Will there be human level artificial intelligence in the next 50 years?

October 13, 2009

Abstract

With computers getting faster and with more and more research being done on general intelligence and advanced AI for things like piloting unmanned aerial vehicles and monitoring CCTV, the possibility that we may be able to create a computer as intelligent as humans has become a topic of much discussion. This essay will explore the advances made in this field and show why the creation of computers as intelligent as humans is unlikely to occur in the next 50 years.

1 Introduction

Artificial intelligence (AI) is the branch of computer science which attempts to provide machines with humanlike reasoning and language-processing capabilities [3]. AI can be divided into two sub-areas: weak AI and strong AI. Weak AI, also known as 'applied AI' or 'narrow AI', is not intended to match or exceed the capabilities of human beings. Its currently used in games, internet search technologies etc. Strong AI, on the other hand, has no cases of implementation as yet. This is the intelligence of a machine that can successfully perform any intellectual task that a human being can, also referred to as 'artificial general intelligence' or 'human level AI'. Humans have been (for a long time since the early 1950s) and still are working hard on the field of AI because it has almost limitless potential which could help humans in many ways. After years of research and experimentation, AI has brought us robotic assisted surgery, safer cars and more. Humans, therefore, cannot deny the fact that AI is improving their standard of living in the current information age. However, this fact has ignited discussions about the possibility and danger that humans may be able to create human level AI that rivals their own and could replace humans in many ways and in some cases even make humans redundant.

This essay will explore the current state of research relating to human level AI in terms of theory, software and hardware and show that despite the extent of progress made in researching this field, we are unlikely to see human level of artificial intelligence in the next 50 years.

2 Evaluating human level AI

Every program needs a test to determine whether it has the required quality. A test can also reveal what areas or requirements are (after testing) or should be (before testing) covered by a program. If a machine has a human-level of artificial intelligence, it should be able to 'think' because that

[1] This essay was submitted for assessment in the subject 433-343 Professional Issues in Computing, in the Department of Computer Science and Software Engineering, The University of Melbourne, 2009.

is basically what humans do and also know what they need to do to complete their tasks such as decision-making [10].

In 1950, Alan Turing (Wikipedia) proposed a 'Turing Test', described in his paper 'Computing Machinery and Intelligence', to determine whether machines can think. In the test, a machine has to communicate with human participants (where these human participants will not know who they are communicating with, hence everyone is anonymous and isolated) by exchanging text only; the machine will only pass the test if it is able to convince other participants that all participants (including the machine) are humans [11].

AI Research Areas

High-level perception
Commonsense reasoning
Natural language
Speech processing
Gesture processing
Planning & counterplanning
Cognitive modeling
Plan recognition
Soft real-time response
Reactive behavior
Teamwork
Scheduling
Path planning
Spatial reasoning
Temporal reasoning
Opponent modeling
Learning
Knowledge acquisition

Figure 1: Areas in human level AI [8]

To date this is the only widely accepted test for Strong AI but not everyone is convinced of its reliability.

First, the ability to communicate with humans using text only does not cover all the areas that one expects in human-level of intelligence e.g. body language, intonation etc. The test does not cover areas such as the sensory-motor system e.g. co-ordination of movement in learning to carry out a task as participants are all anonymous and isolated, nor does it capture the process of learning as it is not practical to run the test for a very long time and see what the machine has learned from the conversation (see figure 1).

Second, the test does not show whether the machine has the required level of intelligence to carry out the tasks/jobs that humans do. However, there are those who consider that a machine goes through some 'thought process' if it can fool a trained human into thinking that they are communicating with another human. The level of reasoning of the machine, therefore, should be equal to or greater than that of a trained human. Moreover, in order to fool a trained human, a machine would need to understand human emotions and be able to come up with responses that only a human being would make but this is not always the case.

This line of thought is based on examples like the Eliza AI created by Joseph Weizenbaum in 1966 and the PARRY AI created by Kenneth Colby in 1972 which have fooled trained humans into thinking they were interacting with humans rather than a program. These examples demonstrate that by keeping the conversation alive it is possible to pass the Turing Test and not have human-level of artificial intelligence.

Hence, more discussion and research is required to develop suitable tests or testing tools to determine whether a machine can think and whether it possesses human-level of intelligence. Other important factors that need to be considered in this discussion and research are:

- the need to have a consensus on the definition of 'human level intelligence' to eliminate ambiguity.
- legal and moral status of a machine and how society will react towards a machine with human level AI etc. [11] (For a thorough discussion of the test and the controversies surrounding it, see the online article [4]).

3 The hybrid intelligence approach to achieve human level AI

In the quest for human level AI, it has become widely accepted that the most likely way human level AI will be achieved is through a combination of different artificial intelligence algorithms and technologies because it produces more immediate results. This approach begins with a brain function or learned ability to model - like object detection or sentence recognition - then a computer is programmed to replicate that function. A human brain is made up of several parts, so (in terms of functionality), it can be divided into discrete parts and each part can be programmed separately. Scientists believe that putting these parts together can create an intelligent entity even though each part is not intelligent on its own. This approach is known as the hybrid intelligence approach because even though the implementation maybe very different from how the human brain accomplishes the same task each part is modelled on a function of the human brain.

There are two main reasons why human level AI is unlikely to eventuate through this method in the next 50 years.

Firstly, weak AI has been the main focus of research and development for all these years. Secondly, the task of combining different AI algorithms and technologies to achieve human level AI itself is a big challenge requiring huge resources and appropriate advances in science.

3.1 Focus on weak AI (a brief history and research developments)

Research into AI started in 1956 when Alan Newell and Herbert Simon demonstrated a reasoning program called the Logic Theorist. Simon claimed that it was the first program based on logic, capable of thinking non-numerically. Logic Theorist was able to prove most of the theorems in chapter 2 of Russell and Whitehead's Principia Mathematica [12, 17] significantly advancing the research in AI.

This early success was followed up with the General Problem Solver. This program was designed to imitate human problem-solving protocols. It changed the order in which it considered some goals and possible actions but it could only handle a limited class of puzzles. During the same year a program written by Arthur Samuel was demonstrated on television playing a very strong game of Checkers which actually learned and improved over time. Then in 1963, the SAINT program was able to solve closed form calculus integration problems that are typical of a first year college course. In 1968, the ANALOGY program solved geometric analogy problems that appear in IQ tests. In 1969, the DENDRAL program was created to solve the problem of inferring molecular structure from the information provided by a mass spectrometer which was an early example of AI using domain specific knowledge.

During the year 1980s, expert systems became the craze with the first commercial expert system, R1, coming online at the Digital Equipment Corporation. By 1986, it was saving the company an estimated $40 million a year. By the year 1988, DEC's AI group had created 40 expert systems that used domain specific knowledge to accomplish various tasks. Another company Du Pont had 100 expert systems in use and at least another 500 in development saving the company an estimated $10 million a year. Almost every large corporation had its own AI group and a lot of very ambitious projects were under way. An example of one was the fifth generation project to create intelligent computers running Prolog. But in the years that followed many companies suffered as they failed to deliver. During these years, neural networks saw resurgence in interest because of the advent of the back propagation algorithm that allowed neural networks to be trained [12].

If we look at the past and fast forward till now, we can see that the area of strong AI has not grown a great deal despite the breakthroughs in AI mentioned above. Only weak AI - which is focused more on building tools for helping humans in their work rather than on replacing them - has been progressing rapidly. So, although, taking the hybrid intelligence approach to achieving human level AI may be possible in theory it is undeniable that more research and experimenting are required to develop more 'smart' algorithms and technologies to make this approach even plausible.

3.2 Strong AI and its challenges

Computer programs like Deep Blue and Alvinn may lead one to believe that we are on the brink of or have already created programs exhibiting some form of intelligence. However, examination of the underlying implementation mechanisms reveals that these programs are not able to demonstrate intelligence outside their chosen domain and do not scale very well. Alvinn, for instance, is able to navigate a car autonomously across the United States, but it is implemented as a large neural network trained on specific data and cannot adapt and learn as it navigates the car. Neural networks now have many algorithms that are vastly superior to the back propagation algorithm like gradient descent, Rprop and more. However, in Alvinns case, these new developments in neural networks will not enable Alvinn to adapt and learn.

IBMs Deep Blue (chess computer) is another perfect example of an AI that appears intelligent but has a very narrow and specialized domain. In 1977, it beat the world's best chess player Garry Kasparov but behind this AI is just an incredibly sophisticated and specialised search algorithm designed only for playing chess and nothing more. It also turns out that the game of chess is very suited for the type of problems that AI can solve and are good at.

While a program can predict chess movements, AI development has not reached a point where it can replicate, pre-empt simple human actions/movements like e.g. visual recognition of objects etc. because of the high degree of unpredictability/uncertainty involved.

Perhaps the answer lies in combining all these different technologies and algorithms together like neural networks, different expert systems, large databases of general knowledge and many other algorithms that have been discovered over the years to combine their strengths and balance out their weaknesses. This has in fact been done by Honda on the robot ASIMO. ASIMO is able to walk, dance, run, hold different objects, talk, understand certain words, and even understand the purpose of some objects that it has never seen before. It is not the first such robot to be built, in fact there have been 11 such robots before it. ASIMO seems to be more outstanding because it is very humanlike and can do far more than any other robot previously made. However, on close inspection, ASIMO actually has severe flaws. For example, ASIMO can certainly walk in places that its creators intended it to but as soon as something unexpected happens something its creators had not foreseen ASIMO fails miserably. Its creators will certainly learn from this mistake and try to rectify the problem by either adding rules into ASIMO's database or teaching it how to deal with this particular situation. But it does not seem like ASIMO has any real intelligence but only has programmed behaviour for a wide variety of situations. In fact one of the creators of ASIMO said that ASIMO is not smarter than a retarded cockroach [7].

Another good example of programs that give the illusion of intelligence that have even on some occasions fooled humans in the Turing Test are commonly referred to as chatbots or chatterbots as commonly referred to on the internet. A Chatbot is merely a massive database with questions and answers. The database is huge and filled with a lot of the responses that are exactly what a real person would have said, so, as a result chatbots seem very intelligent and humanlike with short conversations. However, it becomes immediately apparent that you are talking to a computer program if you chat long enough with one of these chatbots. That is because it cannot reason and come up with new things - something we call creativity and imagination. Also, such a program becomes increasingly difficult to develop to a higher level of communicability because there can be an infinite number of responses as the conversation progresses. It seems impossible to create a program with all possible responses so that it will communicate like a human.

It seems that even though research and development into AI is advancing at an incredible pace, we are not seeing anything that is even remotely artificially intelligent. What we are seeing more of are merely illusions of intelligence designed to seem like a human. One may ask why both ASIMO and chatbots rely heavily on searching algorithms and databases of knowledge but not any learning algorithms that can adapt and learn with time. The answer is: mainly because it is a huge challenge for scientists and engineers to combine different AI algorithms and technologies.

4 Reverse engineering a brain to achieve human level AI

There is another possible way to achieve human level AI and that is by reverse engineering the way the brain works and simulating it in a machine. This approach seems to be the most complicated but it seems to have the greatest chance for success in theory as it mimics how the brain performs a function e.g. from how the brain senses an event to how it reacts to the event (it is not just developing functions for particular tasks like the hybrid intelligence approach). From this point of view, there have been numerous studies carried out on the human and animal brain. Unlike the previous approach, this approach is more likely of achieving human level AI if given enough time and computing power. In theory, this approach seems very promising but the reality is very different and it is very unlikely that human level AI will be achieved in the next 50 years by using this approach due to two main obstacles. The two greatest obstacles that need to be overcome are: obtaining an accurate model of the human or animal brain down to the very last minute detail and having enough computing power to simulate the entire brain on a computer.

4.1 An accurate model of human brain

Many projects have been run in order to gain a better understanding of how the human brain works. Currently the most advanced project for this purpose is known as Blue Brain which is an attempt to study how neurons in the brain interact with one another. The Blue Brain projects initial goal is to accurately simulate the neo-cortical column of a rat brain [9] (as a rats brain only has 10000 neurons while a humans brain contains 60000 neurons). A neocortical column is the basic computational unit of the cortex and it is responsible for functions like processing the colour red or detecting pressure on a patch of skin. The basic structure of this column remains the same from mice to humans, which makes it perfect to assess how difficult it would be to reverse engineer the brain and simulate it on a computer.

According to Markram, [6] one of the scientists working on the blue brain project, the problem is that neuroscience is still woefully incomplete. The scientists (of this project) don't know how billions of discrete cells weave themselves into functional networks or even how the simple neuron (just a sheath of porous membrane) works. All these things remain a total mystery and would need to be solved before a model of the brain can be accurately simulated on a machine. Blue Brain is an attempt to reverse engineer the brain, to explore how it functions and to serve as a tool for neuroscientists and medical researchers. They have no intention to re-create the actual physical structure of the brain and it is not an AI project. Although we may one day achieve insights into the basic nature of intelligence and consciousness using this tool, Blue Brain only focuses on creating a physiological simulation for biomedical applications and thus it is not even close to emulating the real brain for AI. In fact, according to the creator of Blue Brain, creating an exact simulation of the brain for AI will be very difficult because every molecule in the brain is a powerful computer and we would need to simulate the structure and function of trillions upon trillions of molecules as well as all the rules that govern how they interact. One would literally need computers that are trillions of times bigger and faster than anything existing today.

Apart from Blue Brain project, there have been other attempts at understanding the brain. A biomedical engineer named Thomas DeMarse [13] grew a collection of 25,000 rat brain neurons that he connected to a computer via 60 electrodes. This project is meant to bring us closer to understanding how the human brain works but we are no closer now than we were before. In fact, this project is just a typical example of countless other projects that are attempting to figure out how the human brain works. The task is just incredibly difficult and requires a lot of time for research and experimentation, which makes the whole problem of understanding how the human brain works very unlikely in the next 50 years.

4.2 Computing power in the near future

Current estimates of how powerful a computer would need to be comparable to the human brain - fall short of what the Blue Brain project has shown to be necessary for simulating the brain. Based on Blue Brain estimates, a computer as powerful as the brain needs to be able to perform more than 38 thousand trillion operations per second and hold about 3584 terabytes of memory [5]. This is 413 times faster than the world's fastest supercomputer today.

According to Moore's Law, a long-term trend in history of computing hardware, the number of transistors placed inexpensively on an integrated circuit doubles approximately every two years. So, according to Moores Law we should have a computer as powerful as the human brain in about 20 years. However, some argue that Moores Law is not a law, but merely an observed regularity which cannot continue at infinitum [1]. That is, silicon technology hasnt got the capacity to double in capacity every two years and the prediction is that by the year 2014, the rising cost of chip manufacturing equipment will make it economically unfeasible to do volume production of chips with feature sizes smaller than 18 nm although further advances in shrinking processes can be achieved beyond 18 nm. So, even with conservative estimates of how much computing power is needed to simulate the brain, it is undeniable that we are lacking the computing power to do so and this is unlikely to change dramatically in the next 50 years.

4.3 Recreating a human brain on a machine

Despite the fact that this approach seems promising in theory, a critical question needs to be asked: Can the brain be recreated on computer? The answer, according to scientists, is yes in some ways, but mostly no. The brain performs many analogue operations which cannot be performed by machines and in many cases it achieves hybrid digital-analogue computing. Most importantly, the most significant feature of the brain that makes it different from machines is that the brain is constantly changing. If the resistors and capacitors in a machine start changing, then it would immediately malfunction, whereas in the brain such equivalent properties change constantly on the time scales of milliseconds to years. The brain is more like a dynamically morphing machine. Hence, we are still far from understanding the rules that govern the brain genetically and environmentally driven self-organization in response to external stimulus. Currently we have no analogue digital hybrid computers nor is there any research in the area. All these problems make the whole idea of accurately emulating the brain on a computer in the next 50 years merely a dream.

5 Conclusion

As can be seen from the discussion above, much has been achieved in the field of AI, but mainly in the field of weak AI. We are benefiting from this research everywhere; from manufacturing to the search for new drugs in medicine and even in our cars. However, we are far from developing something intelligent that can perform human processes/actions such as visual recognition, decision making due to change in stimuli etc. Hollywood has portrayed strong AI as something that can be easily achieved and as something that is imminent fuelled by those researchers claiming that we are 20 years from developing strong AI. For example, Herbert Simon wrote in 1965 that: 'machines will be capable, within twenty years, of doing any work a man can do' [2] but nothing can be further from the truth as this essay has shown. These sorts of claims will continue to be made but the fact of the matter is developing strong AI is an extremely difficult task and we have not made much headway. While developing a strong AI could be desirable for certain situations like space exploration, it is unnecessary for our everyday lives because we are far more energy efficient and can make 'copies' of ourselves very efficiently. Things that are monotonous and repetitious are already handled by robots with weak AI. So, neither is it technologically feasible to develop strong AI nor is there an urgent need to do so.

References

[1] Nick Bostrom. When machines outsmart humans. *Futures*, 35(7):759 – 764, 2003.

[2] Daniel Crevier. *AI: The Tumultuous Search for Artificial Intelligence*. BasicBooks, first edition, 1993.

[3] Bolton David. Definition of ai (artificial intelligence). `http://cplus.about.com/od/introductiontoprogramming/g/aidefn.htm`, 2009.

[4] David Dowe Graham Oppy. The turing test. `http://plato.stanford.edu/entries/turing-test/`, March 2008.

[5] Larry Greenemeier. Computers have a lot to learn from the human brain, engineers say. `http://www.scientificamerican.com/blog/60-second-science/post.cfm?id=computers-have-a-lot-to-learn-from-2009-03-10`, March 2009.

[6] Lehrer Jonah. Out of the blue. `http://seedmagazine.com/content/article/out_of_the_blue/`, March 2008.

[7] Michio Kaku. Michio kaku on artificial intelligence. `http://www.youtube.com/watch?v=PW8rgKLPHMg`, May 2007.

[8] John E. Laird and Michael van Lent. Human-level ai's killer application: Interactive computer games. In *AAAI/IAAI*, pages 1171–1178, 2000.

[9] Henry Markram. `http://bluebrain.epfl.ch/`, 2009.

[10] Nils J. Nilsson. Human-level artificial intelligence? be serious! *AI Magazine*, 26(4):68–75, 2005.

[11] Graham Oppy and David Dowe. The turing test. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Fall 2008 edition, 2008.

[12] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, second edition, 2003.

[13] Lovgren Stefan. ”brain” in dish flies simulated fighter jet. `http://news.nationalgeographic.com/news/2004/11/1119_041119_brain_petri_dish.html`, November 2004.