John Wyatt examines artificial intelligence and the future of healthcare

THE MACHINE WILL SEE YOU NOW

key points

- Artificial intelligence has the potential to benefit healthcare in developing countries, however, maximising shareholder value seems to be one of the driving forces.
- Christians need to affirm what it means to be human, made in the image of God, and what it means to be machine.
- AI and robotics will bring about many ethical questions, which Christian healthcare professionals should start a conversation about.

Does it matter if 'compassion' or 'friendship' is simulated, a product of clever programming? Who is harmed?

As a child growing up in the 1960s I was an avid reader of science-based predictions of the future. By the year 2000 hover-cars, colonies on the moon, free power through nuclear fusion and robot workers were all forecast. The problem for humankind would be how to fill the endless hours of leisure time once limitless resources and energy were on tap.

Sadly reality turned out to be somewhat different. Yet nobody in the 1960s predicted mobile phones, personal computers, or the all-pervasive nature, power and global spread of the Internet, let alone virtual reality and cloud computing.

What we do know is that computer technology has been doubling in power every 18-24 months since the 1960s and this rate of improvement is predicted to continue. Massive investments amounting to trillions of dollars are currently going into AI and robotics, not just in the US but in Japan, China, India, Europe and across the world. AI is already having a significant impact on our lives, but much of this is invisible. Behind the scenes they are supervising our computer searches, selecting the adverts we see online, trading in financial markets, flying commercial aircraft and translating books.

The fundamental driving force behind these remarkable changes is *laissez faire* capitalism – maximising shareholder value. Intelligent automation optimises speed and productivity whilst minimising expense. AIs can work 24/7 without getting tired, they don't demand pay raises, are less likely to make mistakes and their function doesn't deteriorate with repetition. Rather, they are constantly learning and improving on the job – increasing accuracy and efficiency. Above all, unlike human workers, they are infinitely scalable. Once you have one intelligent machine performing a task effectively you can very rapidly expand. So although the speed with which automation will enter healthcare is debatable, the ultimate direction seems clear – the inexorable logic of the market economy will ultimately triumph.

It is said that computer giant IBM has invested over a billion dollars in healthcare applications for their powerful AI system Watson. The system is capable of extracting and analysing information from free text: thousands of unmodified patient records, genomics databases and the entire scientific and medical literature. The system is currently being used to assist experienced oncologists in the USA, providing diagnosis and treatment options. Watson is constantly learning and improving its accuracy from collaboration with experienced clinicians. One physician was quoted as saying: 'One of the amazing revelations was how much like a learned colleague the system can be'. The Watson system 'understands' natural language and speech and uses the context to determine meaning.

IBM claim that a system that has been trained with experienced US oncologists can be used by a junior doctor working in a developing country to obtain the same degree of accuracy in diagnosis and treatment decisions. This kind of technology has enormous potential in increasing access to expert knowledge and advice across the world. The technology will provide a 'democratisation' of expert medical knowledge that was previously the domain of a few highly-paid specialists. But vast sums have been invested in the development of these AI systems. They have high commercial value in privatised health systems and will concentrate economic power in a small number of extremely wealthy high-tech companies.

DeepMind, a British AI company owned by Google, is collaborating with a number of UK hospitals developing algorithms to interpret head and neck scans at UCL Hospital and retinal images

at Moorfields Eye Hospital. The system learns how to identify potential abnormalities within the scans, and how to recommend the right course of action to a clinician. As AIs become pervasive within healthcare systems, issues of legal control, privacy, copyright and responsibility for malpractice are likely to become more problematic.

The ubiquity of smartphones together with tiny smart sensors and cloud computing will enable sophisticated personal health data tracking. It's claimed that this technology, combined with Big Data and cloud-based expert systems, will allow early diagnosis and continuous monitoring of many medical conditions. For instance, tissue glucose readings can be obtained and analysed in real-time in the cloud, giving real-time warning of hypoglycaemic risk and individualised advice on appropriate management. Some believe that in future many medical consultations will take place using smartphones, home sensors and AI systems with access to NHS-wide data and continuous machine learning from user inputs. The AI could decide what questions to ask, and then use a Bayesian statistical approach to give a probabilistic diagnosis and recommend therapies. Human clinicians may only be involved if the systems cannot solve the problem.

In the field of caring for patients it is very likely that different forms of robotics and AIs will increasingly be seen as providing human-like companionship. Systems are being developed to recognise human emotions using powerful face and speech recognition software and to respond in real-time to these emotions. These systems can be virtual – existing purely as an avatar (a human-like form on a screen), or a disembodied voice like Amazon's Alexa. But they may also be in a physical and embodied form, for instance as a 'cute' child-like robot.

It seems likely that companionship systems will be promoted as technological solutions for 'caring', providing 24 hour supervision and 'friendship' for the elderly, the disabled, babies and infants, those with mental health problems and maybe ordinary people who feel lonely or isolated. Some have emphasised that social robots are meant to partner with humans and should be designed to 'support human empowerment'. When used correctly, it is argued that social robots can even be a catalyst for human-human interaction.

Christian responses

Much of this may seem like science fiction and it is certainly true that it is difficult to distinguish between hype and reality. But there's a clear drive for increasing automation and behind these developments there are deep philosophical and cultural processes. In particular, there seems to be a progressive blurring and merging of concepts in what it means to be human and what it means to be a machine.

Blurring human and machine

Modern academic disciplines, such as cognitive psychology and computational neuroscience, use advances in AI as a means of understanding how the human brain works. The fundamental concept, put at its crudest, is that the brain is a 'computer made out of meat'. The more we understand how computers work, the more we can understand how the human brain works. This approach has been strikingly successful, leading to major advances in brain science and cognitive psychology.

Yet, we are increasingly trying to comprehend what it means to be an intelligent machine in the light of our human experience. Our strong and inbuilt tendency is to *anthropomorphise* intelligent machines – attributing human-like characteristics to them. We instinctively think of them as looking out at the world as we do, having a 'self' with intentions and goals. One troubling aspect of anthropomorphism is that it is not under conscious control; our response is instantaneous and deeply emotionally engaging. Some months ago I was with a group of senior church leaders who were debating the theological implications of AI and robotics. We visited a computer lab in which a group of small child-like robots were active – speaking, waving and moving around on the floor. Instantly the atmosphere changed – people waved back, laughing and responding, and a senior bishop got down on his hands and knees and started engaging delightedly with the robots, as though they were precious and vulnerable children.

The irony is that our very humanity makes us open and vulnerable to manipulation by human-like machines. The aim of many AI and robotics designers is to encourage anthropomorphism – not so much in making the physical appearance indistinguishable from human, but in simulating characteristics such as emotional intelligence and responsiveness, memory, humour, and a sense of personal identity.

1. Understanding and critiquing of modern technology

A common understanding of technology sees it merely as a neutral tool, like a hammer which can be used equally for good or evil. But in reality the power, widespread reach and hiddenness of advanced technology in our lives means that it changes and manipulates the world we see. Technology generates a 'reality distortion field'.

On the one hand modern technology can be seen as a fulfilment of the creation mandates given to the first humans by the Creator: 'Be fruitful… fill the earth and subdue it'.[1] As such we need to celebrate the extraordinary achievements and promise of digital technology in modern healthcare. But we can't be naïve about the hidden power plays, and the invasion of privacy. So before we accept new and powerful technological innovations in medical practice, perhaps we need to ask detailed questions about transparency, vested interests, privacy issues, and potential dehumanising consequences. As CS Lewis put it: 'Man's power over nature turns out to be power exerted by some men over men'.[2]

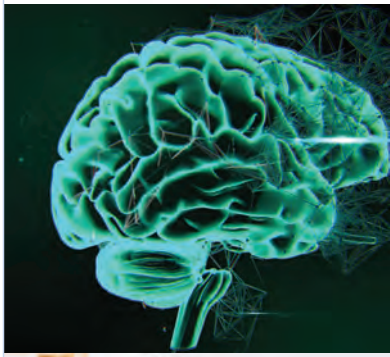

We need to celebrate the extraordinary achievements and promise of digital technology in modern healthcare. But we can't be naïve about the hidden power plays.

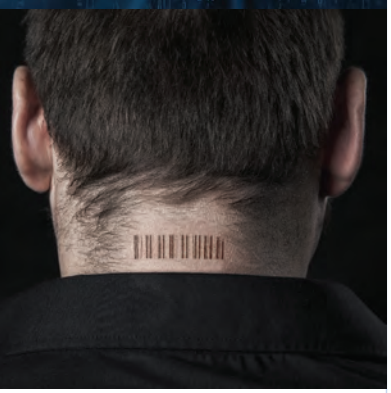

The early chapters of Genesis see both the positive life-enhancing impact of technology in the flowering of metal-working and instrument-making, and its malign counterpart in the profound and mysterious story of the tower of Babel. Perhaps this ancient story has a fresh relevance in a world where computing technology is providing a new and powerful global language: '*The Lord said, behold they are one people and they all have one language and this is only the beginning of what they will do*'.[3]

2. Resisting the conceptual blurring between our created, embodied humanity and intelligent machines

It seems inevitable that AI technology will become increasingly effective at simulating aspects of human intellectual, emotional and relational behaviour. Technology will be capable of providing physical and virtual companions, colleagues, teachers, therapists, carers and playmates. But this will raise complex and troubling issues. Think of an elderly person with dementia feeling lonely, confused and abandoned. She finds a wonderful new companion who makes her feel loved and cared for: a strange friend who seems compassionate, thoughtful, wise and fun to be with. Her mood improves and she becomes more interactive and engaged with her environment. Simultaneously, the robot companion is covertly recording and analysing all of the patient's behaviour and sending it to a central control facility. Does it matter if 'compassion' or 'friendship' is simulated, a product of clever programming? Who is harmed? If simulated companions can allow elderly people to remain in their own homes rather than being admitted to a care facility, wouldn't this be acceptable? What ethical values should be implanted in autonomous caring systems?

It also seems inevitable that AI, virtual reality and robotic technology will enable people to act out sexual and violent fantasies. Should an adult with paedophilia be allowed to abuse a child robot? After all, who is harmed? Is it possible to torture a robot? To what extent am I damaging my own humanity when I mistreat a human-like machine?

Behind these developments lies a conceptual and emotional blurring between the human person and the intelligent machine. It clearly is true that there are aspects of our humanity, including our thinking processes, which are *machine-like*. But to understand ourselves as though we are machines is a new and subtle form of idolatry. It is to worship the products of human ingenuity in place of the Creator. In Christian thinking human beings are unique in the cosmos because they are created in God's image, as embodied reflections of another reality. And the goodness of our embodied humanity is vindicated and reinforced in the miracles of the incarnation and resurrection, when God himself takes on our humanity and is raised as a physical, recognisable and touchable human being.

In the ancient words of the Nicene Creed, the Church Fathers developed the profound understanding that Christ was *begotten* not *made*. Christ was not part of the creation; he was the only begotten Son of the Father. That which we *make* is a product of our will and is ours to control. Children we *give birth* to are a gift, a product of our nature, and equal to us in dignity and significance. However sophisticated the machines we develop, they are still products of our will to be controlled and used for good.

3. We need to develop resilience to the dehumanising and manipulative possibilities of technology

We are so immersed in technology that it is almost impossible to comprehend its all-pervasive nature and influence on our lives and on our practice as health professionals. Whilst we look forward to increasingly powerful diagnostic, therapeutic and caring opportunities, I think it will be important for Christians to develop techniques of resilience and resistance to the dehumanising and manipulative possibilities of technology in the world of healthcare. Perhaps we will need to develop arguments in favour of real human carers rather than simulated ones, real human relationships in favour of simulated compassion and real experiences in place of virtual reality. At the very least Christian people need to start having the conversation as to how we might respond to these startling developments.

And in the face of apocalyptic fears about a future world in which malign technology has taken over, we need to remind ourselves that God remains the supreme Lord of history. The biblical narrative from original creation to the new creation is still underway. Human artefacts and technology have a part to play in his purposes, but ultimately it is God himself who will bring in the reality of the New Jerusalem, seen in the final chapters of Revelation, and in the final destruction of Babylon, the home of all deceptions, idolatries and counterfeits.

In the face of apocalyptic fears about a future world in which malign technology has taken over, we need to remind ourselves that God remains the supreme Lord of history.

John Wyatt *is Emeritus Professor of Neonatal Paediatrics at UCL and Senior Researcher at the Faraday Institute for Science and Religion, University of Cambridge.*

He is currently leading a multidisciplinary research project into the social, philosophical and religious implications of advances in artificial intelligence and robotics, based at the Faraday Institute.

references

1. Genesis 1:28
2. Lewis CS. *The Abolition of Man*. Oxford: OUP, 1943
3. Genesis 11:6