WHAT IS ARTIFICIAL INTELLIGENCE (AI)?

It depends who you ask.

Back in the 1950s, the fathers of the field Minsky and McCarthy, described artificial intelligence as any task performed by a program or a machine that, if a human carried out the same activity, we would say the human had to apply intelligence to accomplish the task.
That obviously is a fairly broad definition, which is why you will sometimes see arguments over whether something is truly AI or not.

AI systems will typically demonstrate at least some of the following behaviours associated with human intelligence: planning, learning, reasoning, problem solving, knowledge representation, perception, motion, and manipulation and, to a lesser extent, social intelligence and creativity.

WHAT ARE THE USES FOR AI?

AI is ubiquitous today, used to recommend what you should buy next online, to understand what you say to virtual assistants such as Amazon's Alexa and Apple's Siri, to recognise who and what is in a photo, to spot spam, or detect credit card fraud.

WHAT ARE THE DIFFERENT TYPES OF AI?

At a very high level artificial intelligence can be split into two broad types: narrow AI and general AI.

Narrow AI is what we see all around us in computers today: intelligent systems that have been taught or learned how to carry out specific tasks without being explicitly programmed how to do so.

This type of machine intelligence is evident in the speech and language recognition of the Siri virtual assistant on the Apple iPhone, in the vision-recognition systems on self-driving cars,

in the recommendation engines that suggest products you might like based on what you bought in the past. Unlike humans, these systems can only learn or be taught how to do specific tasks, which is why they are called narrow AI.

WHAT CAN NARROW AI DO?

There are a vast number of emerging applications for narrow AI: interpreting video feeds from drones carrying out visual inspections of infrastructure such as oil pipelines, organizing personal and business calendars, responding to simple customer-service queries, co-ordinating with other intelligent systems to carry out tasks like booking a hotel at a suitable time and location, helping radiologists to spot potential tumors in X-rays, flagging inappropriate content online, detecting wear and tear in elevators from data gathered by IoT devices, the list goes on and on.

WHAT CAN GENERAL AI DO?

Artificial general intelligence is very different, and is the type of adaptable intellect found in humans, a flexible form of intelligence capable of learning how to carry out vastly different tasks, anything from haircutting to building spreadsheets, or to reason about a wide variety of topics based on its accumulated experience. This is the sort of AI more commonly seen in movies, the likes of HAL in *2001* or Skynet in *The Terminator*, but which doesn't exist today and AI experts are fiercely divided over how soon it will become a reality.

A survey conducted among four groups of experts in 2012/13 by AI researchers Vincent C Müller and philosopher Nick Bostrom reported a 50 percent chance that Artificial General Intelligence (AGI) would be developed between 2040 and 2050, rising to 90 percent by 2075. The group went even further, predicting that so-called ' superintelligence' -- which Bostrom defines as "any intellect that greatly exceeds the cognitive performance of humans in virtually all domains of interest" -- was expected some 30 years after the achievement of AGI.
That said, some AI experts believe such projections are wildly optimistic given our limited understanding of the human brain, and believe that AGI is still centuries away.

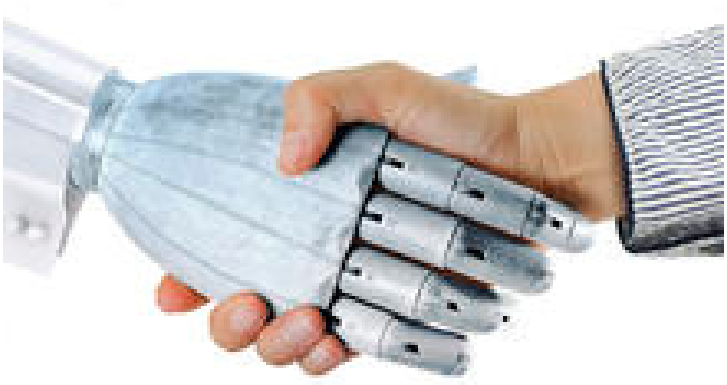

WHAT IS MACHINE LEARNING?

There is a broad body of research in AI, much of which feeds into and complements each other.

Currently enjoying something of a resurgence, machine learning is where a computer system is fed large amounts of data, which it then uses to learn how to carry out a specific task, such as understanding speech or captioning a photograph.

WHAT ARE NEURAL NETWORKS?

Key to the process of machine learning are neural networks. These are brain-inspired networks of interconnected layers of algorithms, called neurons, that feed data into each other, and which can be trained to carry out specific tasks by modifying the importance attributed to input data as it passes between the layers. During training of these neural networks, the weights attached to different inputs will continue to be varied until the output from the neural network is very close to what is desired, at which point the network will have 'learned' how to carry out a particular task.

A subset of machine learning is deep learning, where neural networks are expanded into sprawling networks with a huge number of layers that are trained using massive amounts of data. It is these deep neural networks that have fuelled the current leap forward in the ability of computers to carry out task like speech recognition and computer vision.

There are various types of neural networks, with different strengths and weaknesses. Recurrent neural networks are a type of neural net particularly well suited to language processing and speech recognition, while convolutional neural networks are more commonly used in image recognition. The design of neural networks is also evolving, with researchers recently refining a more effective form of deep neural network called long short-term memory or LSTM, allowing it to operate fast enough to be used in on-demand systems like Google Translate.

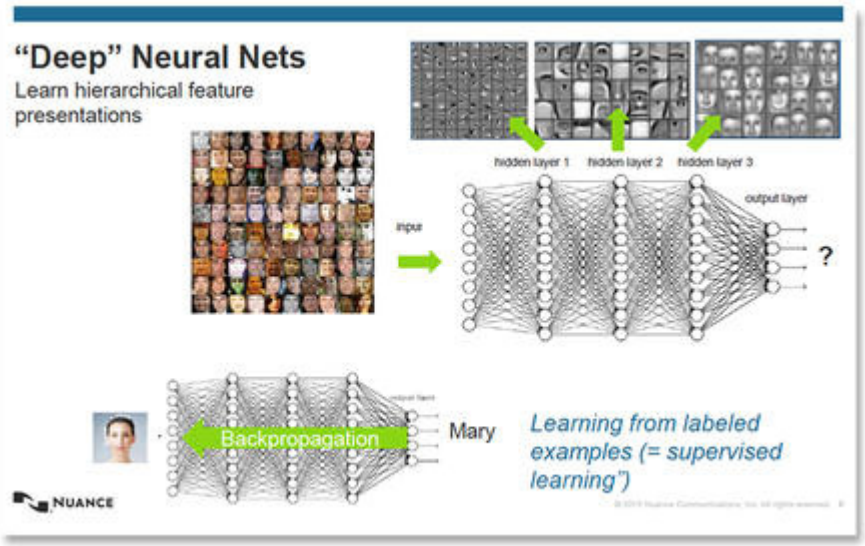

Another area of AI research is evolutionary computation, which borrows from Darwin's theory of natural selection, and sees genetic algorithms undergo random mutations and combinations between generations in an attempt to evolve the optimal solution to a given problem.

This approach has even been used to help design AI models, effectively using AI to help build AI. This use of evolutionary algorithms to optimize neural networks is called neuroevolution, and could have an important role to play in helping design efficient AI as the use of intelligent systems becomes more prevalent, particularly as demand for data scientists often outstrips supply. The technique was recently showcased by Uber AI Labs, which released papers on using genetic algorithms to train deep neural networks for reinforcement learning problems.

Finally there are expert systems, where computers are programmed with rules that allow them to take a series of decisions based on a large number of inputs, allowing that machine to mimic the behaviour of a human expert in a specific domain. An example of these knowledge-based systems might be, for example, an autopilot system flying a plane.

WHAT IS FUELING THE RESURGENCE IN AI?

The biggest breakthroughs for AI research in recent years have been in the field of machine learning, in particular within the field of deep learning.

This has been driven in part by the easy availability of data, but even more so by an explosion in parallel computing power in recent years, during which time the use of GPU clusters to train machine-learning systems has become more prevalent.

Not only do these clusters offer vastly more powerful systems for training machine-learning models, but they are now widely available as cloud services over the internet. Over time the major tech firms, the likes of Google and Microsoft, have moved to using specialised chips tailored to both running, and more recently training, machine-learning models.

.