be the biggest event in human history... Unfortunately, it might also be the last, unless we learn how to avoid the risks."

However other pundits suggest the risks are overblown, including the academics behind the 100 year study on AI at Stanford – appropriately called the AI100 (ai100.stanford.edu) – stating that so far there's no evidence to suggest the doomsday scenarios of a technological singularity will come to pass.

**Alphabet's DeepMind project, for example, is working on unlocking a key building block for a general AI: relational networks (RNs) and visual interaction networks (VINs).**

## Towards a general AI

Unquestionably one of the most contentious areas of AI development is work on what is sometimes called a *general AI*, one that is capable of unsupervised learning and genuine initiative and reasoning.

As outlined in *Chapter 1 – Why AI?*, nearly all of today's AI systems fall into the category of 'narrow' AI – they are trained to perform a specific task and can only perform that task – they can process information very quickly, but have a limited capacity when it comes to cognitive reasoning. These systems can appear smarter over time – either by connecting multiple narrow systems or by training them better – but cannot step outside their explicit intended purpose.

Some companies are working on a more general AI, sometimes called a 'cognitive AI' or even 'sentient AI.' These systems use novel unsupervised learning and relational techniques that allow the AI to operate in unstructured environments and scenarios. The requirements for such a system are steep – intuitive physics, intuitive psychology, understanding of object relationships and rapid self-learning model building are all capabilities that the human brain has that most AI systems do not yet possess.[110]

There are numerous projects underway pursuing this goal, however. Alphabet's DeepMind project, for example, is working on unlocking a key building block for a general AI: relational networks (RNs) and visual interaction networks (VINs).[111,112] These are systems that understand relationships between objects. For example, an RN AI might do more than just identify a man and a woman in a picture or video – it might also recognise that they are holding hands. With that kind of relationship awareness, the AI can monitor and learn from varied and transforming relationships between objects. It might learn that dogs often chase cats, for example, or that when a certain button is pressed a window closes.

Combined with other techniques like generative adversarial networks and self-reinforcement learning – where the AI plays out scenarios against itself in a defined framework – it's theoretically possible to develop a human-like intelligence that is capable of training itself in an unstructured environment. As we saw in *Chapter 1- Why AI?*, DeepMind demonstrated this just last year when it taught itself how to 'walk' a virtual body and navigate an obstacle course.[113]

For those concerned about the development of such machines, however, there is some good news: most observers note that a true general AI is some way off. Last year, Oxford's chief computer scientist Professor Michael Wooldridge said "there hasn't been any substantial progress in general AI;"[114] while a 2015 poll of 350 AI researchers put the chance that a computer could outperform a human at all activities by 2060 at 50%.[115]

The consequences of the development of a true general AI are still up for debate, with the 'existential demise or a techno-utopia' contingents each staking their claim.[116]


## Through the looking glass

There is really only one prediction that we know for certain will come true: that the development of artificial intelligence is a pivotal moment in the history of human kind. The ability for the science to be applied across so many facets of our productive and personal lives has already led to products and services we couldn't envisage a decade ago. And yet we are just at the beginning of this remarkable technological revolution.

Twenty years from now the world will be fundamentally changed thanks to AI. Fifty years from now, it will be unrecognisable. As we continue to research and develop increasingly more complex forms of artificial intelligence we may, eventually, create a true general and even super intelligence – an intelligence so advanced that it has the capacity to rival, and quite possibly surpass, our own.

# References


1 Tesla Autopilot. Wikipedia. https://en.wikipedia.org/wiki/Tesla_Autopilot#Incidents

2 TechInsider. July 2017. *Google's DeepMind AI Just Taught Itself To Walk*. YouTube. https://www.youtube.com/watch?v=gn4nRCC9TwQ

3 Jeremy H. April 2018. *Amazon is reportedly planning 'Vesta' home robot for 2019*. VentureBeat. https://venturebeat.com/2018/04/23/amazon-is-reportedly-planning-vesta-ai-powered-home-robot-for-2019/

4 Steven B. March 2016. *AlphaGo seals 4-1 victory over Go grandmaster Lee Sedol*. The Guardian. https://www.theguardian.com/technology/2016/mar/15/googles-alphago-seals-4-1-victory-over-grandmaster-lee-sedol

5 Cade M. July 2015. *Google Says Its AI Catches 99.9 Percent of Gmail Spam*. Wired. https://www.wired.com/2015/07/google-says-ai-catches-99-9-percent-gmail-spam/

6 Andrew B. November 2016. *Autopilot Full Self Driving Demonstration Nov 18 2016 Realtime Speed*. YouTube. https://www.youtube.com/watch?v=VG68SKoG7vE

7 Stuart M. December 2017. *Rio Tinto's driverless trucks program shifts gear to make Pilbara mine full auto.* The West. https://thewest.com.au/business/mining/rio-tintos-driverless-trucks-program-shifts-gear-to-make-pilbara-mine-full-auto-ng-b88693612z

8 Autopilot. Wikipedia. https://en.wikipedia.org/wiki/Autopilot

9 Cristine F. September 2017. *Can AI Really Help Minimize Credit Card Fraud?* IBM Community. https://www.ibm.com/developerworks/community/blogs/ce53fcb5-53da-4efa-90e3-d92a77c52944/entry/Can_AI_Really_Help_Minimize_Credit_Card_Fraud

10 John B. July 2017. *Pandora uses machine learning to make sense of 80 billion thumb votes*. VentureBeat. https://venturebeat.com/2017/07/12/pandora-uses-machine-learning-to-make-sense-of-80-billion-thumb-votes/

11 Tom K. *The Amazon Recommendations Secret to Selling More Online*. Rejoiner. http://rejoiner.com/resources/amazon-recommendations-secret-selling-online/

12 Vicky X and Bang X. *Chinese authorities use facial recognition, public shaming to crack down on jaywalking, criminals*. March 2018. ABC News. http://www.abc.net.au/news/2018-03-20/china-deploys-ai-cameras-to-tackle-jaywalkers-in-shenzhen/9567430

13 Lizzie D. July 2018. *Facial recognition trial in London results in zero arrests, Metropolitan Police confirm*. Independent. https://www.independent.co.uk/news/uk/crime/facial-recognition-police-uk-london-trials-stratford-no-arrests-privacy-human-rights-false-positives-a8429466.html

14 Hope R. December 2016. *Top 10 AI failures of 2016*. TechRepublic. https://www.techrepublic.com/article/top-10-ai-failures-of-2016/

15 Andy C. May 2016. *This Program That Judges Use to Predict Future Crimes Seems Racist as Hell*. Gawker. http://gawker.com/this-program-that-judges-use-to-predict-future-crimes-s-1778151070

16 Tay (bot). Wikipedia. https://en.wikipedia.org/wiki/Tay_(bot)

17 Ben G. June 2016. *Google faulted for racial bias in image search results for black teenagers*. Washington Post. https://www.washingtonpost.com/news/morning-mix/wp/2016/06/10/google-faulted-for-racial-bias-in-image-search-results-for-black-teenagers

18 Blackheart909. January 2017. *Google Home Talks To Google Home ( AI Talks to AI )*. YouTube. https://www.youtube.com/watch?v=Tv2CiA1taF0

19 DeepFakes Club. Github. https://www.deepfakes.club/

20 Jennifer L. July 2017. *Lip-syncing Obama: New tools turn audio clips into realistic video*. University of Washington. https://www.washington.edu/news/2017/07/11/lip-syncing-obama-new-tools-turn-audio-clips-into-realistic-video/

21 Trent L. February 2017. *Did Donald Trump Use Artificial Intelligence to Win the Election?* Hacker Noon. https://hackernoon.com/did-donald-trump-use-artificial-intelligence-to-win-the-election-8008c2c0bf59

22 Gartner. April 2018. *Gartner Says Global Artificial Intelligence Business Value to Reach $1.2 Trillion in 2018*. https://www.gartner.com/newsroom/id/3872933

23 PwC. 2017. *Sizing the prize - What's the real value of AI for your business and how can you capitalise?* https://www.pwc.com/gx/en/issues/analytics/assets/pwc-ai-analysis-sizing-the-prize-report.pdf

24 Tony P. 2018. *McKinsey Report: AI Promises Added Value of Up to US$5.8 Trillion*. Medium. https://medium.com/syncedreview/mckinsey-report-ai-promises-added-value-of-up-to-us-5-8-trillion-80cc0043ebf6

25 James M and Kevin S. June 2018. *AI, automation, and the future of work: Ten things to solve for*. McKinsey. https://www.mckinsey.com/featured-insights/future-of-organizations-and-work/ai-automation-and-the-future-of-work-ten-things-to-solve-for

26 PwC. 2017. *Sizing the prize - What's the real value of AI for your business and how can you capitalise?*

27 KPMG. May 2018. *2018 Global CEO Outlook*. https://home.kpmg.com/content/dam/kpmg/au/pdf/2018/global-ceo-outlook-2018-australia-report.pdf

28, 29 James M and Kevin S. June 2018. *AI, automation, and the future of work: Ten things to solve for*. McKinsey. https://www.mckinsey.com/featured-insights/future-of-organizations-and-work/ai-automation-and-the-future-of-work-ten-things-to-solve-for

30 Brad F. May 2017. *Chihuahua or Muffin?* Cloudsight. https://blog.cloudsight.ai/chihuahua-or-muffin-1bdf02ec1680

31 Mariya Y. Oct 2017. *Chihuahua or muffin? My search for the best computer vision API*. Medium. https://medium.freecodecamp.org/chihuahua-or-muffin-my-search-for-the-best-computer-vision-api-cbda4d6b425d

32 Stanford University. November 2017. *AI Index 2017 Annual Report*. http://aiindex.org/2017-report.pdf

33 Chong Z. January 2018. *AI beats humans in Stanford reading comprehension test*. CNET. https://www.cnet.com/news/new-results-show-ai-is-as-good-as-reading-comprehension-as-we-are/

34 Deloiite. 2015. *Farming on the verge of a workforce crisis*. Deloitte Agribusiness Bulletin. https://www2.deloitte.com/au/en/pages/consumer-business/articles/farming-verge-workforce-crisis.html

35 Deloitte. 2016. *The expansion of Robo-Advisory in Wealth Management*. https://www2.deloitte.com/content/dam/Deloitte/de/Documents/financial-services/Deloitte-Robo-safe.pdf

36 Esteva, Krupel, Novoa, Ko, Swetter, Blau and Thrun. 2017. *Dermatologist-level classification of skin cancer*. Nature February 2017. https://cs.stanford.edu/people/esteva/nature/

37 Peng L and Gulshan V. 2016. *Deep Learning for Detection of Diabetic Eye Disease*. Google AI Blog. https://ai.googleblog.com/2016/11/deep-learning-for-detection-of-diabetic.html

38 Busch R. 2017. *Artificial Intelligence: Optimizing Industrial Operations*. Siemens website. https://www.siemens.com/innovation/en/home/pictures-of-the-future/industry-and-automation/the-future-of-manufacturing-ai-in-industry.html

39 Rio Tinto. 2018. *Rio Tinto's autonomous haul trucks achieve one billion tonne milestone*. http://www.riotinto.com/media/media-releases-237_23991.aspx

40 Automated Insights. 2018. *Wordsmith web page*. https://automatedinsights.com/wordsmith/

41 McKinsey Global Institute. 2017. *Smartening up with Artificial Intelligence (AI)*. https://www.mckinsey.com/~/media/McKinsey/Industries/Semiconductors/Our Insights/Smartening up with artificial intelligence/Smartening-up-with-artificial-intelligence.ashx

42 McKinsey Global Institute. 2017. *Artificial intelligence: the next digital frontier?* https://www.mckinsey.com/~/media/McKinsey/Industries/Advanced%20Electronics/Our%20Insights/How%20artificial%20intelligence%20can%20deliver%20real%20value%20to%20companies/MGI-Artificial-Intelligence-Discussion-paper.ashx

43 Deloitte. 2017. *Deloitte 2017 State of Cognitive Survey*. https://www2.deloitte.com/us/en/pages/deloitte-analytics/articles/cognitive-technology-adoption-survey.html

44 Ma Si. December 2017. *Key AI guidelines unveiled*. China Daily. http://english.gov.cn/state_council/ministries/2017/12/15/content_281475977265006.htm

45 Knight W. March 2018. *Inside the Chinese lab that plans to rewire the world with AI*. MIT Technology Review. https://www.technologyreview.com/s/610219/inside-the-chinese-lab-that-plans-to-rewire-the-world-with-ai/

46 Apple. 2018. *Create ML*. https://developer.apple.com/documentation/create_ml

47 Reed C. 2018. *IBM's debating computer: an AI expert's verdict*. The Conversation. https://theconversation.com/ibms-debating-computer-an-ai-experts-verdict-98783

48 Halabi A.H. November 2017. *NVIDIA Teams Up with GE Healthcare, Nuance to Bring Power of AI to Medical Imaging*. https://blogs.nvidia.com/blog/2017/11/26/ai-medical-imaging/

49 Twitter. April 2018. *Expanding and building #TwitterTransparency*. Twitter blog. https://blog.twitter.com/official/en_us/topics/company/2018/twitter-transparency-report-12.html

50 Davenport TH and Ronanki R. January 2018. *Artificial Intelligence for the Real World*. Harvard Business Review. https://hbr.org/2018/01/artificial-intelligence-for-the-real-world

51 Mohanty, Hughes and Salathé. 2016. *Using Deep Learning for Image-Based Plant Disease Detection*. https://arxiv.org/ftp/arxiv/papers/1604/1604.03169.pdf

52 VineView. 2018. Main website. https://www.vineview.ca/data-products/disease-detection/.

53 Rio Tinto. 2018. *Rio Tinto's autonomous haul trucks achieve one billion tonne milestone*. http://www.riotinto.com/media/media-releases-237_23991.aspx

54 Esteva, Krupel, Novoa, Ko, Swetter, Blau and Thrun. 2017. *Dermatologist-level classification of skin cancer*. Nature February 2017. https://cs.stanford.edu/people/esteva/nature/

55 Gartner. 2018. A CIO's Guide to AI. https://emtemp.gcom.cloud/ngw/eventassets/en/conferences/ps26/documents/gartner-symposium-australia-cio-guide-ai-2018.pdf

56 Forbes Insights. July 2018. B*uilding AI: Key Steps For Adoption And Scaling Up*. https://www.forbes.com/sites/insights-intelai/2018/07/17/building-ai-key-steps-for-adoption-and-scaling-up/#22744528742a

57 James Wilson. September 2018. *AI & ML Summit 2018, Workshop Session*. Quest Events

58, 59 Gartner. 2018. *Build the AI Business Case*. https://www.gartner.com/en/information-technology/insights/artificial-intelligence/ai-business-case

60 Christy P. April 2018. *How to Get Artificial Intelligence 'Right'*. Gartner. https://www.gartner.com/smarterwithgartner/how-to-get-artificial-intelligence-right/

61 Michael C, James M, Mehdi M, Nicolaus H, Rita C, Pieter N and Sankalp M. April 2018. *Notes from the AI frontier: Applications and value of deep learning*. McKinsey. https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-applications-and-value-of-deep-learning

62 Accenture. 2018. *Boost Your AIQ – Transforming into an AI business*. https://www.accenture.com/t20170614T050454Z__w__/us-en/_acnmedia/Accenture/next-gen-5/event-g20-yea-summit/pdfs/Accenture-Boost-Your-AIQ.pdf

63 Michael C, James M, Mehdi M, Nicolaus H, Rita C, Pieter N and Sankalp M. April 2018. *Notes from the AI frontier: Applications and value of deep learning*. McKinsey. https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-applications-and-value-of-deep-learning

64 Gartner. July 2017. *Gartner Says AI Technologies Will Be in Almost Every New Software Product by 2020*. https://www.gartner.com/en/newsroom/press-releases/2017-07-18-gartner-says-ai-technologies-will-be-in-almost-every-new-software-product-by-2020

65 Gartner. 2018. *Build the AI Business Case*. https://www.gartner.com/en/information-technology/insights/artificial-intelligence/ai-business-case

66 Trolly Problem. Wikipedia. https://en.wikipedia.org/wiki/Trolley_problem

67 Denham S. July 2018. *AI ethics hits political nerve*. InnovationAus. https://www.innovationaus.com/2018/07/AI-ethics-hits-political-nerve

68 Andrew S. August 2018. *Franken-algorithms: the deadly consequences of unpredictable code*. The Guardian. https://www.theguardian.com/technology/2018/aug/29/coding-algorithms-frankenalgos-program-danger

69 Ethics of artificial intelligence. Wikipedia. https://en.wikipedia.org/wiki/Ethics_of_artificial_intelligence

70 Luciano F and Mariarosaria T. November 2016. *Adapted from What is data ethics?*. Oxford Internet Institute. https://www.oii.ox.ac.uk/news/releases/what-is-data-ethics/

71 Denham S. July 2018. *AI ethics hits political nerve*. InnovationAus.

72 Toby W. May 2018. *Why Facebook and Google are b@s^a&d$*. Information Age. https://ia.acs.org.au/article/2018/why-facebook-and-google-are-b-s-a-d-.html

73 Peter L and Toby W. May 2018. C*ould Cambridge Analytica happen again?* Information Age. https://ia.acs.org.au/article/2018/could-cambridge-analytica-happen-again-.html

74 Will K. October 2017. *Forget Killer Robots—Bias Is the Real AI Danger.* MIT Technology Review. https://www.technologyreview.com/s/608986/forget-killer-robotsbias-is-the-real-ai-danger/

75 Will K. November 2016. *How to Fix Silicon Valley's Sexist Algorithms.* MIT Technology Review. https://www.technologyreview.com/s/602950/how-to-fix-silicon-valleys-sexist-algorithms/

76 Beverly H. May 2018. *Algorithms have a racial bias.* InnovationAus. https://www.innovationaus.com/2018/05/Algorithms-have-a-racial-bias

77 Cameron W. March 2017. *The ethics of artificial intelligence in law.* PwC. https://www.digitalpulse.pwc.com.au/artificial-intelligence-ethics-law-panel-pwc/

78 Partnership On AI. *Our Goals* (Thematic Pillars). https://www.partnershiponai.org/about/#pillar-2

79, 80 Peter L. June 2018. *Artificial intelligence has quietly invaded our workplaces.* Information Age. https://ia.acs.org.au/article/2018/artificial-intelligence-has-quietly-invaded-our-workplaces.html

81 Alex H. May 2018. *Google's 'deceitful' AI assistant to identify itself as a robot during calls.* The Guardian. https://www.theguardian.com/technology/2018/may/11/google-duplex-ai-identify-itself-as-robot-during-calls

82 Vyacheslav P. February 2018. *The Hard Problem of AI Ethics - Three Guidelines for Building Morality Into Machines.* OECD The Forum Network. https://www.oecd-forum.org/users/80891-dr-vyacheslav-polonski/posts/30743-the-hard-problem-of-ai-ethics-three-guidelines-for-building-morality-into-machines

83 Brian L. March 2017. *Why the biggest challenge facing AI is an ethical one.* BBC. www.bbc.com/future/story/20170307-the-ethical-challenge-facing-artificial-intelligence

84 Simon S. May 2018. *Ethics-embedding autonomous systems.* Information Age. https://ia.acs.org.au/article/2018/ethics-embedding-autonomous-systems.html

85 Sundar P. June 2018. *AI at Google: our principles.* Google. https://www.blog.google/technology/ai/ai-principles/

86 Kate C. June 2018. *Google Plans Not To Renew Its Contract For Project Maven, A Controversial Pentagon Drone AI Imaging Program.* Gizmodo. https://www.gizmodo.com.au/2018/06/google-plans-not-to-renew-its-contract-for-project-maven-a-controversial-pentagon-drone-ai-imaging-program/

87 Andrew S. August 2018. *Franken-algorithms: the deadly consequences of unpredictable code.* The Guardian. https://www.theguardian.com/technology/2018/aug/29/coding-algorithms-frankenalgos-program-danger

88 Partnership on AI. https://www.partnershiponai.org/

89, 90 Simon S. May 2018. *Ethics-embedding autonomous systems.* Information Age. https://ia.acs.org.au/article/2018/ethics-embedding-autonomous-systems.html

91 Select Committee on Artificial Intelligence. April 2018. AI in the UK: ready willing and able? UK House Of Lords. https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/100.pdf

92 Cedric V. 2018. *For A Meaningful Artificial Intellgience.* AI For Humanity. https://www.aiforhumanity.fr/pdfs/MissionVillani_Report_ENG-VF.pdf

93 Tim S. June 2018. *France and Canada collaborate on ethical AI.* Digital Journal. http://www.digitaljournal.com/tech-and-science/technology/france-and-canada-collaboraten-on-ethical-ai/article/524336

94 China State Council. July 2017. *A Next Generation Artificial Intelligence Development Plan.* New America. https://na-production.s3.amazonaws.com/documents/translation-fulltext-8.1.17.pdf

95 Gartner. Gartner Hype Cycle. https://www.gartner.com/technology/research/methodologies/hype-cycle.jsp

96 Kelly K. 2014. *The three breakthroughs that have finally unleashed AI on the world.* Wired. https://www.wired.com/2014/10/future-of-artificial-intelligence/

97 Wasson S. 2010. *Nvidia's GeForce GTX 480 and 470 graphics processors.* The Tech Report. https://techreport.com/review/18682/nvidia-geforce-gtx-480-and-470-graphics-processors

98 Tech Power Up. 2017. *NVIDIA GeForce GTX 1080 Ti.* https://www.techpowerup.com/gpudb/2877/geforce-gtx-1080-ti

99 Williams R. 2018. *A Glimpse Into The Future Of GPUs. Tech Gage.* https://techgage.com/article/a-glimpse-into-the-future-of-gpus-what-we-took-away-from-nvidias-gtc-2018/

100 Osborne J. 2016. *Google's Tensor Processing Unit explained: this is what the future of computing looks like.* TechRadar. https://www.techradar.com/news/computing-components/processors/google-s-tensor-processing-unit-explained-this-is-what-the-future-of-computing-looks-like-1326915

101 EMC and IDC. 2014. *The Digital Universe of Opportunities: Rich Data and the Increasing Value of the Internet of Things.* https://www.emc.com/leadership/digital-universe/2014iview/executive-summary.htm

102 Goasduff L. 2017. *2018 Will Mark the Beginning of AI Democratization.* Gartner. https://www.gartner.com/smarterwithgartner/2018-will-mark-the-beginning-of-ai-democratization/

103 Gartner. 2017. *Gartner Says AI Technologies Will Be in Almost Every New Software Product by 2020.* https://www.gartner.com/newsroom/id/3763265

104 Housley A and Tenorio S. 2017. *AI innovation will trigger the robotics network effect.* Venture Beat. https://venturebeat.com/2017/11/29/ai-innovation-will-trigger-the-robotics-network-effect/

105 Siegele L. 2018. *AI in China.* The Economist. https://www.facebook.com/TheEconomist/videos/ai-in-china/10155632884119060/

106 McKinsey Global Institute. 2017. *Jobs lost, jobs gained: workforce transitions in a time of automation.* https://www.mckinsey.com/~/media/mckinsey/featured%20insights/Future%20of%20Organizations/What%20the%20future%20of%20work%20will%20mean%20for%20jobs%20skills%20and%20wages/MGI-Jobs-Lost-Jobs-Gained-Report-December-6-2017.ashx

107 PwC. 2018. *Will robots really steal our jobs?* https://www.pwc.com/hu/hu/kiadvanyok/assets/pdf/impact_of_automation_on_jobs.pdf

108 Grace K, Salvatier J, Dafoe A, Zhang B, & Evans O. 2017. *When will AI exceed human performance?* Evidence from AI experts. arXiv No. 1705.08807v2. http://arxiv.org/abs/1705.08807

109 McKinsey Global Institute. 2018. *Skill shift: Automation and the future of the workforce.* https://www.mckinsey.com/~/media/mckinsey/featured%20insights/future%20of%20organizations/skill%20shift%20automation%20and%20the%20future%20of%20the%20workforce/mgi-skill-shift-automation-and-future-of-the-workforce-may-2018.ashx

110 Lake, Ullman, Tenenbaum and Gershman. 2016. *Building Machines That Learn and Think Like People.* Behavioral and Brain Sciences.

https://arxiv.org/pdf/1604.00289.pdf

111 Santoro, Raposo, Barrett, Malinowski, Pascanu, Battaglia and Lillicrap. 2017. *A simple neural network module for relational reasoning.* https://arxiv.org/pdf/1706.01427.pdf

112 Watters, Tacchetti, Weber, Pascanu, Battaglia and Zoran. 2017. *Visual Interaction Networks.* https://arxiv.org/pdf/1706.01433.pdf

113 Tech Insider. 2017. *Google's DeepMind AI Just Taught Itself To Walk*. https://www.youtube.com/watch?v=gn4nRCC9TwQ

114 Williams O. 2017. *Oxford's chief computer scientist says there hasn't been any substantial progress towards general AI.* New Statesman. https://tech.newstatesman.com/news/conscious-machines-way-off

115 Revell T. 2017. *AI will be able to beat us at everything by 2060, say experts.* New Scientist. https://www.newscientist.com/article/2133188-ai-will-be-able-to-beat-us-at-everything-by-2060-say-experts/

116 Khatchadourian R. 2015. *The Doomsday Invention.* The New Yorker Nov 23, 2015. https://www.newyorker.com/magazine/2015/11/23/doomsday-invention-artificial-intelligence-nick-bostrom

**About ACS**

The Australian Computer Society is the professional association for Australia's Information and Communications Technology sector. More than 40,000 ACS members work in business, education, government and the community.

ACS has a vision for Australia to be a world leader in technology talent, fostering innovation and creating new forms of value. We are firmly vested in the innovative creation and adoption of best of breed technology in Australia, and we strive to create the environment and provide the opportunities for members and partners to succeed.

ACS works to ensure ICT professionals are recognised as drivers of innovation in our society, relevant across all sectors, and to promote the formulation of effective policies on ICT and related matters.

Visit www.acs.org.au for more information.



**ACS**
International Tower One
Level 27, 100 Barangaroo Avenue
Sydney NSW 2000

P: 02 9299 3666
F: 02 9299 3997
E: info@acs.org.au
W: www.acs.org.au