Introduction: A computer technology perspective

Peter Robinson

What hath man wrought! The headline above David Lawrence's editorial[1] for the United States News in the immediate aftermath of the destruction of Hiroshima and Nagasaki by atomic bombs could almost apply to the world of artificial intelligence. What have we done? The same question could be asked of the technologists who pioneered the industrial revolution late in the eighteenth century. Their technology has allowed humankind to inflict climate change on the whole planet. As we see these effects 200 years later, we might be tempted to ask, "What have they done?"

The timescales are very different. It has taken 200 years to see the effects of increasing carbon dioxide concentrations in the atmosphere, and the changes are so slow that they can easily be overlooked or even ignored. The consequences of atomic warfare were immediately apparent, so the public reaction was more forceful. Artificial Intelligence has developed at a slower rate. It is 70 years since Maurice Wilkes and his team built the world's first practical computer, 55 years since John McCarthy coined the term 'Artificial Intelligence', 40 years since the Japanese launched their fifth generation computer project, 25 years since IBM's Deep Blue beat the world champion at chess, and 10 years since products using AI entered the domestic market. Has the progress been so slow that we are falling into a trap every bit as serious as man-made climate change?

Artificial intelligence, atomic power and the industrial revolution are just the latest examples of technologies that pose these questions. They go back to prehistoric times, with ploughshares being beaten into swords, wheeled carts being turned into chariots, and fire used for destruction as well as warmth. The fallen nature of mankind makes it only too easy to find unhelpful uses of new technologies. The eminent mathematician GH Hardy wrote that: 'I have never done anything 'useful'. No discovery of mine has made, or is likely to make, directly or indirectly, for good or ill, the least difference to the amenity of the world.'[2]

Hardy could not have been more wrong. His work in number theory, the purest of pure mathematics, is central to all modern cryptography used to secure communications in banking and commerce, by the military to control weapons, and in billions of mobile phone conversations every day. Few technologies are devised with malicious intent, but most technologies can be turned to malicious use. So it is with artificial intelligence.

Before embarking on the chapters that follow, it would be worth exploring the distinctions between information technology, artificial intelligence and robots. Information technology combines computing and communications to automate operations that could be undertaken by people. Computing makes them faster and possibly more accurate, while electronic communication allows them to be geographically dispersed. Banking is an obvious example. Computers automate the work of bank clerks writing in ledgers, and communication allows banking services to be provided remotely, perhaps through automatic cash machines. Satellite navigation uses computing and communications to automate the determination of a location that would previously have been undertaken using a sextant and chronometer. Neither is an example of 'artificial intelligence' in its strict sense.

---

[1] David Lawrence, founder of the United States News, writing on 17 August 1945. https://www.usnews.com/news/special-reports/the-manhattan-project/articles/2015/09/28/editorial-from-1945-what-hath-man-wrought

[2] GH Hardy, Sadlerian Professor of Mathematics in the University of Cambridge, in *A Mathematician's Apology*, Cambridge University Press, 1940.

Artificial intelligence is simply the display by a machine of any cognitive process that we would expect to be undertaken by a person. It has become more widespread in the past 10 years with dramatic increases in the processing power and memory storage of computers, combined with data that can be used to train machine learning systems. The combination of these three allows systems to be built that identify patterns in data and use them to make predictions about the real world in new contexts. An important characteristic is that the systems rely on probabilistic modelling so their outputs identify what is likely rather than anything absolute. Of course, this means that they should be treated with some caution. Banking systems have also grown to use these techniques. When a withdrawal is made from an automatic cash machine that is for an unusual amount or in an unusual location, the conflict with the system's model of a particular customer can be used to prevent fraud. The transaction may be genuine but the calculation shows it to be unlikely, which can give rise to difficulties when it is refused.

The term 'artificial intelligence' is widely used for marketing, scaremongering or, perhaps, simply to describe something that we think computers can't do yet but might do at some point in the indeterminate future. Working computer technology instead has names like language processing, speech understanding, computer vision and so on.

Practical applications of artificial intelligence are now mainly implemented using forms of machine learning, the statistical analysis of large bodies of data to characterise mathematical models that can then predict a person's probable response to novel stimuli. The analysis is computationally demanding, but computers have doubled in speed every two years from 700 instructions/second on the EDSAC in 1949 to 700,000,000,000 instructions/second on a modern (2021) desktop computer while the cost has decreased at a similar rate. Memory density has increased, and its cost decreased even more steeply over the same period. Finally, the growing use of computers in every aspect of our daily life has allowed the collection of the data needed to construct the models. Even social media have played their part, collecting information about their users that can be used to predict their behaviour.

This processing is simply mathematical modelling, but it allows computers to simulate aspects of human understanding and behaviour. Many people have confused this simulation with emerging sentience and speculate that that the machines are exhibiting nascent intelligence akin to that in humans. Taken to an extreme, this leads to the idea of 'artificial general intelligence' where the machines evolve faster than humans and become the dominant species. As the physicist Stephen Hawking wrote[3]:

> *The development of full artificial intelligence could spell the end of the human race. Once humans develop artificial intelligence, it will take off on its own and redesign itself at an ever-increasing rate. Humans, who are limited by slow biological evolution, couldn't compete and would be superseded.*

As we shall see in the following chapters , the theme of machines surpassing their human inventors has run through science fiction from ancient Jewish Golem mythologies, through Talos protecting Crete several centuries BC, Mary Shelley's *Frankenstein* in 1818, Karel Čapek's *Rossum's Universal Robots* in 1921, and the proliferation of novels accompanying the 20th century scientific revolution. This takes us to our third item: robots. A recurrent theme is of humanoid robots made to serve mankind turning on their creators. There is a fascination with machines made in the image of men, physically humanoid and inherently malicious. The character of these robots perhaps says more

[3] Stephen Hawking quoted by Rory Cellan-Jones in a BBC report from December 2014. https://www.bbc.co.uk/news/technology-30290540

about their creators than it does about technology. The fact is that most robots today are simply machines undertaking mechanical tasks that require strength, precision or attention to detail in a repetitive process.

Vint Cerf, widely regarded as "the father of the Internet", helpfully characterises robots more widely[4]:

> *In most formulations, robots have the ability to manipulate and affect the real world. Examples include robots that assemble cars (or at least parts of them). Less facile robots might be devices that fill cans with food or bottles with liquid and then close them up. The most primitive robots might not normally even be considered robots in normal parlance. One example is a temperature control for a home heating system that relies on a piece of bi-metal material that expands differentially causing a circuit to be closed or opened depending on the ambient temperature.*
>
> *I would like to posit, however, that the notion of robot could usefully be expanded to include programs that perform functions, ingest input and produce output that has a perceptible effect.*

Even with this broader definition, Cerf is not concerned at the prospect of sentient robots overthrowing the human race[5]. However, he is worried that flawed software could pose a real threat to humans. "If there are bugs in the software and some device is operating autonomously with regard to that software, the bugs can cause bad things to happen." Programmers' carelessness or incompetence has already led to catastrophic accidents. These programs are the robots that pose a genuine threat.

The Boeing 737 MAX debacle presents an interesting case study. Boeing introduced the 737 passenger jet in 1967. Over the following 50 years the design was modified to expand its capacity, fit more modern engines and improve its control systems, while retaining the the principal features of the airframe. The 737 MAX was introduced in 2017 as a fourth generation of an enormously successful design. It used a new type of engine to improve fuel economy but, unfortunately, this was too large to replace the earlier engines directly. They new engines had to be mounted differently which rendered the aircraft inherently unstable. Boeing solved this problem by adopting a technique used in the design of military aircraft where computers maintain a correct attitude. The Maneuvering Characteristics Augmentation System (MCAS) detected that the plane was pulling up and forced the pilots' controls down[6]. Poor software engineering practices allowed a faulty reading from a single sensor to override the pilot without warning. This only became apparent after two fatal accidents leading to the death of 346 people.

---

[4] Vint Cerf, President of the American Association for Computing Machinery, writing in the January 2013 issue of the Communications of the ACM.
https://cacm.acm.org/magazines/2013/1/158758-whats-a-robot/

[5] Nextgov report on an address by Vint Cerf at the Italian embassy in April 2016.
https://www.nextgov.com/emerging-tech/2016/04/vint-cerf-buggy-software-greater-threat-rogue-robots/127808/

[6] Michael Laris writing in The Washington Post, June 2019.
https://www.washingtonpost.com/local/trafficandcommuting/changes-to-flawed-boeing-737-max-were-kept-from-pilots-defazio-says/2019/06/19/553522f0-92bc-11e9-aadb-74e6b2b46f6a_story.html

Poor programming led to instability in automatic stock trading systems which exacerbated the 2007-08 financial crisis in the wake of concerns about bad debts affecting bank solvency[7].

These two examples also illustrate a second factor alongside flawed software that "can cause bad things to happen": motivation. Both Boeing and the banks were motivated by profit. Boeing wanted to retain the same pilot certification for the 737 MAX while fitting more efficient engines. This required them to retain the same basic airframe despite the new engines rendering it inherently unstable. Stability would be maintained by the automatic software. Unfortunately, indicators of failure in the sensors used by the software was a chargeable option on the aircraft, and not all airlines bought it. In the same way, stock market players were motivated by the additional margins that could be earned by automatic high frequency trading.

It is important to point out that aeroplane control and automatic trading only rely marginally on artificial intelligence. However, artificial intelligence lies at the heart of social media. The companies' motivation is financial, the sale of advertising. That requires users to continue clicking through links on their sites, which is achieved by tailoring the presentation to engage each individual user. That is achieved by inferring their interests and offering them more of the same. (Incidentally, that is why social media are so effective in amplifying people's political prejudices.) Finally, the systems encourage users to share information with their friends, so recruiting them into an ever-expanding audience for advertisements. The financial motivation is met by increasing advertising, engagement and growth[8]. Given that aim, they are very competent.

Lethal autonomous weapon systems also raise questions of motivation and the gravity of their effects raises questions of competence. They increasingly rely on artificial intelligence, delegating decisions on targeting and even tactics to computer systems. Legal decisions, medical diagnosis and medical interventions also rely on artificial intelligence. Civilian as well as military decisions may be delegated to computers. These are serious matters of liberty and even life. In her chapter on *Technology and humanity*, Noreen Herzfeld usefully cites Amish questions directed at any new technology: Does it provide tangible benefits? How does it affect our relationships?

Professions are paid well because their practitioners are expected to be well trained and to use their skills wisely. The rewards are substantial not least because the responsibilities are also substantial. Software engineering in general and artificial intelligence in particular are no different. Their practitioners are expected to hold to higher ethical standards of motivation and competence.

I conclude by returning to David Lawrence. The headline for his editorial was a deliberate play on Balaam's declaration, "What hath God wrought!" in Numbers 23:23. A modern translation would be, "See what God has done!" The context is that God's will cannot be deflected by human endeavours. But it remains the case that humans are still perfectly capable of inflicting pain on themselves by trying to deflect God's will.

Nuclear power is a clear example of a technology that has the potential to do great good as well as great harm. Nine years after the bombing of Hiroshima and Nagasaki, Lewis Straus, then Chairman of the United States Atomic Energy Commission, declared that "[i]t is not too much to expect that

---

[7] United States Senate Permanent Subcommittee on Investigations report on *Wall Street and the Financial Crisis* in January 2011.
http://hsgac.senate.gov/public/_files/Financial_Crisis/FinancialCrisisReport.pdf

[8] *The Social Dilemma*, a drama documentary released by Netflix in 2020.
https://www.netflix.com/gb/title/81254224

our children will enjoy in their homes electrical energy too cheap to meter."[9] That might have been hyperbole but nuclear power has undoubtedly been a great blessing, despite the dreadful spectre of nuclear war. The more powerful the technology, the more wisdom is required in its use.

So it is with computer technology in general and artificial intelligence in particular. Both have the potential to do great good and great harm. The same is true of robots. They serve well by undertaking dangerous, repetitive or delicate tasks and, possibly, in health care (although the risks of abrogating human responsibility to machines is a concern. But their ability to deceive could lead to great harm. Computing professionals will be at the forefront of those who make such choices, guided by their ethical principles of motivation and competence. Many other professionals will be making similar choices in their specialities. The rôle of the Christian professional is to apply those principles in pursuing God's will.

A number of writers in this volume highlight the fact that God created mankind in his image to rule over the earth and to enjoy relationships with him and with each other. We must be careful when we delegate that rule to automatic computer systems, or when we forget the importance of those relationships.

As David Lawrence's editorial concluded:

> *For at last it has been demonstrated to all of us that only by following His guidance in our daily conduct as individuals and as nations can we hope to fulfil our true mission as the children of God on earth. It is the only road left now—the road of mutual forbearance. It is the way to survival and human happiness.*

January 2021

[9] Lewis Strauss, addressing the National Association of Science Writers on 16 September 1954. https://www.nrc.gov/docs/ML1613/ML16131A120.pdf