Evolving scope of AI: From concept to computation

Dr. B. Yegnanarayana,
INSA Senior Scientist, IIIT-H and
Distinguished Professor, IIT Hyderabad
KID: 20200206

Most disciplines in engineering have evolved based on some fundamental principles, theorems and laws, like the strength of materials, electromagnetics and thermodynamics. But the discipline of artificial intelligence (AI) is evolving without any foundation. We are struggling to define the basis or theory to justify the area as it exists today. To appreciate this, we will briefly recall the history and evolution of AI over the past seven decades. We will discuss current notions of AI and highlight what is missing, for the future generations to ponder. The objective of this article is to dispel the misunderstanding and anxiety, especially among the youngsters, that they are missing something if they do not catch up with it now. In my opinion, by the time they are prime in their careers, either the scope of AI would have changed completely from what it is now, or new area(s) may emerge, making our current thinking of AI irrelevant, as it happened to most of us in our careers. For example, I am not sure how many of us now feel the relevance of most of the material in the books on AI written during the period 1960 to 1990 (eg. Books by Patrick Winston, N.J.Nilsson, Elaine Rich, S.Russell).

The term AI was coined in 1956 at Dartmouth workshop by four eminent mathematicians, John McCarthy (Dartmouth College), Marvin Minsky (MIT), N. Rochester (IBM), and C.E.Shannon (Bell Labs). The organizer John McCarthy said that the deliberations at the workshop "to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence in principle be so precisely described that a machine can be made to simulate it". The goal was to understand human intelligence and learning, and not to build a machine to learn and display intelligence the way human does. At that time there was never a focus on building powerful machines for AI.

In an effort to demonstrate by simulation our limited understanding of human way of doing things, it was necessary to develop interpretation of the manipulation of bits in Von Neumann architecture for symbolic operations, instead of numerical operations. The symbolic operations enabled people to simulate the logical inference, which was assumed to be a human trait.

The way people play some of the games is supposed to be a reflection of their intelligence. Hence attention was directed to simulation of games to demonstrate the intellectual activity of human beings through machines. Games were chosen for demonstration of the intelligent behaviour, as it was easy to represent them in terms of states and operators on a machine. Game playing was formulated as a search problem, using heuristics for pruning the search. AI was associated with the heuristics in search. Many real world problems were mapped as search problems, such as speech understanding systems (SUS) and image understanding systems (IUS). But slowly it was realized that mapping a real world problem as a search problem is the real intelligent part of the problem, which human beings were doing with their accumulated knowledge.

Continued…

It was felt that acquisition, representation and invocation of knowledge was the key for simulating the intelligent behaviour of human beings on a machine. Knowledge-based expert systems were developed under the fifth generation computer systems project in early 1980s. But soon it was realized that knowledge of an expert cannot be articulated to extract and represent in the form of rules for an engineer to simulate on a machine.

The AI researchers were desperately looking for an alternative to the knowledge-based systems. They thought they found a way in the emergence of Parallel and Distributed Processing (PDP) volumes by Rumelhart and Mclelland in 1986, which attempt to exploit our limited understanding of the structure and function of the biological neural networks (BNN) to realize computation models that can be used to simulate the intelligent tasks. Several different types of architectures, called artificial neural network (ANN) models, were developed to demonstrate specific pattern processing tasks that reflect some aspects of human learning and intelligence.

But very soon the field ended up in training complex ANN models, resulting in the emergence of three broad categories of learning tasks, namely, supervised, unsupervised and reinforcement learning. The focus shifted to capturing the implicit relations in large volume of data. With the availability of computing power and data, the attention was more on developing algorithms to capture the pattern information in the distribution of the data. Machine learning took the centre stage, and the original AI goal of simulation of learning and intelligence behaviour of human beings has been relegated to the background, along with the idea of building ANN models for specific pattern recognition tasks.

The availability of large volumes of data and huge computing power with storage has increased by several orders of magnitude in the past decade. This enabled people to develop large complex nonlinear models to capture the implicit patterns or relations or mappings in the input data. The complexity of these models is mostly in terms of number of parameters defining the model. A few varieties of models, such as CNN, VAE, GAN and LSTM, are conceived to cater to different types of data and problems. All these models are grouped under the generic name of deep neural networks (DNN), with associated deep learning for training. Note that DNN is simply a nonlinear computing model. It has nothing to do with the BNN in terms of structure or function. The deep learning refers to adjusting the parameters of the model. Deep learning has no significance either of learning or of the depth of learning. It is also interesting to note that deep learning is not even a task like supervised, unsupervised and reinforcement learning, all of which have no links with any particular architecture or model.

DNN with associated deep learning is a powerful computational model, which can be exploited for several sophisticated tasks for which large volume of data is available for training the model. One must acknowledge the power of these models in addressing a variety of practical problems to obtain meaningful predictions from the data.

Most of these tasks are currently interpreted as AI tasks, although no concept of AI, such as learning or intelligence is involved in it. It is also apparent that there is hardly any innovation possible for understanding human learning and intelligence, as envisaged by the pioneers of AI.

Continued…

Those familiar with the evolution of the scope of AI over the decades are amused at the current trends in describing AI as a superset of machine learning, which in turn is viewed as a superset of deep learning. Careful observation indicates that machine learning involves a set of algorithms for executing the tasks which are directly linked with data-driven problems. Likewise, DNN is a nonlinear model with large number of parameters, which can be determined only with large volume of data and huge computing power. There is tremendous potential for these models with many commercial benefits. But calling them as AI problems, and trying even to come up with explanation (explainable AI) for the intermediate stages of computation is far from the way humans do these tasks.

Once the applications using these models become a commonplace, it is likely that the attention may switch back to understanding the human way of doing things, especially in the domains of learning, knowledge and intelligence. It is important to see the complementary roles of human intelligence and machine intelligence. While human intelligence involves pattern processing using the BNN, machine intelligence involves data processing using computational models. Human intelligence is concept-driven, whereas machine intelligence is computation driven. The dichotomy can be expressed in several ways. My explanation is that humans process the data first and then represent it in their memory, whereas machines represent the data first (samples, pixels, numbers, etc) and then process. Some of the obvious attributes of human intelligence are selective attention mechanism, stability-plasticity, continuously reconfigurable BNN architecture, and dealing with variety of data situations. The more we make an attempt to understand these issues, the less we will know on how to describe them for implementation or simulation on a machine.

In this AI journey over decades, a few scientists and philosophers have been alerting on the hype being created on several occasions. Marvin Minsky's books on Perceptrons in 1969 and 1988, and Hubert Dreyfus's books on What computers (still) can't do in 1972 and 1994, are eye-openers in this context.

Current AI is technology-driven, and not concept-driven. While we wait for clearer perspective of AI to emerge in the light of the current developments, it is necessary to explore and exploit the potential of the evolving technology for many practical applications. It may not be wise for us to freeze the scope of AI in terms of these developments, as the scientific community is unable to guess what the next wave would be like in this evolution.

In conclusion, it is apt to quote C.E.Shannon about the limitations of the current computing models for displaying intelligence.

"Efficient machines for such purposes as pattern recognition, language translation and so on, may require a different type of computer than any we have today. It is my feeling that this will be a computer whose natural operation is in terms of patterns, concepts and vague similarities, rather than sequential operations on ten-digit numbers". Claude E. Shannon