The BRAIN

BROUGHT TO YOU BY THE QUEENSLAND BRAIN INSTITUTE

Issue Four

INTELLIGENT MACHINES

How our brains inspire AI

What is intelligence?

Future AI: should we be worried?

Will robots outsmart us?

Ethics: the blurring lines between humans and machines

THE UNIVERSITY OF QUEENSLAND AUSTRALIA

Publisher
Professor Pankaj Sah

Editor
Carolyn Barry

Editorial assistance
Karen McGhee

Writers
Dr Alan Woodruff
Carolyn Barry
Zoe McDonald

Scientific advisers
Professor Geoffrey Goodhill
Professor Mandyam Srinivasan
Professor Justin Marshall
Dr Peter Stratton
Dr Rodrigo Suárez

Designer
Katharine McKinnon

Marketing and Communications
Bronwyn Adams

The Queensland Brain Institute is located at The University of Queensland, St Lucia, QLD, 4072

communications@qbi.uq.edu.au

The Brain Series: Intelligent Machines was produced by the Queensland Brain Institute, The University of Queensland, January 2019

The University of Queensland
AUSTRALIA
CREATE CHANGE

Special Report: Intelligent Machines

DOWNLOAD, READ AND SHARE

this special report on
your computer or smartphone.
Go to qbi.uq.edu.au/intelligentmachines

FACEBOOK
facebook.com/QldBrainInstitute

TWITTER
twitter.com/QldBrainInst

INSTAGRAM
instagram.com/QldBrainInstitute

CONTENTS

A MESSAGE FROM

Professor Pankaj Sah

DIRECTOR, QUEENSLAND BRAIN INSTITUTE

PATRICK HAMILTON

THE BRAIN: YOUR INTELLIGENT MACHINE

The brain is the most powerful machine currently in existence, so it's little wonder it is a central source of inspiration for artificial intelligence (AI) and robotic technology. With unparalleled efficiency and the ability to learn and adapt, it has formed the blueprint of much research in the fields of AI and robotics.

Over millions of years our brains have increased in size and complexity, evolving from basic nerve structures to the sophisticated organs we have today. Our brain is not just a machine though; it is the reason we can think, move, love, remember, socialise. In **Chapter 1** we look at why we have a brain, and its evolution.

As you'll discover in **Chapter 2**, the links between neuroscience and AI are much deeper than would appear. The brain has been—and continues to be—the inspiration for developing intelligent machines, with many features of AI built on principles of how the brain functions. Take, for example, the rapidly growing interest in 'deep neural networks': this technology is based loosely on the way the brain is organised, in layers of networks that 'talk' to each other to decode information.

We also discuss the possibility of intelligent robots, which—unlike computers or AI assistants that sit on a desk—would have to intelligently interact with their environments. What robots lack in accurate perception, and the ability to learn and interact with their environment, the brain can offer. Rbotics researchers are working with neuroscientists to build intelligent robots. We delve into these issues in **Chapter 3**.

In **Chapter 4**, we discuss advances that move beyond inspiration, and look at connecting brains and machines. While this may sound like science fiction, it's already happening in thought-controlled prosthetic limbs and electrical brain implants, and even devices that track your brain activity.

With any technology, there are serious ethical issues to consider when merging biology with machines, as we cover in **Chapter 5**, including how it may redefine what we consider as human, and what levels of robot-human merging we should accept. **Chapter 6** considers what a future with this technology would look like—what concerns we should raise, and what exciting innovations are on the horizon.

1 THE POWER OF OUR BRAINS

WHY WE HAVE A BRAIN

The brain is the ultimate machine, more complex than just about anything we have ever studied. There is much we don't know about the brain, including exactly how it processes all of the information we receive.

The brain—composed of ~100 billion cells known as neurons—is the central command of our nervous system, and is what we use to interact with the world. Sensory neurons relay information about the world around us. Motor neurons allow us to move around our environment.

The brain integrates sensory information, tests it against our stored memories and model of the world, and decides when and how to act; it's also responsible for predicting possible future scenarios and giving meaning and emotional context to our daily experiences.

Our brains weren't always so complex. Over millions of years, brains have evolved from simple networks to become the intelligent, complex and efficient machines that we have today.

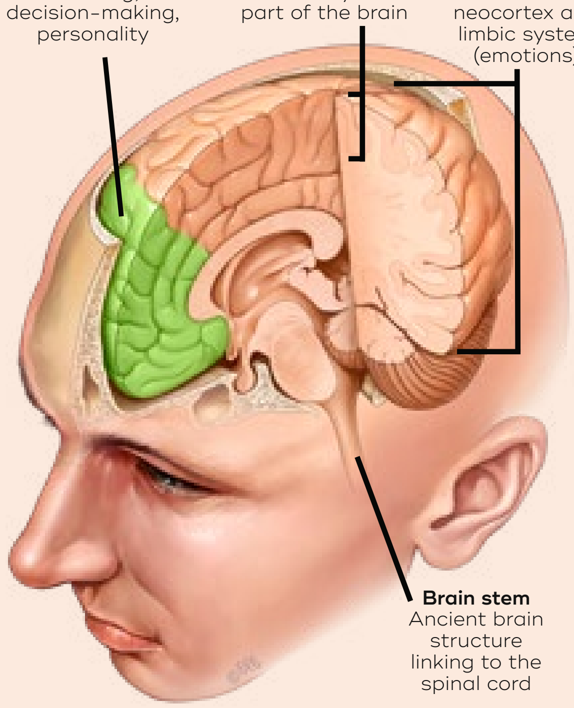

CREDIT: LEVENT EFE

THE EVOLVING BRAIN

When multicellular organisms evolved from single-celled ones over 1.9 billion years ago, they found a way to communicate by using chemical and electrical messages. As nervous systems evolved, communication between neurons became faster, more precise, and increasingly complex.

The first nervous systems were likely 'nerve nets', which have no central command or brain. They are still seen today in simple organisms such as hydras and jellyfish. As evolution continued, neurons began grouping together around the mouth and sensory organs of free-swimming animals. This allowed signals to be processed and transformed, rather than just transmitted from one part of an organism to another—which enabled more complex behaviours such as hunting for prey, escaping from predators or finding a mate.

Brains have evolved from simple networks to become the intelligent, complex and efficient machines that we have today

Around 500 million years ago, the first vertebrates developed basic versions of a pattern of connections broadly shared by all species today. As four-legged vertebrates colonised land over 350 million years ago, the complexity of their brains, in particular the forebrain (see above), increased to be able to process information about their surroundings and remember and learn from experience.

With the appearance of mammals about 200 million years ago, the forebrain further developed in size and complexity. Connections between the left and right sides of the brain were further expanded in placental mammals by evolution of the corpus callosum—the main bundle of nerve cells linking both hemispheres. A six-layered structure evolved, known as the neocortex—the wrinkly outer region of our brains, important for complex tasks like abstract thinking and planning. The neocortex is considered to be the key to our intelligence, including our ability to reason, plan ahead, communicate, store memories and solve complex problems. This intelligence, and the ability of our brains to handle more than just basic survival, has allowed humans to dominate the Earth.

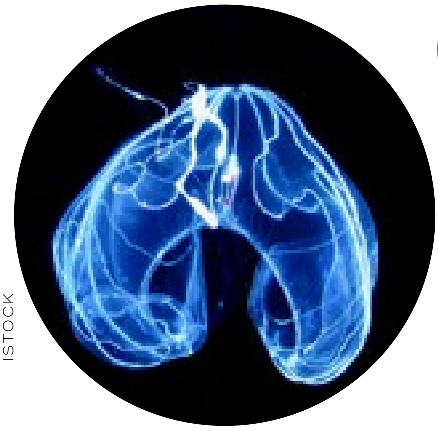

ISTOCK

Single-celled organisms (above) communicated by chemical cues. The comb jelly (left), like modern jellyfish and hydra, has a nerve net but no centralised brain.

Evolution of the brain

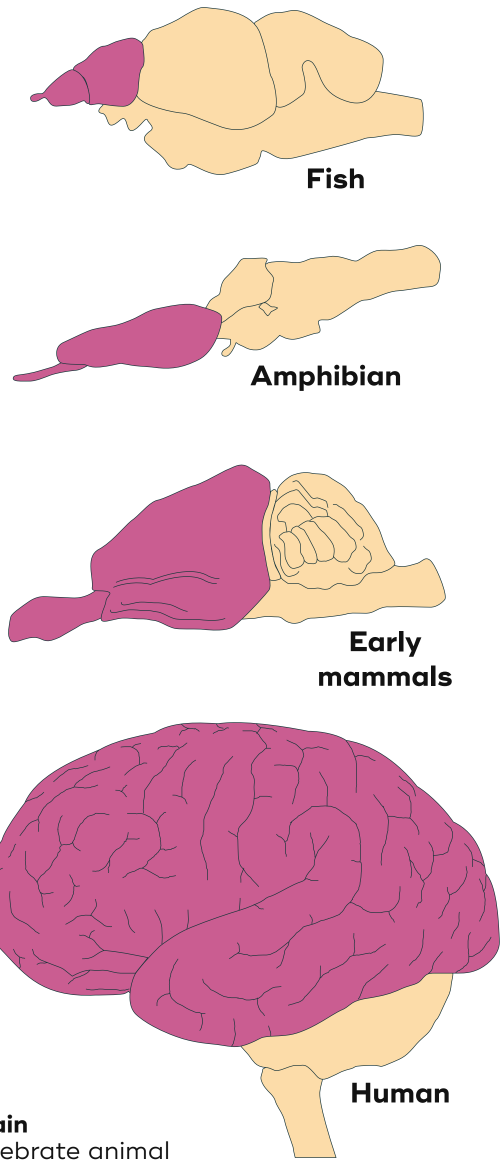

Forebrain
As vertebrate animal brains evolved, so did their forebrains, responsible for functions like planning, remembering and thinking.

3.9 BILLION YEARS AGO

First organisms appear. This included the ability to sense the world and react accordingly. Basic chemical cues likely allowed them to communicate with other single-celled organisms by using similar molecules to those used by our own neurons.

1.9 BILLION YEARS AGO

Ability to produce and use electrical signals develops in eukaryote cells (cells with a true nucleus), which are like our cells and those of other multicellular animals.

600 MILLION YEARS AGO

The earliest nerve nets appear in jellyfish-like creatures; light sensors, which function like early eyes, evolve in flatworms.

500 MILLION YEARS AGO

First vertebrates with complex brains appear, likely associated with the evolution of free-swimming in animals.

350 MILLION YEARS AGO

Four-legged vertebrates, with the ability to breathe air, colonise land and develop a complex forebrain, critical for processing information for cognitive and sensorimotor functions.

200 MILLION YEARS AGO

The six-layered neocortex (wrinkly outer region of the brain) originates as the first egg-laying mammals appear.

120 MILLION YEARS AGO

Placental mammals evolve the corpus callosum as the main link between the brain's hemispheres, which allows a further expansion of left-right connections.

6–2 MILLION YEARS AGO

Early hominids (human ancestors) appear, and as they evolve bigger bodies, develop social interaction, begin to walk upright and make simple tools, they develop slightly larger and more complex brains to process and store information.

2 MILLION–800,000 YEARS AGO

As early humans spread globally, they encounter new environments and new challenges, which further increases brain size.

800,000–200,000 YEARS AGO

Human brain size—particularly the neocortex—increases rapidly; the modern human brain is three times the size of the earliest human brain.

UNDERSTANDING THE BRAIN: A BRIEF HISTORY

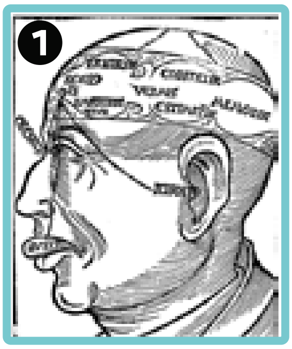

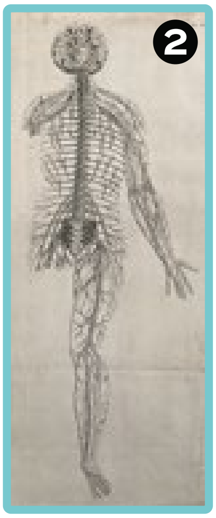

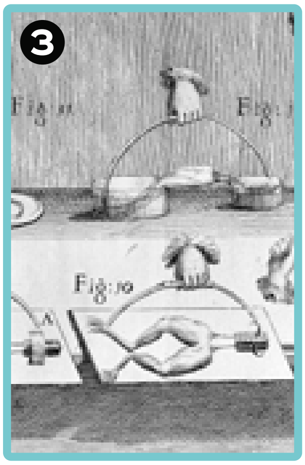

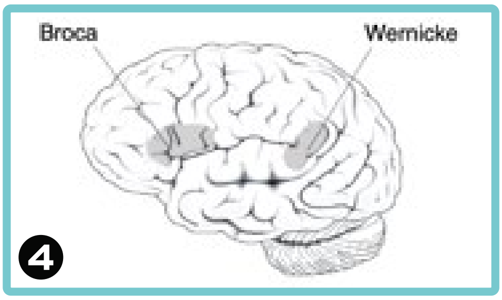

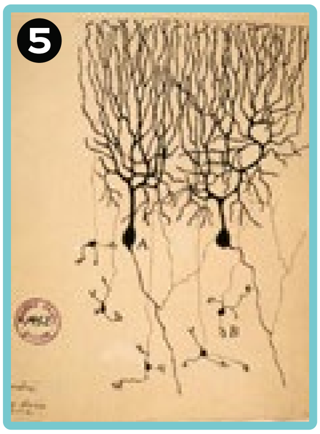

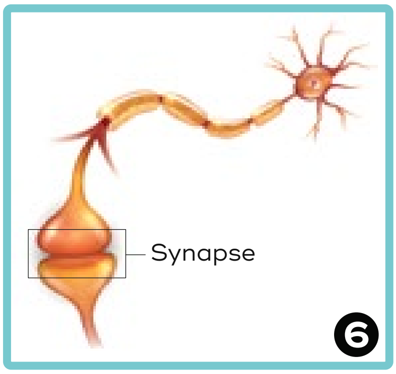

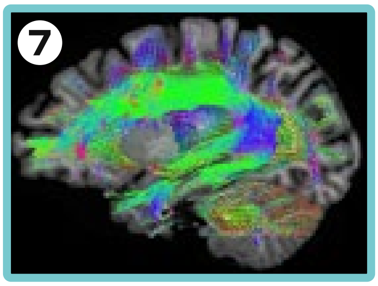

IMAGE CREDITS 1. THE NOBLE EXPERIENCE OF THE VIRTUOUS HANDY WARKE OF SURGERY, LONDRES, 1525 2, 3. WIKIMEDIA/ WELLCOME IMAGES 4. WIKIMEDIA COMMONS 5. INSTITUTO CAJAL/WIKIMEDIA COMMONS 6. ISTOCK 7. QBI

1 THE BRAIN AS A RADIATOR

In 335 BC, Greek philosopher Aristotle thought the brain was simply a radiator that kept the all-important heart from overheating. Around 170 BC, Roman physician Galen suggested the brain's four ventricles (fluid-filled cavities) were the seat of complex thought, and determined personality and bodily functions. This was one of the first suggestions that the brain was where our memory, personality and thinking reside.

2 FIRST NERVOUS SYSTEM SKETCH

In the 16th century, Belgian anatomist Andreas Vesalius created a highly detailed map of the nervous system and argued against the ventricles as the site of brain functions. We now know he was correct: the ventricles are filled with the cerebrospinal fluid that nourishes brain cells and cushions the brain against physical impact.

3 ELECTRICITY FIRES THE NERVES

In 1791, in the first suggestion that electrical impulses were important in the nervous system, Italian Luigi Galvani showed that electricity applied to nerves could make muscles contract.

4 THE BRAIN HAS REGIONS

In 1848, American railroad worker Phineas Gage had an iron rod strike his head, passing through his left frontal lobe. He survived, but aspects of his personality changed, suggesting that specific brain regions were important for certain functions. This idea grew stronger after studies in the 1860-70s by physicians Paul Broca and Carl Wernicke showed that specific parts of the brain were dedicated to different components of speech.

5 NEURONS, UNITS OF THE BRAIN

In the early 1900s, anatomists were taking advantage of microscopes and new staining methods to explore the smallest parts of the brain. Neuroanatomists Santiago Ramón y Cajal and Camillo Golgi were awarded the 1906 Nobel Prize for identifying that nerve cells (neurons) are the building blocks of the brain, and showing there are many different types.

6 COMMUNICATION IN THE BRAIN

In 1932 Sir Charles Sherrington and Edgar Adrian won the Nobel Prize for proposing the concept of synapses (junctions between neurons, pictured), which advanced the understanding of the central nervous system; Alan Hodgkin, Andrew Huxley and Australian Sir John Eccles won a Nobel Prize in 1963 for showing how neurons communicate via electrical and chemical signalling.

7 EXPLOSION OF NEUROSCIENCE

From the 1960s there was an explosion of neuroscience research. With rapid advances in technology, and collaboration across fields such as physics and genetics, scientists have made great leaps in understanding the brain, through detailed imaging and mapping of networks (pictured), and deciphering chemical pathways.

CHAPTER

2

ARTIFICIAL INTELLIGENCE

Unlike our own intelligence, artificial intelligence (AI) isn't driven by a survival instinct. Yet human-like robots and machines that aim to out-compete humanity are typically what we conjure up when thinking of AI. The fact is, AI is all around us and we often use it without realising. Every time you ask Siri to do a task, use a Snapchat filter, have junk mail filtered from your inbox, or get a traffic update on Google Maps, AI is at work.

What is Artificial Intelligence?

First, let's be clear on what intelligence is: the ability to acquire knowledge through learning, and to apply that knowledge to solve problems. Artificial intelligence is the use of a program, or algorithm, to perform a task which requires humans to use our intelligence.

Just as we use our brains to learn from new information gathered by our senses, AI learns from information fed to it—either many examples of an image, say, or the rules of a game—and processed according to how it is programmed. When AI is programmed to follow an algorithm, it is called machine learning. AI can handle massive amounts of data that our brains simply cannot cope with.

For the simple task of identifying a cat, humans take into consideration the shape and physical attributes and check this against our knowledge of what a cat is, based on our memories and experience. Importantly, though, our learning naturally involves building abstract representations, meaning we can recognise a cat even if we see only the back legs and a furry tail, or if we see a drawing of a circle for a head and two triangles for ears. For AI to identify a cat, you would feed it millions of cat images and train it to recognise particular pixel patterns that form a cat shape.

Intelligence is the ability to acquire knowledge through learning and to apply that knowledge to solve problems.

This is exactly what Google did in 2012, using a type of machine learning known as a deep neural network (inspired by the brain) to develop a program that learned to identify images that contained cats. The program was not fed the rules that cats have four legs, a tail, two ears and so on, but because the images were labelled as containing or not containing cats, the program was able to learn the visual concept of a cat. When given a new image, the program could label it as cat-containing or not. In other words, this Google AI acquired information from the images, learned to identify cats, and could then apply rules to solve the question of which new images contained cats. But unlike our brains, the AI doesn't really *know* what a cat is, or understand it can have abstract representations or is an object that changes shape.

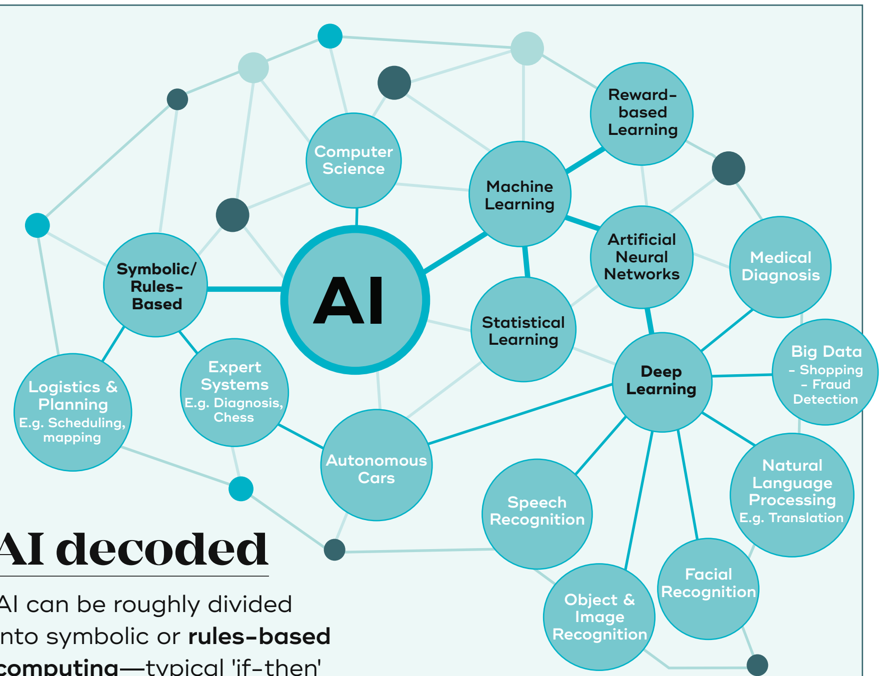

AI decoded

AI can be roughly divided into symbolic or **rules-based computing**—typical 'if-then' statements—and **machine learning**, where AI is programmed with an algorithm that can modify itself as it gains new information. There are different ways machine learning can be employed: deep learning by using artificial neural networks, reward-based learning, and learning based on statistical principles. Many AI programs are used for analysing and processing images or speech or extracting information from them. Deep learning is often used for making predictions, like medical diagnoses or possible credit card fraud.

What AI is good at

What AI is not good at

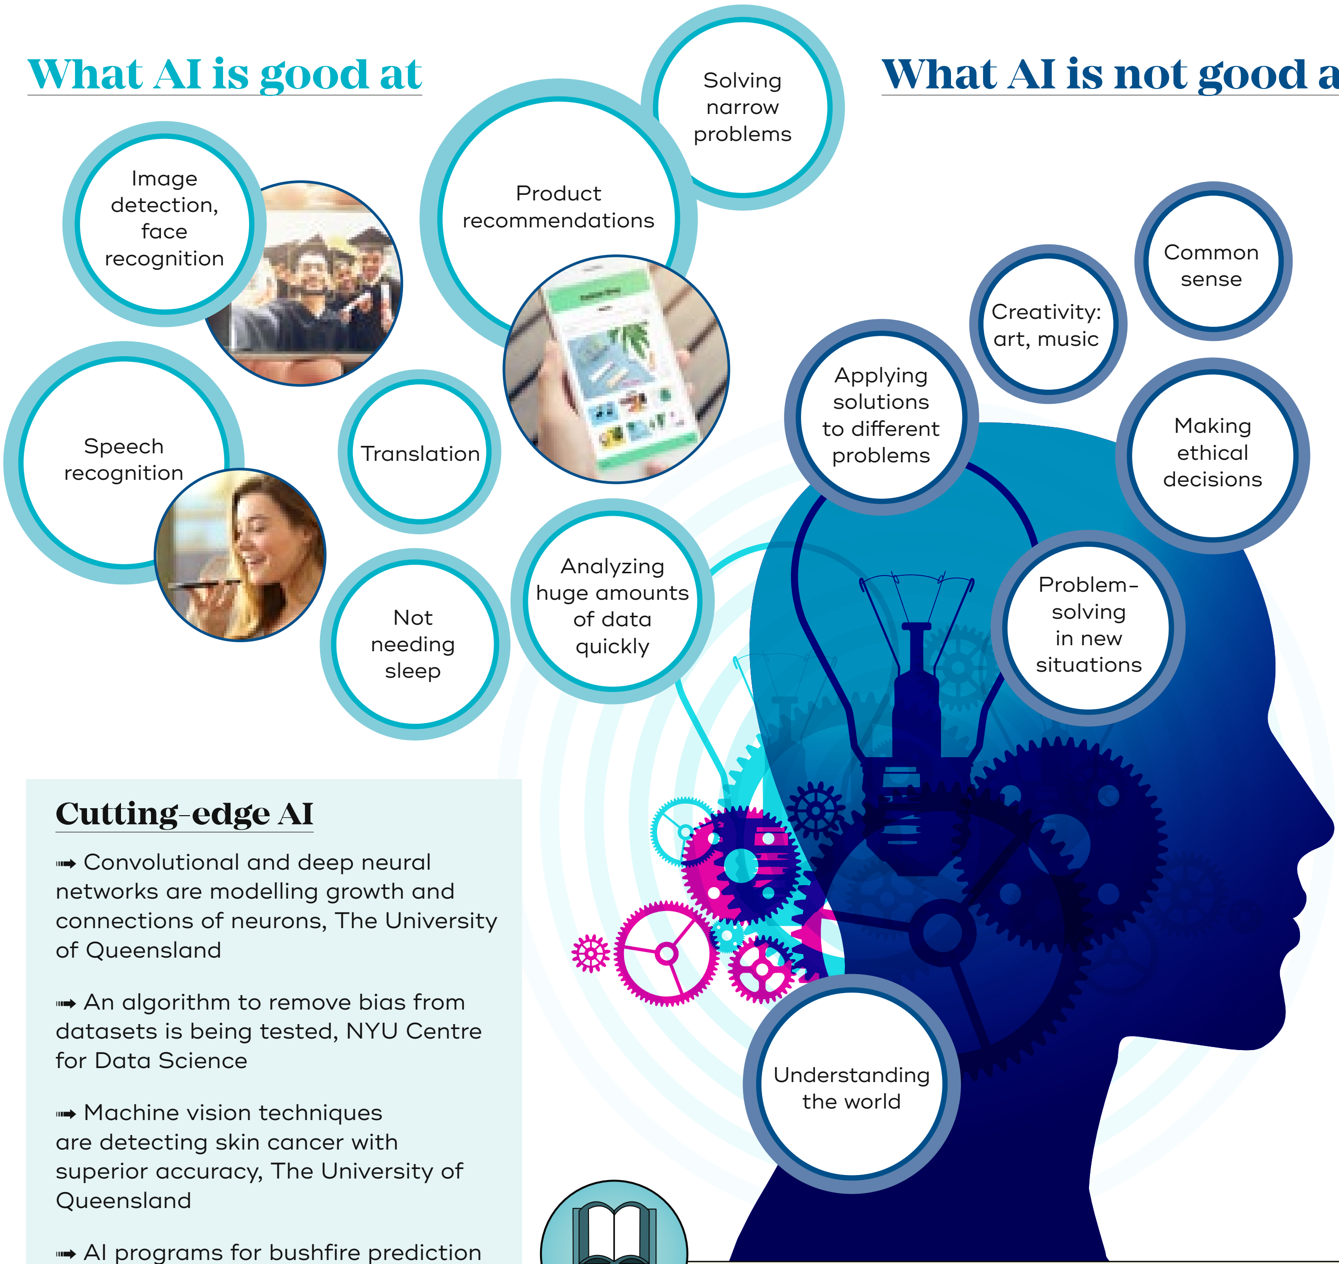

GETTY/ISTOCK

Cutting-edge AI

➟ Convolutional and deep neural networks are modelling growth and connections of neurons, The University of Queensland

➟ An algorithm to remove bias from datasets is being tested, NYU Centre for Data Science

➟ Machine vision techniques are detecting skin cancer with superior accuracy, The University of Queensland

➟ AI programs for bushfire prediction and analysis, CSIRO Data61

➟ Machine learning and computer vision are streamlining pathology diagnostics, The University of Queensland

➟ Computer vision is helping to understand superbugs, to eradicate them without antibiotics, The University of Queensland

➟ A tetris-like AI program is helping to detect breast cancer tumours, University of Adelaide

➟ Generative adversarial networks (GANs) are being researched for unsupervised learning in AI, MIT

Find out more:
qbi.uq.edu.au/intelligentmachines

A-Z

DEFINITIONS

Machine learning
Computer-based learning achieved by following an algorithm (a set of instructions or rules) to maximise the chance of a prediction being correct. Any program that improves its performance based on new inputs is an example of machine learning. It may involve artificial neural networks, but doesn't have to.

Artificial neural network
A computer-based system to process information, loosely modelled on the brain's architecture. It typically involves units, or 'neurons' connected to one another, comparable to how the brain's neurons communicate. Artificial neural networks are highly simplified, abstract forms of brain networks.

Deep learning
A type of machine learning that uses 'deep' neural networks. The word 'deep' refers to the many layers in the neural network. Deep learning is how most of today's artificial intelligence programs operate.

Neural
Refers to anything relating to neurons, or the nervous system.

History of AI

Modern AI began in the 1950s with the view to solving complex mathematical problems and creating 'thinking machines'. From the outset, there were two competing approaches. One used formal rules to manipulate symbols, a logic-based approach not at all based on biology. This became known as 'good old-fashioned artificial intelligence': GOFAI. The other camp took inspiration from how the brain works and created 'artificial neural networks' loosely inspired by our brains. These still had to be trained using certain procedures to solve problems.

In the first 20 years, GOFAI was the more successful approach, leading to much hype and significant government funding. But in real-world settings GOFAI didn't achieve its outcomes. Artificial neural networks also struggled, and in the 1970s funding dried up, research slowed and the AI community shrank. In the 1980s, improvements were made in both the rules-based GOFAI systems and biologically-inspired neural networks. Previously difficult problems became achievable and AI seemed promising once again. However, the hope and hype exceeded reality, and by the 1990s AI research again diminished.

The latest surge of interest comes off the back of the power of deep learning, a type of biologically-inspired neural network that harnesses the huge amounts of data now available, and the massive computational power and speed of today's computers.

With enormous data sets, modern AI neural networks can often exceed human performance in many tasks, including pattern recognition and playing games like Go, previously very difficult for AI. Importantly, these systems can learn from experience, unlike GOFAI.

AI's ubiquity might now *appear* like it's not far off reaching human-level intelligence. But AI needs massive amounts of data to learn, unlike our brains, which can learn from a single experience.

Some researchers argue that for AI to advance further, more needs to be understood about the basic principles of how our brains function, and the kinds of biological shortcuts our brains take to complete tasks.

...AI needs massive amounts of data to learn, unlike our brains...

TWO TYPES OF AI

Logic-based rules system 'Good Old-Fashioned AI'	Brain-inspired neural networks

1642
First mechanical calculating machine built by French mathematician and inventor Blaise Pascal.

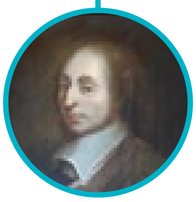

1837
First design for a programmable machine, by Charles Babbage and Ada Lovelace.

1943
Foundations of neural networks established by Warren McCulloch and Walter Pitts, drawing parallels between the brain and computing machines.

1950
Alan Turing introduces a test—the Turing test—as a way of testing a machine's intelligence.

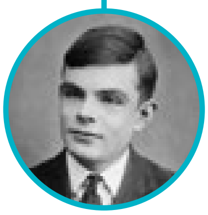

1955
'Artificial intelligence' is coined during a conference devoted to the topic.

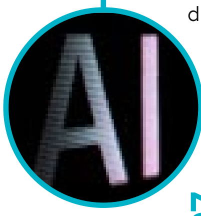

1965
ELIZA, a natural language program, is created. ELIZA handles dialogue on any topic; similar in concept to today's chatbots.

1980s
Edward Feigenbaum creates expert systems which emulate decisions of human experts.

1997
Computer program Deep Blue beats world chess champion Garry Kasparov.

2002
iRobot launches Roomba, an autonomous vacuum cleaner that avoids obstacles.

2009
Google builds the first self-driving car to handle urban conditions.

2011
IBM's Watson defeats champions of US game show Jeopardy!

2011–2014
Personal assistants like Siri, Google Now, Cortana use speech recognition to answer questions and perform simple tasks.

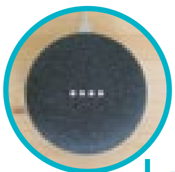

2014
Ian Goodfellow comes up with Generative Adversarial Networks (GAN).

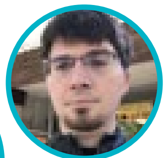

2016
AlphaGo beats professional Go player Lee Sedol 4-1.

2018
Most universities have courses in Artificial Intelligence.

IMAGE CREDITS: 1, 2, 3, 5, 9, 10, 11 WIKIMEDIA COMMONS; 4, 6, 7 ISTOCK; 8 TWITTER

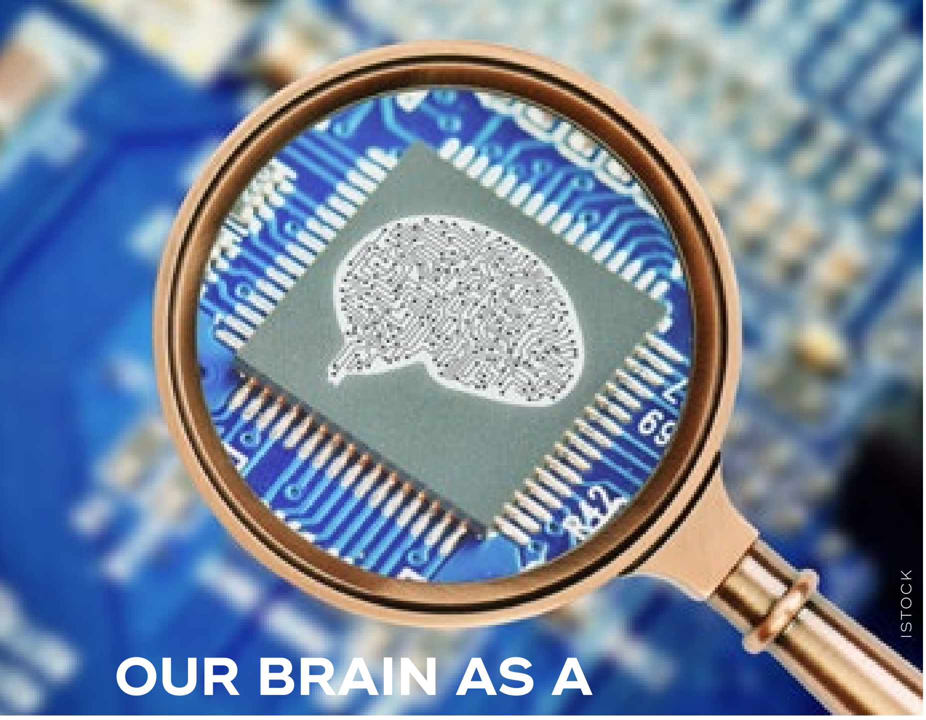

ISTOCK

OUR BRAIN AS A SOURCE OF INSPIRATION

THE HUMAN BRAIN PROVIDES MUCH INSPIRATION FOR THE DEEP NEURAL NETWORKS THAT POWER TODAY'S AI.

Scientists know how powerful and efficient our brains are, and if artificial intelligence is to match or even approach human intelligence, then it makes sense to be inspired by nature.

Most current AI is built to learn by using artificial neural networks, which emulate many structural aspects of how neurons are organised in the brain.

Neuroscientists still don't know exactly how our brains process all of the information we take in, or decide what's important to learn. But based on studies, even some from the 1950s, we know that information from our senses travels up and down different layers of processing in the brain.

For example, when looking at a cat, your eye detects the image through the retina. Information about that image is then transferred to the thalamus, a part of the brain important for relaying sensory information and regulating sleep and consciousness. The signal then travels sequentially through multiple areas of the neocortex. At each level, different features of the cat image are processed (see illustration, p9). This all happens in a third of a second, as we recognise the shape in the image as a cat.

Artificial deep networks operate in a similar way. The 'deep' part of the name refers to the fact they can have many layers, which is one reason we need powerful computers to make them work.

ARTIFICIAL NEURAL NETWORKS

Another similarity of AI with our brain relates to neurons—the cells of the brain. The equivalent components in deep neural networks are called 'units'. Like neurons, these units are connected to each other, providing a way for information to move between layers.

And like neurons, the strength of connections between artificial deep network units can change. The more a group of connected neurons is used, the stronger that path becomes. The less it's used, the weaker it becomes. In the brain, changes in strength occur because of a process called plasticity—the ability of the brain to adapt or respond to repeated stimulation—which underlies learning.

Deep networks also learn by adjusting the strength of connections between units. After the network processes an input image (i.e. a picture of a cat), its output is checked, and, if it made a mistake (e.g. it didn't detect the cat), the connections are re-adjusted so that it will be better at recognising the cat next time. Over time, if the network is trained on a sufficient number of images, it learns to find cats even in pictures it has never seen before.

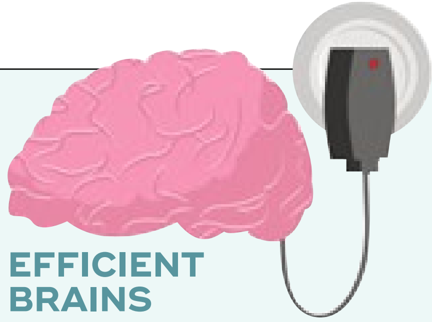

EFFICIENT BRAINS

An adult brain weighs about 1.3kg, and while that's about 2% of the overall body weight, it uses about 20% of the body's resting energy to power its ~100 billion neurons. Our brains are incredibly efficient, though, compared to computers, even with all the sleep they require. The brain needs about 15 Watts of power to perform basic functions; compare this to IBM's Watson supercomputer, which needed 90,000W to defeat Jeopardy! champions in 2011.

ISTOCK

Reward-based learning

In 2015, researchers at DeepMind (now part of Google's Alphabet) created a deep neural network that learned to play Atari video games. The program wasn't told the basic rules or purpose of the game, which was to shoot down spaceships. Instead, it was designed only to increase its 'reward' in the form of a higher game score. Through trial and error, the AI learned to play those games. This is a different type of learning than detecting images of cats.

In biology, reinforcement- or reward-based learning, is a fundamental part of how the brain works. It's a similar way to how you'd train a dog with treats to perform a desired action.

"The human brain is the only existing proof that the sort of general intelligence we're trying to build is even possible," Demis Hassabis, CEO of DeepMind said in an interview with The Verge, "so we think it's worth putting the effort in to try and understand how it achieves these capabilities."

ISTOCK

Limits of today's AI

- Enormous data and computational requirements
- Like a black box: Difficult to know how or why a deep learning AI program produced a particular outcome
- Difficulty understanding irony, humour, sarcasm or metaphors
- Bias: AI performance is based on data that contains human biases
- Interpreting tone and intentions from speech
- The need for human supervision to set the rules and goals
- Difficulty creating or understanding meaningful art or music, novels or poetry

How brains and AI process images

Both the brain and deep neural networks process information in stages. For visual processing in the brain, early stages (V1, V2) detect simple features like edges (yellow) and contours (light orange). Later stages (V4, IT) represent increasingly complex features (dark orange and red). This all occurs in about a third of a second.

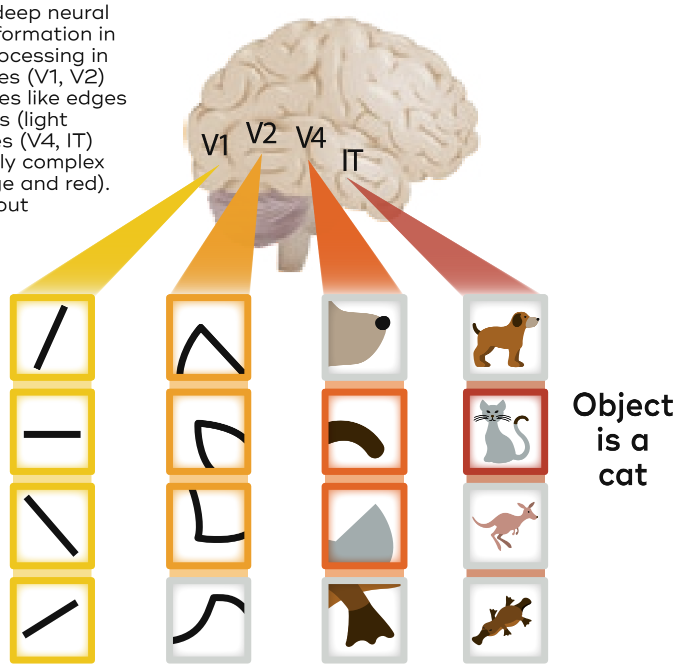

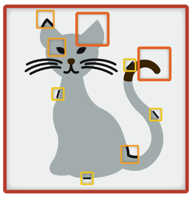

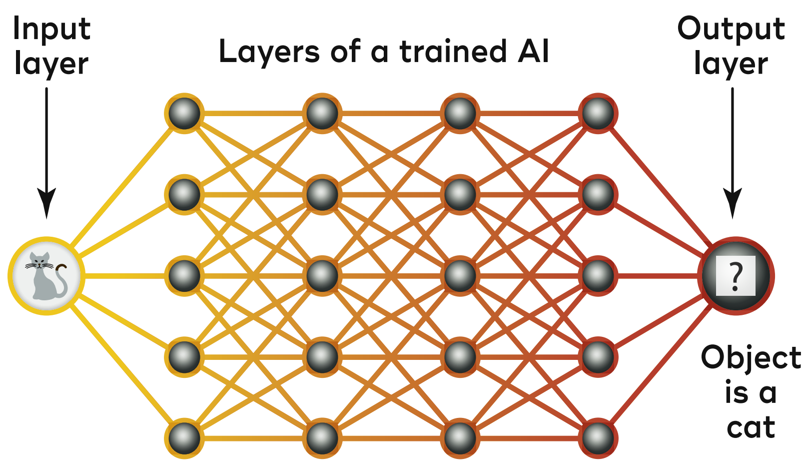

QBI/NICK VALMAS

Above: Deep networks also learn to represent features in different stages. In this way the layers of a neural network emulate the different parts of the visual cortex of our brains. The deep network AI learns from pixel information of images with labels (e.g. images labelled "cat" or "not cat"). A trained AI like this one would then scan a new image, and go through multiple stages, processing increasingly complex visual information until it came out with a result, in this case, the result is that the object is a cat.

A–Z

DEFINITIONS

Neuron Also called a nerve cell. It is the basic unit of our nervous system.

Reinforcement or reward-based learning
Using rewards or penalty to guide learning. With the aim of maximising its reward and minimising mistakes, the network learns an approach that meets its goal.

IS AI SMARTER THAN A HUMAN BRAIN?

The answer to that question is...sometimes.

At some tasks—such as playing Atari video games, chess or Go (see p10), or identifying cats in YouTube videos—today's deep networks can closely match or even outperform humans. Some types of image processing and pattern detection are now performed better by AI, such as knowing whether chest X-rays show signs of pneumonia. But, if you took the incredible AlphaGo network and asked it to identify cat videos on YouTube, you'd see it fail. Even the Atari-playing AI could only perform well at one task at a time: once it was trained on a new game, it forgot everything it knew about the previous one.

In other words, AI right now is extremely task-specific.

The quest for artificial general intelligence, capable of producing the range of human cognitive abilities, is considerably more difficult. Although scientific and tech influencers such as the late Professor Stephen Hawking and Elon Musk have talked up the dangers of super-intelligent AI, people working in the AI field, like DeepMind CEO Demis Hassabis, are far more restrained. The AI experts will tell you we are a long way from any such thing.

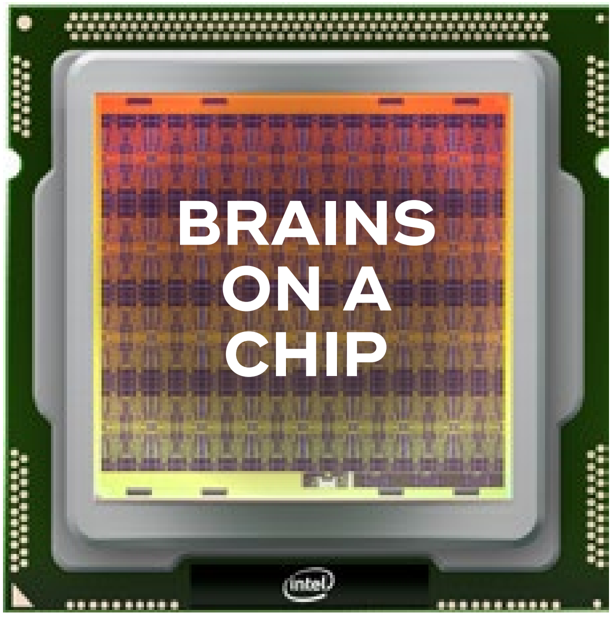

INTEL

Today's AI uses the brain as inspiration for software that runs on traditional computers (or supercomputers). But scientists are working on another option to model the brain's hardware by building a neural network on a microchip.

Called neuromorphic computing ('in the form of a nervous system'), in theory, the chips offer a completely different and more energy-efficient way to build a computer. Traditional computers are power-hungry because information is shuttled back and forth between the central processing unit (CPU) and the memory storage. In neuromorphic chips, units process and store information as an integrated part of the same operation, just like neurons do via synapses (the connection points between neurons). And the strength of those connections can grow or fade, as they do in our brains. This would let neuromorphic chips undertake parallel processing, similar to the brain, and would make some computing tasks like image processing (e.g. detecting cats) much faster and less energy intensive, by about 1000-fold. For pure calculation, though, a normal computer would still outperform a neuromorphic chip – just like it outperforms a brain.

As it is, neuromorphic chips aren't so useful…yet. "Just because we know how to build hardware that simulates brain components doesn't mean we know how to make use of it," says Dr Peter Stratton, a computer scientist working at the Queensland Brain Institute.

"That's because we don't really know how the brain uses its own hardware! Until we understand more about how a brain computes things, the full potential of neuromorphic chips isn't likely to be realised."

AI BECOMES A BOARD GAME MASTER

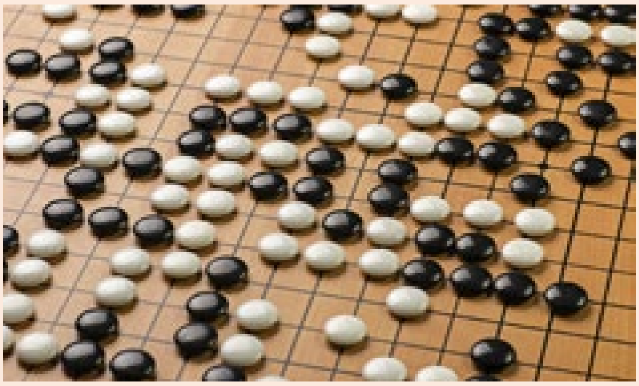

ISTOCK

The game of Go is a board game for two players invented in China more than 2000 years ago. Just like chess, it's a game of strategy for which the best players require extraordinary intellect and insight. It's considered the world's most complex and longest continuously played board game—and a good candidate for testing AI.

AlphaGo was developed by Google's British AI company DeepMind Technologies to beat a human in the game of Go. In 2016, at a much publicised machine-human showdown in Seoul, South Korea, AlphaGo convincingly defeated legendary human grandmaster Lee Se-Dol, by four games to one.

Just 18 months later, a newer version of AlphaGo, named AlphaGo Zero, thrashed its predecessor 100 games to 0. Improvements in AI are to be expected, but the different way AlphaGo Zero was 'trained' is particularly interesting. AlphaGo learned by studying 30 million moves from human experts and then played against itself to improve even further. It was a process that took many months.

AlphaGo Zero, however, was simply given the rules of the game and then played against itself 5 million times, all within the space of just three days. AlphaGo Zero trained itself by relying on deep reinforcement learning. A 'reward signal' was calculated for every possible move, based on the probability of winning from the new game position. Because the program was built to maximise its reward signal, it chose the move that most increased its probability of winning.

The success of AlphaGo Zero has shown that 'pure reinforcement learning', unguided by any human data, is possible. This offers the tantalising prospect that AI might be useful in developing entirely new solutions to problems, uncorrupted by any human biases.

3

ROBOTIC FUTURE

Our brains exist to help us survive. Organisms that don't have a brain, like bacteria or jellyfish, can still find food and evade predators. But with a brain, survival is more likely because we can remember where food is found, due to the hippocampus; plan a route that avoids danger, thanks to the prefrontal cortex; and prioritise memories with emotional significance, courtesy of the amygdala. These cognitive abilities wouldn't be so helpful, though, if we couldn't move. So, without having a body to move around and manipulate its environment, could AI ever be as intelligent as a human? Some scientists believe intelligence, as we know it, is so intricately connected with our bodies that building robots is the surest way to create truly intelligent machines.

ROBOT REALITY

CAN A ROBOT MAKE A CUP OF COFFEE IN YOUR KITCHEN? THIS BASIC TEST IS YET TO BE PASSED

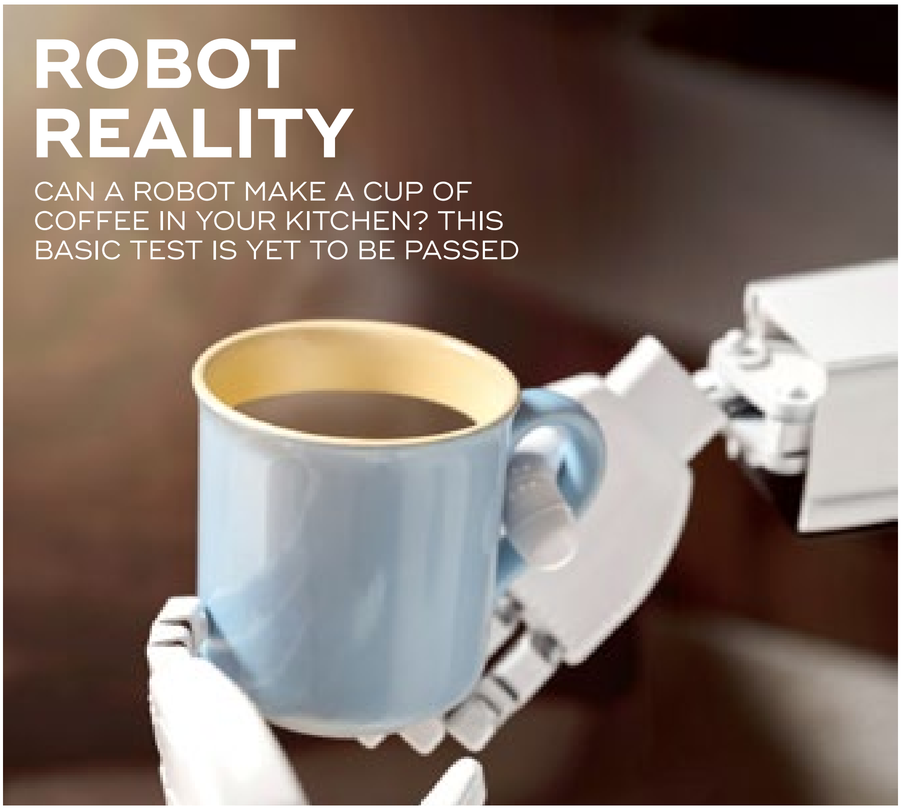

ISTOCK

Robots are machines that can automatically carry out a complex series of actions, particularly those programmed with sophisticated computer algorithms. Computers have been growing exponentially in power and speed ever since their invention more than 70 years ago. Today even a smartphone can perform more than a billion calculations per second. Our biggest supercomputers are capable of 100,000 trillion calculations every second – equivalent to tens of millions of smartphones working simultaneously.

Why then are computers still mostly boxes with screens and maybe keyboards attached? Why don't we have computer-controlled robot servants that can, for example, bring us breakfast and clean our houses? Modern robots can be agile enough to perform such tasks – the Atlas robot from Boston Dynamics is an example. So, physical limitations are not the issue. The problem is our current best robots, and the computers powering their 'brains', are bad at making sense of the world.

Computers excel at performing precise calculations at lightning speed, but struggle with simple real-world tasks, like turning on a tap, folding a towel or finding and picking up an object. A well-known benchmark in robotics research is the 'coffee test': can a robot enter a house it hasn't seen before and make a cup of coffee? This simple-sounding task includes many steps that are difficult, even impossible, for current robots. They would need to know what any kitchen could look like and be able to find it in any house. Then they'd need to open cupboards, find the coffee, identify and pick up a spoon from a drawer, find the kettle and sink, position the kettle to fill it with the right amount of water by operating the taps, and so on.

Such tasks are beyond the capabilities of current technology. The real world is too complex and variable to program a robot to successfully complete most of these steps. The solution is to build robots that can learn.

... our current best robots, and the computers powering their 'brains', are bad at making sense of the world.

ISTOCK

Robots that learn

WE'RE STILL A LONG WAY FROM INTELLIGENT ROBOTS

Industrial robots programmed for repetitive movements, to build cars for example, have been around for decades, but cannot learn. Newer warehouse robots, like those used by retail giant Amazon to pick up, store and deliver stock, can operate in environments with people and other obstacles. Unlike their industrial predecessors, these are mobile and can deal with some uncertainty about what they'll encounter while performing tasks, although all their actions are still programmed, not learned.

Scientists are now attempting to build robots that can learn by example. These robots often have sensors for vision and touch, for example, and manipulators such as simple hands or grippers. The idea is to teach the robot by guiding its manipulators to perform tasks you need, like picking up an item from a conveyor belt and placing it in a box, or collecting litter.

Basic manipulation tasks, however, are challenging for a robot to copy. Something that seems obvious to us, like the difference between an empty chip packet, a large leaf or small purse, can be difficult for a computer to comprehend, due to infinite possibilities of shapes, colours and sizes (see image below).

Having a body can help. We humans can pick up an object, turn it over, look inside and even smell it. We use our senses and bodies in flexible ways that allow us to accumulate evidence and make clear decisions when situations are ambiguous.

How can you teach a robot this? Our intelligence (and common sense) is the culmination of years of experiences, each of which adds a little more to our knowledge and ability to interpret the world. This is why scientists believe that to build truly intelligent machines, we'll need to equip them with bodies like ours and let them experience the world and mentally 'grow up' as children do.

For a taste of the current state-of-the-art domestic robots: qbi.uq.edu.au/best-robots

Something that seems obvious to us...can be difficult for a computer to comprehend...

Is this a leaf or a purse? We humans can easily tell, by using our hands, and senses to spot the difference. A robot would have great difficulty with the same task.

DIANA ULANOVA

ROBOT BRAINS:
GENERAL ARTIFICIAL INTELLIGENCE

WILL ROBOTS SURPASS HUMANS AND TAKE OVER THE WORLD? NOT LIKELY.

Despite impressive feats like deciphering human speech and beating the Go world champion (see p10), the current best AIs can perform only specific tasks: the AI that can understand speech for example, cannot play Go, and vice versa.

One big AI goal is to build artificial general intelligence, or AGI, which is good at most things, like humans are. It's what we'll need to create a helpful house robot, one that can do dishes, mow lawns and even read stories to children.

Controlling a body in the real world, however, is a far more complex problem than any AI can currently solve. The world is too varied and unpredictable to train an AI on everything it could possibly encounter. Our brains work most things out for us and an intelligent robot would need to do this as well.

Current AI programs also need massive amounts of training data, whereas our brains can often work things out from a handful of examples. We do this through trial and error and by applying what we already know to new situations—a form of creativity beyond current AI.

Possibly the biggest difference is that brains are adept at thinking ahead: planning. We can foresee consequences of our actions and mentally trial different options before doing anything. Currently, AI lives in the present or makes plans based on very structured rules. The Go-winning AI didn't make plans; it made the best moves it knew based on the current patterns of stones on the board and what

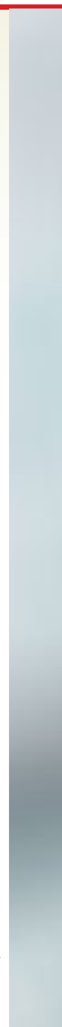

ISTOCK, GETTY

happened with similar past patterns. Current AI is good at sensing—deciphering speech, recognising objects and detecting board game patterns—but useless at abstract thinking.

The reason is that AI is input driven. It calculates the best output for a given pattern of input. But brains, especially complex ones like ours, are mostly internally driven and don't need input to keep going.

We don't know much about this internal activity, but its purpose is most likely to consolidate memories by revising past occurrences, make future plans, and look for shortcuts in procedures and links between concepts—those "AHA! I get it!" moments. AI is challenged by such things.

In fact, thinking may be closely related to planning ahead for our next movements. It's possible that having a body is an important step in creating an AI that thinks.

AI can't learn about or understand the world as we do, as there are too many possibilities and contexts. And that means we're likely still decades from helpful house robots taking over our jobs...or fully autonomous robots taking over the world!

Current AI programs also need massive amounts of training data, whereas our brains can often work things out from a handful of examples.

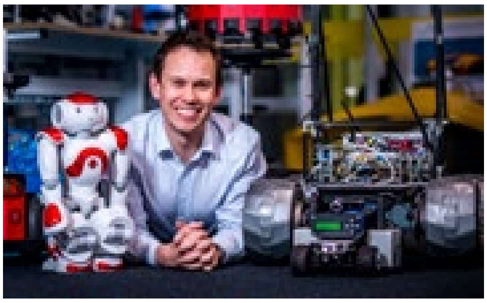

PROFILE

MICHAEL MILFORD

Robotics Professor Michael Milford's work is intrinsically linked with neuroscience. The Queensland University of Technology engineer draws inspiration from biology and has collaborated with neuroscientists at QBI to study how rats navigate, and if this can be applied in robots.

One of the greatest issues facing brain-inspired robotics, he says, is that technology for studying the brain isn't yet informative enough to really help robotics' researchers, often leaving them to make wild, albeit educated, guesses. "From a computer science perspective, it's frustrating that we still have very limited observability of the brain," he says. "In my area, we create networks that mimic the navigation processes in the brain, and neuroscientists can only observe maybe a couple of hundred of these navigation cells at any one time, when there's hundreds of thousands. It's like peeking into the brain through a tiny hole and only seeing a little bit."

Drawing inspiration from neuroscience is not simply about mimicry, he says. "The perfect [neuro-robot] product, in terms of fulfilling the consumer's need, would not necessarily be a perfect replication of the brain, because we don't do everything perfectly," he explains. "There are lots of examples where technology is actually better than parts of biology. Look, for example, at cameras: 10 years ago, the human eye was far superior to any camera you could get. Now, some cameras are comparable to the human eye." Likewise, technology doesn't need to sleep.

He predicts roboticists and neuroscientists will work together more closely. "The brain is the ultimate machine, and our technology that builds on top of it won't be identical, but it will draw a lot of inspiration from it."

4
MERGING BRAINS AND MACHINES

The brain inspires some of the most exciting computer technologies of our time, but the relationship isn't just one-way. Computers—including the deep-learning algorithms used in today's AI—can also help us understand how the brain works. The brain–machine relationship goes much deeper than just serving as a research tool or inspiration. In this chapter we'll delve into some of the incredible medical treatments and other developments emerging in the field of brain-machine technologies (see p18) which aims to create a seamless human-machine interface by combining engineering and neuroscience.

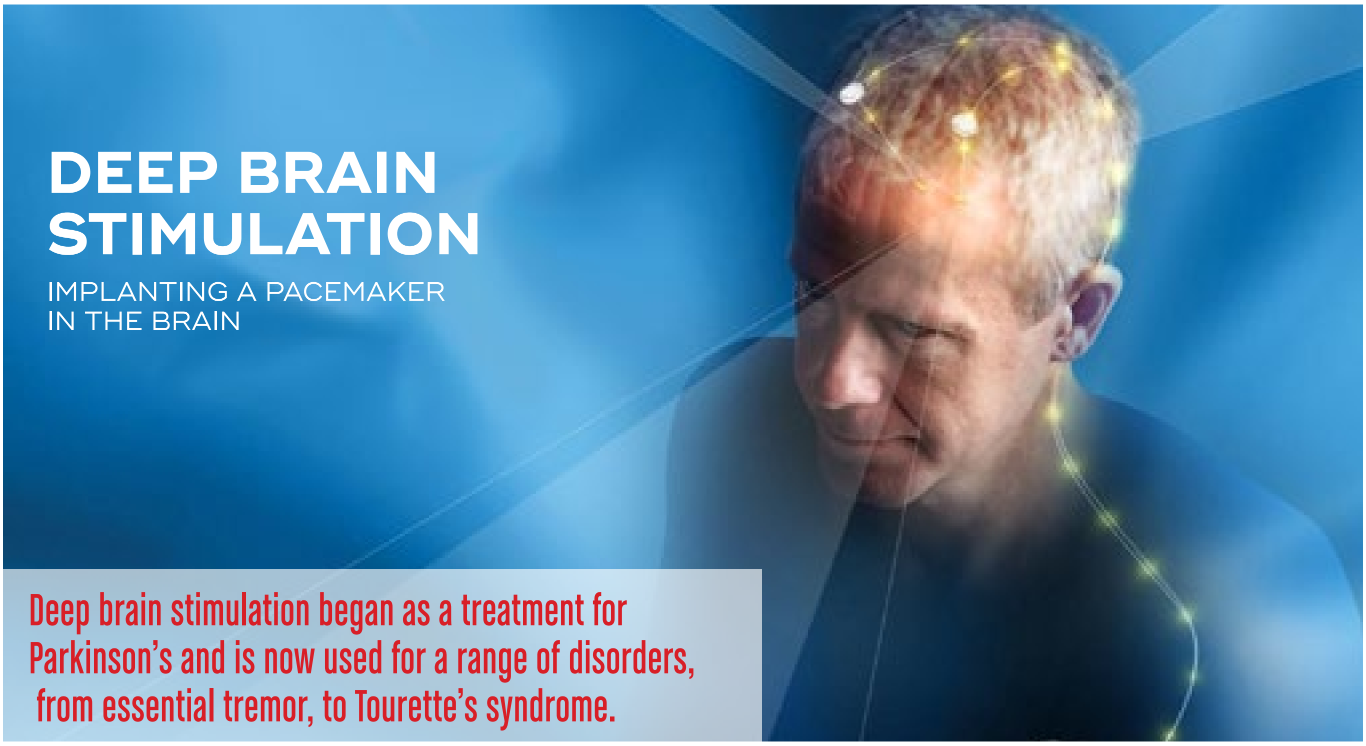

MEDTRONIC

Deep brain stimulation (DBS) is the use of electrodes surgically inserted into target regions deep in the brain, to stimulate neurons. This reduces or stops unwanted symptoms of disorders, such as the tremors of Parkinson's disease. Likened to a pacemaker for the brain, the device is programmed to deliver regular stimulation, as a way of keeping misfiring neurons on the right track.

Which brain structure is targeted depends on the neurological disorder and the specific symptoms being treated. Deep brain stimulation began as a treatment for Parkinson's and is now used for a range of movement disorders, from essential tremor, to dystonia. Clinicians and researchers see huge promise for the technique in treating a range of other conditions, like Tourette's or OCD. Doctors can also use the electrodes to record direct activity of the brain and learn more about how it functions.

THE FUTURE OF DEEP BRAIN STIMULATION

Despite the success of deep brain stimulation, it's not fully understood how or why it works, which makes it difficult to improve the technology.

One problem with current devices is that the stimulation is always on and always delivered in exactly the same pattern, at least until the patient has a check-up with their neurologist. The constant use drains the chest-implanted battery that powers the electrode in the brain and may also produce side effects. An ideal solution would be to create a device that automatically adjusts the level of stimulation based on the person's symptoms. For example, the brain of a patient with epilepsy can exhibit tell-tale signs of an oncoming seizure, which could be used to trigger stimulation that may stop its progression. Once the seizure activity ends, the stimulation would switch off. The key is being able to detect and respond to abnormal activity, which is not an easy task. Some experts believe this will be easier for epilepsy than other conditions because the misfiring of neurons is more defined.

Another challenge is the placement of electrodes. Because some brain regions affected by neurological disorders are extremely small, placing an electrode in the right location can be very difficult. Neurosurgeons have reference atlases for a patient to help guide placement, but every brain is slightly different. Doctors use many sources of information to confirm the best location, including: the type and amount of electrical activity of neurons; a pre-surgery MRI scan of the patient's brain; and the path of the electrode within a patient's brain.

DEEP BRAIN STIMULATION USES

CONDITION	BRAIN REGION TARGET	STATUS
Parkinson's disease	Subthalamic nucleus	Approved for use
Essential tremor	Thalamus	Approved for use
Dystonia	Pallidum	Approved for use
Obsessive-compulsive disorder	Nucleus accumbens	Clinical trial
Major depressive disorder	Multiple	Clinical trial
Epilepsy	Multiple	Clinical trial

DEFINITION

Brain-computer interface
Also called brain–machine interfaces. These devices use decoded brain activity to control an external device, such as a mouse cursor or prosthetic limb.

Electrode
A device used in deep brain stimulation to measure or stimulate electrical brain activity. Any instrument used for recording brain electricity (eg. EEG or electroencephalography) or stimulation requires electrodes.

PROFILE

PETER SILBURN & TERRY COYNE

Terry Coyne

Peter Silburn

Professor Peter Silburn and Associate Professor Terry Coyne head one of the country's most experienced deep brain stimulation teams and are also research leaders at QBI's Asia Pacific Centre for Neuromodulation, where they are developing new and improved ways of placing deep brain stimulation electrodes for the best outcome. Together with their highly specialised team, the pair has performed 1000 deep brain stimulation surgeries in the past 16 years. "We've learned a lot from stimulating the deep brain of humans," Prof Silburn says.

Although deep brain stimulation is used mostly for movement disorders such as Parkinson's disease and essential tremor, Professor Silburn's work is now expanding in other directions. "What we learnt from Parkinson's and tremor disorders, and also from Tourette's, really opened our eyes," he says. "We could settle down the tics and uncontrollable movements, but we also noticed that things like anxiety, obsession and compulsion settled.

"Now we can potentially treat people with psychiatric illness, and that's where our team is now progressing. We're looking at obsessive-compulsive disorder (OCD), anorexia nervosa, and other disorders."

Professor Silburn is currently running a clinical trial that uses deep brain stimulation to help people with severe, treatment-resistant OCD. "When the drugs don't work, and the psychological therapies don't work, that's when deep brain stimulation comes into the picture," he explains.

The effects can be dramatic. "One of the most satisfying things in my career is you see people with 20, 25 years of problems just dissolve in front of your eyes. The expression in people's eyes is really quite moving, to be honest. It's extremely rewarding. There's no doubt about that."

"When the drugs don't work, and the psychological therapies don't work, that's when deep brain stimulation comes into the picture."

MIND-CONTROLLED MACHINES

British science fiction writer and futurist Arthur C. Clarke once said, “Any sufficiently advanced technology is indistinguishable from magic”. It seems prophetic in the face of recent advancements in brain–computer interfaces, which appear to border on the magical.

In 2012, for example, a woman paralysed from the neck down used nothing but her thoughts (and a robotic arm) to bring a bottle filled with coffee to her mouth. At the start of the 2014 FIFA World Cup in Brazil, a young paraplegic man used his thoughts and a robotic exoskeleton to kick a soccer ball. And in 2017, ‘locked-in’ patients, unable to move a muscle or even blink, were able to communicate with doctors via their brain activity.

What connects these extraordinary examples is the brain–computer interface. This works by relaying brain activity to computers, which translate it into actions. In the case of the paralysed woman, 96 electrodes recorded her brain activity over many hours as she trained herself to think “move the arm in direction X”. Eventually, researchers could decipher the patterns of brain activity that produced different movement commands. These patterns could then be translated to a digital signal sent from a computer to the robotic arm.

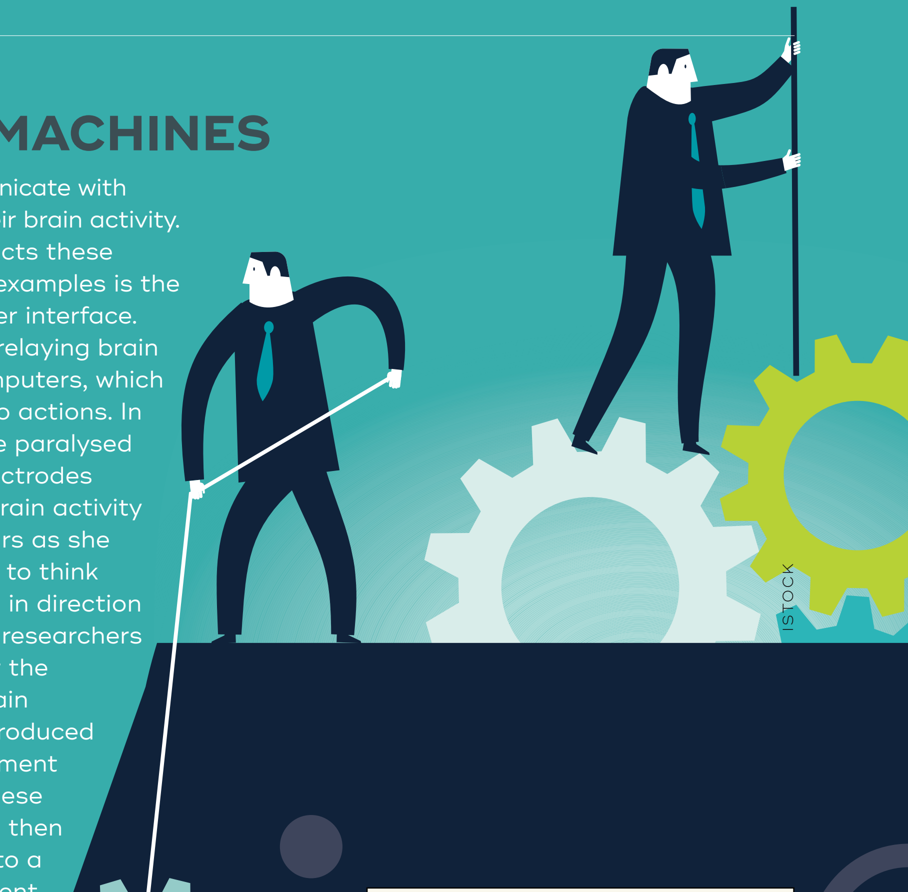

ISTOCK

Eavesdropping on neurons

TO LISTEN TO NEURONS, YOU HAVE TO GET INSIDE THE BRAIN

Given adult brains have about 100 billion neurons, it is very difficult to home in on the ones responsible for a certain action. To get closer to neurons, scientists have developed brain-computer interfaces that need to be inserted into the brain, to sit alongside neurons and ‘listen’ in on their activity. Researchers have been able to isolate and listen to about 200 individual neurons at a time, similar to eavesdropping on 200 conversations, rather than hearing an indistinct crowd murmur (see p18).

ISTOCK

Neuroscientists and engineers are working to increase this to 1000, even 10,000 neurons. The more neurons we can listen to, the more communication we can decode. Another goal is to make electrodes more tolerable to the brain. Long-term implantation causes the brain to scar where electrodes are located, which makes the results less effective.

MACHINES THAT BYPASS THE BRAIN

Devices that stimulate muscles could help those with damaged spinal cords

Today’s invasive brain-computer interfaces are highly experimental, and because they need to be surgically implanted, their use is limited to people with severe medical conditions. Spinal cord injury has, so far, been the most common target, the idea being to restore movement to legs or arms by bypassing spinal cord damage: instead of a brain command going through the spinal cord to muscles in the limbs, it is decoded by a computer that then sends a command to a limb. The limb may be a prosthetic device or the patient’s own arm or leg. While still in the early days, the technology has been successful. In 2017 a man with quadriplegia successfully manipulated his own arm to feed himself using a brain-controlled system that electrically stimulated his arm muscles, in an arrangement known as functional electrical stimulation.

Wearable brain-computer interfaces

SLIPPING ON AN ELECTRODE CAP COULD HELP GUIDE OUR ATTENTION AND PERFORMANCE

ISTOCK

A less invasive way to eavesdrop on neurons is by using something called EEG (electroencephalography) to measure brain activity through the scalp. Typically, this is done by placing a cap with electrodes on the head, hooking up the wires to a computer and recording the electrical signals of the brain. Whereas implanted devices allow researchers to measure the activity of individual cells, EEG signals represent the entire activity of thousands of neurons, so their level of detail is much lower. Instead of listening to individual conversations among a crowd of people watching a game of tennis, say, it's like you're outside the stadium only able to hear the collective oohs and ahhs or murmuring of the crowd. Still, this general information about brain activity can be useful for medical therapies, education and even gaming.

REHABILITATION

Stroke survivors with severe movement problems may be unable to physically practice rehabilitation tasks, but they may be able to imagine making those movements. This process called 'motor imagery' activates some of the same brain networks used in actual movement. Then, using neurofeedback (see Definitions), an EEG program can provide real-time feedback on motor imagery. For example, a program that detects the brain signals of a movement such as closing a fist could send a signal to stimulate forearm muscles and cause the fist to close. This feedback helps patients link intended movements to successful ones, promoting new connections in the brain to compensate for damaged areas.

LEARNING

Another area of interest is using brain-computer interfaces to improve learning, in particular, how attention can be measured from brain activity, and how that might be used to keep students' minds focused.

Using an EEG, leading QBI cognitive neuroscientist Professor Jason Mattingley and colleagues from The University of Queensland have converted attention signals into an easy-to-understand visualisation on a computer monitor, a form of neurofeedback. In one study, subjects used this visualisation to regulate their own attention levels. QBI psychology researcher, Dr David Painter, who developed the technology, says it's advanced enough to be tested in applied settings, such as classrooms or workplaces.

VIDEO GAMES

Video games are a good platform to test the control of actions with thoughts, as many require only a few commands or buttons; a major limitation with brain-computer interfaces is the number of signals that can be reliably decoded and distinguished. QBI's Dr Painter says this is just one of several hurdles. "The decoding, or so-called mind-reading technology, is already here," he says. "The challenge for developers is to invent natural and useful means for interaction—the interface itself." Also important, he says, is what advances this technology offers: the unique experiences it could create. "Developers worldwide are solving this problem as we speak." This means mind-controlled video games could be the next big advance in gaming.

MEMORY MANIPULATION

A possible use for brain-computer interfaces is to enhance brain function. Studies have shown that electrically stimulating the brain can boost memory. Most of these take a general approach, trying to enhance all memories using the same stimulation pattern. But one 2018 study used stimulation to improve a particular memory. Scientists recorded the brain activity of epilepsy patients with implanted electrodes as they played a memory game. Later, when their brains were stimulated with the recorded pattern, they scored about 35% better on the memory test.

A-Z

DEFINITIONS

Neurotechnology
Engineered hardware that connects with the nervous system. Neurotechnologies can be input devices that alter brain activity (eg. deep brain stimulation electrodes) or output devices that record brain activity (eg. EEG devices). Prosthetics such as the cochlear implant and robotic arms are also neurotechnologies.

Neurofeedback
Feedback given to a person about their own brain activity, generally to help them self-regulate or train aspects of their own brain function. For example, brain signatures of attention levels can be detected and converted into a visual scale, which the user can learn to modulate themselves.

Machines that read minds

BLURRING THE LINE BETWEEN BRAINS AND COMPUTERS

If we can control machines with our brains, does that mean machines can read our minds? Not really. Using brain signals to control a device is very different to a computer knowing what we're thinking. Nevertheless, some experiments have shown that our brain activity can be used to digitally recreate images we have seen. Although impressive, these machines can mostly only interpret what we're seeing—a long way from knowing our thoughts and intentions.

This 'mind-reading' technology works by first showing people hundreds of images while their brains are scanned in a particular type of MRI machine. A computer program, often based on a deep neural network, learns that particular brain activity patterns occur when certain image features are displayed. Then, when a person in the MRI machine is shown a new image, the computer program can decipher, at some level, what the person is seeing.

In 2013, Japanese researchers used a similar approach to deduce some basic content from people's dreams. Our dreams are highly visual, making them suitable for the type of decoding described here. Still, like other 'mind-reading' feats, the level of detail and accuracy is low and restricted to visual features. There's no evidence we are close to true mind-reading, such as gauging what a person is thinking, intending or desiring.

...our brain activity can be used to digitally recreate images we have seen.

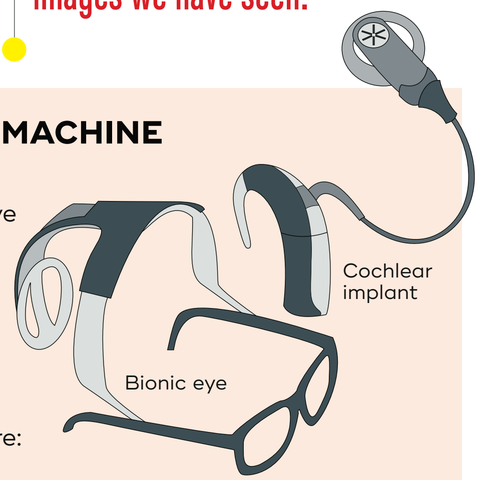

CUTTING-EDGE BRAIN-MACHINE TECHNOLOGIES

There are a host of other impressive technologies either currently available or being researched and trialled for eventual public use. Together, they show the promise of merging our brains and nervous systems with machines, while also illustrating how close we already are:

- ➟ The cochlear implant, invented by the Australian Professor Graeme Clark AC in the 1970s, has helped hundreds of thousands of people worldwide to hear.
- ➟ Researchers at Melbourne's Bionics Institute and others at UNSW are at the forefront of research to create different versions of bionic vision, to restore some sight in people with degenerative eye diseases.
- ➟ A clinical trial is underway in Australia for a spinal cord stimulator to relieve chronic pain, potentially providing a drug-free way to treat this common complaint.
- ➟ Avoiding the need for brain surgery, an Australian team has created a device that can be delivered to the brain's motor cortex via a vein in the neck. Once in the brain, the 'Stentrode' can record nerve signals from within a blood vessel, with the potential to assist paralysed patients.
- ➟ In the US, the Department of Defense research agency DARPA is making impressive progress on a bionic arm that connects directly to a person's limb. It is controllable by linking electrical signals from muscles of the remaining stump to manipulations of the prosthetic arm, giving the person mobility.


Professor Geoff Goodhill

QBI

DECODING THE LANGUAGE OF NEURONS

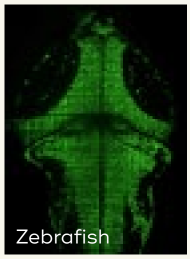
Zebrafish

Our 100 billion or so neurons communicate in a very complex language that scientists have not yet been able to translate. By understanding this language, called the neural code, we will be able to better communicate with neurons and build better brain-computer interfaces.

Computational neuroscientists like QBI's Professor Geoff Goodhill are starting to crack this code, using the tiny zebrafish as a model. These 4mm-long transparent fish larvae are ideal subjects because scientists can see the simultaneous electrical activity of neurons – up to the 100,000 in the entire fish's brain!

Professor Goodhill is studying how brain activity patterns in the zebrafish represent stimuli in their environment, and how these patterns change as the fish develops. His team recently discovered that the sensory environment of the fish is very important for its brain activity to develop correctly. For example, fish reared in the dark had different brain activity patterns to those in normal light, and were worse at catching prey when the lights were back on. Researchers believe the basic principles of neural coding—how information is stored and transmitted by brain activity patterns—apply across all animals.

CHAPTER

5

ETHICS

ISTOCK

Neuroscience is merging with technology in ways that will have a huge impact on society. This won't be limited to improvements in just health or brain function. There will also be profound ethical challenges and possibly even a redefinition of what it means to be human. So, what are some of the ethical issues created by merging machines with our brains—the organs that define us as humans and as individuals?

- IDENTITY
- RESPONSIBILITY
- MORALITY
- BRAIN ENHANCEMENT
- PRIVACY
- BIAS
- DEFINITION OF HUMAN

ISTOCK

DISCUSSION QUESTION

How much is a person responsible for their actions if their brain is changed by drugs or devices?

IDENTITY AND RESPONSIBILITY

Today, we are the 'agents' of our own actions – meaning we are in control. But technologies that alter our brain activity have the potential to blur that line. For example, what is someone's responsibility if they commit an out-of-character crime while being stimulated by such a device? Would the answers be any different for somebody on antidepressants or other medications, which also affect brain activity?

Identity could be another issue with devices that can change our patterns of brain activity. One anecdotal report described a patient with a brain stimulation device who sometimes wondered "who he was". Reports like this are extremely rare (and a loss of identity isn't rare for people with brain degeneration, suggesting an alternative explanation for his feelings); nevertheless, they remind us that changing our brain activity can, at least in theory, change our sense of self.

Taking the pulse of the public can help guide ethically fraught research. For example, one survey of students found that people generally don't like the idea of altering their personality traits. They were willing to improve cognitive abilities such as attention, alertness and memorisation, but empathy and kindness were off limits—perhaps because they shape a person's emotions, identity and sense of self.

In raising these issues, the point isn't to question the value of brain stimulation devices as medical therapies—they've proven themselves as safe, effective treatments, and patients are delighted to have their quality of life restored. The point is that as the technology progresses, we need to make sure that the legal and ethical guidelines keep pace.

How human do we want robots to be?

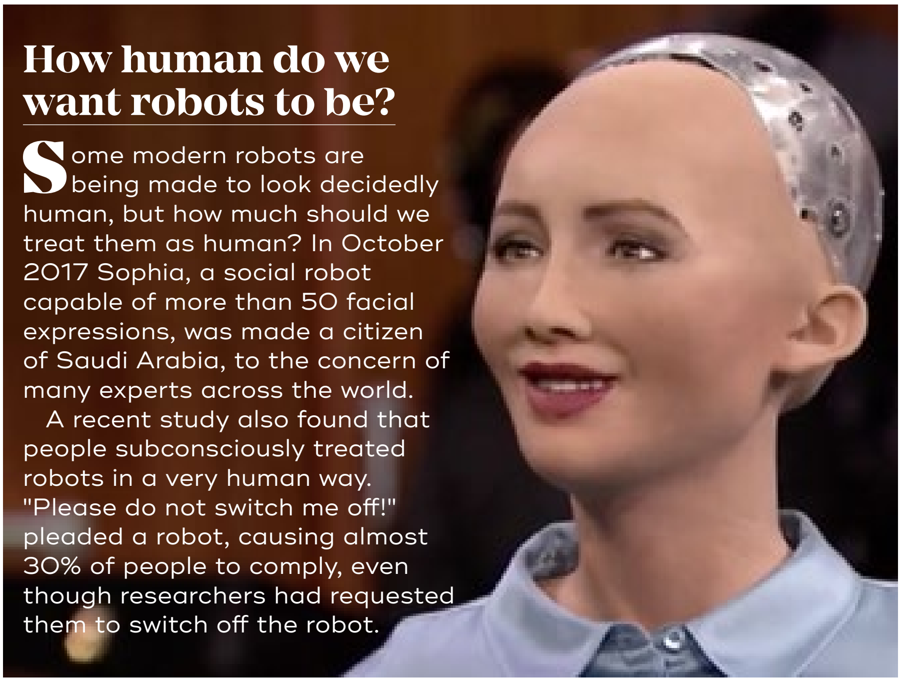

HANSON ROBOTICS

Some modern robots are being made to look decidedly human, but how much should we treat them as human? In October 2017 Sophia, a social robot capable of more than 50 facial expressions, was made a citizen of Saudi Arabia, to the concern of many experts across the world.

A recent study also found that people subconsciously treated robots in a very human way. "Please do not switch me off!" pleaded a robot, causing almost 30% of people to comply, even though researchers had requested them to switch off the robot.

Brain boosting

Enhancing brain function is certain to be one goal of future research, and the military may be the best example of who might want this. The United States defence research program DARPA is already investing heavily in brain-computer interfaces that could one day boost combat readiness, performance and recovery of military personnel.

But is cognitive enhancement something we should allow? The truth is some people already take steps to achieve it; 'smart' drugs like Ritalin and modafinil can help focus and extend attention; Prozac alters mood to ward off depression and anxiety.

The prospect of cognitive enhancement raises issues of equality and fairness—who should have access to these enhancements, and would they be limited to those who can afford it? Would a high score on a test be fair with the use of brain enhancement? Similar issues regarding performance-enhancing drugs are confronted in professional sports.

PRIVACY

Today's biggest technology companies collect huge amounts of personal information because they can sell it for commercial gain. This means that if or when your brain activity can be recorded with a wearable EEG headset (see p17), companies would find huge value in accessing your brain-based information.

Imagine, in the future, if you merely thought about buying a new Smart TV while wearing an EEG headset, and that information was relayed to a big online retailer. The retailer could use AI programs to automatically contact you with their latest specials on Smart TVs. You wouldn't even have to act on your thoughts by typing it into a search engine; your head-mounted device would record the activity pattern caused by thinking 'Smart TV', and commercial operators could act.

There are huge privacy concerns around who would or should have access to your brain activity. For example, what if a health insurance company wanted to buy your brain activity, which could indicate if you had a mental health disorder?

These are important issues to consider as technology improves and brain data becomes more readily available.

?

DISCUSSION QUESTION

Who should have access to your brain activity recordings? How could that be protected?

ISTOCK

Morality

The goal of robot researchers is to build robots with general intelligence (run by AI), guided by a sense of morality. But whose morals should a robot take? What principles should guide robot decisions? What happens when human input is purposely dishonourable?

For example, less than one day after its release in March 2016 to engage in 'conversational understanding' with the general public, Microsoft's AI-based chatbot Tay became a Hitler-loving racist and conspiracy theorist based on the interactions it had on Twitter.

Teaching general moral principles to robots and letting them deduce appropriate decisions won't work all the time either; there will always be exceptions to the rule or ambiguous situations. The other option is to have the robot learn through experience, just as humans do, but perhaps with ethicists guiding it.

But what moral code should the robot learn? There are many, though these are just some: the Golden Rule (do unto others as you would have them do unto you); utilitarianism (act for the greater good); the categorical imperative (actions bound by moral duty to do what's right, regardless of outcome).

Humans still often disagree on what the right moral decisions are even within a single country or culture.

?

DISCUSSION QUESTION

What moral code should AI follow?

AUTONOMOUS CARS: SHOULD THEY BE ON THE ROADS?

ISTOCK

With testing in progress from some big global companies, autonomous cars appear to be the next big thing. But they've also emerged full-force into the debate around AI and robotics, particularly after the March 2018 death of a woman in Arizona, US after she was hit by a self-driving Uber car.

This incident raises the obvious issue: do we allow autonomous vehicles on the road when inevitable accidents will happen?

Driverless vehicles are powered by machine learning, which relies on predictability to form rules and actions for the car. But people are inherently unpredictable. We make mistakes. We often behave erratically or ambiguously, which makes it almost impossible to teach machines what actions to predict and the appropriate responses to make. Just as with human drivers, self-driving cars will have accidents. As a society, we have to decide if that is acceptable, and who is at fault.

While self-driving cars may be touted as an inevitable future, there are countless obstacles that will keep them in test phase for a while.

ISTOCK

Bias

Can AI be biased? Algorithms programmed for AI are coded in a very logical way, with data fed into them and clear outcomes defined. The primary role of AI is to find patterns in huge data sets, so it would appear that they are impartial machines.

But the rules they operate by are encoded into the programs, and the data fed to them come from humans, who are not without flaws and biases. AI then, reflects the biases of their human creators.

In the US, for example, an AI-powered 'judge' predicted that the likelihood of criminal re-offence was significantly higher among African-Americans compared to Caucasian defendants. And yet when both groups were tracked for the following two years, the rates of re-offence were the same; African-American defendants had been wrongly stereotyped as more likely to commit a future offence. AI researchers are becoming more aware of the bias problem and are working to overcome it. IBM says more than 180 different human biases have been defined and it is working to address them in AI programs.

DISCUSSION QUESTION

What sort of biases could be unknowingly fed into AI algorithms? How could you offset them?

AI researchers are becoming more aware of the bias problem and are working to overcome it.

PROFILE

MARY WALKER

Rapidly expanding developments in neurotechnology could provide hope to many people. But they'll come with ethical issues that will be difficult to navigate, warns Dr Mary Walker, a bioethicist based at Melbourne's Monash University.

SUPPLIED

Dr Walker worries about the impact that technology may have on our identity and autonomy. For example, there have been a couple of reports of character or personality changes in specific cases where brain stimulation has been used. Dr Walker notes, however, that due to variations in reports of identity change it has been difficult to create a solid ethical framework addressing such issues.

"Sometimes people will say it changed their character and that's a good thing," she says. "Other people will say it's...something they experienced as aversive. So, it's really hard to pick out any particular patterns."

It's important, Dr Walker adds, to weigh up aversive experiences that may challenge someone's identity against the benefits that deep brain stimulation and similar technologies can provide.

"At the moment, implants are really used for therapeutic purposes," she explains. "They've got implications, not just in terms of character changes, but also in reducing motor symptoms in people with, say, Parkinson's."

However, using technology to enhance function, instead of overcoming impairments, shifts the technical debate. "At the moment, the people being treated often have important clinical problems, and the technology is being used to overcome them." Dr Walker says. "If you take that out of the picture, then the equation of cost-and-benefit changes significantly."

She is also concerned that these technologies will become "fads" that people will have difficulty keeping up with.

"There are problems around obsolescence—particularly something like a neural implant, which you may not be able to [easily] replace," Dr Walker says.

"So, if you spend a lot of money getting an implant or something to make you smarter, and five years later the technology improves, then everyone gets the new, improved implant and they're going to be smarter than you and you might not be able to catch up. So, there's some interesting questions around that as well."

CHAPTER

6 THE FUTURE

ANTICIPATING THE ISSUES AHEAD

What is next for the technologies covered in this magazine? The answer will be driven by the pace, direction and funding of research, the potential good that could result and, equally importantly, the ethical issues they raise. As with research into areas of medical science like genetic engineering and human stem cell technologies, both of which have profound potential benefits, society needs to decide how artificial intelligence should be regulated. And those decisions should have inputs from people with diverse perspectives and expertise, including the general public, lawmakers, scientists, ethicists and medical professionals.

The discoveries scientists make can certainly have great positive impacts on our quality of life, but unintended outcomes are also possible. So it's important to consider the many possible uses and motivations for developing these technologies.

Cyborgs: future fact or science fiction?

The term cyborg, (cybernetics + organism) was first used to describe a human-machine hybrid more than 50 years ago by Austrian-born, University of Melbourne-educated computer engineer Manfred Clynes. Writing in the journal Aeronautics, Clynes and American psychiatrist Nathan Kline explained that: “The purpose of the Cyborg…is to provide an organisational system in which… problems are taken care of automatically and unconsciously, leaving man free to explore, to create, to think, and to feel.” In other words, cyborgs are meant to liberate us from the drudgery of routine tasks, letting us concentrate on more creative, more human endeavours.

Today’s trends indicate that in the medium- to long-term future, a machine-enhanced humanity is possible. Although we are not there yet, at a very small scale, we are already becoming a mix of human and machine, in the form of deep brain stimulation devices, prosthetic limbs, cochlear implants and artificial pacemakers. These aren’t quite cyborg elements because they restore function rather than enhance us beyond normal human limitations. What's more, the most complex part of our bodies, our brains, will be the hardest part to augment.

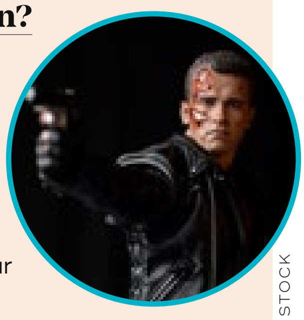

ISTOCK

But we aren’t far away. Several research groups are already using devices to improve learning and memory, with some success (see p17). And as technology becomes better at treating disease and dysfunction, and as we learn more about how the brain operates, there’s little doubt that human-machine hybrids are something we’ll confront.

NEAR FUTURE

- Better product recommendations
- More AI helping with medical diagnostics and personalised medicine
- Brain-computer implants, bionic eyes, arms, and legs and mind-controlled exoskeletons will become more advanced, and more available
- Chatbots will be commonly used by companies for customer interaction

MEDIUM FUTURE

- Brain-computer interfaces will be increasingly able to enhance memory function, especially for people with traumatic or invasive memories.
- Brain-to-brain interfaces will become viable between humans—sending telepathic messages not bound by location.
- Automation of many jobs
- Personal assistants like Alexa and Siri will be commonplace and more advanced
- Autonomous cars on suburban roads

DISTANT FUTURE

- Cyborgs will exist
- Artificial general intelligence (very distant future…)
- Robot servants or carers

ISTOCK

An intelligent future

Modern computing's high-powered technology, along with the emergence of strikingly human-like robots, would make you believe intelligent machines are not far off taking over our jobs and our homes. But the reality is, while AI programs are sometimes better than humans at specific tasks—like pattern recognition or playing games—they're still a long way from achieving the kind of general intelligence that we humans possess.

The mammalian brain took millions of years to develop its complex yet efficient abilities, and despite the power of modern computers, AI is not likely to match its potential for decades. There is still much that even 'simple' animal brains can tell us about how information is processed and what intelligence is. For instance, many animals—like octopus, squid and even bees—display intelligence, despite not having a forebrain.

These animals can outperform AI, in terms of how quickly they can master multiple tasks, with limited time and input for learning. "Typically, deep learning networks need millions of samples before they get the hang of it," says neuroengineer Professor Srini Srinivasan, from UQ's Queensland Brain Institute. "A bee lives for barely a month, and it's got to learn everything within that time," he says. "In the first week of its life it learns to forage for food, recognise species of food-bearing flowers, and navigate back home without getting lost."

AI is not even close to producing the kind of authentic, creative works that humans can conjure.

The joke is on AI: Will creativity distinguish humans from machines?

AI programs have been developed to do very human things, like write novels, compose music and create art, but can they take a joke—or even understand one? AI is more likely to be the punchline. Attempts by AI programs to write prose are often decidedly funny (to humans).

Take this example from an AI program that was fed all seven Harry Potter books:

"Leathery sheets of rain lashed at Harry's ghost as he walked across the grounds towards the castle. Ron was standing there and doing a kind of frenzied tap dance. He saw Harry and immediately began to eat Hermione's family."

Music, novels, poetry and visual art need nuanced cultural and historical understanding to be meaningful. AI is not even close to producing the kind of authentic, creative works that humans can conjure. In fact, AI cannot grasp humour, sarcasm, irony or abstract ideas.

"Producing a song with AI requires you to feed it thousands of rhymes and sound patterns," says Professor Srinivasan. "You don't have the organic process that humans have to create music, where we take inspiration from a life event or something you see."

Before AI can achieve the kind of creativity that humans can, says Professor Srinivasan, we first need to figure out the basics of how AI even works.

"No one fully understands how AI works and why it works so well," says Professor Srinivasan. We just feed in the data and turn the crank to get an answer—it's a wonderful but very mechanical process."

The answer may lie in knowing more about how the brain works, he says. "Understanding the biological short-cuts that brains use to perform the same tasks as AI can give us insights that can help to create better machine learning algorithms."

Our scientists are unlocking the mysteries of the brain to understand and treat disease, improve learning and memory, and inspire technology. Partner with us.

partners@qbi.uq.edu.au

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA
CREATE CHANGE

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

CREATE CHANGE

w: www.qbi.uq.edu.au / e: qbi@uq.edu.au / p: +61 7 3346 6300