Artificial Intelligence

1

Life with Artificial Intelligence

Prabir K. Pal

Former Outstanding Scientist, Division of Remote Handling & Robotics (DRHR), Bhabha Atomic Research Centre, Mumbai 400085, INDIA
Former Professor, Homi Bhabha National Institute, Mumbai 400094, INDIA

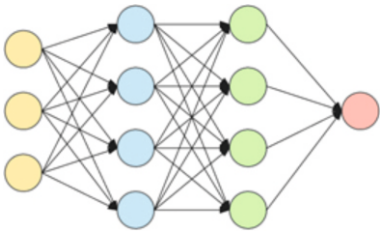

A Feedforward Neural Network

ABSTRACT

As the Artificial Intelligence (AI) technology grows in power, the question arises, can it ever give rise to an artificial mind? A mind which not only remembers facts and solves hard problems, but is conscious about its existence, and has a will of its own? Can such a mind surpass human intelligence? What can be the consequence of these mind boggling possibilities on our society? The present article provides a brief overview of AI as it stands today, and explores what lies ahead as this technology grows big and influences all aspects of our life.

KEYWORDS: *Artificial intelligence, Neural network, Superintelligent, Expert systems, Artificial general intelligence, Deep learning*

Introduction

Soon after the modern digital computer was invented, and the notion of a central processing unit fetching and executing instructions from a program stored in the memory became clear, quite a few visionary minds started speculating on the possibility of programming a human mind on such a machine. The idea was far too ambitious at that point of time, and the way forward was not at all clear. Nevertheless, the seed was sown. Alan Turing even laid down a procedure for testing such a machine, if ever it comes up. He suggested that a person engage in conversation with the artificial mind program and a human being separately using a keyboard, and if he is unable to say for sure which one is the human being and which one is the computer, then only the mind program can claim success.

Intelligence through Search

Human mind has diverse capabilities. Which feature do we pick-up to start building an artificial mind? Shaw, Newell and Simon started looking into ways in which a computer program could possibly solve a problem like proving a theorem of Geometry, or playing a game like checker, or even chess. They demonstrated that solving such problems boils down to searching through a maze of possible decisions to reach a desired goal. The sequence of decisions, that leads to the goal, forms the solution. Often the search space is potentially infinite. So there are strategies to search selectively through this space exploiting any prior knowledge about the nature of the problem. This was the first major breakthrough of Artificial Intelligence, which eventually led to the IBM Deep Blue Computer beating Grandmaster Kasparov in a game of chess in 1997. Even before that, in the eighties, there was an avalanche of Expert Systems in various domains, primarily triggered by techniques to capture domain knowledge in the form of facts and rules and use search to find an appropriate sequence of application of these rules to arrive at a solution.

Computers are extremely fast and efficient in certain matters, like numerical computations, storing and retrieving information to and from a database, etc. Humans fare very poorly in these tasks. However, there are certain other tasks that humans do effortlessly, like seeing and recognising objects in their environment, communicating in spoken and written languages etc., which the computers find very hard to perform correctly. Getting computers to carry out these tasks as effectively as humans do, or even better, has been the principal aim of Artificial Intelligence. Modelling these processes by appropriate facts and rules and searching for the right sequence of their application on input data has broadly been the approach to recognise an object, or a spoken word. These methods have improved considerably over time, but have not been able to reach the point where they can be used in daily life.

**Author for Correspondence: Prabir K. Pal*
E-mail: drhroff@barc.gov.in

Neural Network

There was a parallel effort for some time to build a computing system out of a network of artificial neurons. Each artificial neuron is a simple computing element whose output is a simple function of its input. Neurons form the building blocks of brains of living beings, so there was a strong suspicion that human-like capabilities may be easier to achieve using neuron like computing elements. The major difference in this case is of course the fact that we do not try to build any model for the process. We allow the network of neurons to learn to emulate the process by watching its input and output data under various circumstances. Thus we achieve the functionality of the process without knowing how it actually functions. In effect, we do not need the facts and rules any more. We only need input output data – and a lot of them covering almost all possible situations the process might encounter. The neural network deciphers the process and learns to produce the desired output for any input. Here the input may be an image data, while the object it represents can be the output. Similarly, a stretch of speech data can be the input, while the word it represents can be the output.

Artificial neurons can be connected and organised in many ways to exhibit diverse behaviours. One of them, called a Layered Feed forward Neural Network, in which the neurons are connected in several layers from input to output (Fig.1), has found many applications, and has been studied extensively to demonstrate their ability to emulate a process by observing its

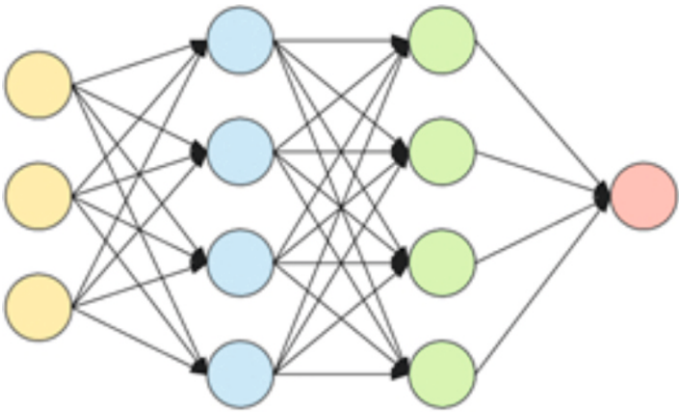

Fig.1: A Feedforward Neural Network of 4 layers with 3 neurons in the input layer, and a single neuron as an output.

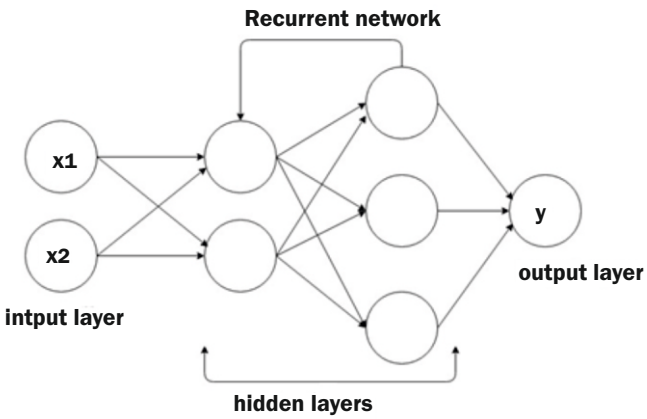

Fig.3: A Recurrent Neural Network (RNN) feeds a part of the output as input to a part of a preceding layer for the next cycle of computation, thus introducing a notion of memory.

input and output data. They were typically arranged in 3 to 4 layers, as a higher number of layers proved almost impossible to train. Because of the limitations on the number of layers, such neural nets could not emulate very complex processes. In 2004 however, ways to train deep neural networks with 30 or more layers of neurons were invented, resulting in the development of many successful applications of this technology in areas like object recognition, speech processing, natural language processing etc.

Image and speech have this property that they have a huge data to be presented as input, in which however only neighbouring data may be related in some ways. For processing such data efficiently, a Convolutional Neural Network (CNN) was devised, which permits only local connectivity among neurons in consecutive layers to identify certain basic patterns in different parts of the input data. This process progresses through several layers, ultimately creating a concise and abstract description of the input data. This is then fed as input to a deep neural network to produce the desired output. The entire Deep CNN so formed (Fig.2) was utilized with great success to identify objects in images and words in speeches.

There are some data that have an inherent notion of sequence in them, e.g., in a natural language, in which words appear in certain sequences to form meaningful sentences. A Recurrent Neural Network (RNN) was devised to process such data. In this, at any stage of computation, a part of the output is fed back and presented as a part of the input to one of the preceding layers of the Neural Network for the next stage of computation (Fig.3). This introduces a memory-like quality in the network. A word is no more interpreted in isolation. Instead, it is interpreted in the light of the words that have preceded it. Deep RNNs and their variants proved to be very useful in understanding and synthesising sentences and also in translating from one language to another. Translation of-course requires the network to get trained by reading many such documents which are published simultaneously in different languages.

Artificial General Intelligence

The recent spurt in AI applications have mainly been triggered by the advent of Deep Learning techniques. However effective they are, the old knowledge based search algorithms are not discarded either and are utilized freely in association with deep learning wherever appropriate. Some of the major achievements of AI in recent years have been the Google Translator which translates from one language to another (2006), IBM's Watson computer which unseated Jeopardy champion (2011), Deep CNN-based AlexNet which won Image Net Competition (2012) for recognizing images, Google voice for Automatic Speech Recognition (2015), Google Assistant /Alexa which answers your questions (2016),Google's AlphaGo which beat human Go champion (2016), and IBM's Debater which provided a spirited fight against human debater (2019). However intelligent these programs are, they operate well only within their domains. They are no substitute for the general intelligence exhibited by a human mind in varied circumstances. That is why there is a quest now for building Artificial General Intelligence (AGI) and to that end, it is pertinent to figure out how the human brain functions.

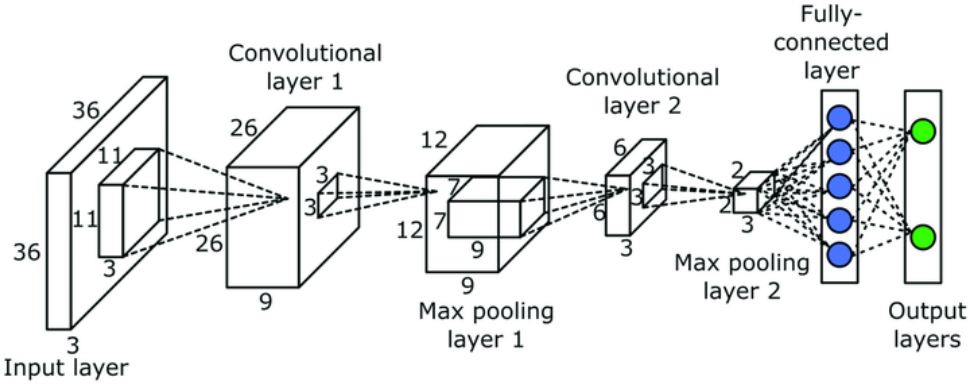

Fig.2: A Deep CNN comprises several stages of convolution and pooling, and finally a Deep Neural Network for training.

There are about 100 billion neurons in a human brain. They are arranged in clusters and these clusters are networked in layers. They receive signals through sensory nerves and these signals get interpreted by these neuronal layers to form a model of the environment. The outermost layer of human brain, called Cortex, has many regions in it corresponding to different functions of the brain. One of the regions recognises faces or objects and brings in corresponding emotions. Another region understands and forms sentences, an ability unique to humans. One creates an image of the body parts and helps maintain balance. Another does arithmetic, chooses words, creates metaphors and abstractions. Human qualities like ambition, vision, ethics and self-dignity are the prerogatives of one such region. There are also Mirror Neurons which enable us to feel for others, a quality that is possibly the origin of human culture. These cortex regions and their functions have been identified by associating abnormal behaviours exhibited by patients with regions of their brains that are injured, isolated, or dead, as indicated by MRI scans. It seems, all aspects of the personality of a human being are generated by the neuron clusters of the brain. This provides a mechanistic view of the human mind. And if mind is really a machine, it may not be impossible to simulate it.

One approach to emulate the human mind is to identify the components and processes of the human brain in sufficient details and simulate them on a machine. For that purpose, one has to get thin slices of the brain scanned to trace all neurons and their inter-connections in 3D format. This is easier said than done and at present it is possible only for brains of very small organisms with a few hundred neurons. This process must be mechanised so that it can be scaled up first for a mouse, then for a monkey, and ultimately for a human brain. An artificial neuronal model of the brain can be tried for creating AGI. If necessary, each cell or even each molecule of the brain may have to be simulated. But that is a gargantuan task, and far beyond the capacity of present day computers. Although this approach appears quite promising, AGI may also be achieved independently by means other than emulation of the brain.

Superintelligence

It is quite likely that at some point, computers will get smarter than humans. Smart computers will design even smarter computers with improved AI programs. Soon this may snowball into an intelligence explosion. Science, Technology, Defence, Business, International Relations, everything will undergo a rapid transformation with the advent of superintelligence. Even creative jobs will be performed by computers – faster and better. That will result in a life without jobs for humans! Even if everybody is provided for, this will bring in a major change in the way we live. Shall we become irrelevant and extinct as we lose our role in leading civilisation?

A Superintelligent System will take many decisions and actions to pursue the goals set by its creators. Even if some of them hurt us, we may not be in a position to restrain it from acting them out. Some experts feel that a Superintelligent System cannot be trusted and must be caged in a closed environment away from the Internet and all wireless communications. It may be provided with necessary information in digital storage medium, and its advice may be taken down in printed form, so that the machine does not get any opportunity to initiate actions or influence individuals to attain its goal. That sounds like a contradiction, since the whole purpose of building superintelligent machines is to obtain solutions of difficult problems faced by the human civilisation. Much of-course depends on the way Superintelligence is actually realized.

Conclusion

Max Tegmark in his book 'Life 3.0' describes three stages of our evolution. In the first stage, biological evolution from micro-organisms to humans took place over millions of years. Since this was a natural process, it was extremely slow. In the next stage, human civilisation quickened the pace of evolution of our lifestyle and culture through acquiring, storing and sharing of knowledge. We are at present going through that stage of evolution. In the penultimate stage, he predicts, superintelligent robots will take over humans and spread out to farthest corners of the Universe. We can mark them as our descendants.

That is of course only a prophecy. It is almost impossible to predict far into the future, as so many things change meanwhile. Whether humans will continue to have a role, or the machines will completely takeover is a big question that cannot be answered so easily. Possibility of superintelligence is strong, but how it may be realised is not clear at the moment. The future will unfold on its own. Hope, the journey will be beneficial and enjoyable to humanity.