

Artificial Intelligence and Machine Learning – Is it real or just hype?

A brief history of Artificial Intelligence (AI)

It is tempting to think of AI as a modern concept, but a short search of “Google” soon reveals that AI has roots in much older civilisations. Early Greek mythology contains concepts of AI and intelligent machines. Aristotle, in the 5th century, toyed with logic and reasoning systems. Throughout history, literature has recorded instances where AI concepts have been discussed and, in some cases, attempts made to build intelligent machines.

From the 1200’s rumours abound of “Talking Heads”, Roger Bacon and Albert the Great are both rumoured to have owned such artifacts. The Knights Templar order was destroyed in 1307, by Philip IV of France, accused of worshipping a brass talking head, called Baphomet, able to answer any question put to it. Cervantes, in Don Quixote, makes reference to a head, supported by a tube, that spoke.

As time progressed, science and technology developed and mechanical devices, such as clocks, became common. Mechanical toys were developed, Vaucanson’s mechanical duck and Von Kempelen’s mechanical chess player. In 1642, Pascal created an adding machine which, in 1673, was improved by Leibniz to include multiplication and division. In the 19th century, George Boole developed binary algebra which is the basis of today’s computer logic systems. Charles Babbage and Ada Byron, who the programming language ‘ADA’ was named after, worked on programmable mechanical calculators (the analytical engine was derived from these). Science fiction stories introduced us to robots and automatons, science and technology seemed to be making these ideas real but the computer age gave these AI concepts real substance and the possibility that these artificially intelligent entities could be created.

What is Artificial Intelligence?

Intelligence is defined as “The ability to acquire and apply knowledge and skills” and artificial usually refers to something that is “made or produced by human beings rather than occurring naturally, especially as a copy of something natural” and may be characterised as being false, fake, simulated or unnatural.

Artificial intelligence, a term usually associated with a computer-based system, can therefore be thought of as unreal, or simulated, intelligence. Physically, a computer system consists of a series of mechanical, electrical and electronic components which, integrated together, form an electronic system able to perform many thousands of small operations and process large amounts of data very rapidly. To be able to perform any useful functions this system, or hardware, needs a software programme. Without this software a computer is an inanimate object. It has no life, intelligence or purpose and certainly, cannot think for itself.

A software programme contains a series of commands which instructs a computer, or system how to obtain data (input), how to manipulate it (process) and what to do with it (output). Because a computer system is able to perform billions of operations per second and control trillions of pieces of data it appears to be intelligent but, in fact, is just responding to its software programme and so displays artificial intelligence.

Deep Blue versus Garry Kasparov

In the late 1990s, a pair of six-game chess matches, between world chess champion Garry Kasparov and an IBM supercomputer called Deep Blue were arranged. The first match, played in Philadelphia in 1996, was won by Kasparov. The second, took place in 1997 in New York City, was won by Deep Blue. The 1997 match was the first defeat of a reigning world chess champion by a computer under tournament conditions.

It was generally held as proof that computers were now more intelligent than humans. In reality this did not demonstrate that a machine has intelligence rather it showed that a machine's software can simulate learning.

Machine Learning

Machine learning is a part of the Artificial Intelligence discipline where the machine appears to learn without being explicitly programmed. It relies on software algorithms to look for patterns in data. Much like statistics, if data is available on a particular subject, careful analysis of that data will highlight patterns within the data and allow for predictions about likely outcome given particular conditions, to be made about the subject. The larger the data samples, the greater the accuracy of predictions which can be achieved.

Suppose a car manufacturer wants to streamline its production and stop producing cars of less popular colours. It decides to carry out a survey to determine what the most popular colours for cars on the road are.

They pay someone to stand on road at a junction and record the colour of the first 1000 cars that pass. The results show that black, white, red and blue cars are all equally popular at 25%. The survey is inconclusive as the data set was so small and only gathered from one location. To be of relevance a much larger data set is required, one which uses multiple locations and requires the use of a computer system.

A network of cameras is established on, motorway bridges and major road junctions and, connected to a central computer system, which has software which can distinguish a range of different colours.

Now the system can observe thousands of cars and store the information in a data base. As the system runs it counts the numbers of each colour car, if it encounters a new coloured car the software creates a new location to store and count that colour car.

When the sampling is complete the system can be queried and will provide details of the number of cars in each colour group and also the most popular for each location. To an observer the system appears to have intelligence and also the ability to learn, since it can expand its data base to include new colours. The system can only do this if the colours it encounters have already been defined in its software.

Can a computer really play chess?

If we think about a chess match there are, potentially, thousands of moves available.

But there are only 20 possible moves to open. Any of the pawns to row 3 or 4 or one of the two knights to A3 or C3 or F3 or H3.

A computer software programme, designed for playing chess, if able to access data from thousands of actual games, would be able to select which move to make. After each move the computer would systematically look at each move available and assess the options available to its opponent. Drawing on data from real games, and their outcomes based on statistics, it would select a move.

This is similar to the process a human would go through in deciding the next move and is actually planning several moves ahead, although would not have access to as much historic data.

A computer appears to display artificial intelligence and also, since it is able to store data from each new game, would also appear to be able to learn.

But can a computer “think” about what move to make and actually play chess or does it just simulate it? Of course, it cannot it is inanimate and not a sentient being hence, through its ability to access vast amounts of historic data very rapidly, it displays artificial intelligence.

A dog is a sentient creature but does a dog display intelligence? Many dogs will give you their paw if prompted. Is this a display of intelligence, or is it simply a conditioned response that may result in a treat?

Machine learning is essentially a means of achieving artificial intelligence and, while it may appear that the machine is exhibiting intelligence, it will always remain artificial.

There are several different approaches used to implement machine learning:

- Decision tree learning
- Inductive logic programming
- Clustering
- Reinforcement learning
- Bayesian networks

Deep learning, probably the best-known approach to machine learning, is based on the functions and structure of the brain and the interconnections of neurons.

Neural Networks

An artificial neural network (ANN), usually referred to as a neural network (NN), is a computer system which is designed to emulate the workings of the brain. In an ANN there are discrete layers of “neurons” which have connections to other neurons. Each of these layers are assigned a specific feature to learn, these could be colour, shape, edges, curved or straight lines etc. Through data analysis, connections are made in this network and patterns established. As an example, a computer with a neural network, is loaded with thousands of images of cats. By comparison with these retained images the computer can “learn” what a cat looks like and should then be able to select, from new images presented to it, which ones contain cats or not. Difficulties exist with this method if, for instance an image contained a white cat against a white background unless there is clear outline distinction on which to base the decision. However, this method can work well for detecting differences between two images such as in a crack detection system for structures.

Facial recognition apps are now common on mobile phones as a security measure. Comparing a face against a known image is a relatively straight forward process but trying to identify an unknown person by comparison with millions of images is much more challenging.

In the US, the National Institute of Standards and Technology (NIST) Facial Recognition Vendor Tests (FRVT) found that one industry leading facial recognition algorithm achieved an error rate of only 0.1% with good quality images with the subject looking directly at the camera but, when filmed “in the wild” with the subject moving and subject to light, shadow and moving normally, the error rate increased to 9.3%. A further problem was noted that if the images, being used as a comparison, were not recent the error rate rose by a factor of 10 due to the ageing process and facial features altering. Tests carried out at an airport boarding gate (a relatively controlled environment) achieved a 94.4% accuracy whilst a test at a sports venue returned accuracies between 36% and 84%. Although one top algorithm achieved an 87% accuracy the median was only 40% which explains why there is so much concern with the use of these current technologies as evidence by police and security agencies.

AI and Internet of Things (IoT)

AI and IoT are linked together in the same manner our body and brains are connected. Our bodies sense the state and condition of the environment we are in by touch, smell, sight hearing and taste and transmit this data to our brains via our nervous system. Our brains interpret that data and transform it into a format which we can understand and use. Sensors, in an IoT system, gather data and input it into a computer where software manipulates and interprets the data and produces useful outputs.

Uses of AI and ML

There are many areas where this machine learning will be of benefit to humans such as medical diagnosis. Vast amounts of medical data can be stored and processed quickly enabling rapid assessment of a patient's symptoms and conditions likely to develop and most beneficial treatments.

Similarly, with engineering, a gas turbine engine fitted with a range of sensors which can monitor temperature, vibration, oil pressure, speed and oil condition can be monitored by a software programme. This programme can monitor the state of each main bearing and, based on historical data, can predict likely failures before they occur. This predictive maintenance (PM) enables a more efficient maintenance program. A traditional maintenance system relies on a time-based periodic condition system, where a system is stripped down and inspected with components being replaced based on their observed condition as necessary. PM allows for a more focused and efficient maintenance programme where the actual condition of a system is monitored constantly and components changed prior to failure. This system also allows for a more efficient logistics system where components, displaying signs of imminent failure can be sourced ahead of the event. These system should lead to much more safe, efficient and cost effective maintenance programmes.

For more information, please contact andre.rose@imca-int.com.