

Move 78

The Soviet Union still existed throughout my formative years, along with a vague undefinable miasma of possible war that generated films such as *Red Dawn* and *Rocky IV*. My favourite film from that era was called *War Games*, in which, by modern day standards, laughably primitive computers came close to annihilating the world with only a teenage hacker able to prevent an accidental nuclear holocaust. The enemy in that film was an artificial intelligence programme called Joshua rather than the Soviets, with the hero preventing disaster in the end through convincing the computer that there was no possible winner and hence the “only possible winning move is not to play”. Other films expanded the concept of artificial intelligence turning against us, with the *Terminator* franchise, as well as *I, Robot*, based on the book by Isaac Asimov. The scenarios that all of these works illustrated were universally horrible, but for me at least there was absolutely nothing to fear in real life as my own experience of computing told me that it was trouble enough for a desktop to perform basic functions correctly to have enough spare time to plot the destruction of humanity.

Last year I noted the fanfare surrounding Pearse Keane’s efforts to get Google’s DeepMind to process OCT scans and learn from them, but confess I didn’t really believe it would amount to anything as I was still having difficulty sending and receiving emails and more complicated tasks such as attempting to view an actual scan could result in irreversible freezing of the screen. All nice in theory but nothing good would come of any artificial intelligence creations as the things were still much too basic and humans much too complex. All that changed when I attended Congress this year in Liverpool and heard Pearse speak in person about the technology itself. He explained in detail what advances had been made but one story in particular stood out for its sheer jaw dropping implications. The story of a game called Go.

I had previously thought of this game as being the pastime of Chinese citizens who didn’t have the advantage of proper entertainment, whereby white or black stones are placed on intersections on a grid with the simple aim of dominating the board area and encircling the other player. Pearse explained it was the most complicated game yet invented by humans with more possible game permutations than there were atoms in the universe. Various numbers came up on screen which, beyond a certain point,

meant nothing to me, though they did make chess out to be the equivalent of pin the tail on the donkey. While Deep Blue had beaten Kasparov 20 years ago at chess, Go was so complicated humans reigned supreme. Until last year. A competition between a Korean master called Lee Sedol and Google DeepMind’s AlphaGo was arranged where five matches would be played. AlphaGo won the first three matches with one move in the second match, move 37, being so wonderful and brilliant that for the first time the game was teaching the best Go players in the world new techniques. The machine had learnt so much that its intelligence was surpassing that of humanity’s best. Lee Sedol won the fourth match with another famous move, move 78, which was in fact rather a bad move but its badness somehow confused the brilliant yet artificial opponent such that it couldn’t work out what to do and lost the match. But it learnt from this and every single match since between man and machine has been lost.

Pearse eloquently described move 78 as the last gasp of brilliance from humanity and showed videos of DeepMind overnight smashing Atari video games by learning hitherto unknown tricks and Google self-driving cars navigating American streets, while the legally required designated human driver sat uselessly inside doing nothing. We were invited to think of a future whereby DeepMind can interpret OCT scans and angiograms and pretty much any piece of medical data and use complicated and hungrily self-learning self-altering algorithms, not only to diagnose disease, but to answer clinical questions that humans had not even had the gumption to think about, let alone ask. The end of us having to triage increasing numbers of referrals for possible wet age-related macular degeneration (AMD). The end of us having to sift through all the patient’s scans going back years to work out whether another injection was needed or not. The end of having to work out the most suitable drug or the regimen needed for every condition based on painstaking attention to journal articles and attendance at conferences such as the one I was currently at. Apparently, this technology could already identify skin cancers better than board certified dermatologists.

One of Karl Marx’s most famous lines is that for life to be satisfying we must “see ourselves in our work”, otherwise we become alienated drones and working loses all meaning beyond being purely economic. There were only two possibilities I could see; the most likely perhaps was for this

experiment with DeepMind technology to fail with all the associated sadness that would bring, but I saw success as being much more hideous. If this beast could think and develop for itself with the aim of doing our jobs better than us then we become technicians, drones, with no real input into anything. We become alienated from our work. I don’t want to be a designated driver for a driverless car, uselessly sitting by a computer in clinic infinitely better than me at my job and getting smarter all the time. Our role was suggested to be the human face of technology and we could communicate these decisions and plans with the patient, but that if anything made it sound worse. Like a deconstructed eye clinic liaison officer giving advice to a patient about a condition they have never treated themselves trying to second guess the decisions made by the ‘oh so clever’ computer ophthalmologist and when faced with a question they cannot answer shrugging apologetically and turning to an array of learnt communication skills instead of actual ophthalmology.

If this plan actually works it means the end of ophthalmology. The end of medicine, the end of all professions and the end of any true thinking satisfying human work. Perhaps the end of humanity itself. All because DeepMind might be much better than us at playing the game that is ophthalmology, and perhaps the game of life. Whilst Lee Sedol’s move 78 was superficially a bad move it did end up with him winning the game. Perhaps our move 78 is also a superficially bad move in that we triage the referrals ourselves rather than let a computer do it for us. We should listen to Joshua’s conclusion at the end of *War Games*. We should at least consider it. The only winning move might be not to play.

SECTION EDITOR



Gwyn Samuel Williams,

Consultant Ophthalmologist, Singleton Hospital, Swansea.

E: gwynwilliams@doctors.org.uk