

How do you think that artificial intelligence will influence our lives in ten years time?

N.C.C. Brown (neil@twistedsquare.com)

June 18, 2003

1 Introduction

To answer this question, artificial intelligence must first be defined, to give a better understanding of what is being discussed. Once that is clear, there are two factors to consider in this question – firstly what advances will be made in AI in the next ten years, and secondly how artificial intelligence will be applied in the next ten years. To answer the first point the history of artificial intelligence should be considered, so that the progress made so far can help in predicting the future (as the saying goes, “in order to see the future, you must first look to the past”). To answer the second we will look at previous predictions and also reason about likely application areas. Finally the question will be concluded upon.

2 Definition

Artificial Intelligence has many different definitions. A handful are given here:

1. The branch of computer science concerned with making computers behave like humans. [1]
2. Artificial intelligence is the study of ideas to bring into being machines that respond to stimulation consistent with traditional responses from humans, given the human capacity for contemplation, judgment and intention. Each such machine should engage in critical appraisal and selection of differing opinions within itself. Produced by human skill and labor, these machines should conduct themselves in agreement with life, spirit and sensitivity, though in reality, they are imitations. [2]
3. Simply stated, artificial intelligence, AI, is the study of the abilities for computers to perform tasks which currently are better done by humans. [3]
4. In the minds of most people (especially the general public) the field of artificial intelligence is or should be going about the business of producing intelligence in digital computers. [4]

5. The science and engineering of making intelligent machines, especially intelligent computer programs. [5]

From my research on the matter, these definitions offer a representative cross-section of the definitions of Artificial Intelligence (AI). The first two definitions are, to my mind, very blinkered in their views on the (long-term) potential for AI. They limit the aims of AI development to imitating humans – it does not seem entirely inconceivable that computers could evolve a form of intelligence that is better or at least different from our own.

The third definition is not limited in the same sense, but instead focuses on the benefit of AI to humans. From a philosophical standpoint this is also blinkered – it is only concerned with the uses of AI for our own uses, not for the sake of developing AI because of its philosophical possibilities, but it does not limit the potential of AI like the previous definitions and also details the motivation of most researchers for studying AI.

The fourth definition concerns the public’s view of artificial intelligence. It is simplistic but nicely sums up AI at its most basic – intelligent machines. The fifth definition is from a question and answer page written by John McCarthy, a veteran AI researcher who invented the LISP language [6]. Like the previous definition, it is short, simple and to the point.

McCarthy also points out that “AI does not have to confine itself to methods that are biologically observable... On the one hand, we can learn something about how to make machines solve problems by observing other people or just by observing our own methods. On the other hand, most work in AI involves studying the problems the world presents to intelligence rather than studying people or animals.” [5]. This is important – while AI often seeks to surpass the intelligence of humans, it does not have to do so by emulating our brains. AI is far wider than just neural networks.

It is also useful at this juncture to define intelligence, or rather to point out that intelligence is incredibly difficult to define. There is an almost intuitive understanding of the meaning of the word that cannot be easily put into words. One common observation on the matter which is useful to note is that “intelligence is what you use when you do not know what to do” [7].

3 History of AI Research

The serious study of AI first began following the invention of modern computers during World War II [8]. Many researchers began work in the area, including Alan Turing – considered by many to the grandfather of Computer Science. AI blossomed through the decades only to provide disappointment during the 1980’s. As an AI veteran, Marvin Minsky, said: “these seemingly intelligent programs simply make dumb decisions faster”.

AI could be said to have captured the public’s imagination more than any other branch of Computer Science, in fact the idea of AI – in an abstract philosophical sense – has been of interest since before computers were invented. The idea of an artificially created

intelligence has been featured in myths through the ages, from the Ancient Greeks through to modern science fiction films. As one observer notes, “Though fast-growing, AI has never actually caught up with all the expectation imposed on it” [8].

It is without question that this expectation of AI has been increased by the wealth of science fiction surrounding the matter that has been produced throughout the decades. Isaac Asimov produced robot science fiction stories in the 1950’s concerning intelligent sentient robots, Arthur C. Clarke gave us the intelligent computer HAL in his seminal 2001: A Space Odyssey in the 1960’s, George Lucas’s ever-popular Star Wars films featured two charismatic intelligent robots C3P0 and R2-D2 in the 1970’s, Ridley Scott’s Blade Runner brought Philip K. Dick’s ruminations on what separates humans and artificial intelligence to life (and a much wider audience) in the 1980’s before imaginings of the dangers of artificial intelligence were shown in James Cameron’s Terminator films and most recently Spielberg and Kubrick’s epic A.I. considered our potential relationship with intelligent machines.

Ignoring the fiction, the facts must be considered – particularly what AI research has achieved since it first began. In the classical definition of AI, attempting to produce human-like intelligence, there has been precious little in the way of results. Much research has been done over the past fifty years but there is no clear indication that our machines are becoming intelligent. Minsky’s quote that the programs “simply make dumb decisions faster” is particularly telling. Computing power has increased exponentially and hence so has the speed of computers, enabling more and more complex number-crunching tasks to be performed, but there are few signs of a program that could actually be considered intelligent.

A focal point of interest in AI has always been the Turing Test. Invented by Alan Turing, this test was based on the argument that “if the machine could successfully pretend to be human to a knowledgeable observer then you certainly should consider it intelligent”. There have been various uncertified claims of passing the test but despite being considered by many to the holy grail of AI, researchers in the field have been sceptical and claimed it is irrelevant.

This was exemplified by the recent debacle [9] surrounding the Loebner Prize, offered to anyone who can produce a computer program capable of passing the Turing Test. Minsky himself calls the prize “obnoxious and stupid” and has offered a reward to anyone who can convince the organiser, Hugh Loebner, to withdraw the prize. There is an argument beneath this matter that the test is of no use to real AI research (the researchers claim), but the fact is that regardless of what the researchers says, the test is yet to be passed (Loebner points out). Is it a bad benchmark or are the researchers just trying to cover up their lack of progress towards reaching it? Whoever is right, little effort has recently been applied by AI researchers in attempts to pass the Turing Test.

In recent years AI researchers seem to have begun to realise that they are not satisfying the ultimate aim of truly intelligent machines, and some attempts have been made to shift the goalposts, as it were: “by the original definition of AI, AI has basically been failing all these years... maybe the most honest definition right now is that AI is ‘advanced computing techniques’ ” [4]. One AI researcher wrote: “If we mean AI to be a realization of real human

intelligence in the machine, its current state may be considered primitive... However, when AI is looked at as advanced computing, it can be seen as much more” [10]. This makes the point that while AI research is not at all close to human levels of intelligence, the uses of clever techniques and algorithms discovered during computer science research are wide-ranging.

So do the spin-off benefits of AI research justify the continuation of aiming for a seemingly impossible goal? Would the technological benefits of the space race, such as CAT scans and cordless power tools [11], have justified continuing the Apollo program even if the moon had never been reached? These are questions for the AI researchers to consider but AI research will doubtless continue regardless – it has been going for the past fifty years without achieving its goal, and it seems unlikely that the same will not be true for the next fifty years.

4 Current Applications of Artificial Intelligence

The potential practical uses for AI, namely intelligent machines, can be broken down into two categories – intelligent devices (e.g. fridges, vacuum cleaners, wristwatches) and intelligent personal computers. It is interesting how easily the term intelligent is applied to machines – it is done in a manner that suggests it is always applied relative to our current expectations, especially for devices.

A fridge that scans barcodes from the food you show to its barcode reader and orders a replacement is deemed intelligent [12]. Is it really intelligent though? All it does is scan a barcode and place a food order. Scanning barcodes is neither new nor intelligent, the same applies to internet shopping. It seems this device was considered to be intelligent because it does something new (or rather combines technologies in a new way) and potentially useful. There are many more examples of the term intelligent being used freely, often in contexts with extremely tenuous links to what is really considered intelligence. An intelligent device is even defined by some as “any type of equipment, instrument, or machine that has its own computing capability” [13].

It is clear that devices are termed intelligent too easily, according to the understanding of intelligence cultivated here, so it must now be considered whether any devices are intelligent. The observation listed earlier, that intelligence “is what you use when you do not know what to do” is useful here. This means that an intelligent device/computer algorithm is one that can still provide a good solution to a problem even when the problem falls outside its pre-programmed knowledge domain.

Neural networks (specifically those for recognising patterns) are an example of a computer algorithm that can operate well while considering inputs (patterns) outside of its existing knowledge (i.e. the patterns it used to learn). A neural network is intelligent as once it has learned to recognise a category of object well it can then recognise that objects it has never seen before also belong in this category.

It is difficult to discern between intelligent and non-intelligent computer programs, although in reality it is not as clear-cut as intelligent and not intelligent. We have agreed

that neural networks are intelligent, as data mining would seem to be and it is also clear that an average program, such as Microsoft Word, is not. It is simply a pre-programmed set of responses to various inputs.

The public, especially science fiction writers/enthusiasts, have long associated speech producing/interpreting capabilities in computers and intelligent computers. This is symbolic of the AI desire to replicate human behaviour in computers and also our own inherent recognition that fully understanding speech represents intelligence. A machine that could understand the full nuances of human language would surely deserve the right to be called intelligent. The near future promises nothing such as this though.

Machines that can speak are neither new nor (in their current form) intelligent. Rudimentary automated speaking machines can be found in train stations and in large companies' telephone exchanges all over the country. Machines for voice/speech recognition are slowly developing but are not very intelligent (they usually use neural networks which it has already been decided are intelligent, so they have that at least) especially considering the human potential for speech recognition and processing. It remains a very useful and appropriate area for AI research.

Advanced computer techniques are becoming more widespread in their use for an ever increasing number of applications, but intelligent computer programs, in the strict sense of intelligence, remain in their infancy in terms of their complexity and potential.

5 Future Applications of Artificial Intelligence

It seems to be a favourite pastime of most technological "gurus" to predict the future developments in computing and electronics, and AI is far from an exception in this field. Fuelled perhaps by humanity's seemingly innate interest in artificially creating intelligence, or maybe the science fiction that was born of that desire, amateurs and professional scientists alike love to predict that humanity is about to create highly intelligent (and possibly sentient) computers.

If this "future impatience" [14] can be avoided then likely future uses of AI can be considered. Speech recognition is a popular area of research that can benefit from AI, and is a growth area that is not very far from becoming a mainstream computer business area. Speech recognition in call centres is widely predicted to be a multi-billion dollar industry within the next ten years [15] even if the keyboard is still many years more away from being redundant.

Digressing into the endless potential applications of artificial intelligence to try and determine the future for AI would be a large and ultimately uninformative task. Predicting possible niches for the use of AI in order to determine the overall future trend would be a large and inefficient task. A step away from the details must be taken in order for the whole picture to be considered.

One factor that many people do not consider when making their predictions for the future uses of computers is that in many cases the non-computerised method of doing things is quicker, easier, cheaper or perhaps all three. Computers do not control the lights

in our houses because it is easier just to press the switch. Our recipes are not available on a computer screen on the front of the cooker because it is just as easy to pick up the cookbook from beside the oven. The list goes on, and the principle applies to AI as well as general computer use.

Computing is an area that can be hard to predict though. Thomas Watson, chairman of IBM famously said in 1943: “I think there is a world market for maybe five computers”. The interesting thing about this quote is that it was fairly accurate, based on the state of computing at the time. The subsequent developments of silicon and semiconductors were not foreseen in the early days of computers. “There is no reason for any individual to have a computer in his home,” Ken Olson, president of Digital Equipment, stated in 1977. This was perhaps slightly more short-sighted but it shows how difficult the field is to predict.

It seems highly unlikely that AI will suddenly make gigantic leaps in its application during the next ten years and produce intelligent talking computers for light conversation and witty repartee with an amazing cognitive and learning capacity. Its development will be much slower than that. Advanced computing techniques will continue to be developed and integrated into the current uses for computers, and AI will gradually develop and be applied to more areas in the current applications of computers, but it is noteworthy that the applications of AI(a discipline that tends to require much computer power) will remain a subset of the applications of personal computers until miniaturisation has advanced (and costs have reduced) to a point where embedded microchips can cope with current and future AI algorithms.

It should be pointed out that practical applications of AI research are not the only use of AI. It has been pointed out that greater understanding of how to implement human-like learning and decision making in computers could potentially help treat mental disorders in humans. This is of course different from using AI to help diagnose patients, another oft-cited use of AI.

6 Summary and Conclusions

It has been ascertained that research into AI can bring about many potential benefits. This is peripheral to the question however, which asks how artificial intelligence itself will affect our lives. It has also been determined that few current computer technologies can actually be termed intelligent, therefore there is little current evidence of systems that can be called artificially intelligent.

It is often claimed that the next big breakthrough in AI is imminent. With the track record in research as the only impartial and easily comprehensible guide it is difficult to believe that this is the case. The harsh view of AI research in relation to its goal of achieving true artificial intelligence is that it has achieved little so far. Computing is a difficult area to predict but the best guess must be that although AI will develop in the next ten years, it seems unlikely that it will take a quantum leap forwards.

The other way in which AI might become more influential over the next ten years is in applying existing techniques (and their developments) to new application areas or by

applying them to greater effect in current uses of AI. Speech recognition is a good example of the latter, and will likely be the most important and noteworthy area of computing that AI will benefit. New application areas for AI are very hard to predict, but HCI (Human-Computer Interaction) seems like a strong candidate for the application of AI. We much prefer dealing with intelligent beings to using “dumb” computers.

My answer to the question posed is that artificial intelligence will not influence our lives much more in ten years time than it does now. Artificial intelligence has a long way to go yet before it can achieve its goals, but the discoveries made by research in this area justify its continuation. There are intelligent techniques that have been developed though and it is doubtless that these will continue to be developed and similar new discoveries made. However, true intelligence in machines appears to be, for now at least, beyond our reach. Only time will tell whether this remains to be so.

References

- [1] Definition of Artificial Intelligence,
<http://3dgraphics.about.com/library/glossary/bldef-ai.htm>
- [2] Definition of Artificial Intelligence,
<http://privacy.cs.cmu.edu/people/sweeney/aidef.html>
- [3] What Is Artificial Intelligence?,
<http://www.cwrl.utexas.edu/~tonya/cyberpunk/projects/ai/definition.html>
- [4] A comp.ai.philosophy FAQ,
<http://www.dontveter.com/caipfaq/ai.html>
- [5] What is Artificial Intelligence: Basic Questions,
<http://www-formal.stanford.edu/jmc/whatisai/node1.html>
- [6] ALU: Lisp History,
<http://www.lisp.org/table/history.htm>
- [7] Fridges, Elephants, and the Meaning of Autonomy and Intelligence
<http://138.100.76.200/~sanz/ISIC-2000.pdf>
- [8] The History of AI,
<http://www.generation5.org/aihistory.shtml>
- [9] Salon.com – Artificial Stupidity,
http://archive.salon.com/tech/feature/2003/02/26/loebner_part_one/index.html
- [10] Toshinori Munakata, Guest Editor, Communications of the ACM, March, 1994
- [11] Apollo's Contributions to America, NASA,
<http://www.sti.nasa.gov/tto/apollo.htm>

- [12] Business: The Economy Guru predicts web surge, BBC News,
http://news.bbc.co.uk/1/hi/business/the_economy/521065.stm
- [13] intelligent device – a whatis definition
http://whatis.techtarget.com/definition/0,,sid9_gci812508,00.html
- [14] Plocktau: The Portfolio of Caleb John Clark
http://www.plocktau.com/portfolio/publications/web_review/wr_lion.html
- [15] PCs and Speech: A Rocky Marriage
http://www.businessweek.com/magazine/content/02_36/b3798078.htm