Faith and Science Collaborative Research Forum
Redeeming Technology 2017-18

The Business of Artificial Intelligence

Keith C.C. Chan
Department of Computing
The Hong Kong Polytechnic University
Hung Hom, Kowloon
HONG KONG
Email: keith.chan@polyu.edu.hk

Synopsis

Artificial Intelligence (AI) is concerned with making computers perform tasks that otherwise can only be done with human intelligence. Such tasks include learning, problem solving, knowledge processing, uncertainty management, natural language understanding, vision and speech recognition, etc. Since AI became an academic discipline in 1956, it has gone through several periods of boom and bust. It began with a boom period with a lot of attention and funding for research. However, this period ended in the mid-seventies when AI ran into several fundamental problems relating to the rather limited computing power at that time. Computing hardware was just not powerful enough to address the combinatorial explosion problems caused by the need to consider too many possible solutions when dealing with more realistic applications. This led to funding for AI research being reduced and eventually withdrawn.

It took until the early eighties that AI experienced its boom years again when such AI techniques as expert systems and fuzzy logic, etc., had matured and were considered potentially useful. The Japanese government's announcement of a *Fifth Generation Computer Systems* (FCGS) project to develop a massively-parallel computing platform for AI added fuel to the heat and AI received a lot of attention once again. Unfortunately, the recovery period was relatively short-lived. By the early nineties, it was concluded that AI was not able to meet commercial needs and the highly parallel computer architecture that the Japanese researchers were trying to build was surpassed in speed by less specialized hardware developed by Intel and Sun Microsystems.

At the end of the first decade of the 21st century, important factors that affect the development of AI started to emerge and AI entered another boom period again. The most notable

such factor is the advances in computer hardware and communication. Directly and indirectly as a result, machine learning techniques have greatly improved in accuracy and efficiency, to the extent that AI can now solve many real-world problems better than human beings.

Machine Learning is concerned with making computers learn by programming them with how, rather than what, to learn. In other words, computers are programmed with learning procedures rather than to obtain specific learning outcomes. The learning procedures that they are programmed with can either take a *symbolic* or a *connectionist* approach. Methods that adopt the symbolic approach are developed based on the assumption that many aspects of intelligence can be achieved by manipulating symbols. The symbolic approach spawned a lot of research in Cognitive Sciences and had led to significant advances in the understanding of cognition. In terms of application, the symbolic approach had led to such successes as IBM's Deep Blue which won the match again the world chess champion in 1997. There has been much development in machine learning based on the symbolic approach for the last several years due to advances in computer hardware and big data infrastructure. New machine learning algorithms that make use of Apache Hadoop, which is the open source framework for distributed processing, to process very large data sets, have been used to successfully learn useful patterns in very large data sets for prediction or other applications.

Despite the development of these symbolic machine learning algorithms, when it comes to even very simple tasks related to perception, such as face recognition or speech recognition, etc., the symbolic approaches do not usually work well. This is because the logic for performing such tasks is not known and programming computers to perform such tasks as symbol manipulation is not easy. It is for this reason that the connectionist approach is proposed. Such non-symbolic approach does not manipulate a symbolic representation to find solutions to problems. Instead, as human beings can perform such perception tasks relatively easily, the connectionist approach attempts to simulate neuro-biological processes of the human brain. Based on the process that human neurons take to process action potentials, a mathematical model called the artificial neural network (ANN) has been proposed. The most remarkable thing about the ANN is that one does not need to know any logic about human perception to program a computer to perform such tasks. What we need to do instead is to feed a ANN with many images or years' worth of speech samples to train it to allow it to figure out for itself how desired objects, words, or sentences should be recognized.

Unlike machine learning methods that adopt the symbolic approach which can be made to easily explain why a certain conclusion is reached and what the reasoning steps are, the reasoning process of those methods that are based on the connectionist approach are not understood. Nevertheless, the availability of powerful computer hardware, including the latest GPU development and high-speed communication, have driven significant ANN progresses in image and video processing, text analysis, and even speech and facial recognition, etc. Despite the black-box, over the past four years, there have been quantum leaps in the development of a wide range of everyday technologies that are developed based on ANN. For example, the speech-recognition functions on our smartphones work much better than they used to. We can now just talk to such devices as Amazon's Alexa, Apple's Siri, Microsoft's Cortana, or the many voice-responsive features of Google. Machine translation and other forms of language processing have also become far more convincing. Google Translate now renders spoken sentences in one language into spoken sentences in another for 32 pairs of languages, while offering text translations for 103. Other than translation, there have been advances in image recognition. There are applications that let users search or automatically organize collections of photos with no identifying tags. You can ask to be shown, say, all the ones that have dogs in them, or snow, or even something fairly abstract like hugs. Companies are expected to roll out products that can generate sentence-long descriptions for the photos in seconds. Better image recognition is expected to unleash improvements in robotics, autonomous drones, and, self-driving cars.

Programmers can only expose ANN to terabytes of data to train it, and then allow the computer to figure out for itself how to recognize desired objects, words, or sentences. Such a black-box approach is why the learning outcomes are unpredictable and this has caused concern and controversies among tech celebrities. The most notable among them are Prof. Stephen Hawking and Elon Musk. At the end of 2014, Prof. Stephen Hawking told the BBC that full AI could spell the end for mankind. According to him, "Humans, who are limited by slow biological evolution couldn't compete, and would be superseded." Other than Prof. Hawking, Elon Musk, CEO of SpaceX and Tesla, also believes it's highly likely that AI will be a threat to people. has also warned of AI being of fundamental risk to the existence of human civilization…" Hawking, Musk and other AI scientists signed an open letter to caution people the risks of AI and at the same time, promise to ensure AI research benefits humanity.

We may not face imminent extinction but many scientists have been concerned about the intended and unintended side effects of AI. They worry about accidental creation of a super-intelligent AI and forget to program it with a conscience and they also worry about people actively trying to create AI for immoral, criminal or malicious purposes. Given that AI has much accelerated, it will not only create impacts on the society but also on our Christian faith.

One of the most notable issues among those that are relevant to the Christian faith is the recent creation of a new religion based on an AI god, the others are related to the possibility of dealing with a robot that has a soul and wanting to become a Christian. There have also been discussions related to AI being the Anti-Christ of the last days, and what machine consciousness means for Christianity. When the goal of AI is to develop computers that exhibit human intelligence, there has been discussion among Christians and non-Christians about what it means to be human and whether or not the gift of salvation can be given to a robot if AI achieves consciousness. The effort to differentiate human from AI can weaken our view of humanity when we create definitions of thinking, emotions or creativity. The question of whether or not AI can worship becomes relevant to some people. In addition, one may also ask how thinking about AI may help us respect the fullness of each other's personhood and how technology development can be shaped by the full diversity of human experience. In addition, it may be worth asking where in our lives have automated technologies deepened or extended our worship and how AI is shaping ideas of sin, justice, freedom, and forgiveness. Can Christianity clearly imagine a human identity that is not diminished if machines are able to match or exceed our abilities in producing, creating, and serving?

While these issues are theological in nature, there are also more practical issues related to the use of AI to communicate our faith. Can AI find use in counselling and spiritual care. Can AI transform care for mental and spiritual well-being? Is it possible for large-scale data to be collected on people's intimate thoughts/actions, interpret people's emotional and spiritual conditions, systems that coordinate counselling and emergency response given that Facebook uses AI to assess users' risk of suicide from what they write, say or how our friends respond. Can AI be used to reduce costs of mental health counselling by intelligently assembling personalised therapies. Can AI be used to measure spiritual well-being by tracking people's mobile phones and communications, monitor and intervene on a wide range of spiritual and moral issues. If so, we have to ask questions related to what theologies of privacy and accountability could guide an era

where our intimate spiritual, emotional, and prayer lives are observed and judged by automated systems? What might Christian teaching on personal transformation offer a society that increasingly relates to each other based on predictive modeling? Just as Christians have pioneered crisis telephone lines, how can Christian ministries use AI systems to heal the brokenhearted?

The development of humanoid robots also raises issues which include, for example, what actions we should take when robots become subjects of fantasies about power and sex and if it is possible that humanoid robots may cause human to turn away from human contact. In addition to these practical impacts, there are other potential social impacts which include mass unemployment due to many jobs being replaced by AI. As a result, this may cause societal damage to people's purpose and sense of self-worth and Christians may have to be prepared to cope with the consequences. While AI can take over jobs, they may also transform jobs. The question that needs to be addressed is then how AI will transform work, creativity, and purpose.

Machine learning based on predictive analytics on personal data, news and social media analytics may affect ranking for donation, thereby transforming charity, humanitarian relief, and international development, and also deciding who should or should not be eligible for loans, humanitarian response, etc. Decisions can be made without understanding of fairness, and with limited accountability. news and social media analytics affect ranking for donation. AI trading based on "fake news". Financial access to poorest people based on credit score prediction. How much of our prayers and generosity are shaped by news and social media trends? What priorities might Christians offer to an AI that decides how to allocate resources? How might AI help churches & charities predict and respond to emerging needs? What ideas of fairness might Christian thought bring to the governance of AI systems?

Advances in computer hardware and in deep learning have made surveillance and law enforcement system very effective. Human behavior can now be observed easily so that alleged wrongdoing can be detected and predictions can be made for law enforcements. AI systems have been used to direct police patrols, prompt investigations of domestic violence. In fact, predictive sentencing has been adopted by judges in the US as AI learn from historical data. The concern, of course, is related to whether or not unjust discrimination may be reproduced. When human justice is delegated to AI, the question becomes what Christians can offer a world that expects machines to predict our moral futures.

My research into AI help me discover new understandings that reflect wisdom and creativity of our Creator. I am amazed at how the human brain functions and how humans reason with intelligence, I appreciate more and more the psalmist's exclamation, "I will praise You, for I am fearfully and wonderfully made; Marvelous are Your works, And that my soul knows very well." (Ps. 139:14). My research applying AI and machine learning to bioinformatics and computational biology allow me to understand why the DNA has been described as the Language of God and why Bill Gates of Microsoft says that "DNA is like a computer program, but far, far more advanced than any software we've ever created." According to Prof. Alan Perils, the winner of the first Turing Award, "A year spent in artificial intelligence is enough to make one believe in God." If you a Christian, you are invited to put some effort into understanding AI.