# When information is imperfect

A NEW ALGORITHM COULD PUSH THE BOUNDARIES OF WHAT COMPUTERS CAN DO, WRITES **S ANANTHANARAYAN**

In *The Honeymoon Machine*, a 1961 Hollywood farce, the computer aboard a US naval ship off the coast of Venice is used to predict outcomes at roulette. The National Bank was heading for bankruptcy when the role of the ship was detected and the US Government stepped in to avert a diplomatic crisis.

From an artificial intelligence viewpoint, predicting results at roulette is still outside the computers' capacity. With progress in the mathematics of how strategies are planned and methods of data analyses, however, computers are now doing better in other areas, which were thought to be outside their reach. Matej Moravcík, Martin Schmid, Neil Burch, Viliam Lisý, Dustin Morrill, Nolan Bard, Trevor Davis, Kevin Waugh, Michael Johanson and Michael Bowling, from the University of Alberta at Edmonton and the Charles University and the Czech Technical University at Prague report in the journal, Science, that yet another conceptual milestone has been passed, on the route of computers doing the things we thought only people can.

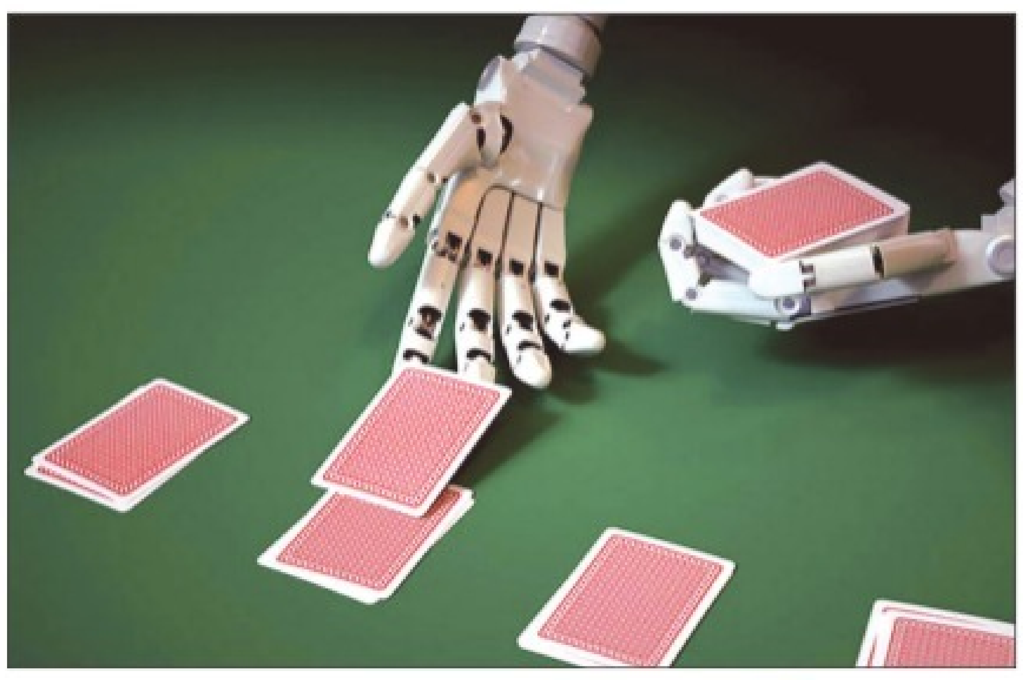

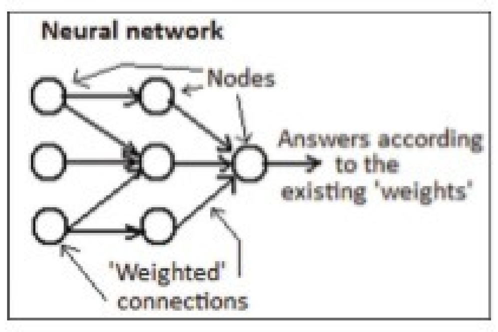



*The Honeymoon Machine*

One area where computers have proved themselves is in playing games of skill against human adversaries. A typical example is in chess, where each player has a vast choice of permissible moves, but only a few can lead to victory or avoid defeat. A method that many beginning human players try to use is to imagine the possible moves and counter-moves and then the answers, and the next lot of counter-moves, and so on. This, however, does not lead the player very far and, with practice, she switches to generic strategies. The computer, in contrast, has the capacity to analyse a huge number of possible exchanges. This ability alone, however, could not help computers, which still lost out to the experience of a good human player.

Computers then used a different tack, more like what a human player does -- to look for patterns in board positions and select strategies from experience, making use of libraries and learning and getting better as the game progressed. Given the parameters of a situation, there is mathematical theory that can work out the strategy that is most likely to succeed, whatever the moves of the opponent, especially when the opponent uses the best possible strategy. Given the ability to analyse trends of data and to evaluate strategies, computers have got very good at chess, and a fine example is the hard time that IBM's Deep Blue computer gave grand master Gary Kasporov. Computers have similarly done very well in different kinds of board games, like Checkers, Backgammon and recently, the ancient Chinese game of GO.

What these games have in common is that, with different degrees of complexity, all the data for generating moves that strengthen players' positions is there for both or all players to see. It is in these circumstances, of "perfect information", that the algorithms, which help computers play these games, have worked. Another kind of game, which may be more akin to real life, is where all the data is not available to all players -- a typical case is card games, where some of the cards known to one player are not known to others. How to programme computers to play this category of "imperfect information" games has been a challenge for over 60 years, say the authors of the paper in the journal. The authors then describe "DeepStack", a computer procedure that can deal with this category of games, specifically a variation of poker, called "Heads up, no limit Texas hold'em".

One method that computers use is to classify trends in data, which may represent images, for recognising patterns, or abstractions like the value of chessboard positions. This method seeks to fit the data into a formula, called a "hypothesis" and then to test the formula with untried situations. For example, a value, like the price of a house, may depend on a combination of factors like the area in square feet, the number of bathrooms, the number of floors, the area of the garden, the distance to the supermarket, the Metro station and so on. The computer tries out different weightages, or multiplying factors that attribute importance to the factors, to create an algebraic expression, based on the value of the factors, that works out the value of the house.

The method that is used to go from one set of multiplying factors to a better set is to change the factors in such a way that it minimises the sum of the differences between the values given by the hypothesis and the actual values. This difference is known as the cost function, and the values used are called the training values. Using these methods, which have grown in complexity and sophistication, a powerful computer can take in the data of the pixels of the screen of a digital camera and make out if the image is of house, a car, a dog or man or a woman. While it is routine for these programmes to compare fingerprints, they can now even recognise faces.

These programmes are now run the way the human brain works, which is, not by maintaining an exhaustive library with which to compare new data, but by strengthening the paths that lead to specific responses. In this way, networks of nerve cells in the brain go through a process of learning every time differently strengthened, random responses lead to correct or incorrect results. This is the way a human baby, for instance, learns the use of language long before she hears about the rules of grammar. Computers now mimic this process by creating layers of software "nodes" that receive inputs from the nodes in an earlier layer and then pass signals to the nodes in a subsequent layer. Whether the answer that the final layer outputs is correct or not then controls how multipliers along the path are increased or decreased, till paths that lead to correct answers become more likely.

As we just said, these methods have been successful in automating complex, adversarial problems, exemplified by games of strategy, but of the kind where the information available is there for all players to use. The other kind –"imperfect information" games -- has additional features of each player not only using the extra information she has but also of trying to mislead other players into forming incorrect impressions of the assets the first player holds. And the game of poker, where players make bets on cards that only they know, against cards that only the other players know, consists of "bluffing, of little tactics of deception, of asking yourself what the other man is going to think…" in the words of Von Neuman, pioneer of strategy theory and computing, that the journal paper quotes.

In games like poker, the deception is practiced by both sides and "how our opponents' actions reveal their (private) information depends upon their knowledge of our private information and how our actions reveal it," the authors say in the paper. Solving imperfect information games then involves this kind of "recursive" reasoning, where one thing depends on another, which itself depends on the first thing, and so on. The probability of a line of play being chosen by a player depends on the history of the game, the private information of the player, which could be revealed by the game history, and again the player's suspicions about others' private information revealed by their own past play history. The factor to be minimised, for improving strategy, then moves from the cost function, to a value called the counterfactual regret, which measures the difference between the utility of a method of play chosen and the value that was possible.

The DeepStack algorithm was then tried out in actual practice, playing 3,000 games each against eleven professional poker players. The result was that DeepStack could beat ten of these eleven pros "by a statistically significant margin". While this is an achievement, the authors emphasise that the DeepStack algorithm represents a paradigm shift in the approach to large, sequential, imperfect information games.

"With many real-world problems involving information asymmetry, DeepStack also has implications for seeing powerful Artificial Intelligence applied more in settings that do not fit the perfect information assumption. The abstraction paradigm for handling imperfect information has shown promise in applications like defending strategic resources and robust decision making as needed for medical treatment recommendations. DeepStack's continual resolving paradigm will hopefully open up many more possibilities," the authors say.

THE WRITER CAN BE CONTACTED AT response@simplescience.in

# CONTROL ELEMENTS

**TAPAN KUMAR MAITRA** EXPLAINS THE PROCESS OF ATTENUATION IN DNA

All the regulatory mechanisms so far control the initiation of transcription. Prokaryotes also employ some regulatory mechanisms that operate after the initiation step. A classic example was discovered when Charles Yanofsky and his colleagues found that the trp operon of E coli has a novel type of regulatory site located between the promoter/operator and the operon's first gene, trpE. This stretch of DNA, called the leader sequence (or I), is transcribed to produce a leader mRNA segment, 162 nucleotides long, located at the end of the polycistronic *trp* mRNA.

Analysis of trp operson transcripts made under various conditions revealed that, as expected, the full-length, polycistronic trp mRNA is transcribed when tryptophan is scarce. This allows the enzymes of the tryptophan bio-synthetic pathway to be synthesised, and hence the pathway can produce more tryptophan. On the other hand, when tryptophan is plentiful, the genes coding for the enzymes of the tryptophan pathway are not transcribed, also as expected. An unexpected result, however, was that the DNA corresponding to most of the leader sequence is transcribed under such conditions. Based on these findings, Yanofsky suggested that the leader sequence contains a control region that is sensitive to tryptophan levels. This control sequence somehow determines not whether trp operon transcription can begin, but whether it will continue to completion. The effect of this control element was called attenuation because of its role in attenuating, or reducing, the synthesis of mRNA.

How the attenuation mechanism works, it must start with a closer look at the trp operon leader mRNA segment. This leader has two unusual features that enable it to play a regulatory role. First, in contrast to the non-translated leader sequences typically encountered at the end of mRNA molecules, a portion of the trp leader sequence is translated, forming a leader peptide 14 amino acids long. Within the mRNA sequence coding for this peptide are two adjacent codons for the amino acid tryptophan; these will prove important. Second, the trp leader mRNA also contains four segments (labelled regions 1, 2, 3, and 4) whose nucleotides can base-pair with each other to form several distinctive hairpin loop structures. The region comprising regions 3 and 4 plus an adjacent string of eight U nucleotides is called the terminator. When base pairing between regions 3 and 4 creates a hairpin loop, it acts as a transcription termination signal.

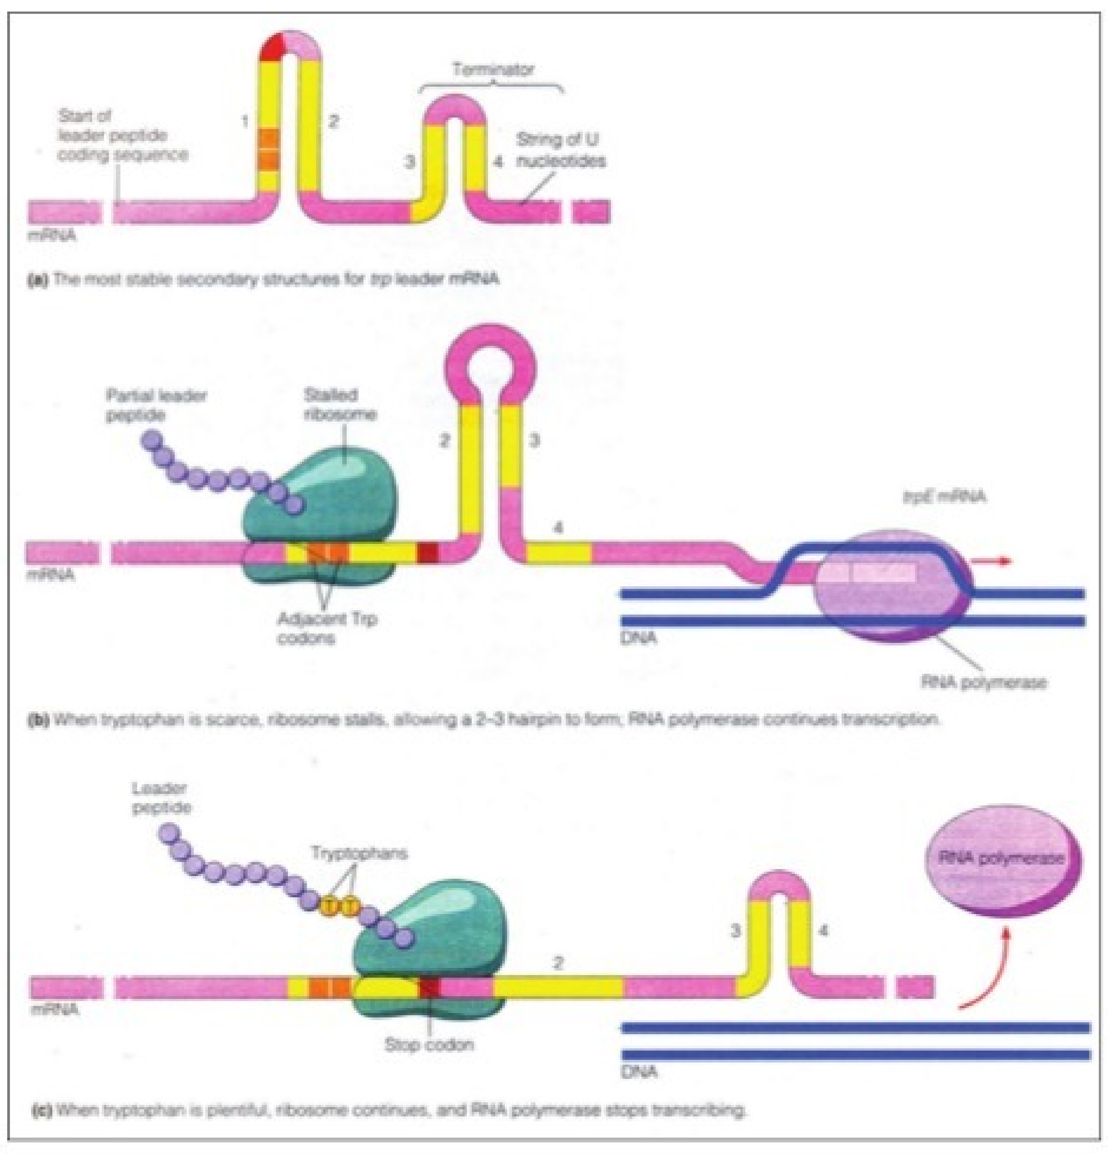


Attenuation depends upon the ability of regions 1 and 2 and regions 3 and 4 of the trp leader sequence to base-pair, forming hairpin secondary structures. The 3-4 hairpin structure acts as a transcription termination signal; as soon as it forms, the RNA and the RNA polymerase are released from the DNA. (b) During periods of tryptophan scarcity, a ribosome translating the coding sequence for the leader peptide may stall when it encounters the two tryptophan (Trp) codons because of the shortage of tryptophan-carrying tRNA molecules. Because a stalled ribosome at this site blocks region 1, a 1-2 hairpin cannot form, and an alternative, 2-3 hairpin is created. The 2-3 base pairing prevents formation of the 3-4 transcription termination hairpin, and therefore RNA polymerase can move on to transcribe the entire operon. (c) When tryptophan is readily available, a ribosome can complete translation of the leader peptide without stalling. As it pauses at the stop codon, it blocks region 2, preventing it from base pairing. As a result, the 3-4 structure forms and terminates transcription near the end of the leader sequence.

As Yanofsky's experiments suggested, translation of the leader RNA plays a crucial role in the attenuation mechanism. A ribosome attaches to its first binding site on the *trp* mRNA as soon as the site appears, and from there it follows close behind the RNA polymerase. When tryptophan levels are low, the concentration of tryptophanyl tRNA (tRNA molecules carrying tryptophan) is also low. Thus, when the ribosome arrives at the tryptophan codons of the leader RNA, it stalls briefly, awaiting the arrival of tryptophanyl tRNA. The stalled ribosome blocks region 1, allowing an alternative hairpin structure to form by pairing regions 2 and 3. When region 3 is tied up in this way, it cannot pair with region 4 to create a termination structure, and so the RNA polymerase continues, eventually producing a complete mRNA transcript of the trp operon. Ribosomes use this mRNA to synthesise the tryptophan pathway enzymes, and production of tryptophan therefore increases.

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

# *Star Wars*, a thousand light years away

A TEAM OF RESEARCHERS IN THE UK HAVE DISCOVERED A PLANETARY SYSTEM, WHICH HAS DEBRIS ORBITING A DOUBLE STAR

Evidence of planetary debris surrounding a double sun in a *Star Wars*-like system has been found by a team of researchers from the UK.

The study, published recently in the journal *Nature Astronomy*, has found the remains of shattered asteroids orbiting a double sun consisting of a white dwarf star and a brown dwarf star roughly 1,000 light years away in a system called SDSS 1557. The debris appears to be rocky suggesting that terrestrial planets, such as Tatooine, the home planet of Luke Skywalker in the film series *Star Wars*, might exist in the system. To date, all exoplanets discovered in orbit around double stars are gas giants, similar to Jupiter, and are thought to form in the icy regions of their systems.

In contrast to the carbon-rich icy material found in other double star systems, the planetary material identified in the SDSS 1557 system has a high metal content, including silicon and magnesium. These elements were identified as the debris flowed from its orbit onto the surface of the star, polluting it temporarily with at least 1.1 trillion tons of matter, equating it to an asteroid at least four km in size.

The discovery came as a complete surprise, as the team assumed the dusty white dwarf was a single star but co-author Steven Parsons from the University of Sheffield, who is an expert in double star (or binary) systems, noticed the tell-tale signs. "We know of thousands of binaries similar to SDSS 1557 but this is the first time we've seen asteroid debris and pollution.

"The brown dwarf was effectively hidden by the dust until we looked with the right instrument but when we observed SDSS 1557 in detail, we recognised the brown dwarf's subtle gravitational pull on the white dwarf."

Lead author of the study, Jay Farihi from University College London's department of physics and astronomy, said, "Building rocky planets around two suns is a challenge because the gravity of both stars can push and pull tremendously, preventing bits of rock and dust from sticking together and growing into full-fledged planets.

"With the discovery of asteroid debris in the SDSS 1557 system, we see clear signatures of rocky planet assembly via large asteroids

that formed, helping us understand how rocky exoplanets are made in double star systems."

In the solar system, the asteroid belt contains the leftover building blocks for the terrestrial planets Mercury, Venus, Earth, and Mars, so planetary scientists study the asteroids to gain a better understanding of how rocky and potentially habitable, planets are formed. The same approach was used by the team to study the SDSS 1557 system as any planets within it cannot yet be detected directly but the debris is spread in a large belt around the double stars, which is a much larger target for analysis.

The team studied the binary system and the chemical composition of the debris by measuring the absorption of different wavelengths of light or "spectra", using the Gemini Observatory South telescope and the European Southern Observatory Very Large Telescope, both located in Chile.

Co-author Professor Boris Gänsicke from the University of Warwick analysed these data and found they all told a consistent and compelling story. "Any metals we see in the white dwarf will disappear within a few weeks, and sink down into the interior, unless the debris is continuously flowing onto the star.

"We'll be looking at SDSS 1557 next with Hubble, to conclusively show the dust is made of rock rather than ice."

## PLUS POINTS

### Living on Mars

If living on Mars is of the realm of sci-fi fantasy, one Nasa scientist has an idea, which could make it possible in a matter of years.

Speaking at the Planetary Science Vision 2050 Workshop at the Nasa headquarters in Washington, one scientist presented the extraordinary idea to put a magnetic shield around Mars to restore its atmosphere, which eventually could make it habitable.

In a talk, Nasa's Planetary Science Division director James Green set out how the organisation could be in a position to carry out daily science and exploration on Mars. The workshop was aimed to discuss ambitious space projects that could be implemented or at least started by 2050. Speakers included a range of experts on space technology, which set out their vision for what planetary science may look like in the future.

Green said that launching a "magnetic shield" to a stable orbit between Mars and the sun could shield the planet from high energy solar particles. In the past, Mars had a significant amount of water before the planet lost between 80 and 90 per cent of its atmosphere over its lifetime.

The shield would consist of a large dipole, which is a close electric circuit powerful enough to generate an artificial magnetic field, *Popular Mechanics* reports. The shield would allow Mars to slowly restore its atmosphere.

Green's modelling of the shield found that the structure could enable Mars to build up half the atmospheric pressure of the earth in a matter of years. The shield would protect the planet from solar winds and the greenhouse effect would start to heat the planet and eventually melt the ice under its poles.

"Perhaps one-seventh of the ancient ocean could return to Mars. The solar system is ours, let's take it," he said.

CHLOE FARAND/THE INDEPENDENT

### Image from up above

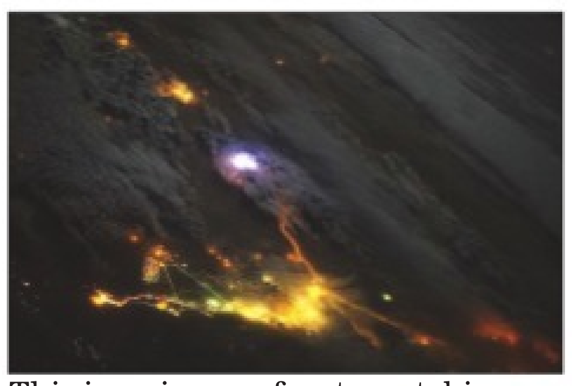

This is an image of a storm taking place over West Asia captured by an astronaut aboard the International Space Station. The station recently received two new instruments — the Stratospheric Aerosol and Gas Experiment III and the Lightning Imaging Sensor.

Sage III monitors the condition of the ozone layer, which covers an area in the stratosphere 16km to 48km above Earth and protects the planet from the Sun's harmful ultraviolet radiation, while the LIS records the time, energy output and location of lightning events around the world, day and night. LIS will improve coverage of lightning events over the oceans and in the Northern Hemisphere during its summer months.

THE STRAITS TIMES/ANN

### Asteroid exploration

China plans to conduct at least one asteroid exploration mission before 2025, a senior space scientist said recently.

Ye Peijian, a leading scientist on deep-space exploration at the China Academy of Space Technology and a national political adviser, told *China Daily* that the asteroid mission will be carried out after the country's first Mars expedition in 2020 and before 2025.

"The detailed schedule and the target asteroid have yet to be determined, but we are working on them. We want to explore asteroids because their resources will be important to mankind's development in future," he said.

The country has also decided to send a probe to Mars to take samples and return to Earth around 2030, Ye said. Additionally, Chinese scientists are conducting preliminary research on a Jupiter mission, he said.

CHINA DAILY/ANN