

When the Machine gets Smarter!

*Summoning the Demon –
The Last Invention?*

Dan Remenyi PhD

dan.remenyi@gmail.com

(graphics free version for download)

© Dan Remenyi, November 2018.

The challenges inherent in smart machines

- There is an increasing belief that AI is already becoming one of the most powerful technologies at humanity's disposal. It is becoming ubiquitous. However although it is incomplete, progress is being made rapidly. Full AI or AGI (Artificial General Intelligence) is coming.
- This has led to a series of warnings by individuals such as Bill Gates, Sam Harris, Elon Musk and even Stephen Hawking before he died.
- These warnings forecast AI getting out of control. Exactly how this could happen is not at all clear. It is not likely that the scenarios used by Hollywood to entertain us involving killer robots will come to pass. But there are other challenges relating to how the technologies which are becoming available can seriously impact society and in a negative way.
- The two most important issues revolve around what has been referred to as "*hacking the human*" and the *re-conceptualising of the world of work*.

A Bit of History – AI as Expert systems

- In the early days what passed for AI were Expert Systems which were rules-based processing systems that matched input to an established database.
- These were/are the Mr Memory (The 39 Steps) of the current era. This has been taken to a high level by Watson, the winner of the Jeopardy TV Programme. Imagine the intellectual capital (?) represented by that system.
- Within the parameters of a given system this technology is most efficient.
- It can be highly effective provided there are no changes to the parameters.
- But what happens when the conditions/requirements/parameters change? Enter machine learning!

Will a machine ever really be smarter? Is super-intelligence a real possibility?

- Clearly the first question requires the careful definition of smarter, but the likelihood is YES. As academics we have to say nothing is impossible.
- At a simple level humanity has a long history of being able to teach animals “tricks”. There is now a history of teaching machines to perform amazing feats in the form of automation. Machines have now been *taught to learn* and some machines have been taught to do extraordinary tasks and this trend appears to be accelerating.
- What is called a robot is often simply automation. Some argue that to be classified as a robot it is necessary to be autonomous, sense the environment, make a decision and change activity if circumstances require. Also a physical presence. Thus computer “bots” are not really robots.
- Connecting learning and doing is a challenge. So far “*doing*” has been highly controlled. Although more than a million people die in road accidents every year, when one person is killed by a driverless car it makes headlines across the world.

Learning machines

- We arrive at a point where we may not need to populate the database with the “knowledge” needed for expert or rule-based systems.
- Now the machines can learn – largely on their own. Machines can improve themselves? But within what parameters? The collection of knowledge per se does not imply any resultant action.
- Does this mean that a machine might improve itself up to creating its own consciousness? Ignoramus. Miracles do appear to happen! But is consciousness an improvement in the eye of the machine? Would a machine want to be conscious? What could it mean to say that a machine might “want to be” anything?
- HAL has arrived?..... maybe in a sense HAL has been superseded.
- The fear of the big names is that we are in danger of becoming redundant/obsolete. (?)
- If knowledge is simply the accumulation of “know how” then the answer is yes. But we can be swamped with knowledge and this can reduce action.

The conundrum of values

- If machines require super-intelligence, and we don't really know what that means, then the conclusions which they could reach in solving some of humanity's problems might be completely unacceptable. It is clear that being rational is not the only way in which humans address problems. In fact it is probable that rationality is relatively low among our problem-solving activities.
- We cannot conclusively resolve this issue which might turn out to be like those suggested by Musk, Gates, Harris etc.
- On the other hand this fear might be a mirage.
- But there are other big issues.

Hacking the human - Enter Big Data

- The idea of hacking the human became visible as a result of the exposure of the work done by Cambridge Analytica. Although there is no hard evidence in the form that would be accepted by academic scholars there is a strong suggestion that the algorithms used by Cambridge Analytica were instrumental in securing the 2016 USA presidential election for the present incumbent.
- The suggestion made is that by the use of Big Data together with appropriate algorithms the computer can get to know an individual better than almost anyone else including that individual him/herself. This allowed highly focused targeting in the campaign.
- Do we believe this? The jury is out.

The availability of Big Data for the purposes of “knowing” individuals

- Living in the Western world and participating in its society as a resident, a worker, a student, a tourist, a patient - in fact anyone other than a very young child or a very elderly individual, it is impossible not to have some degree of a digital footprint.
- Most individuals freely offer a digital footprint through the use of one or other social media software product.
- Any online purchasing or selling provides more material for the digital footprint.
- Every time anyone indicates a like or dislike another important element of the digital footprint is supplied.
- The suggestion is that the digital foot print is the most authentic representation of an individual. Is this true?

One of the dangers

- We are asked to believe that AI algorithms applied to Big Data resources will massively increase our ability to manipulate individuals, not only commercially but also politically. This is what is being referred to as “hacking”. It could be described as digital social engineering.
- This is clearly a hypothesis for which there has been no evidence which would stand up to any rigorous examination.
- However prima facie there is no obvious reason to outrightly discount this hypothesis.

Our history and our future

- We have generally overestimated the amount of choice we have. This is not only individual choice but also the choices faced by our community. It is clearly true that technology will allow us to collect vast amounts of data, and algorithms will be able to perform the most intricate types of calculations. But it has never been true that the past uniquely dictates the future.
- The past can give us very useful information with which we can make assessments about what is likely to happen, but it is complete folly to believe that there is a direct connection between the two.
- But how much of a “direct” connection is needed?

Reservations about the hypothesis

- Although the hypothesis cannot be discarded out of hand there is reason to believe that more is being claimed for this technology than can be reliably delivered and this is perhaps well expressed by an extract from GK Chesterton:-

“The real trouble with this world of ours is not that it is an unreasonable world, nor even that it is a reasonable one. The commonest kind of trouble is that it is nearly reasonable, but not quite. Life is not an illogicality; yet it is a trap for logicians. It looks just a little more mathematical and regular than it is; its exactitude is obvious, but its inexactitude is hidden; its wildness lies in wait”.

Chesterton, G.K, Orthodoxy – The romance of faith. Doubleday Dell Publishing Group, New York, 1990.

Minimising the digital footprint

- It has been argued that the only way of minimising the digital footprint is to utilise the products and tools currently being used for the dark web which has been traditionally seen as the locus of crime.
- There are clearly going to be many issues to be resolved in this respect as our need for privacy and the need to participate in the Internet society further develops.
- Knowledge managers will certainly have a role in monitoring and explaining this to everyone concerned.

Re-conceptualising of the world of work

- Virtually no job is immune from being transferred in whole or in part to a computer.
- Previously computers largely caused a shift in work from relatively unskilled jobs to more skilled work, but now AI will see professional work be severely impacted by computers as AI is focused at knowledge workers.
- We cannot forecast the impact of AI and robotics on the labour force but there is reason to suspect that within the 5 to 10 year timeframe there will be a very severe impact on the number of jobs available.

Jobs and the future

- Is every job vulnerable to automation?
- No. But an awful lot are.

Universal Basic Income

- UBI sometimes also referred to as Unconditional Basic Income is based on the notion that our society will not need the labour of all citizens in order to provide adequate goods and services for everyone.
- If this is the case then the argument goes that we will need to provide a reasonable level of physical well-being for that part of the population which is not required in the production of these goods and services. Clearly this would make sense in a world where automation (or if you will AI and robots) are able to deliver. However we do not appear to be anywhere near this position, yet.
- Interestingly enough UBI has been supported by a number of high-profile individuals from the high-tech industry. It is not simply a spin-off from traditional socialist thinking.

When the machine gets smarter!

- Will it make sense to have an activity called knowledge management?
- Will the amount of intellectual capital simply be a reflection of the amount of money which an organisation is prepared to spend on its computers. Computers at this level are really expensive.
- Will the knowledge manager become the chief human manipulator in the organisation?
- How will we be able to prevent this from happening?
- Do we need an ethical code to be in place before this whole thing get out of hand?

The re-conceptualisation of the world of work

- Are we really prepared to pay a Unconditional Basic Income to individuals who make no material contribution to the society?
- Are we prepared to make this payment adequate so that these individuals do not live in penury?
- Are we able to we envisage our lives so that work does not define who we are?

Where is the end?

- We do not know and we cannot know..... but
- It is pretty clear that the arrival of the smart machine is going to facilitate the production of a new type of society in which we will have to question many of our values.

Let's go back to knowledge

- Knowledge is a slippery word.
- The Latin word **Scio** means to know and we have taken Scio as the root of science.
- Thus we have placed an emphasis on the word **know** and thus knowledge to have some sort of quality of an artefact.
- In short knowledge is something that we create and which exists outside of us in the same way as the book does.
- But knowing is never enough. Knowing without doing is seldom of any value. While our minds are of great importance to us even when there is zero action such as while we sleep, in society our thoughts are only of value if they result in action.

And the action has got to be appropriate

- There are many stories of the inappropriate suggestions made by computers in former days.

"The owl was the wisest of animals. A centipede with 99 sore feet came to him seeking advice. 'Walk for two weeks one inch above the ground; the air under your feet and the lack of pressure will cure you,' said the owl. 'How am I to do that?' asked the centipede. 'I have solved your conceptual problem, do not bother me with the trivia concerning implementation,' replied the owl." Shubik M, A Game-Theoretic Approach to Political Economy, The MIT Press, USA, 1988.

- Many of our problems occur because of our attitude towards personal freedom which is not always rational.

Uncovering knowledge

“Nothing can have value without being an object of utility”.

Karl Marx *Capital Part II* Chapter 3, 1867-1883.

“Knowledge is knowledge for some purpose. The validity of the knowledge depends on the validity of the purpose”.

E H Carr, *What is History?*, p22, Penguin Books, 1967.

“One may understand the cosmos, but never the ego; the self is more distant than any star.

G. K. Chesterton *Orthodoxy*, “The Logic of Elfland” (1908).

Knowledge is.....

- Dewey thought that ideas and beliefs are the same as hands instruments for coping. An idea has no greater metaphysical stature than, say, a fork. When your fork proves inadequate to the task of eating soup, it makes little sense to argue about whether there is something inherent in the nature of forks or something inherent in the nature of soup that accounts for the failure. You just reach for a spoon Knowledge is not a copy of something that exists independently of its being known, *“it is an instrument or organ of successful action”*.

Dewey, cited in *The Metaphysical Club*, Lewis Menand, P 361, Flamingo, London, 2002

- This definition begs a number of very important fundamental questions one of which is where is the boundary between knowledge and other things including information?

Following Dewey - Pragmatist

- Maybe what the Pragmatists are saying is that acquiring information is a strategy of being prepared. Exercising one's knowledge come into play when one does something with it.
- Any particular parcel of what we now tend to call knowledge only has meaning when it is actioned. Before it is actioned it is information or maybe even data.
- It is the connection of data/information with action which delivers something we can say has been the result of the application of knowledge.
- The idea of knowledge without it being located in someone is empty.

Encouraging the use of knowledge

- One of the issues which has to be faced is the fact that the application of knowledge is not necessarily neutral. It may be argued that knowledge itself may not be value free.
- Knowledge can be used for good and for bad in the same way as the Pragmatists saw a fork.
- Before we can simply applaud the acquisition and the use of knowledge it is necessary to ask the question *“For what purposes is knowledge to be used?”*
- In general it is difficult to see that the search for AI has asked this question. It is probable that this question is unanswerable from where we stand at the moment.
- The problem is that there is so much at stake in that the apparent potential of AI is vast and it is not at all clear that if we achieve a state of proper AI that this artificial intelligence will be put to use in a way which will be conducive to the furthering of values.

Is control the issue?

- Ultimately this seems to be a control issue and it has been argued that if the world creates AI then it will be possible to switch it off. However this is almost certainly a highly naive statement which can be seen by the thought experiment of asking the question or making the statement "Let's switch the Internet off!"
- The way ICT is progressing AI systems may end up being contained within one or a small number of microchips and may exist independently around the world.
- When we think of how difficult it is to contain nuclear power, any attempt to control the proliferation of AI could well be impossible.

Perhaps this is not the question we should be asking?

- Do we want machines to be smarter? If the answer to this question is yes then we should ask "*for what reason?*"
- What are the ethical implications of such machines?
- What might be the unanticipated consequences (of course we cannot definitively answer this question)? If computer code can learn then can smart programs/machines create/build smarter programs/machines?
- What do we do when interests collide and values clash? Is this *Summoning the Demon?*
- Do we plug out the machine or does the machine plug us out?
- The wise scientist in the Hollywood film Jurassic Park makes the comment that scientists were too busy working out how they could produce the monsters to bother to ask the question as to whether they should be doing it at all. So the central question is, *Do we really need AI and if so why?*

The last words of Alan Turing

- *“We can only see a short distance ahead, but we can see plenty there that needs to be done”.*
A. M. Turing (1950) Computing Machinery and Intelligence. *Mind* 49: 433-460.
- *“What we can achieve has always been governed by our imagination. Our imagination is a principal driver of our progress. Our imagination is usually tuned to look at **“the bright side of things”**. It might now be time to look at what could also happen which we may not like”.*
Remarks by John Searle
- Can a machine think?.....Thinking is a biological process. A human is in a sense a machine (a number of parts which work in collaboration to achieve an objective) and therefore the answer to this question is in a sense YES.
- Can an artificial machine think?.....We do **not** know how the brain does it but in theory we can learn this and if we do learn then **there is no reason why not**. The brain is a causal mechanism that produces consciousness by complex and still imperfectly understood neural biological processes. *Will we ever know how these process work? We do not know.*
- Consciousness.....a common sense definition.... It consists of all those states of feeling, awareness which feels like it is best described as being conscious. We don't know what the term artificial feeling might mean so we cannot really talk about artificial consciousness.

KM as a Management Paradigm

- We value knowledge because it is at the base of all our activities and it is the ultimate source of all our aims and ambitions
- KM has been remarkably successful. It was originally thought to be a fashion.
- KM has evolved and IC has flourished.
- The importance of these subjects continues to grow in the period when we have to face up to the fact that machines will contain more and more “knowledge” and may be able to, or perhaps most probably will be able to out manoeuvre humans intellectually.
- How do we position KM/IC in a world which is increasingly pre-occupied with finding every possible application for AI?
- AI was once defined as a machine (hardware and software) which could perform (at least conversationally) as well as a human. However there were far too many definitional problems with this understanding and this led to the splitting of AI into a number of subsets including expert systems, computer vision, voice recognition etc.

A last word

- There was a piece in a recent Sunday Times (Business Section page 10) arguing that '*Knowledge doesn't help us*' and that more and more tasks are best handled by technology / AI.
- For example, we know that we should change computer passwords weekly and go to the gym to keep fit. But many do not do it. AI would do a better job of replacing passwords. But that is rather different to teaching and learning.