Benjamin Rosman gives us an introductory overview

It is difficult to imagine what the modern world would look like without computers. We have come to rely on them for everything, from handling our financial transactions to maintaining our social lives.

Having computers assist in all these tasks requires a way of giving them instructions. For this we traditionally use programming, where programs are written as sequences of steps that should be followed by a computer to perform some function. The programmer here has to think carefully about any situation that the computer may find itself in, and explain to it how that problem should be dealt with.

Unfortunately, there are many problems that are just too hard to solve in this way. Think of the challenge of detecting if there is a face in an image. To do this, a programmer would describe an algorithm – or set of instructions – to identify two eyes, placed above a nose, which in turn is above a mouth. But how should the computer recognise these eyes, noses and mouths? And what happens if the face is seen from the side so only one eye is visible, or the image is upside down? And how do we make sure this works for anyone, of any age or race, even if they have hats, glasses or beards?

Machine learning is an area of artificial intelligence that focuses on learning patterns from data, rather than being programmed by humans.

The insight of machine learning is that even though solutions to problems like these are very difficult for a human to specify, it is easy to collect examples of images that do and don't contain faces. If this can be done, a computer should be able to write the program or algorithm itself using these examples. It does this by trying to learn the rules that would distinguish the examples of 'faces' that it has been given from the examples of 'not faces', using the underlying data such as the pixels in the images.

This is a specific case of machine learning, known as **supervised learning**, because the algorithm is supervised by a human, training it with a set of examples with faces

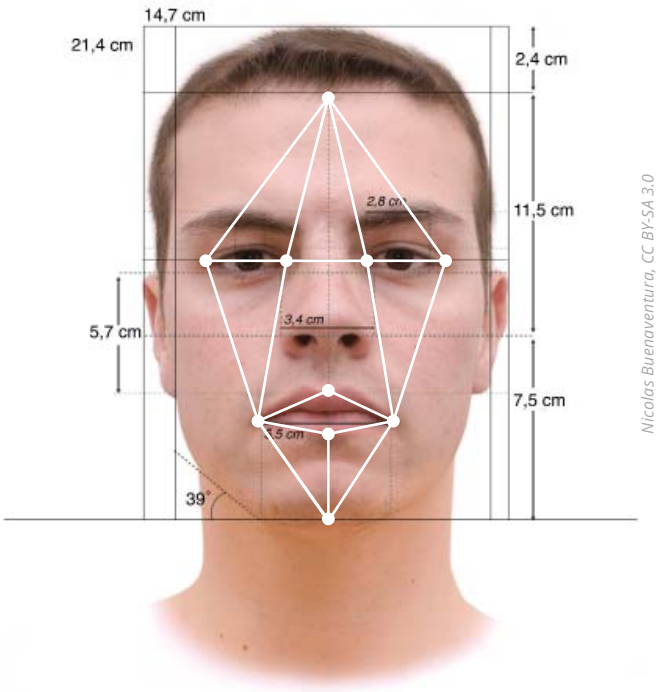

Nicolas Buenaventura, CC BY-SA 3.0

Machine learning is used in biometric facial recognition software, which analyses nodal points on the face.

(the positive examples) and a set of examples without faces (the negative examples). Technically, the algorithm tries to optimise, or find the best set of mathematical rules, to be able to divide the positive and negative examples as well as the human did.

This same basic methodology can be used anywhere that a human can provide samples of what the machine is required to learn. This approach has been used for everything from handwriting recognition, to detecting fraudulent bank transactions, performing medical diagnoses from x-rays, and recognising and translating speech.

Ukufunda kwemishini kusewolunye uphiko lwe artificial intelligence olumayelana nokuthi imishini izifundele yona ngokwayo iphethini yedatha ngaphandle kokuthi kungenelele noma kuhlelwe ngabantu.

Translation by Zamantimande Kunene

Machine learning is at the core of **artificial intelligence**: the science of building intelligent machines. This sits at the intersection of computer science, mathematics and statistics, but also draws on many other fields, such as psychology and neuroscience. Many of the exciting new ideas in machine learning come from knowledge of how the brain works, or how young babies and children learn about their environments.

Another area of machine learning that has seen considerable advances in recent years is **reinforcement learning**. This involves having a computer learn to make sequences of decisions towards some desired goal. For example, a chess-playing algorithm would need to learn to take a whole sequence of actions to win a game, where each move it makes may not by itself be good or bad, but sets the game up for a later win. This kind of algorithm has recently been used to learn to beat expert human players in games as diverse as the board game Go and the computer game StarCraft II. These techniques are also used for robots to learn new skills and for autonomous cars to learn to drive.

With the current frenzy of research in this area, machine learning is touching every sphere of society, from medicine to education, with significant implications for ethics and society. Global advancements here will affect everything from the law to our relationship with work. It is likely that, in the coming years, every profession will have artificial intelligence and machine learning playing a greater role in how it operates, and every person will be influenced by it to a much greater extent, in virtually every aspect of life.

- **Artificial intelligence:** The creation of computers and computer software capable of intelligent behaviour.
- **Machine learning:** The subfield of artificial intelligence involving the scientific study of computational and statistical models and algorithms that can learn from and make predictions on data. Machine learning is often divided into the three areas of supervised, unsupervised and reinforcement learning.
- **Deep learning:** A family of machine learning methods based on artificial neural networks, a class of model very loosely inspired by information processing in the brain.
- **Supervised learning:** A type of machine learning task that involves training on a set of known inputs (data) and outputs (targets) in order to make predictions from new inputs.
- **Unsupervised learning:** An area of machine learning concerned with finding structure and patterns in data, where no targets are specified.
- **Reinforcement learning:** The area of machine learning that studies how an agent, or decision maker, can choose a sequence of actions in some environment so as to maximise a long-term payoff.

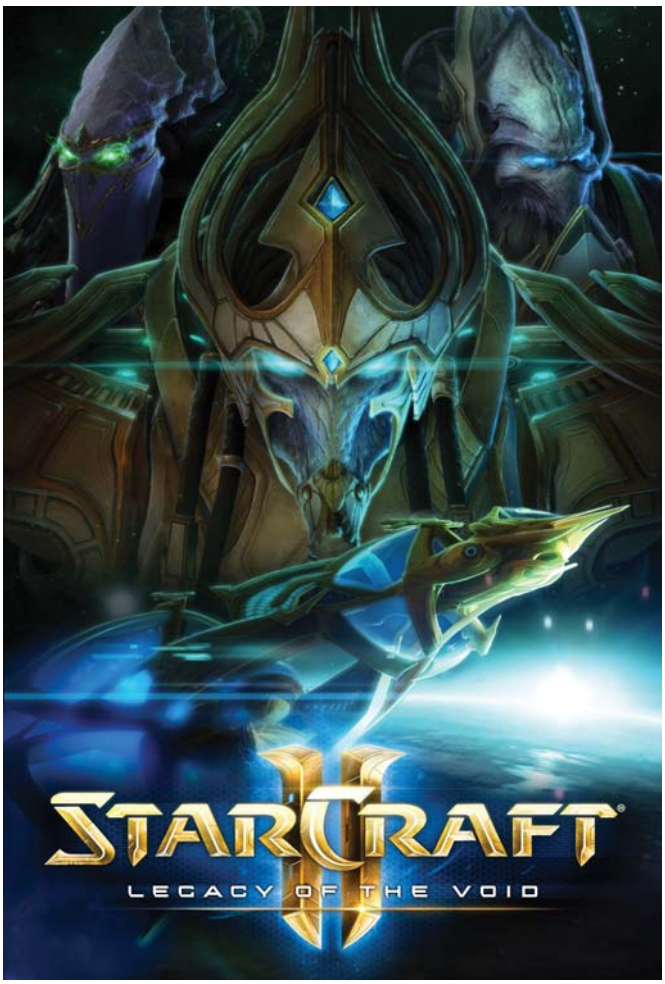

Google's DeepMind artificial intelligence (AI) program, called AphaStar, used machine learning to beat some of the world's top Starcraft II players in January.

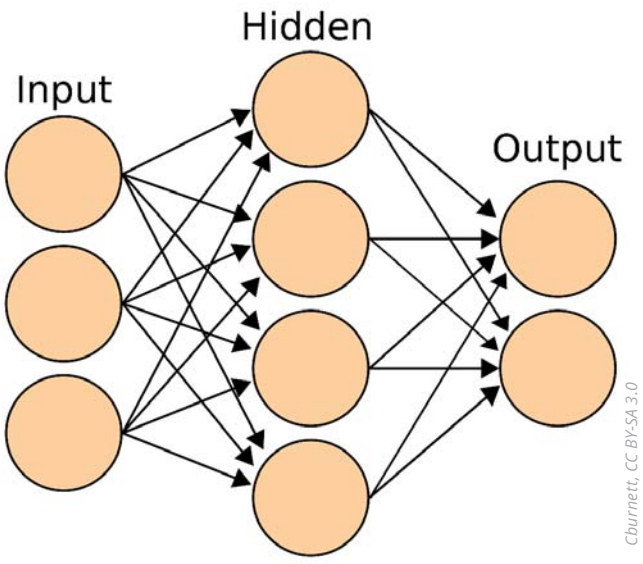

Artificial neural networks are inspired by the neural network of the brain, where nerve calls called neurons transmit information for processing.

Dr Benjamin Rosman is a principal researcher in the Mobile Intelligent Autonomous Systems group at the CSIR, and is also a senior lecturer in the School of Computer Science and Applied Mathematics at the University of the Witwatersrand, where he leads the Robotics, Autonomous Intelligence and Learning (RAIL) laboratory.