



FIRST PRINCIPLES FOR MEDICAL AI

William P Cheshire looks at the emerging use of artificial intelligence (AI) and machine learning in medicine, the great opportunities it presents to improve treatment and care, and the pitfalls and dangers we should be aware of as Christian health professionals.



William P Cheshire
is a neurologist at
Mayo Clinic in
Jacksonville, Florida

The history of medicine is punctuated by novel innovations. Consider the impact of anaesthesia, antibiotics, the hypodermic syringe, the stethoscope, X-ray imaging, MRI, artificial ventilators, pagers, pacemakers, organ transplants, gene sequencing – the list goes on. The first use of each of these inventions revolutionised the diagnosis and treatment of the sick. Technology turns the pages of the book of medical progress.

Prophets of the technosphere predict that the next wave of extraordinary progress lies with artificial intelligence, or 'AI'. This refers to the simulation of human intelligence by machines programmed to mimic human thought patterns such as learning, adapting to new data, problem-solving, executing complex tasks, and even social interaction – without being assisted by humans. Enthusiasm for its

transformative potential has led to predictions that AI will usher in a fourth technological revolution comparable to the neolithic transition to agriculture, the industrial revolution utilising mechanised production, and the digital revolution based on computer processing of digital information.

medical AI

What AI might mean for medicine is a subject of exciting speculation. Inspired by accelerating change, the CEO of Google has predicted that, within the next five years, 'we will shift to a world that is AI-first'.¹ The challenge for physicians will be to calibrate medical AI in such a way as to place the interests of patients first. With that in mind, a responsible approach to making decisions about medical AI must not be limited to technique but should thoughtfully explore the ethics of human interaction with technology.

ways to think about technology

The bioethicist Michael Slesman offers a helpful, ethical framework for thinking through decisions about the use of technology.² He proposes four categories, which he labels as sentimentalism, messianism, pragmatism, and responsibilism. How one navigates these categories and balances their concerns draws from prior assumptions about what it means to be human and a thinking, morally responsible being.

technological sentimentalism

In one corner are those who instinctively question or resist technological innovation. They may feel threatened by technological trends that erode kindness in communication or seem to treat people as things to be manipulated. Finding it increasingly difficult to unplug from electronic connectivity, they feel that technology deprives them of a more natural way of inhabiting the world. They see technology as the source of a host of social problems, and thus novel technologies such as AI as a Pandora's box of potential new problems. Recognising that AI has the potential to magnify biases in data, they worry that AI could worsen unjust disparities and social inequalities, placing the poor and vulnerable at a greater disadvantage. Sentimentalists value being satisfied with what one has and lament the loss of what is nostalgically imagined as a simpler, less technological age. Sentimental rejection of medical AI, however, would deny us its benefits.

technological messianism

At another corner are those who regard technology as intrinsically good and desirable. The potential useful benefits of AI stagger the imagination, especially as medical progress relies increasingly on digital technologies. The enormous amount of data in every person's genome and the accumulating scientific content in medical journals exceed the capacity of the human brain to recall, interpret, or synthesise all there is to know about health and disease. AI promises to augment our ability to assess this 'megadata' and bridge the information gap. An already familiar application of AI is natural language processing to generate clinical notes efficiently from speech. In medical research, machine-learning algorithms are mining vast clinical data sets to detect previously unrecognised disease patterns. Artificial neural networks utilising deep learning are beginning to analyse radiographic, histologic, and morphologic images with stunning diagnostic accuracy. Such advances may lead to earlier detection of disease when it is more easily treatable.

We can trace optimism in technological progress to the European Enlightenment, which sought to improve the human condition by applying rational knowledge and invention to fulfil material needs.

Optimism in technology can overlook unintended bad consequences, and it flies off the rails if taken to the extreme by those who would place ultimate hope in machinery as the saviour of society.

The uncritical adoption of technology sometimes appeals to the technological imperative, which is the philosophical claim that technology in all its achievable forms is inevitable. The technological imperative bows to this supposed inevitability and reduces to the belief that if it can be done, it should be done. Such fatalism is seriously mistaken because it abandons ethics by denying the role of human decision and responsibility.

technological pragmatism

The prevailing attitude toward technology today is pragmatism, which regards usefulness as the ultimate measure of truth and meaning. From a pragmatic perspective, AI is a transactional tool for solving human problems by plugging predicted benefits and risks into a mathematical equation, digesting more and more data to calculate potential consequences far into the future. Would such hyper-pragmatism ultimately deliver human flourishing?

Thinking about thinking technology brings this question to the forefront. Projects to construct a machine intelligence superior to human intelligence presuppose that thought is reducible to matter and its physical causes and random collisions. If intelligence were nothing more than neurones firing and neurochemicals churning in a particular pattern, then, in principle, intelligence equivalent to or surpassing the human mind could be replicated in electronic circuits. Some philosophers go so far as to suggest that if machines can mimic human intelligence, then more powerful machines will eventually exceed it and potentially replace it. The Oxford bioethicist Julian Savulescu speculates that, '*Humans may become extinct...We might have reason to save or create such vastly superior lives, rather than continue the human line*'.³ One rightly shudders to consider the implications of such a view for medicine.

Whereas pragmatism seems objective, its blind spot is the assumption that empirically verifiable facts are the only valid knowledge. If that were true, then ethical problems could be resolved by acquiring a sufficient amount of measurable factual data and subjecting that data to intense machine analysis. In reality, facts are necessary but insufficient for ethical analysis. As the volume of clinical data increases, the variables that can interact increase, and the number of potential outcomes enlarges exponentially. Having more facts is not enough. Ethical decisions depend on considering what is meant by 'useful,' which requires assigning to facts not just numerical value but moral value. Lacking or ignoring knowledge of moral right and wrong, pragmatism and computer programs that emulate it descend into moral relativism.⁴




SCAN FOR MORE

key points



- AI will transform healthcare in many ways, some as yet unforeseen, and differently from other technologies in that it imitates human intelligence.
- While customary ethical approaches to technology can guide us, the author argues that responsible use of AI must affirm and preserve the human aspect of medicine.
- The author challenges us to use this new technology to serve Christ and glorify God in improving the care of patients.



The Robot Will See You Now
edited by John Wyatt and Stephen N Williams, £11



The Robots Are Coming: Us, Them & God
Nigel Cameron
£7

Available online at cmf.org.uk/bookstore



references (accessed 17/02/23)

1. Schultz S. Artificial Intelligence: The Next Industrial Revolution. *Inside HPC*, 5 February, 2018. bit.ly/3S9jP7z
2. Sleasman M. Thinking Through Technology Part III. The Bioethics Podcast, *Center for Bioethics & Human Dignity*, 25 July, 2008. 8:1. bit.ly/3xuPpTX
3. Savulescu J. The human prejudice and the moral status of enhanced beings: what do we owe the gods? In: Savulescu J & Bostrom N, editors. *Human Enhancement*. Oxford: Oxford University Press 2009, p. 244.
4. Hughes A. The folly of scientism. *The New Atlantis*, 2012; 37: 32-50. bit.ly/3l8AeET
5. West J, editor. *The Magician's Twin: C. S. Lewis on Science, Scientism, and Society*. Seattle, WA: Discovery Institute Press, 2012.
6. DeCamp M & Tilburt JC. Why we cannot trust artificial intelligence in medicine. *The Lancet*, 2019; 1(8): E390. [doi.org/10.1016/S2589-7500\(19\)30197-9](https://doi.org/10.1016/S2589-7500(19)30197-9)
7. John 3:16-17
8. Genesis 1:27
9. Philippians 2:5-7
10. John 15:9, Romans 5:8
11. John 10:10
12. Richards J. *The Human Advantage: The Future of American Work in an Age of Smart Machines*. New York: Crown Forum, 2018, 195.
13. Mark 12:31, Luke 10:24-37, Galatians 5:13
14. Hippocrates. *Epidemics*. Circa 400 BC, Book I, Section XI.
15. Isaiah 55:8
16. John 14:6

This article was jointly published with *Ethics & Medicine*, Vol 37, No. 2, 2023. Early access available online at bit.ly/3my0sJX

◀ CS Lewis understood this when he warned that recognising facts while denying the realm of value undermines the authority of the ethical principles needed to justify limits on human manipulation.⁵ Unconstrained pragmatism makes it possible to justify committing any evil, provided a greater benefit is predicted for society at large.

Powerful computational technology, even if supplied with unimaginable quantities of data and outfitted with the superfast computer processing speed of AI, operates within the realm of fact. Technology, while useful, is not, on its own, the final answer to human problems. Virtue is also needed, an aspect of human intelligence that surpasses the realm of factual knowledge where AI operates, not by degree but in kind.

technological responsibility

Technological responsibility realistically accepts both the promise and peril of technology while engaging with it critically and with an awareness of what it means to be human. This perspective recognises both fact and value and distinguishes between the real and the artificial.

Technology itself is morally ambiguous and laden with the values of its designers and implementers. AI is a tool, not an agent; a machine made of matter and not a living being. Although AI might be programmed to imitate human emotion and, in some ways, to deliver care, it cannot truly care about the patient. An AI might be programmed to render responsible decisions that conform with humanly decided values, but it cannot be programmed to be responsible. Applying the attribute of trust to AI, according to physicians Matthew DeCamp and Jon Tilburt, 'is a category error, mistakenly assuming that AI belongs to a category of things that can be trusted'.⁶ AI cannot be said to have the human attributes of voluntary responsibility, moral agency, motives, or character.

Some uses of AI blur the line of moral responsibility. A signature feature of deep learning AI is what is known as its 'black box' problem, in which hidden layers within artificial neural networks are inaccessible to human users, who see only input and output and cannot know how the AI is analysing data. One might accept an AI diagnosis of cancer based on a proven track record for accurate analysis of images. But if an unsupervised AI supplied with clinical data were to advise an elderly patient to forego life-prolonging treatment without giving reasons, one would be justified in querying whether the machine's decision was ethically valid.

Technological responsibility entails a more comprehensive perspective on human nature than is available within the paradigm of pragmatism. Christian healthcare professionals affirm that humans are much more than the sum of their cells.

Every patient is a precious unity of body and spirit who is loved by God.⁷ All people are endowed with unique dignity as image-bearers of the Creator.⁸ Responsible use of technology respects this innate dignity in others, a dignity affirmed by Jesus Christ, who, though being in the form of God, took on humanity.⁹ Unlike AI, which is incapable of compassion, Jesus loves us¹⁰ and came that we might have life abundantly.¹¹

Moral responsibility for medical AI lies with those who design and implement its programming. This is why it matters what moral vision is embedded in the technology. For medicine to retain its moral integrity, AI must remain our tool and not become our master. We must ensure we retain the prerogative to override an AI-generated healthcare decision we believe to be morally wrong and harmful to patients. In harnessing the power of AI, we must be careful not to become so enchanted by its technical charm that we lose sight of the special dignity of our patients, who are bearers of God's image.

Responsible use of AI extends to an awareness of how its use may influence implicit attitudes toward others. Once we habitually converse with AI, a mere machine imitating human emotion, we must guard against thinking of other persons as objects. The philosopher Jay Richards observes, '*The greatest delusion of our age is the paradoxical penchant to deny our own agency while attributing agency to the machines we create*'.¹²

conclusion

The potential medical benefits of AI are promising. If used wisely, AI promises to be a powerful tool to assist Christian healthcare professionals in their calling to heal the sick and to love and serve those in need.¹³ However, caution also is needed to fulfil the Hippocratic imperative to '*first do no harm*'.¹⁴

The eventual entry of superintelligent AI into medicine will challenge healthcare professionals in new ways to choose whom we will serve. When we reach the limit of what we as finite humans can accomplish, will we place our ultimate trust in supercomputers or in Christ? Will we look to and empower technology, potentially allowing it to master us, or will we strive to align our use of technology with God's good purposes? Christians understand that God's thoughts are infinitely higher than anything possible by machine intelligence.¹⁵ Christians believe that Christ, not technology, is the true saviour, and his teachings are our first principles for providing medical care. Christians know that Christ, not machine intelligence, is the way, the truth, and the life,¹⁶ our only reliable source of wisdom, our rescue from disease and death, and our everlasting hope. ◉