Artificial Intelligence – Contemporary Views

In this occasional paper various views are expressed about advances in artificial intelligence since computing pioneer and artificial intelligence (AI) theorist Alan Turing postulated the idea in his 1950 paper, *Computing Machinery and Intelligence*. These articles produced by the BBC mark the centenary of his birth in 1912.

Alan Turing: The experiment that shaped artificial intelligence

Noel Sharkey, professor of artificial intelligence and robotics at the University of Sheffield and co-founder of the International Committee for Robot Arms Control - 21 June 2014 BBC News

“In this fourth article of the BBC series, his influence on AI research and the resulting controversy are explored. Despite advances in computer technology, scientists have not been able to create a "thinking machine" that can pass the Turing Test.”

“Alan Turing was clearly a man ahead of his time. In 1950 at the dawn of computing he was already grappling with the question: "Can machines think?" This was at a time when the first general purpose computers had only just been built.”

“The term artificial intelligence had not even been coined. John McCarthy would come up with the expression in 1956, two years after Alan Turing's untimely death. Yet his ideas proved both to have a profound influence over the new field of AI, and to cause a schism amongst its practitioners.”

Knocking down naysayers

“One of Turing's lasting legacies to AI, not necessarily a good one, is his approach to the problem of thinking machines. He wrote: "I have no very convincing arguments of a positive nature to support my views." Instead, he turned the tables on those who may be sceptical about the idea of machines thinking, unleashing his formidable intellect on a range of possible objections, from religion to consciousness.”

“With so little known about where computing was heading at this time, the approach made sense. He asserted correctly that "conjectures are of great importance since they suggest useful lines of research".

“But 62 years on, now that we have advanced computers to test, it seems wrong that some proponents of AI still demand the onus be put on sceptics to prove the idea of an intelligent machine impossible.”

“The philosopher Bertrand Russell ridiculed this type of situation, likening it to asking a sceptic to disprove there is a china teapot revolving around the sun while insisting the teapot is too small to be revealed.”

“This can be seen as wrong-footing the scientific process of hypothesis testing and evidence collection.”

The Imitation Game

“In fact, Turing well understood the need for empirical evidence, proposing what has become known as the Turing Test to determine if a machine was capable of thinking. The test was an adaptation of a Victorian-style competition called the Imitation Game. It involves secluding a man (A) and woman (B) from an interrogator who has to guess which is which by asking questions and studying written replies. The man aims to fool the interrogator, while the woman tries to help him.”

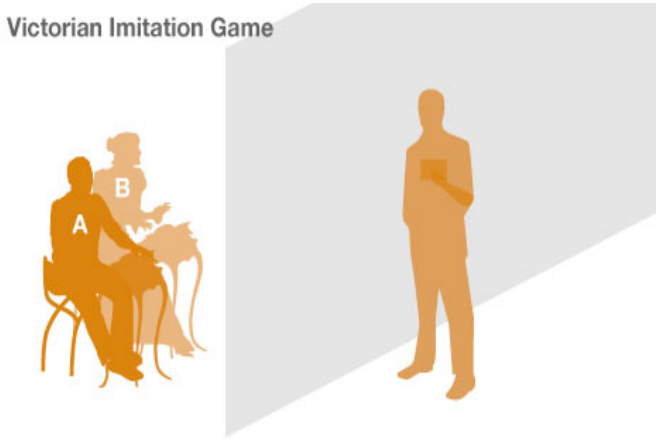

“In the Turing test, a computer program replaces the man and studies whether an interrogator can determine which is computer and which is human; although Turing did not explicitly say that the interrogator should be told that one of the respondents was a computer it seems clear to me from his example questions that this was what he intended.”

“Turing's test replaces the man with a computer running a program designed to deceive the questioner about its true identity. Will he still be able to determine which is the woman?”

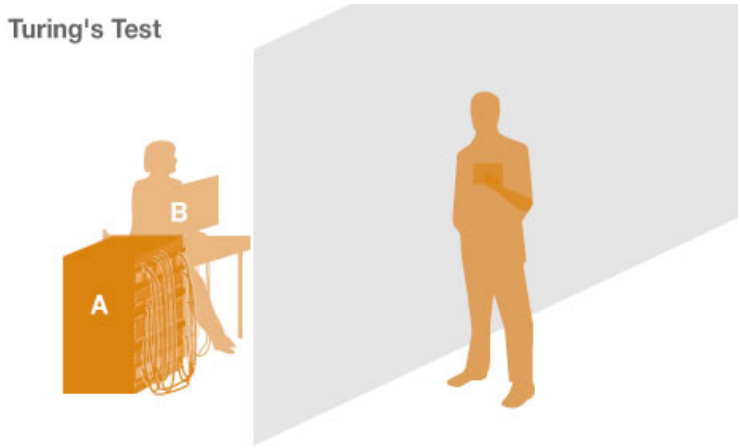

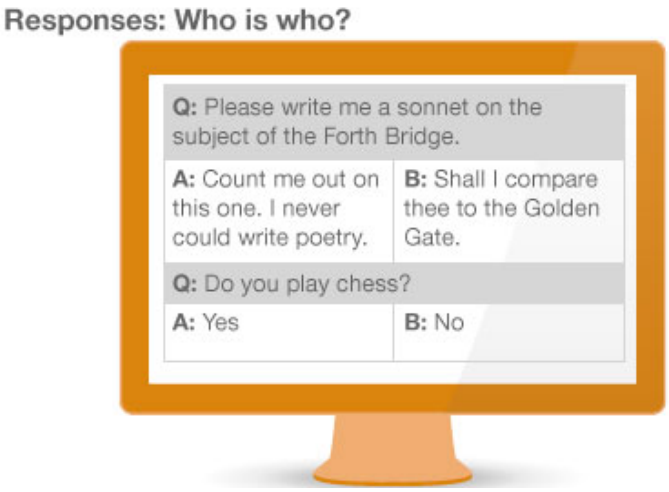

“The idea was that if the person asking the questions could not tell the difference between human and machine, the computer would be considered to be thinking and have artificial intelligence.”

Failing the test

“Turing suggested that by the year 2000 the average interrogator would have less than a 70% chance of making the right decision after five minutes of questioning.”

“My iPhone has more than 500 times the storage capacity he thought would be required and orders of magnitude more processing power, yet passing the test still seems a long way off.”

“In 1990, New York businessman Hugh Loebner set up the annual Loebner Prize competition with a prize of $100,000 (£63,500) to the creator of a machine that could pass the Turing Test. Loebner Prize judges have five minutes to ask questions to determine which respondent is a computer and which a person.”

How the Turing Test inspired AI

“The AI aristocracy strongly supported the contest until it became clear how badly the machines were performing. Now in its twenty-second year, no machine has come even close to winning.”

“Marvin Minsky, one of the fathers of AI, wrote in 1995: "I do hope that someone will volunteer to violate this proscription so that Mr Loebner will indeed revoke his stupid prize, save himself some money, and spare us the horror of this obnoxious and unproductive annual publicity campaign."

“No-one in AI seems to take the failure of the Turing Test as an argument against the possibility of thinking machines. Because Turing talked about a future "imaginable machine", some of the proponents say that we will have them at a future date. But others now argue that the Turing Test simply is not the best way to measure machine intelligence.”

AI moves on

“Despite the failure of machines to deceive us into believing they are human, Turing would be excited by the remarkable progress of AI. It is still flourishing in so many spheres of activity, from robots investigating the progress of climate change to computers running the world's finances. I expect that Turing would have danced for joy in 1997 when Deep Blue defeated world champion Gary Kasparov at chess.”

“I can also imagine him cheering in the wings of the TV game show Jeopardy when in 2011 the program Watson beat the two best human opponents in the history of the American game. It is difficult to tell how any of these achievements would have been possible without the continued inspiration from Turing's original and radical ideas. In my opinion, the Turing Test remains a useful way to chart the progress of AI and I believe that humans will be discussing it for centuries to come.”

Stephen Hawking warns artificial intelligence could end mankind

Rory Cellan-Jones Technology correspondent: BBC News Technology - 2 December 2014

“Professor Stephen Hawking, one of Britain's pre-eminent scientists, has said that efforts to create thinking machines pose a threat to our very existence.”

“He told the BBC:"The development of full artificial intelligence could spell the end of the human race." His warning came in response to a question about a revamp of the technology he uses to communicate which involves a basic form of AI. But others are less gloomy about AI's prospects.”

“The theoretical physicist, who has the motor neurone disease amyotrophic lateral sclerosis (ALS), is using a new system developed by Intel to speak. Machine learning experts from the British company Swiftkey were also involved in its creation. Their technology, already employed as a smartphone keyboard app, learns how the professor thinks and suggests the words he might want to use next.”

“Professor Hawking says the primitive forms of artificial intelligence developed so far have already proved very useful, but he fears the consequences of creating something that can match or surpass humans. Stanley Kubrick's film 2001 and its murderous computer HAL encapsulate many people's fears of how AI could pose a threat to human life.”

"It would take off on its own and re-design itself at an ever increasing rate," he said. Cleverbot is software that is designed to chat like a human would.”

"Humans, who are limited by slow biological evolution, couldn't compete, and would be superseded." But others are less pessimistic. "I believe we will remain in charge of the technology for a decently long time and the potential of it to solve many of the world problems will be realised," said Rollo Carpenter, creator of Cleverbot.”

“Cleverbot's software learns from its past conversations, and has gained high scores in the Turing test, fooling a high proportion of people into believing they are talking to a human.”

Rise of the robots

“Carpenter says we are a long way from having the computing power or developing the algorithms needed to achieve full artificial intelligence, but believes it will come in the next few decades.”

"We cannot quite know what will happen if a machine exceeds our own intelligence, so we can't know if we'll be infinitely helped by it, or ignored by it and sidelined, or conceivably destroyed by it," he says. But he is betting that AI is going to be a positive force.”

“Professor Hawking is not alone in fearing for the future. In the short term, there are concerns that clever machines capable of undertaking tasks done by humans until now will swiftly destroy millions of jobs. In the longer term, the technology entrepreneur Elon Musk has warned that AI is "our biggest existential threat".

Robotic voice

“In his BBC interview, Professor Hawking also talks of the benefits and dangers of the internet. He quotes the director of GCHQ's warning about the net becoming the command centre for terrorists: "More must be done by the internet companies to counter the threat, but the difficulty is to do this without sacrificing freedom and privacy."

“He has, however, been an enthusiastic early adopter of all kinds of communication technologies and is looking forward to being able to write much faster with his new system. But one aspect of his own tech - his computer generated voice - has not changed in the latest update. Professor Hawking concedes that it's slightly robotic, but insists he didn't want a more natural voice. "It has become my trademark, and I wouldn't change it for a more natural voice with a British accent," he said. "I'm told that children who need a computer voice, want one like mine."

Does A1 really threaten the future of the human race?

Rory Cellan-Jones BBC Technology correspondent: BBC News Technology – 5 December 2014

Cellan-Jones, referring to the Stephen Hawking interview, questions “how imminent is the danger and if it is remote, do we still need to worry about the implications of ever smarter machines?”

“My question to Professor Hawking about artificial intelligence comes in the context of the work done by machine learning experts at the British firm Swiftkey, who have helped upgrade his communications system. So I talked to Swiftkey's co-founder and chief executive, Ben Medlock, a computer scientist with a Cambridge doctorate which focuses on how software can understand nuance in language.”

“Ben Medlock told me that Professor Hawking's intervention should be welcomed by anyone working in artificial intelligence: "It's our responsibility to think about all of the consequences good and bad", he told me.”We've had the same debate about atomic power and nanotechnology. With any powerful technology there's always the

dialogue about how do you use it deliver the most benefit and how it can be used to deliver the most harm."

“Medlock says take any speculation that full AI is imminent with a big pinch of salt. He is however sceptical about just how far along the path to full artificial intelligence we are. "If you look at the history of AI, it has been characterised by over-optimism. The founding fathers, including Alan Turing, were overly optimistic about what we'd be able to achieve."

“He points to some successes in single complex tasks, such as using machines to translate foreign languages. But he believes that replicating the processes of the human brain, which is formed by the environment in which it exists, is a far distant prospect: "We dramatically underestimate the complexity of the natural world and the human mind, "he explains.”Take any speculation that full AI is imminent with a big pinch of salt."

“While Ben Medlock is not alone in thinking it's far too early to worry about artificial intelligence putting an end to us all, he and others still see ethical issues around the technology in its current state. Google, which bought the British AI firm DeepMind earlier this year, has gone as far as setting up an ethics committee to examine such issues.”

“DeepMind's founder Demis Hassabis told Newsnight earlier this year that he had only agreed to sell his firm to Google on the basis that his technology would never be used for military purposes. That, of course, will depend in the long-term on Google's ethics committee, and there is no guarantee that the company's owners won't change their approach 50 years from now.”

“The whole question of the use of artificial intelligence in warfare has been addressed this week in a report by two Oxford academics. In a paper called Robo-Wars: The Regulation of Robotic Weapons, they call for guidelines on the use of such weapons in 21st Century warfare.”

"I'm particularly concerned by situations where we remove a human being from the act of killing and war," says Dr Alex Leveringhaus, lead author of the paper. He says you can see artificial intelligence beginning to creep into warfare, with missiles that are not fired at a specific target: "A more sophisticated system could fly into an area and look around for targets and could engage without anyone pressing a button."

“But Dr Leveringhaus, a moral philosopher rather than a computer scientist, is cautious about whether there is anything new about these dilemmas. He points out similar ethical questions have been raised at every stage of automation to the arrival of artillery, allowing the remote killing of enemy soldiers to the removal of humans from manufacturing by mechanisation.”

“Still, he welcomes Stephen Hawking’s intervention.” Dr Leveringhaus contends “We need a societal debate about AI. It’s a matter of degree.”

“And that debate is given added urgency by the sheer pace of technological change. This week the UK government has announced three driverless car pilot projects, and Ben Medlock of Swiftkey sees an ethical issue with autonomous vehicles.”

"Traditionally we have a legal system that deals with a situation where cars have human agents," he explains. "When we have driverless cars we have autonomous agents. You can imagine a scenario when a driverless car has to decide whether to protect the life of someone inside the car or someone outside."

“Those kind of dilemmas are going to emerge in all sorts of areas where smart machines now get to work with little or no human intervention. Stephen Hawking's theory about artificial intelligence making us obsolete may be a distant nightmare, but nagging questions about how much freedom we should give to intelligent gadgets are with us right now.”

Pondering the implications of artificial intelligence

Oxford Dictionaries define artificial intelligence as an ability to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making and translation between languages.

Arguably some of these functions appear to require (only) the application of complex algorithms rather than the use of heuristic methods to discover or learn something by using other means. Examples might include intuition, judgement and conjecture. At present we appear only to have the Turing Test but the a major problem is not establishing robust criteria but applying a burden of proof sufficient to exclude extraneous factors that might skew the findings.

Another method might be used and that is to programme two identical computers with precisely the same sophisticated software but introduce randomized functions and behaviour patterns that do not give a predictable result. The problem then is how can we be sure it replicates human thinking that in itself creates difficulty? We may not think the same, especially where judgement is required and where we need to transcend decision-trees and ordered processes such as ranking or rating.

A stage beyond this is replicating the functioning of the human brain by use of MRI and similar techniques or directly by transplantation. This takes us into a very dark world indeed. That we might shrink in horror does not mean somewhere in the world such experiments are not being conducted or may be in the future.

Should we ignore such sinister thoughts and say it is impossible to create artificial intelligence that can replicate the human brain and achieve an intelligence level way beyond that of humans? How would we know? I am minded of Karl Popper’s theory of falsification in that prevailing wisdom and science holds good until and unless it can be disproved. The jury is still out and likely to remain so for some while yet.

Stuart Sherring
Group leader Bournemouth U3A December 2014