The robot that plays Go

Machines have been taught to learn and trump their trainers

S ANANTHANARAYANAN

Artificial intelligence is the field where computing power gets machines to do things that lie in the domain of thinking. Simple applications are where machines process data to make evaluations or predictions. And then, computers operate machinery, or play simple games. It becomes more complex when computers develop strategy or modify their way of working on the bases of how good their predictions or the results of their decisions have been.

Card games like bridge or poker, where opponents use deception, introduces another level of complexity. But the pinnacle was programming a computer to match grand masters in playing chess. When this bastion fell, the challenge shifted to the game of Go — an ancient Chinese board game that has even greater complexity. This too, was mastered in 2016, when the British Artificial Intelligence company, DeepMind Technologies, created the programme, AlphaGo, which defeated Lee Sedol, the 18-times world champion. David Silver, Julian Schrittwieser, Karen Simonyan, Loannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel and Demis Hassabis, all from DeepMind, now report in the journal Nature, an improvement, named AlphaGo Zero, which achieves "a long-standing goal of AI — an algorithm that self-trains to superhuman proficiency without expensive human-input".

A neural network

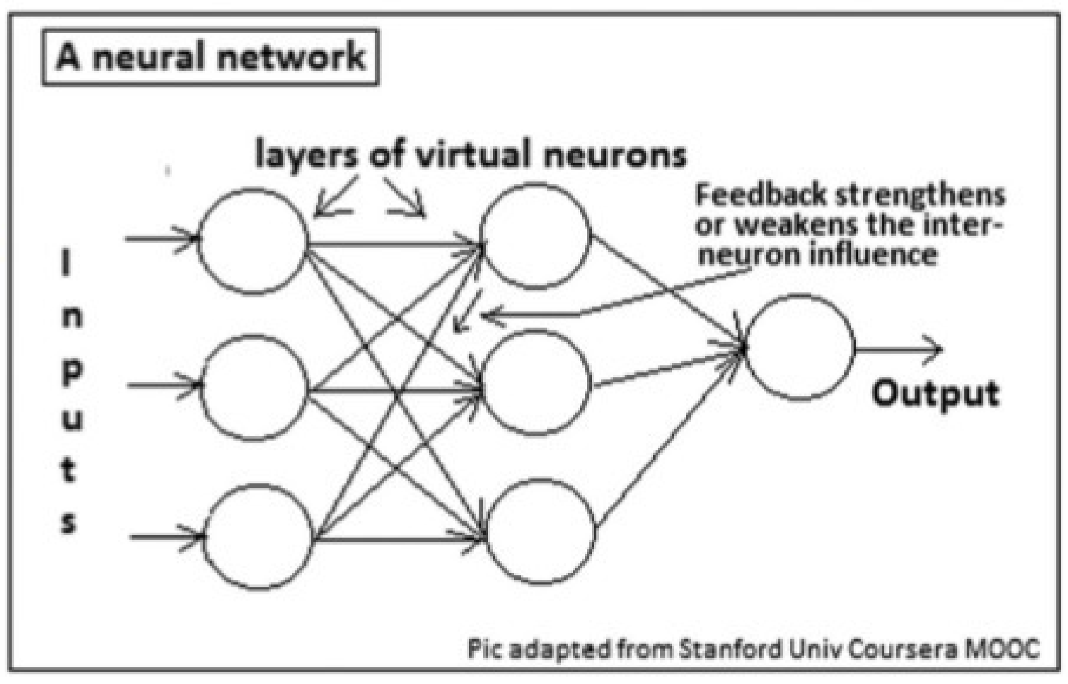

Pic adapted from Stanford Univ Coursera MOOC

This ability, of the algorithm to train itself, is of great importance, as machine learning is now used in a variety of commercial, industrial and scientific fields, all of which face high costs in training the AI machine. Applications include analyses of customer trends, improving supply of products and services, traffic control, health administration, automated diagnosis and drug discovery, voice and image recognition, running driverless motor vehicles in a busy street.

The basic method used is to fit complex but known data to a mathematical formula. The formula is then tested on other sample data and if it qualifies, it could predict unknown data too. While predictions were improved with massive data sets, it was soon realised that the animal brain does even better through a different angle of attack. In playing a game of chess, for instance, a computer could work out all possible moves and counter-moves, by both players, from a given board position, and then choose the move to make.

Human players lack such ability but they still defeat even powerful computers. It was realised that the human brain does not follow the brute force method of the computer, but takes in some features of the chess board position, which may seem to be unrelated, and use these, and experience, to play in a more effective way. The brain charts the different responses of brain cells to the features of the chess board, and the outcomes. Given a set of responses, the responses are strengthened or weakened according to the outcome. Over a series of actual instances, the brain adapts to making more effective responses and continues to learn with experience. This mechanism leads not only to good chess playing but is also the way a child internalises the nuances of a language faster than years of study by scholars.

Computers were now programmed to simulate this architecture by creating virtual neurons, or software that behaved like brain cells. In a simple instance of recognising a just one feature, the feature could be presented to a single virtual neuron. The neuron responds at random, from a set of choices. If the answer is correct, there is feedback that adds to the probability of that response, and if the answer is wrong, the feedback lowers the probability. We can see that this device would soon learn, through a random process, to consistently make the correct response. A brace of artificial neurons that send responses to another set of neurons, and so on, could deal with several inputs with greater complexity. A network like this could learn to identify an image as being that of a car or pedestrian, for instance, and if a pedestrian, whether a man or woman.

In learning to play games, AI systems are designed to take in different features like immediate threats, potential threats, a metric to evaluate positional advantage, devices to evaluate patterns in opponents' moves, relative values of costs to be paid for gains and so on. The system then plans a series of responses to possible future moves of the opponent and could evaluate the options against a data base of past games. Alternately, the system could shadow real games of expert players and evaluate proposed moves by comparison with moves selected by the experts. One strategy to filter moves is to evaluate a move till it is found to be inferior to a previously evaluated one. The evaluation itself may be by traversing the tree that is formed when each move can have different responses, each of which can have several responses, and so on. Or there could be ways of assigning approximate values to moves. A combination of both is the Monte Carlo Tree Search, which employs statistical methods.

AlphaGo, which beat Lee Sedol in four out of five games, employed a pair of neural networks — one to predict the probability of opponent moves and another to evaluate the potential of a board position. The first network was trained with the help of human experts and the second by simulating the game after the given position. These two networks gave rise to a branching inverted tree of moves and countermoves, which helped select the way forward.

AlphaGo Zero, the current version uses a single neural network and a simpler tree search to evaluate moves and positions. The greatest improvement, however, is that the algorithm needs no human input for training. The game, which it plays with itself, starts with simple, untrained, random moves and the positions created are "evaluated". Evaluation results in reinforcing some choices in the random selection and the game is played again. "In each iteration, the performance of the system improves by a small amount, and the quality of the self-play games increases, leading to more and more accurate neural networks and ever stronger versions of AlphaGo Zero," two of the authors say in a press release. The algorithm hence starts to form "tabula rasa" or a 'blank slate", as philosopher John Locke described the mind of an infant, and learns from itself to rapidly become the strongest Go player (and hence teacher) in the world. The method "is no longer constrained by the limits of human," the press release says.

AlphaGo Zero was hence able to beat AlphaGo 100-0 in a trial, with just three days of training on a single machine with four specialised neural network chips against multiple machines and 48 specialised chips. But the real gain is not a better Go player, it is an artificial intelligence strategy that would help machines handle complexities of an interconnected world of rising mechanisation and under pressure to optimise energy and resources.

The writer can be contacted at response@simplescience.in

Game of Go

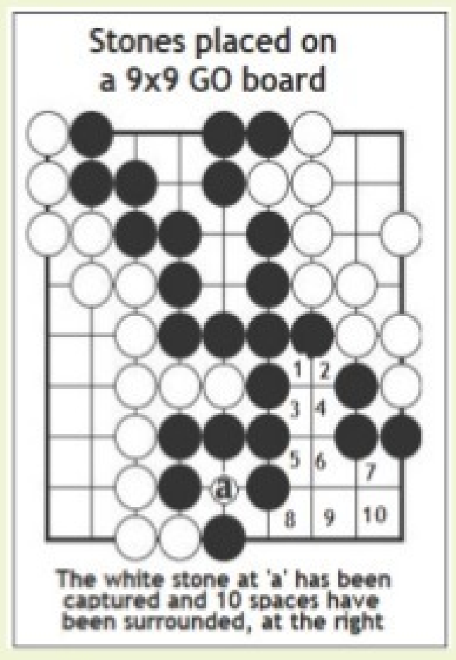

Black and white tokens, called stones, are placed, one in every turn, on the intersections of 9x9, 13x13 or 19x19 lines drawn on a board. The objective is to surround the stones of the opposite colour, which counts as "capture", or to create closed loops of stones, when the number of vacant areas enclosed adds to the player's score. The total of captures and spaces enclosed is the final score.

Players use strategies to maximise the score. For example, a stone may be at risk of capture unless the player occupies the last free intersection. But the player may choose to "sacrifice" the stone for the advantage of progressing a loop in another part of the board. The opponent would then have to choose between the stone offered for capture and preventing the player from stealing a march elsewhere!

The contest rapidly becomes devilishly complex. IBM's Deep Blue succeeded in beating chess world champion, Garry Kasparov in 1996. But it took 20 years before AlphaGo was developed to beat Lee Sedol!

With myriad functions

Here's a look at the different types of membrane proteins

TAPAN KUMAR MAITRA

Membrane proteins have not yielded as well as other proteins to biochemical techniques, mainly because of the problems involved in isolating and purifying hydrophobic proteins in physiologically active form. Procedures such as SDS-polyacrylamide electrophoresis and hydropathy analysis have certainly been useful, as have labeLling techniques involving radioisotopes or fluorescent antibodies.

In the last three decades, however, the study of membrane proteins has been revolutionised by the techniques of molecular biology, especially DNA sequencing and recombinant DNA technology. DNA sequencing makes it possible to deduce the amino acid sequence of a protein, a prerequisite for hydropathy analysis to identify trans-membrane segments. Moreover, sequence comparisons between proteins often reveal evolutionary and functional relationships that might not otherwise have been appreciated. DNA pieces can also be used as probes to identify and isolate sequences that encode related proteins. In addition, the DNA sequence for a particular protein can be altered at specific nucleotide positions to determine the effects of selected changes in the sequence on the activity of the mutant protein for which it codes. What functions do membrane proteins perform?

Some of the proteins in membranes are enzymes, which accounts for the localisation of specific functions to specific membranes. Each of the organelles in a eukaryotic cell is in fact characterised by its own distinctive set of membrane-bound enzymes. Closely related to enzymes in their function are electron transport proteins such as the cytochromes and iron-sulphur proteins that are involved in oxidative processes in mitochondria, chloroplasts, and the plasma membranes of prokaryotic cells.

Other membrane proteins function in solute transport across membranes. These include transport proteins, which facilitate the movement of nutrients such as sugars and amino acids across membranes, and channel proteins, which provide hydrophilic passageways through otherwise hydrophobic membranes. Also in this category are transport ATPases, which use the energy of ATP to pump ions across membranes.

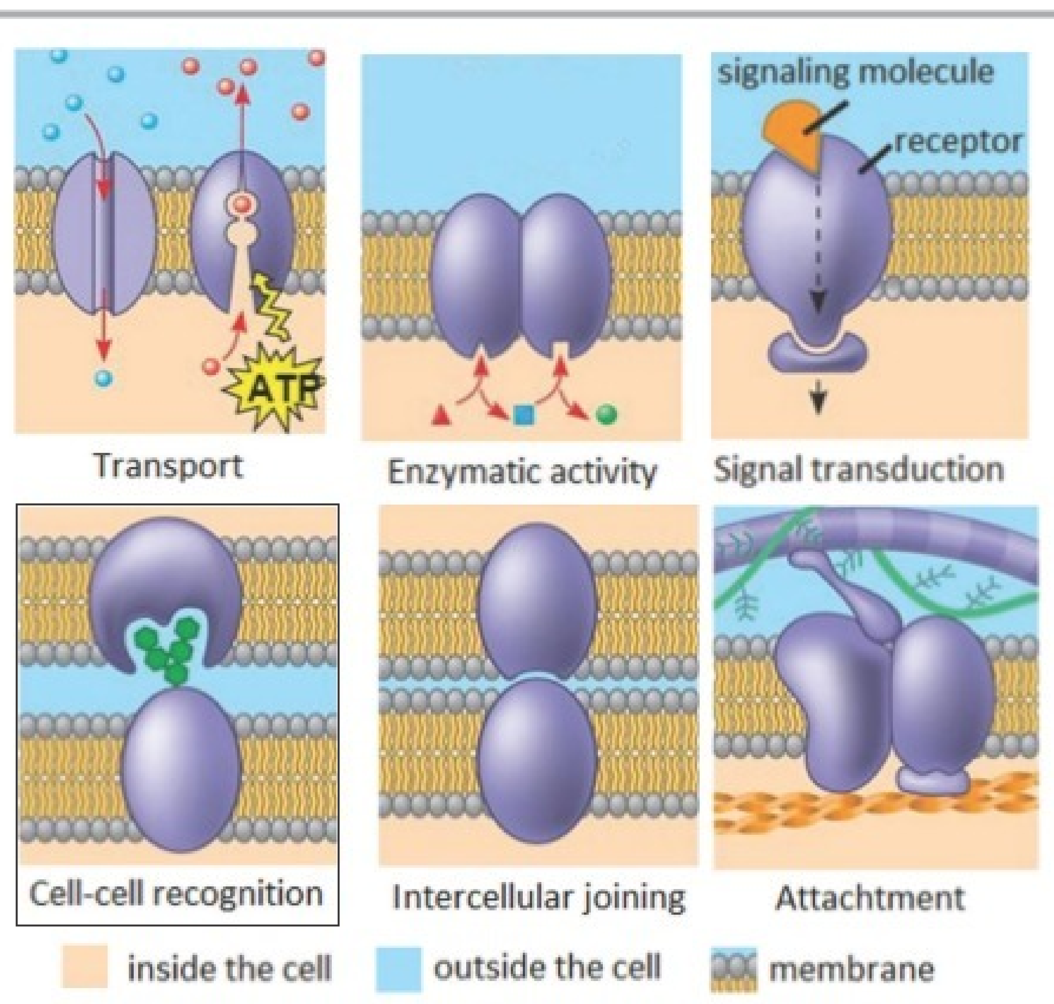

Still other membrane proteins are receptors involved in recognising and mediating the effects of specific chemical signals that impinge on the surface of the cell. Hormones, neurotransmitters, and growth-promoting substances are examples of chemical signals that interact with specific protein receptors on the plasma membrane of target cells.

A final group of membrane-associated proteins are those with structural roles in stabilising and shaping the cell membrane. Examples include spectrin, ankyrin, and band 4.1 protein, the erythrocyte peripheral membrane proteins. A spectrin-based network gives the red blood cell its distinctive biconcave shape and enables the cell to withstand the stress on its membrane as it is forced through narrow capillaries in the circulatory system. Proteins that are structurally homologous to spectrin and spectrin-associated proteins are found just beneath the plasma membrane in many other cell types also, indicating that a cytoskeletal meshwork of peripheral membrane proteins underlies the plasma membrane of many different kinds of cells.

The function of a membrane protein is usually reflected in the way in which the protein is associated with the lipid bilayer. For example, a protein that functions on only one side of a membrane is likely to be a peripheral protein or a lipid-anchored protein. Membrane-bound enzymes that catalyse reactions on only one side of a membrane, such as the ER enzyme glucose phosphatase, are in this category.

In contrast, the tasks of transporting solutes or transmitting signals across a membrane clearly require trans-membrane proteins. Signal transduction is mediated by transmembrane receptors that bind signalling molecules such as hormones on the outside of the plasma membrane and generate signals inside the cell.

The writer is associate professor, head, department of botany, Ananda Mohan College, kolkata, and also fellow, Botanical Society of Bengal, and can be contacted at *tapanmaitra59@yahoo.co.in*

PLUS POINTS

Cool refuges

Tropical rainforests continue to buffer wildlife from extreme temperatures even after logging, a new study has revealed.

Scientists had previously assumed that cutting down trees caused major changes to local climates within tropical forests — something, which would have a devastating effect on the animals living there.

However, new research conducted by the Universities of Sheffield, York and Universiti Malaysia Sabah, shows logged forests on the island of Borneo were thermally indistinguishable from the nearby pristine forest. This is good news for the huge diversity of globally important species that live in logged forests, which may have previously been further destroyed or converted into agricultural land.

The international team of scientists examined the impact that commercial selective logging had on local temperature nine to 12 years after the trees had been chopped down. This type of logging removes a large amount of timber and can be extremely disruptive to rainforest habitat, as is particularly evident in Southeast Asia.

Rebecca Senior, a PhD student from the University of Sheffield's department of animal and plant sciences, led the ground breaking study, which was published last week in the journal, *Global Change Biology*. "Logging activity affects 20 per cent of the world's tropical rainforests and in many places only logged rainforest remains, so it is extremely positive news that even after trees have been logged the forest can continue to support many species of conservation," Senior said.

During the study the researchers travelled to Borneo to establish whether intensive logging activity had altered the availability of cool refuges for animals like the Bornean horned frog, Bornean keeled pit viper, and Wallace's flying frog. "We recorded temperature using a thermal camera and tiny temperature loggers," said Senior, "After nine to 12 years of recovery after logging, the logged forest had a very different structure with fewer large trees and more young saplings. Surprisingly, though, we discovered the average temperature and the availability of cool refuges was comparable with a pristine forest that had never been logged."

Lunar pit stop

A potential shelter has been found on the Moon, which could be a temporary home for future astronauts, says the Japanese space agency. Scientists first spotted it using an orbiter, which saw only a 50-metre hole. But the agency said further exploration using radio waves found a lava tube that stretched for miles.

The potential shelter is thought to have been formed by lava flowing on the Moon as much as 3.5 billion years ago. It sits around a series of volcanic dimes known as the Marius Hills.

The cave, stretching 30 miles over the lunar surface, could be used as protection while astronauts establish a more permanent base. The astronauts inside the cave would also be protected from asteroid impacts — which are more dangerous for the lack of atmosphere — and there would be a more consistent temperature inside the hole.

As well as allowing for further exploration of our closest neighbour, it could also serve as a stop off for further exploration of our solar system. The Kaguya satellite that found the cave has been offline since 2009, when it crashed into the Moon's surface. But scientists continue to find surprises in the data it sent back.

The Independent