VIGYAN 2047 Volume 1 No. 2 November 2024 Pages 8–14

This Year's Physics and Chemistry Nobel is AI Enabled

TV Venkateswaran

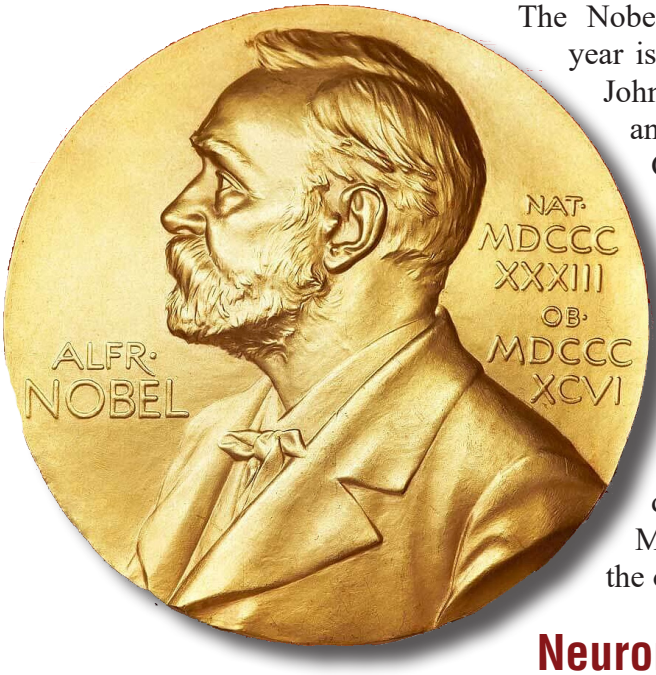

The Nobel Prize in Physics and Chemistry awarded this year is surprisingly related to Artificial Intelligence (AI). John Hopfield at Princeton University in New Jersey and Geoffrey Hinton at the University of Toronto in Canada, who have made significant contributions to this field, have jointly been awarded the Nobel Prize in Physics 2024 "for foundational discoveries and inventions that enable machine learning with artificial neural networks."

David Baker of the University of Washington, Seattle, WA, USA and Howard Hughes Medical Institute, USA, received half of this year's Chemistry Nobel prize for "computational protein design." At the same time, Demis Hassabis and John M. Jumper of Google DeepMind, London, UK, shared the other half for "protein structure prediction".

Neuron-biological and artificial

All our actions are controlled by the brain. Suppose you touch a hot glass and immediately take your hand away. The 'hotness' signals go to the brain and trigger a set of neurons to 'fire'. That is, these neurons send a pulse of signal, called an 'action potential', that travels through the nerves and activates the tissue in your hand; this makes the hand move away from the hot glass.


Geoffrey Hinton


John Hopfield


Gary Ruvkun


Victor Ambros

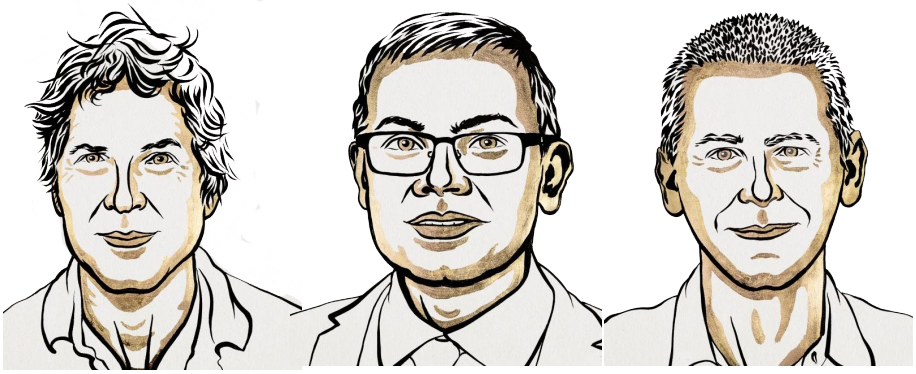

David Baker Demis Hassabis John Jumper

When does the neuron fire? Does it get triggered to all inputs?

Some actions are triggered by a single stimulus. If you inhale a cloud of pollen, you uncontrollably sneeze. But not in the case of itching. You may scratch, but if you are in a social gathering and feel embarrassed, you will hesitate and endure.

There are two inputs: your itch and the awareness of the surroundings. The itch will have an excitatory effect, making you desire to scratch. The ambience will have an inhibiting effect, encouraging you to resist. If you are among your friends, even if it is a social event, you may not hesitate to scratch; nevertheless, if you attend an official occasion, your inhibitions will be stronger. Your action is the result of both excitatory and inhibitory inputs.

In 1943, neurophysiologist Warren McCulloch and logician Walter Pitts postulated that neurons act in this manner in the human brain. Today, this groundbreaking neuron model, the McCulloch-Pitts neuron, depicts neurons as logic gates and neural networks as digital circuitry. According to the framework, the neuron's inputs might be either zero or one. The neuron's output signal is connected to (or a function of) the sum of its inputs. If the input total exceeds or equals a threshold value, the output is one; otherwise, it is zero. Researchers constructed artificial neural networks and artificial intelligence algorithms based on this rudimentary neuron model.

To better understand this, let us elaborate on the itching comparison.

There are four probable scenarios:-

First scenario: It is not itching, and you are in a social event.

Second scenario: It does not itch, but you are in a social event.

Third scenario: It is itching, and you are not in a social gathering.

Fourth scenario: You have itchiness and are at a social function.

To assess scenarios using the McCulloch-Pitts neural model, use the following input signals:

I: Does it itch?

SG: Are you in a social gathering?

If it is itching, the value of X1 is 1; else, 0 (because it is excitatory)

If you are in a social gathering, the value of X2 is –1, else 0 (because it is inhibitory)

The threshold is greater than or equal to 1.

So, the neural network model will look like:

Scenario	I	SG	Sum	Threshold value
1	0	0	0	0
2	0	–1	0	0
3	1	0	1	1
4	1	–1	0	0

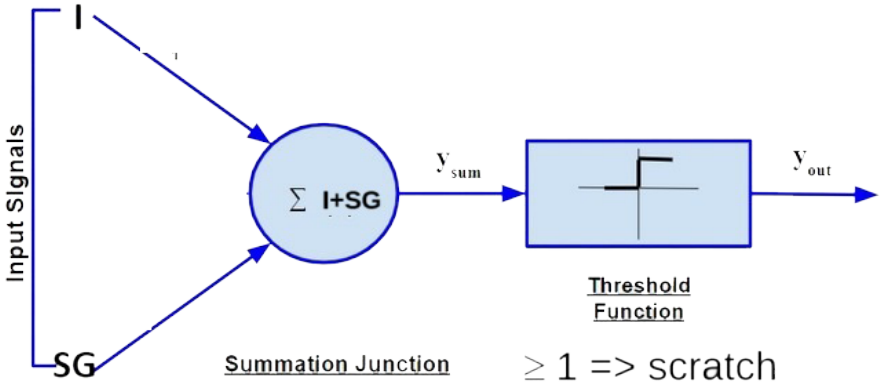

Figure 1: Simple model of the neuron

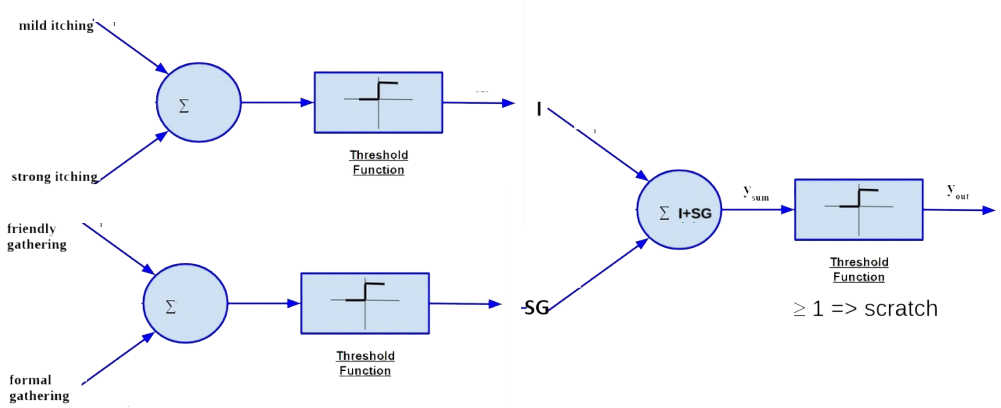

Figure 2: neuronal network

Therefore, we can see that only in case 3 the neuron will 'fire'.

This can be schematically represented as shown in Figure 1.

We can now unite three neurons in decision-making. For example, your itching may be barely noticeable, and you may not feel the need to scratch. This implies that the 'activity' of the neurone that evaluates the amount of itching serves as the input (I) for the neurone that decides whether or not to scratch. Similarly, not all social situations may be humiliating. If you are in the presence of relatives or college acquaintances, you may not feel ashamed and may continue to itch. Thus, the second input (SG) neurone examines the type of social gatherings. This can be represented as shown in Figure 2.

In the preceding example, we considered a basic scenario in which two factors impact a specific summation. In fact, there might be more than two. In addition to the severity of the itch, the setting in which you are situated and the area of your body on which you are scratching will all impact it. Thus, a simplified McCulloch-Pitts neural model may be envisaged as having several inputs, each with a relative weight, and the action is determined by the summing value and the threshold.

Artificial neural networks (ANNs)

Artificial neural networks (ANNs) are intended to function similarly to neurons in the brain. Biological neural networks employ biological neurons, whereas artificial neural networks use digital neurons/perceptrons that are connected in various ways.

In the above example, the neurons are hierarchical, with the top layer sending input to the layer underneath them. However, the output of one neuron may become the feedback input of the layer before it. Recurrent neural networks are a unique type of neural network.

John Hopfield, a physicist, researched the recurrent neural networks using physical principles. In recurrent neural networks, data is sent back and reprocessed several times. He concentrated on the dynamics: what happens to the network over time?

Hopfield network

A Hopfield network is an artificial neural network analogous to how friends recall how you completed a jigsaw puzzle together. You may remember how some parts fit together, but your buddies may recall a joint with other pieces. You can recollect the entire memory together, but none of you can recall the answer. Even if some of you may not recall every detail, you can help each other fill in the gaps by sharing what you remember. The Hopfield network resembles a circle of friends. Each colleague works as a neuron in the network. The memory of how one piece fits into the next is comparable to the patterns the Hopfield network stores. Talking about the answer is like activating the network; everyone brings up what they remember. Even if specific details are missing or confused, this discussion will help you to rebuild your memories of the experiences together.

In conclusion, Hopfield demonstrated how an artificial neural network can be trained to store memories (patterns) so that they may be recovered even when the input is fragmented or noisy. It uses neuron connections (similar to friendship connections) to help retrieve whole memories.

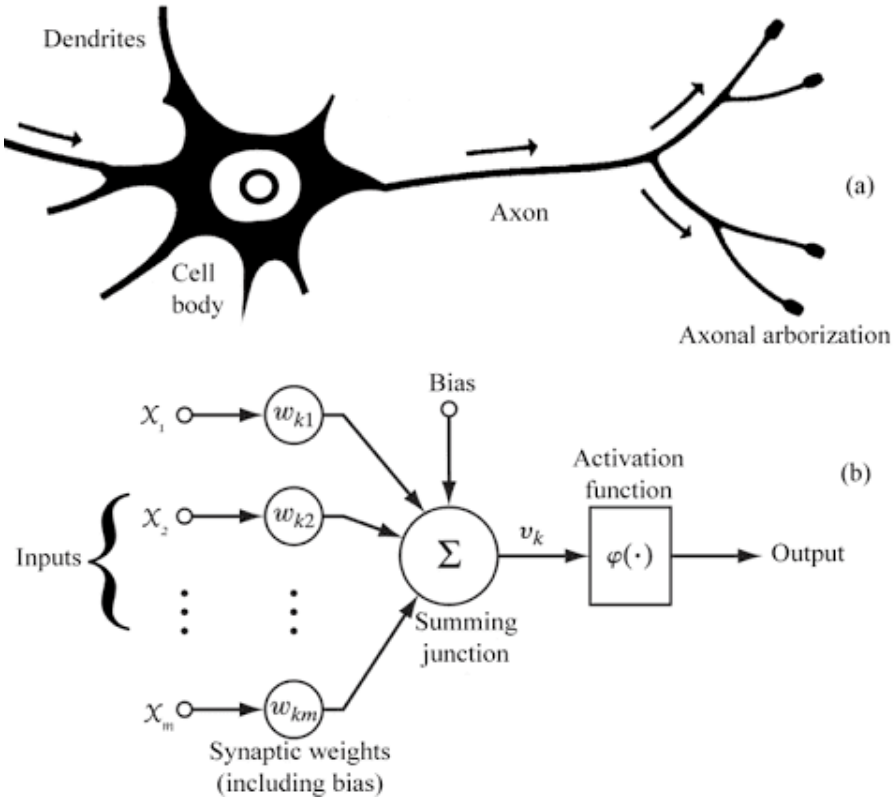

Figure 3: Artificial neurons mimic the function of biological neurons, with input, processing and output.

Backpropagation

In 1986, Geoffrey E. Hinton, David E. Rumelhart, and Ronald J. Williams published the seminal paper 'Learning representations by back-propagating errors'. They popularised this learning approach, allowing them to continually adjust the weights of the network's connections to reduce the difference between the actual and intended outputs.

For example, you scratched yourself in a specific gathering and became really embarrassed. Another time, you resisted, but then you noticed another person in the group scratching an itch. Now you recognise you need to adjust the 'embarrassment weightage' you gave to these gatherings. Learning from experience allows you to change your actions and improve social etiquette.

This is known as backpropagation. You evaluate the preceding steps and make adjustments based on the final outcome. To get artificial neural networks to do fascinating tasks, you must first define the correct weights for the connections between artificial neurons. Backpropagation is required when choosing weights depending on the network's performance on a training dataset.

Training the deep neural network is beneficial. For example, when we type on WhatsApp, words are autocompleted and autocorrected; this is training. It's frustrating at first, but after a while, your Whatsapp is so thoroughly trained that it saves you time typing.

Deep neural network (DNN)

Let us further explore the decision of scratch or not analogy. Perhaps you are at a party where scratching would be really humiliating. However, you are in the restroom. As a result, location in both time and space will be evaluated. The choice can be divided into modules. Where you are now, what social event, whatever portion of your body, the severity of the itch, and so on. These components can function concurrently. Partial shortlisting occurs layer by layer before a final judgement is made at a higher level. If the venue is excessively inhibitive, the module will respond negatively, bringing the computation to a halt. If the itch is intense, the calculation will continue.

The itch event starts the network and concludes with an output layer that recommends whether or not to scratch. This is an illustration of a simplified depiction of deep networks. The massively parallel network of the deep neural network guarantees that tasks are distributed. Each neuron only fires when the specified conditions are satisfied, allowing the group to decide.

The deep neural network was inspired by the human visual cortex's layered architecture, which converts signals from the retina into meaningful perceptions. Hinton and neurologist Terry Sejnowski developed an algorithm that enables deep neural

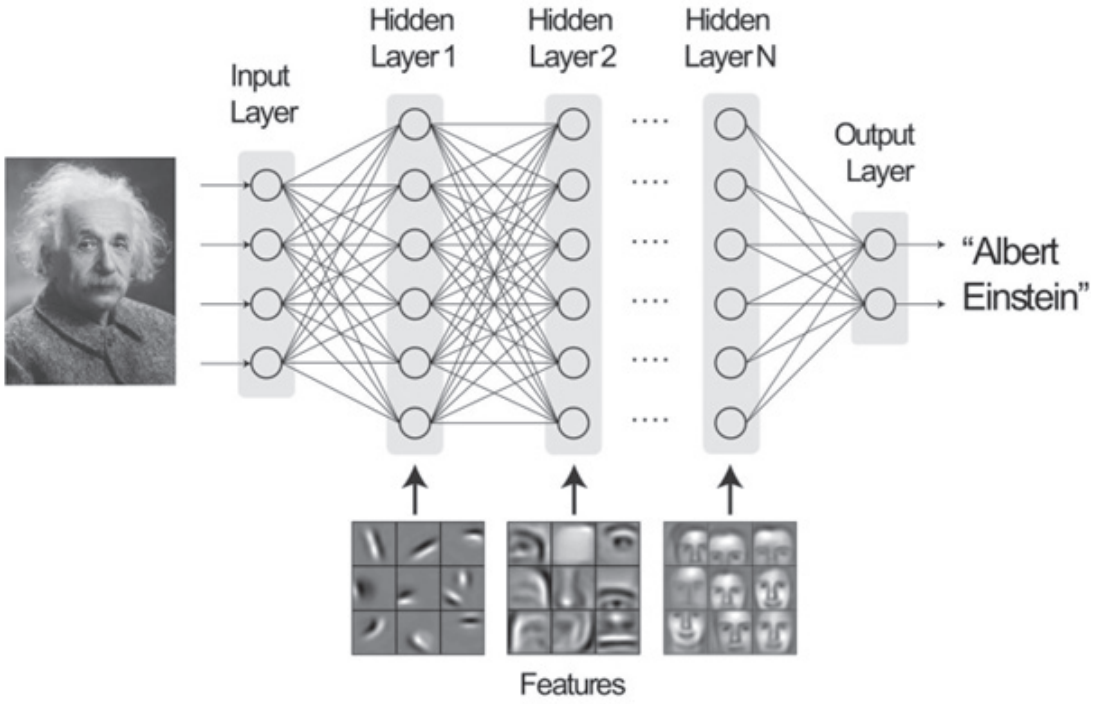

Figure 4: The leftmost layer encodes the input, which is faces. In this example, the rightmost layer produces the result if the photo is of Albert Einstein. Training with millions of labelled trial faces helps to trim and perfect the weights between neurons. Backpropagation alters the connection weights during each trial to obtain the desired output. After enough training, the network matures to the point where each subsequent layer in the neural network learns to recognise and accurately classify progressively complex features (e.g., lips, nose, eyes, and so on) in faces.

networks to form. A deep neural network (DNN) consists of sequence layers of neurons. All network connections were initially given equal strength as if each user had one vote. When two neurons fire simultaneously in response to an input, their connection increases. Conversely, seldom-used connections wilted. In artificial intelligence jargon, the act of strengthening certain connections while weakening others is known as 'training'.

In 2005, Hinton and colleagues created an entirely novel training program based on these findings. Hilton and his colleagues discovered that each layer may be trained individually, similar to how neurons in cat eyes are trained to discriminate between vertical and horizontal lines. The output of the trained layer is sent into the next layer, which is now being trained. They revealed that layer-by-layer training can be significantly more effective.

After being blasted with millions of data points, the initial layer of neurons in facial recognition AI software will learn, for example, how to recognise superficial characteristics like an edge or contour. Once the layer has trained to distinguish these elements correctly, it is passed on to the next layer, which trains itself to detect more complex characteristics like a nose or an ear. Then that layer is passed on to another layer, which trains itself to detect even higher levels of abstraction, and so on, layer by layer—hence the "deep" in deep learning—until the system can reliably recognise highly complex phenomena, such as human faces.

Using AI to learn about protein structure

We know that proteins are chains of amino acids. A garland is created by precisely arranging multiple flowers. At the same time, a chain of amino acids in various patterns produces diverse proteins. Proteins are sometimes called macromolecules to distinguish them from smaller molecules like

glucose and water.

You usually get one piece when you order dosa, but idli comes in two pieces. An idli vada combo typically consists of two idlis and one vada. As is the situation with proteins.

Some proteins are simple, with just one chain of amino acids. One of them is myoglobin, a tiny, brilliant red protein that receives its colour from the iron atom it binds to and is present in muscle cells, giving the meat its red hue. Proteins made up of a single chain of amino acids are called monomers since they act independently. Monomers, like dosas, come in singles.

Human insulin, on the other hand, is made up of two chains that are connected together. Each of these monomers is known as a sub-unit. These proteins are similar to an idli-vada combination in that the end product contains more than one monomer.

When served, dosa has a specific form. A set dosa is distributed in a circular pattern. In contrast, a sada dosa is wrapped into a cylinder, and a podi dosa is folded in a semi-circle. If you order a 'rocket dosa', it is twisted and shaped into a conical form. Similarly, monomers fold and gain form. For example, myoglobin has a spherical form. It serves as a placeholder in the centre for oxygen transfer.

In contrast, collagen is a long protein comprising numerous amino acid chains twisted together like a rope or cable. Collagen's rope-like stiffness strengthens the tendons and ligaments that connect bones and muscles. Porin protein, on the other hand, forms a cylindrical tube with open ends, allowing it to function as a channel for a tiny chemical to diffuse across a cell membrane.

It is the shape that gives the protein its biological property.

Knowing the shape

Knowing whether the end of the pencil is sharp or blunt allows us to predict its behaviour. Similarly, the structure and shape of the protein fold reveal information about how it will function. Histidine is the 98th amino acid found in healthy myoglobin. However, if it is replaced with Tyrosine through mutation. Due to this mutation, the amino acids that appear in the outer layer change, which makes them stick together and form lumps. For example, rice cooked with too much water gets sticky and aggregated. They cannot store oxygen as effectively, resulting in myoglobinopathy muscle disease. Similarly, changing the sixth amino acid in the beta chain of the beta-hemoglobin gene from ordinary glutamic acid to valine causes significant alterations in the protein's behaviour. Healthy red blood cells are typically laddu-shaped, but aberrant ones formed by mutant proteins are sickle-shaped.

AI and protein structure

By the 1990s, artificial intelligence has progressed in image identification and facial recognition. Researchers sought to use AI algorithms to detect trends in protein folding and amino acid sequences.

Demis Hassabis co-founded DeepMind, a company that developed AI applications for popular board games such as chess, shogi, and Go, in 2010. Using his experience in AI pattern recognition, he and his colleagues developed AlphaFold, a neural network-based AI model for predicting protein structure. John Jumper, interested in protein dynamics, joined the Google DeepMind research team in 2017. Jumper and Hassabis co-led a new enterprise that employed the latest neural networks known as transformers to discover patterns in amino acid sequences of the 2,00,000+ proteins whose structures have previously been deciphered. The much-improved version AlphaFold2, released in 2020, was a big success for monomeric proteins.

Designer proteins

David Baker also began utilising AI technologies to forecast the structure. He created a method called Rosetta to predict protein structure. Suddenly, he had a brainwave: instead of attempting to identify the structure from the amino acid sequence, he and his colleagues sought to determine which sequences would generate a particular structure, similar to how different clothing fit a tailor's dummy. He adapted Rosetta to accomplish this purpose and obtained predictions for amino acid sequences for a specific protein structure.

When a structure was provided, Rosetta first recognised its constituent parts before searching a

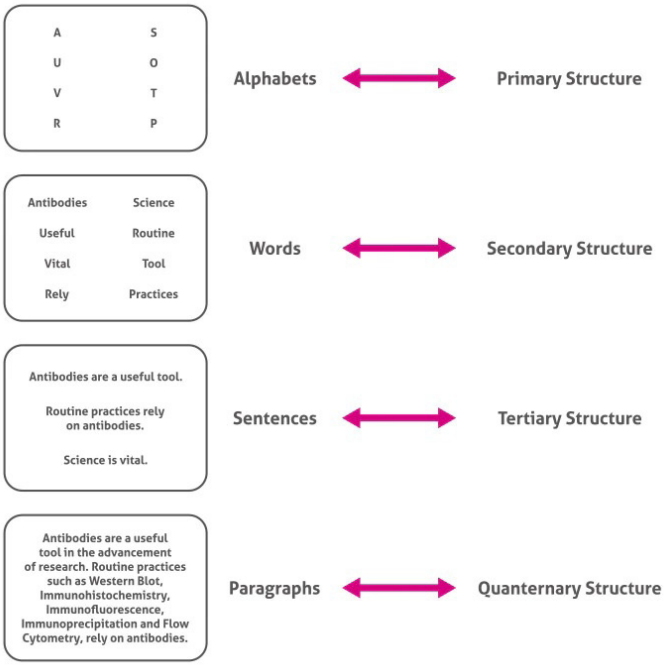

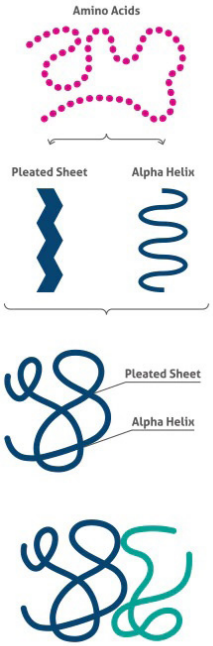

database of all known proteins for a small piece that matched the detected components. The program then determined the amino acid sequences for that component. As a result, the amino acid sequences for the appropriate protein structure were found in stages. His team synthesised the proper DNA sequence for the required amino acid sequence. They put the DNA into bacteria using genetic engineering. They produced artificially designed proteins, known as 'de novo proteins, ' as they do not exist in nature and are created using the structure as a guide. Thus, Baker demonstrated that functional designer proteins, or proteins with desired features, may be made.

Controversy

Some physicists are dissatisfied with the Nobel committee's choice to award what they believe is computer science and not physical sciences. However, how AI has transformed science and our way of life is mind-boggling. AI techniques are invaluable and vital for analysing and making sense of large amounts of data acquired from satellite telescopes to vast genetic information obtained worldwide.

Nevertheless, the criticism against AlphaFold is severe and a significant cause of concern.

Baker has made Rosetta's code available as open source for researchers. Google DeepMind has made the AlphaFold2 code publicly available, and over two million people from 190 countries are using it to detect protein structures and conduct drug development research.

However, in May 2024, Google's DeepMind and Isomorphic Labs (a subsidiary of Google Group's parent firm, Alphabet) disclosed AlphaFold3, which works for complexes, as a closed-source technique, prompting condemnation from scientists worldwide. "The AlphaFold3 case is a cautionary tale. While private money might speed scientific advancement, it must not come at the price of open research," says S Krishnaswamy, a retired professor at Madurai Kamaraj University's School of Biotechnology and also in All-India People's Science Network.

DeepMind did not reveal the whole code or the model's inner workings; instead, it provided a simplified algorithm description and a web server for restricted use. "Companies have a valid interest in protecting their investments, but excessive IP restrictions can hinder scientific progress and limit the societal benefits of innovation", according to him. "The limits of AlphaFold3 in terms of access present ethical problems. Restricted access has the potential to limit AlphaFold3's use for drug development to well-funded universities and enterprises. This might drastically slow the development of life-saving therapies, especially for illnesses affecting impoverished countries," he noted. ◆

Dr. TV Venkateswaran *is a senior scientist and profound science communicator. He can be reached via email at tvv123@gmail.com*