

## Artificial Intelligence: is it finally arriving?

Professor Leslie Smith, Computing Science and Mathematics, University of Stirling.  
9 May 2013.

Artificial intelligence (AI) systems have been promised since the early 1960's. Fifty years on, it is reasonable to ask whether they are about to arrive. In the intervening half-century, what we think of as AI has changed. While the robots from films (HAL from 2001, or Marvin from the Hitch-hiker's Guide to the Galaxy) clearly are intelligent, we have come to no longer regard machines that play games (chess, or go, or backgammon) as intelligent once we have managed to construct them. Another issue is what is meant by human intelligence. There is little or no agreement amongst Psychologists as to what constitutes human intelligence. Turing, one of the forefathers of Computing, devised the Turing Test whereby a machine is to be deemed intelligent if its conversation is indistinguishable from human conversation. Devisers of such machines can now compete for the Loebner prize, awarded for success in this test.

But is this really intelligence? The philosopher Searle suggested in 1980 that it would be possible to devise a machine that worked in a purely mechanistic way, holding a conversation without having any understanding of the conversation itself. Yet surely intelligence, whether artificial or animal, requires understanding. There are two polarised views of how to build an intelligent machine, the top-down view and the bottom-up view. The top-down view suggests that intelligence can be reduced to a clean mathematical abstraction, which can then be implemented in electronics. The bottom-up view is that we need to consider the low-level issues of how animal intelligence is created by the brain. The brains of different vary hugely in size, and there are animals with larger brains than humans, (e.g. sperm whales and elephants), so that size of brain is not everything. For many years now, researchers have developed systems based on brains, using artificial neural networks. These can do pattern completion, and some (simple) forms of prediction, and have found their way into many household devices, such as cameras that detect faces. But although research continues on these brain-inspired systems, they don't seem to be about to gain what we would consider to be intelligence.

A more useful approach might be to consider what we actually want from artificial intelligence. A robotic helper for the aged and infirm? A humanoid robot that can assist us in everyday life? Machines that can perform difficult or dangerous tasks? Back in the early 1980's the Japanese 5<sup>th</sup> generation project aimed to produce a machine that that could solve real (that is, social and political) problems using a top-down approach: clearly this didn't happen. But what is missing? Truly intelligent machines need to be able to understand their environment, and to have a degree of autonomy. I do not believe that intelligence can be separated from an environment: nor do I believe in abstract, disembodied intelligence. This means that intelligence requires interaction with an environment, and if that is the real world, then understanding of that environment really matters. Simple tasks, like walking in

rough terrain, or picking up an object, require understanding – how one picks up a log is different from how one picks up a cup of tea, or an egg. Often it is this everyday knowledge that is particularly difficult to give to a machine. Boston Dynamics (in the US) has developed a sophisticated walking machine, called Big Dog, that can carry loads over rough terrain: but even it still lacks autonomy.

Autonomy is the ability to make decisions independently. Rule-based systems can attempt this, but the world is an unpredictable place, and decisions need to be made rapidly. This unpredictability makes generating a usable set of rules an enormous (and perhaps impossible) task. Systems need to have aims, or goals, but translating these goals into the low level actions that might eventually fulfil the goals is difficult, and needs detailed understanding of the world in which these goals are to be achieved. Unsurprisingly, much of the research in this area has military funding, and there are even groups working on ethics and rules of engagement for autonomous military robots. (It's also worth noting that the current military robots, such as drones are generally not autonomous, but are remotely controlled: this is quite different).

This brings us back to the question of imbuing machines with a human-like intelligence. What is it that makes brains intelligent? Human brains consist of a number of parts: there is the brainstem which enables very fast reactions like reflexes. The midbrain underlies some slightly slower reactions, as well as the very basic integration of the different sensory systems. And then there is the cortex, the part of the brain that gives us our more deliberative capabilities, allowing us to analyse and predict our environment. This seems to consist of many different but similar modules which learn to pick up regularities in their world. The large cortex in humans does seem to be at the root of what we often think of as intelligence, but it does not exist in isolation: it needs the brainstem and midbrain structures that both provide it with appropriate pre-processed information, and enable it to act. Can we build a system with this type of architecture? Perhaps the hardest part is the integration of these very different elements that seem to me to underlie what we think of as intelligence.

But even if we did build this, we would still lack machines that show creativity or insight. At the Artificial General Intelligence conference, in Oxford, UK, in December, Professor Margaret Boden of Sussex University spent an hour's lecture discussing creativity, and essentially concluding that we didn't know what its basis was. Perhaps insight comes from synchronisation across many cortical modules, but this is only a suggestion.

So, is Artificial Intelligence (or rather Artificial General Intelligence, (AGI), to distinguish it from Artificial Intelligence that arises purely from clever programming, and for example plays games like chess) really finally arriving? AGI, with all the connotations of human-like intelligence still seems a long way off. On the other hand we are seeing a gradual increase in "intelligent" semi-autonomous systems that can cope with gradually more complex environments. Both the military and

organisations for caring for the increasing aged population are very interested in research underpinning these types of AGI. To conclude, human-level AGI is not here yet, and not even on the horizon, but systems that use knowledge of their environment as context, and can generate a degree of autonomy from their goals are starting to appear.