International Journal of Innovations in Computing
(ISSN : 2319-8257) Vol. 1 Issue 1

Future of Artificial Intelligence

Deepak Sharma
Gujarat Technological University, Ahmedabad
depaks92@gmail.com

Abstract— ***Artificial intelligence is the science and engineering of creating machinery that can carry out tasks such as perception, reasoning, and learning. Artificial intelligence research has foundered on the issue of representation. When intelligence is approached in an incremental manner, with strict reliance on interfacing to the real world through perception and action, reliance on representation disappears. It is recommended that research and creation of artificial intelligence should continue because it is so important for the safety of the people of our country. In this paper I will give the basic introduction of AI and its stand in the 21st century, its associated applications and certain applications which have not yet been materialized but there is a scope for their further development and worthy application, such futuristic and fictitious applications are just an imagination but with the exponential growth of science and artificial intelligence we can assume that the day is not far when the artificial intelligence will be at its zenith.***

Keywords- AI; Expert System; Turing Test.

I. INTRODUCTION

Artificial Intelligence (AI) is a subfield of computer science that is concerned with the representation, study, and automation of knowledge and intelligent reasoning. One of the main goals of AI is to create computer programs that exhibit "intelligent" behaviour. The AI field is a vast one, encompassing a variety of problem domains, methodologies, and applications.

AI is the field within computer science that seeks to explain and to emulate, through mechanical or computational processes, some or all aspects of human intelligence. Included among these aspects of intelligence are the ability to interact with the environment through sensory means and the ability to make decisions in unforeseen circumstances without human intervention.

The birth of AI is said to have occurred at a conference held at Dartmouth College in the summer of 1956. The conference was organized by **Marvin Minsky**, who later helped found the AI laboratory at the Massachusetts Institute of Technology (MIT) and who is currently at the MIT Media Laboratory. Also present was **John McCarthy**, creator of the **LISP programming** language that is still an important tool for AI research and development. McCarthy himself at that time named the field "artificial intelligence." **Herbert Simon** and **Allen Newell**, who had already implemented an automatic theorem-proving program for the Rand Corporation called the Logic Theorist, were also present. These four--Minsky, McCarthy, Newell, and Simon--are considered the "fathers" of AI. McCarthy went on to found an AI lab at Stanford while Simon and Newell did the same at Carnegie Mellon. At the Dartmouth conference were also assembled a number of other fledgling researchers who had all written programs exhibiting the rudiments of intelligent behaviour, and it was the collective efforts of these individuals that gave impetus to the young field of AI.

II. HISTORY OF ARTIFICIAL INTELLIGENCE

The history of artificial intelligence begins with the four thousand year old wish to craft a copy of a human being out of spirit, or alchemy, or clockwork, or chemistry, or by the infusion of newly discovered energies. In the middle of the 20th century, a handful of scientists explored a new approach to an ancient dream, based on their discoveries in neurology, a new mathematical theory of information, an understanding of control and stability called cybernetics, and above all, by the invention of the digital computer, a machine based on the abstract essence of mathematical reasoning. The field of artificial intelligence research was born at a conference on the campus of Dartmouth College in the summer of 1956. Those who attended would become the leaders of AI research for many decades. Many of them predicted that a machine as intelligent as a human being would exist in no more than a generation and they were given millions of dollars to make this vision come true. Eventually it became obvious that they had grossly underestimated the difficulty of the project. In 1973, in response to the criticism of Sir James Lighthill and ongoing pressure from congress, DARPA and the British_Government stopped funding undirected research into artificial intelligence. Seven years later, the Japanese_Government and American industry would provide AI with billions of dollars, but again the investors would be disappointed and by the late 80s the funding would dry up again. The cycle of boom and bust, of AI winters and summers, continues to the present day. Undaunted, there are those that make extraordinary predictions even now. But, despite the rise and fall of AI in the perceptions of venture capitalists and government bureaucrats, AI has made continuous advances in all areas regardless of the climate, overcoming unexpected obstacles, reorienting priorities in light of new discoveries and riding the crest of the wave of increasing computer power. Progress has been slower than predicted but has continued nonetheless. Artificial intelligence problems that had begun to seem impossible in 1970 have been solved and are now successful commercial products, for example: optical

character recognition, industrial robotics, speech recognition, data mining, and Google's search engine, to name a few. In other areas, such as robotics, tremendous progress has been made. For example, in 1970 the robot Shakey could not reliably cross a room in 8 hours, but by 1995, the VaMP robot car of Mercedes-Benz and Ernst_Dickmanns was driving on the Autobahn in traffic at up to 180 km/h. It remains to be seen when or if an AI system will be built with a human level of intelligence. Alan Turing, in a famous 1950 paper, asked the question "can machines think?" and concluded: "We can only see a short distance ahead, but we can see plenty there that needs to be done."

III. Expert Systems

Expert systems occupy a type of microworld—for example, a model of a ship's hold and its cargo—that is self-contained and relatively uncomplicated. For such AI systems every effort is made to incorporate all the information about some narrow field that an expert (or group of experts) would know, so that a good expert system can often outperform any single human expert. There are many commercial expert systems, including programs for medical diagnosis, chemical analysis, credit authorization, financial management, corporate planning, financial document routing, oil and mineral prospecting, genetic engineering, automobile design and manufacture, camera lens design, computer installation design, airline scheduling, cargo placement, and automatic help services for home computer owners.

The applications of expert systems such as Mycin, The CYC Project are described in further section.

IV. The Turing test

At what point does the behaviour of a machine display intelligence? The answer to this question has raised considerable debate over the definition of intelligence itself. Is a computer capable of beating the world chess champion considered intelligent? Fifty years ago, the answer to this question would most likely have been yes. Today, it is disputed whether or not the behaviour of such a machine is intelligent. One reason for this shift in the definition of intelligence is the massive increase in computational power that has occurred over the past fifty years, allowing the chess problem space to be searched in an almost exhaustive manner.

Two key ingredients are seen as essential to intelligent behaviour: the ability to learn and thereby change one's behaviour over time, and synergy, or the idea that the whole is somehow greater than the sum of its parts.

In 1950 British mathematician Alan Turing proposed a test for intelligence that has, to some extent, withstood the test of time and still serves as a litmus test for intelligent behaviour. Turing proposed that the behaviour of a machine could be considered intelligent if it was indistinguishable from the behaviour of a human. In this imitation game, a human interrogator would hold a dialogue via a terminal with both a human and a computer. If, based solely on the content of the dialogue, the interrogator could not distinguish between the human and the computer, Turing argued that the behaviour of the computer could be assumed to be intelligent. Opponents of this definition of intelligence argue that the Turing Test defines intelligence solely in terms of human intelligence. For example, the ability to carry out complex numerical computation correctly and quickly is something that a computer can do easily but a human cannot. Given that, is it reasonable to use this ability to distinguish between the behaviour of a human and a computer and conclude that the computer is not intelligent?

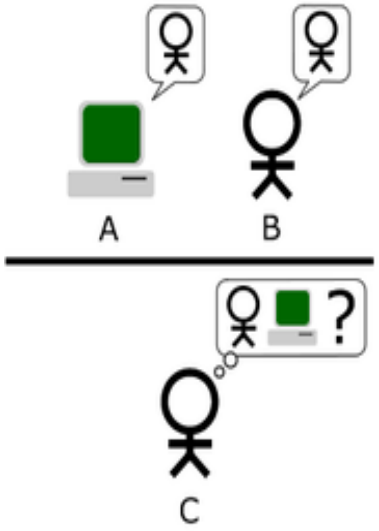

Fig. 1-The Turing test

The Turing test involves three participants: a computer, a human interrogator, and a human foil. The interrogator attempts to determine, by asking questions of the other two participants, which is the computer. All communication is via keyboard and display screen. The interrogator may ask questions as penetrating and wide-ranging as he or she likes, and the computer is permitted to do everything possible to force a wrong identification. (For instance, the computer might answer, “No,” in response to, “Are you a computer?” and might follow a request to multiply one large number by another with a long pause and an incorrect answer.) The foil must help the interrogator to make a correct identification. A number of different people play the roles of interrogator and foil, and, if a sufficient proportion of the interrogators are unable to distinguish the computer from the human being, then (according to proponents of Turing's test) the computer is considered an intelligent, thinking entity.

V. 21ST century AI

Artificial intelligence systems provide a key component in many computer applications that serve the world of business. In fact, AI is so prevalent that many people encounter such applications on a daily basis without even being aware of it.

One of the most ubiquitous uses of AI can be found in network servers that route electronic mail. Expert systems are routinely utilized in the medical field, where they take the place of doctors in assessing the results of tests like

mammograms or electrocardiograms. Neural networks are commonly used by credit card companies, banks, and insurance firms to help detect fraud. These AI systems can, for example, monitor consumer spending habits, detect patterns in the data, and alert the company when uncharacteristic patterns arise. Genetic algorithms serve logistics planning functions in airports, factories, and even military operations, where they are used to help solve incredibly complex resource-allocation problems. And perhaps most familiar, many companies employ AI systems to help monitor calls in their customer service call centers. These systems can analyze the emotional tones of callers' voices or listen for specific words, and route those calls to human supervisors for follow-up attention.

Although computer scientists have thus far failed to create machines that can function with the complex intelligence of human beings, they have succeeded in creating a wide range of AI applications that make people's lives simpler and more convenient.

VI. Applications of AI

Artificial intelligence has been used in a wide range of fields including medical diagnosis, stock trading, robot control, law, scientific discovery and toys. However, many AI applications are not perceived as AI: "A lot of cutting edge AI has filtered into general applications, often without being called AI because once something becomes useful enough and common enough it's not labelled AI anymore.", "Many thousands of AI applications are deeply embedded in the infrastructure of every industry." In the late 90s and early 21st century, AI technology became widely used as elements of larger systems, but the field is rarely credited for these successes.

Typical problems to which AI methods are applied

- Pattern recognition
  - Optical character recognition
  - Handwriting recognition
  - Speech recognition
  - Face recognition
- Artificial Creativit
- Computer vision, Virtual reality and Image processing
- Diagnosis (artificial intelligence)
- Game theory and Strategic planning
- Game artificial intelligence and Computer game bot
- Natural language processing, Translation and Chatterbots
- Non-linear control and Robotics

Other fields in which AI methods are implemented

- Artificial life
- Automated reasoning
- Automation
- Biologically-inspired computing
- Colloquis
- Concept mining
- Data mining
- Knowledge representation
- Semantic Web
- E-mail spam filtering
- Robotics
  - Behavior-based robotics
  - Cognitive
  - Cybernetics
  - Developmental robotics
  - Epigenetic robotics
  - Evolutionary robotics
- Hybrid intelligent system
- Intelligent agent
- Intelligent control
- Litigation

VII. Futuristic and Fictitious applications

In modern science fiction, AI is not necessarily limited by the fundamental problems of perception, knowledge representation, common sense reasoning, or learning. This allows speculation on the technology's potential impact on humanity, meditations on metaphysics or the nature of awareness, and the use of novel plot devices. AI has appeared in fiction as a servant (R2D2), a comrade (Lt. Commander Data), a technology expanding human ability (Ghost in the Shell), a conqueror (With Folded Hands), an exterminator (Terminator, Battlestar Galactica), a manager (Portal (video game)). Some realistic potential consequences of AI investigated in fiction are decreased labour demand, the enhancement of human ability or experience, and a need for redefinition of human identity and basic values (or a threat to existing identity and values). (Uncanny Valley hypothesis) One area of speculation focuses on potential disaster. (AI and Society in fiction) Though in fiction AI are often aware and capable of feeling, the phenomena that allow these experiences are not understood, and as such, there is no theoretical basis for their synthesis. Current theories provide for machines that can replicate or surpass all external human behaviour, but not necessarily human experience. An AI can play any role traditionally assigned to humans in a narrative, such as that of protagonist (Bicentennial Man), antagonist (Terminator, HAL9000), faithful companion (R2D2), or comic relief (C3PO). (See Sentient AI in fiction.) Many portrayals of AI in science fiction deal either with person-like or sentient AI, but the technology of AI appears in many other forms. (See non-sentient AI in fiction.) The inevitability of the integration of AI into human society is also argued by some science/futurist writers such as Kevin Warwick and Hans Moravec.

VIII. Conclusion

In the end I would like to conclude by imposing a question '**Can true AI ever be accomplished?'**, With all that processing power should we make plans for machines to replace humans as the most powerful creature roaming the

earth? Well, certainly not in our lifetime and perhaps never. There are numerous reasons for this. The biggest argument against developing true machine intelligence is the argument of evolution. Machines have not undergone the rigors of survival for millions of years the way humans have. The way we interact, think, respond and adapt are all developmental phases that are critical our intellectual dominance. This process hundreds of millions of years to evolve and the failures along the way are critical to our ultimate intellectual capacity and will be a major hurdle in our efforts to develop a truly artificially intelligent machine.

REFERENCES

[1] G.F. Cooper, E. Herskovits, A Bayesian method for the induction of probabilistic networks from data, Machine Learning 9 (4) (1992) 309–347.

[2] D. Dash, Caveats for causal reasoning, PhD thesis, Intelligent Systems Program, University of Pittsburgh, Pittsburgh, PA, April 2003.

[3] D. Dash, Restructuring dynamic causal systems in equilibrium, in: R.G. Cowell, Z. Ghahramani (Eds.), Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics (AIStats 2005), Society for Artificial Intelligence and Statistics, 2005, pp. 81–88.

[4] J. de Kleer, J. Seely Brown, A qualitative physics based on confluences, Artificial Intelligence 24 (1–3) (December 1984) 7–83.

[5] J. de Kleer, B.C. Williams, Reasoning about multiple faults, in: Proceedings of the 5th National Conference on Artificial Intelligence (AAAI-86), Philadelphia, PA, American Association for Artificial Intelligence, 1986, pp. 132–139.

[6] M.J. Druzdzel, H.A. Simon, Causality in Bayesian belief networks, in: Proceedings of the Ninth Annual Conference on Uncertainty in Artificial Intelligence (UAI-93), San Francisco, CA, Morgan Kaufmann Publishers, Inc., 1993, pp. 3–11.

[7] M.J. Druzdzel, H. van Leijen, Causal reversibility in Bayesian networks, Journal of Experimental and Theoretical Artificial Intelligence 13 (1) (January 2001) 45–62.

[8] A.S. Goldberger, Structural equation methods in the social sciences, Econometrica 40 (6) (November 1972) 979–1001.

[9] C.W.J. Granger, Investigating causal relations by econometric models and cross-spectral methods, Econometrica 37 (1969) 424–438.

[10] T. Haavelmo, The statistical implications of a system of simultaneous equations, Econometrica 11 (1) (January 1943) 1–12.

[11] Y. Iwasaki, H.A. Simon, Causality in device behavior, Artificial Intelligence 29 (1) (July 1986) 3–32.

[12] Y. Iwasaki, H.A. Simon, Theories of causal ordering: Reply to de Kleer and Brown, Artificial Intelligence 29 (1) (July 1986) 63–67.

[13] H. Kiiveri, T.P. Speed, J. Carlin, Recursive causal models, Journal of the Australian Mathematical Society 36 (1984) 30–52.

[14] S.L. Lauritzen, D.J. Spiegelhalter, Local computations with probabilities on graphical structures and their application to expert systems, Journal of the Royal Statistical Society, Series B (Methodological) 50 (2) (1988) 157–224.

[15] G. Kondrack and van Beek. P. A theoretical evaluation of selected back- tracking algorithms. J. Artificial Intelligence, 89:365{387, 1997.

[16] R. Korf. Planning as search: A quantitative approach. J. Artificial Intelligence, 33(1):65-88, September 1987.

[17] N. Kushmerick, S. Hanks, and D. Weld. An Algorithm for Probabilistic Least-Commitment Planning. In Proc. 12th Nat. Conf. AI, 1994.

[18] N. Kushmerick, S. Hanks, and D. Weld. An Algorithm for Probabilistic Planning. J. Artificial Intelligence, 76:239-286, 1995.

[19] D. McAllester. An outlook on truth maintenance. Ai memo 551, MIT AI LAB, 1980.

[20] D. McAllester. Truth maintenance. In Proc. 8th Nat. Conf. AI, pages 1109-1116, 1990.