PART ONE

Introduction

01

Homo technus

AI has and will continue to change many areas of human endeavour. Almost everything we do online is mediated by AI: search through Google; social media, whether Facebook, Instagram, or Twitter; buy something on Amazon; entertain yourself through Netflix – all are mediated by AI. AI now touches almost all global, online services. Perhaps the only online sector that is not yet mediated by AI is learning.

The nature of work is also being shaped by AI, not only in the automation of manufacturing and warehouses, but also in our homes, offices and services. This change in the workplace, in itself, will surely have a profound effect on *what* we learn, *why* we learn and *how* we learn. This is already changing through AI-driven, online learning.

Technological revolutions

Technological revolutions are not new. We as a species have shaped and been shaped by technology, from the first intentional use of stone hand axes to artificial intelligence. There has been a relentless rhythm to this progress.

The problem with most descriptions of this technological progress is that they focus on the *physical* technology itself, stone tools (Stone Age, Neolithic), metals (Bronze Age, Iron Age), age of steam engines, railways, mass production, computers (Industrial Revolutions). We see this when AI for learning is couched in terms of the 'Fourth Industrial Revolution' (Seldon and Abidoye, 2018), which is neither the fourth nor industrial.

A far better lens through which to look at AI in learning is not in terms of industrial revolutions, but cognitive revolutions. It is more revealing to see AI in terms of those revolutions in learning technology, such as language, writing, alphabets, printing, the internet and now AI. Our *physical* technology is underpinned, supported and created by *psychological* technology that enables its very conceptualization, design, development and delivery. The stone axe was imagined, shaped and used by minds. Cave paintings were the product of sophisticated imaginations. Clay tablets, papyrus, manuscripts and the entire technology of writing were a psychological breakthrough that externalized and archived thought for others to access. Printing gave rise to the scientific revolution, the Reformation and the Enlightenment. The internet, more specifically the web, gave us global access to knowledge. Now we have AI, the next technological leap, again a product of pure psychological endeavour.

Technology that enables learning is often overlooked when the history of technology is written. It is all too easy to focus on the physical objects. But without learning technologies, no other technologies would have developed. We are the species that 'learned' faster than the others. Our evolution as a species over the last few million years has been one of learning to adapt. It is this that has given us global dominance, allowing us to walk on the moon and reach out beyond our solar system.

Without the ability to shape stone tools we would never have avoided predators, sought out prey and become that dominant species. Prehistoric technology like pointed axes allowed us to kill, crush, scrape and cut. With bone needles we could dress ourselves efficiently; pots cook and axes chop down trees for fuel.

We are called *Homo sapiens* but our genus '*Homo*' emerged with the appearance of *Homo habilis* (handy man). They were so described because of their association with stone tools, but recent evidence has shown that tools were used by previous species. We are, more accurately, *Homo technus*, the species that uses tools and technology, both physical and, more importantly, psychological.

In addition to physical tools we are the masters of symbolic tools. It is difficult to see *language* as a tool, but if we define technology as something that exists outside of ourselves, that we create to exist outside of our minds and bodies, to enhance us psychologically and physically, then language is a technology. We create sounds that exist separate from us, travel to others across distance to be heard by others. It is the mainstay of communication, whether face-to-face, across the globe by telegraph then telephone and now face-to-face online. Voice underpins all other forms of technology and is being embedded in the powerful and personal mobile devices we have in our pockets, as well as in our homes, as the way of controlling the internet of things.

From Gutenberg to Zuckerberg, language, writing, printing, distribution, communication and sharing came together on the back of the internet in the form of the World Wide Web. The global scale and cumulative effect means it may prove more disruptive in the long term than all that went before. Some compare now to the 1930s, but truer historical comparison would be the 15th and 16th century, when printing shook the world. The internet has unleashed forces we are still struggling to understand. This deep tectonic shift on technology is still in its infancy. The same creative and destructive forces are being unleashed as were with writing and printing, and AI has given it a new impetus.

As smart AI technology emerges, technology challenges and, in some cases, supersedes human competences. We enter another unpredictable phase of technological change. This, some argue, is an existential threat. Whatever it turns out to be, it is certainly changing the very nature of work. For all its dangers, AI will therefore certainly shape, in some form, how we learn. Unlike speech, writing, printing and the internet, this is software that matches, and in some cases even more than matches, us as humans.

Daniel Dennett, philosopher and polymath, in *From Bacteria to Bach and Back*, attempts a synthesis of human evolution and AI (Dennett, 2017). Just as the Darwinian evolution of life over nearly 4 billion years has been 'competence without comprehension', the result of blind design, what Dawkins called the 'blind watchmaker', an invisible process that drove biological evolution, so cultural evolution and AI is often competence without comprehension (Dawkins, 1996). His vision, which has gained some traction in cognitive science, is that the brain is a prediction machine, constantly modelling forward. He also sees cultural evolution as the invasion or infection of the brain by memes, primarily words. These informational memes, like Darwinian evolution, also show competence without comprehension and fitness in the sense of being advantageous to themselves.

His hope is that machines will open up 'notorious pedagogical bottlenecks', even 'imagination prostheses' working with and for us to solve big problems. We must recognize that the future is only partly, yet largely, in our control. Let our artificial intelligences depend on us even 'as we become more warily dependent on them'.

Technology comes in revolutionary waves that have disruptive effects. It is combinatorial and cumulative building upon previous revolutions, not something separate from us but ultimately a dialectic or an accommodation, between it and we humans. We must also recognize that technology is almost always a double-edged sword that needs to be overseen and controlled, so that technology for good overcomes technology for evil. At heart these technologies define what we must learn and how we learn. They transform learning. We are *Homo technus*.

Culture

To understand AI in learning, one also has to dig deep into our cultural history, be almost archaeological, to uncover the historical paths that gave rise to AI. Eliezer Yudkowsky is right to warn us that 'by far the greatest danger of Artificial Intelligence is that people conclude too early that they understand it' (Yudkowsky, 2015). This means understanding where it came from and how we got here.

AI has its origins in Greek philosophy and mathematics. It has also been interpreted for centuries by *culture*: poems, plays and novels. In addition to these ancient and older origins, a more recent art form, the movies, has almost defined AI in the modern mind.

Culture can illuminate, but also mislead, and there is no more misleading cultural forms than those that deal with AI. AI has been shown to us in Western culture largely through dystopian theatre, literature, then movies, with endless re-treads of the Prometheus (Frankenstein) myth. This has distorted thinking around AI for learning, but we do have something to learn from its presentation in culture, and in movies in particular.

From Aeschylus's *Prometheus Bound* (1961), to Mary Shelley's *Frankenstein: The Modern Prometheus* (1990), the creation of a monstrous force took hold of the popular imagination, a myth fed straight into the movies in the 20th century. Asimov's novels have provided the famous Three Laws of Robotics in his short story, *Runaround* (1942), and there is a slew of modern novels about AI that have started to appear. Typical is Ian McEwan's *Machines Like Me* (2019), still stuck in the Mary Shelley Frankenstein myth, with Turing as the gratuitous Frankenstein.

Over the last 100 years, from *Metropolis* (1927) through the lens of movies, AI has largely been portrayed as dystopian and evil. AI has, in film, reflected our fears, often representing the fear of technology but also of the 'other', whatever that 'other' was at the time – the Cold War, crime, violence, helplessness, corporate greed, climate change and so on. There have been glimpses of a more sophisticated and subtler dynamic around AI, in *Blade Runner* and more recently a rush of movies around AI, as it takes hold in our lives through the internet.

There are several movie themes that have shaped the common perception of AI, primarily as robots. AI will lead to robots that will turn on us and kill us all. AI will take over the internet and kill us all. AI will fool us into thinking it is good but it is bad. This is similar to the popularity of child characters in horror movies, where our creations, our children, become our worst nightmares. These are variations of the Prometheus myth.

Technology is always ahead of cultural commentary. It will always be thus. Only now, over these last few years, as AI becomes operative in many domains, is it receiving subtler cultural appreciation and critiques, rather than robot fantasies.

Philosophy and mathematics

Technology is not a 'black box', something separate from us. It has shaped our evolution, shaped our progress, shaped our thinking and it will shape our future. There is a complex dialectic between our species and technology that is far more multifaceted than the simplistic 'it's about people not technology' trope one constantly hears.

Descartes (2013) saw the body as a machine; others see Leibniz (1989) as the true progenitor of AI, with his theory that language mirrors thought and a universal language may be written that manipulates symbols representing concepts and ideas using logic to simulate reasoning. Note also that Descartes and Leibniz made significant contributions to mathematics, in algebra, geometry and calculus, influencing AI in other, more purely mathematical ways.

In the 20th century, Sartre in *Being and Nothingness* (1956) and Heidegger in *The Question Concerning Technology* (1954) explore the place technology has in our very being. But it was Turing's speculations, modified by the likes of Searle (1980) and Dennett, that shaped speculative thinking on AI.

To understand why AI has real potential in learning, we also need to understand its two and half millennia gestation period, through *mathematics*.

The modern era of AI started in 1956, at the famous Dartmouth conference convened by John McCarthy and Marvin Minsky, whose aim was to 'to proceed on the basis of the conjecture that every aspect of *learning* [my emphasis] or any other features of intelligence can in principle be so precisely described that a machine can be made to simulate it'.

Note the emphasis on learning, the realization of Turing's vision of machines that could learn like children to become competent. There were successes, like Arthur Samuel's checkers software, but the promise was never realized and the first AI winter arrived in the 1960s. The early 1980s saw a resurgence of interest in expert systems but it got bogged down in rules-based reasoning and a second winter came. Only when probability and statistics were literally introduced into the equations, to give us deep learning, did we succeed in translation, speech and image recognition. AI, with exponential growth in processing power, data and powerful devices, can now deliver on that early promise.

From Euclid onwards, maths and algorithms were laying the foundations for what we now call 'artificial intelligence'. So AI has not sprung up out of nowhere; it has had long gestation period, 2,500 years of mathematics, logic, probability, statistics, algorithmic progress and machine learning.

We can revisit and implement that 1956 objective, with a focus on how AI can be used to not only learn itself but accelerate our learning. AI makes us rethink learning. It holds the possibility that we do not have to learn some old knowledge and skills; it may help us learn new knowledge and skills, even improve the process and speed of learning. Stuart Russell, a major figure in AI, rightly claims in *Human Compatible* that, 'With AI tutors, the potential of each child, no matter how poor can be realised. The cost per child would be negligible and that child would live a far richer and more productive life' (Russell, 2019). I think he's right. Even if he's only partly right, this is the right direction of travel. What greater social good than AI helping us to learn?

Learning technology

Our final look at AI is through the history and development of AI in learning technology. Mechanical devices have been used to teach and learn for a long time. Behaviourism spawned many such efforts. Sidney Pressey (Petrina, 2004) showed his 'automatic teacher' at an American Psychology Association meeting in 1924. Skinner's teaching machines took positive reinforcement as their learning principle. The problem with these and much of the simple behaviourist modelling is the lack of data and knowledge about what the learner was actually thinking and why. Adaptive AI goes much further with individual and adaptive data that guides the user, much like a GPS or satnav, through the learning experience. The experiences themselves can also adapt through machine learning.

The go-to paper for much of this is Benjamin Bloom's 'The 2 Sigma Problem' (1984), where he compared the lecture, formative feedback lecture and one-to-one tuition. Taking the straight lecture as the mean, he found an 84 per cent increase in mastery above the mean for a formative approach to teaching and an astonishing 98 per cent increase in mastery for one-to-one tuition.

Intelligent tutoring was attempted for decades, with limited success, most notably in PLATO (https://en.wikipedia.org/wiki/PLATO_(computer_system)). These systems started to improve as interface design, hardware capabilities and more sophisticated pedagogies emerged. These were heavily influenced by the idea of expert systems and knowledge management in the 1980s and 1990s, but were still hampered by limited hardware and, just as significantly, limited software design. These were simple 'hard-coded' conditional response systems or rule-based systems that trundled through 'rules' to capture insights and therefore decisions about what to teach next to the learner.

It is often thought that all the real progress was made in universities and education, but in practice it was in organizational learning, especially corporates, that established computer-based training (CBT) as a standard feature of the learning and development budget. Companies grew to serve the real needs in this market and learning management systems emerged to manage the learners and learning content, along with standards, such as SCORM, for passing data from learning experiences back to a database. There was little interest in certification. This was all about serving actual organizational needs in what were starting to become fast-moving businesses.

Alongside this emerged the multimedia age, with an increased range of media (audio, images and video) available through hardware and software advances. Laserdiscs could store tens of thousands of images and hours of audio and video. They were read-only devices but led to a much richer use of visual imagery in learning. CD-I and many other consumer devices were also tried but failed. Again, they all tended to have limited memory and speed. The hardware was always a rate-limiting factor.

Many of the multimedia, CBT and LMS businesses started in the 1980s and 1990s are still around today,. The sector has grown year on year, steadily taking on learning tasks that were previously thought beyond the reach of computers.

But one significant event led to increased growth in the market – the internet. At first, with slow dial-up, there was little to be gained, but as bandwidth and consumer adoption exploded, new opportunities were available for networked and more sophisticated pedagogic approaches. This gave us online access to Google, YouTube, Wikipedia, open educational resources, MOOCs and cloud-based services.

The development of online learning proceeded in fits and starts. Exponential growth in hardware meant faster, smaller, more powerful and cheaper machines. Interactive design adopted heuristics such as blended learning, the flipped classroom, scenario-based learning and simulations. The internet then gave us the power to network and cloud computing. Interface design, always dependent on hardware capabilities, developed from commands and menus to window metaphors to touch screen and now gesture and voice.

Post-internet, the ecosystem morphed, alongside these advances, into the online learning market. This has thrived and most large organizations have learning management systems, buy catalogues of content, commission content, use webinars and encourage social learning. More recently, as mobile devices became ubiquitous, mobile learning, or m-learning, became viable, putting the potential for learning into one's pocket. Most online learning content is now 'responsive' to desktops, laptops, tablets and mobile devices.

Beyond this we have the enhancement of real-world experience through augmented reality. The democratization of experience through VR is now with us, as are massive open online courses (MOOCs). All of these are, to some degree, already delivered through the smart use of algorithms *implicit* in the hardware and software. All of these offer opportunities for AI to become more pedagogically *explicit* and greatly enhance these learning channels

Then, as the devices became more powerful and data more plentiful, AI had a surge on the back of new techniques, especially in natural language processing (NLP) and machine learning. This grew out of the existence of the internet, then began to shape it.

The first great triumph in AI-led learning was in search, in its many forms – the ability to find things out, quickly and accurately. It was here that AI took centre stage in the learning ecosystem, opening up the future to irreversible, pedagogic shifts. As Google became the way we found things – including finding the answers to things, scholarly articles, places on maps, images, videos – huge productivity gains were made. In research alone, months of getting to and from libraries, walking up and down shelves of books and journals, were reduced to seconds.

Buying books, no matter how obscure, became quick and easy, accelerated by an AI-driven recommendation engine. This had a deep effect on AI on the internet as data mining, and a whole host of AI techniques, were brought to bear on the presentation of content and ads. Social media became a global phenomenon and many found these incredibly useful in terms of learning and CPD, all mediated by AI. Facebook timelines and Twitter feeds were suddenly algorithm driven.

AI-driven face recognition developed quickly as did voice to text and text to voice. NLP led to advances in search, translation and a whole host of new advances in interface design. AI also, sadly, made an appearance in plagiarism checking. It may seem odd that the only place you are likely to see AI in learning in a university is in checking on whether students cheat.

Despite all of these advances, online learning still remains stubbornly primitive, almost behavioural in approach. Even the gamification of content can be doggedly Pavlovian, with its use of external stimuli, such as scores, badges and leader boards, to motivate learners. The issue that remains is to what degree AI-driven learning can replicate an actual one-to-one tutor. What held such systems back were the limitations of the hardware, software, tools and a paucity of data. To have any real impact on reproducing the process of actual teaching, we had to wait for huge increases in processing power, memory and data. This has proved significant for those working in AI in learning, as data is the fuel that allows algorithms to function more effectively. But there are other advances that are less data hungry, in interfaces, content creation, curation, personalization, assessment and the consolidation of learning.

In organizational learning the rise of learning experience platforms (LXP) is an enabling development. Learning management systems (LMS) have their uses, especially in large organizations where enterprise software is necessary. But it has also become a bind for those that want to move on to use AI, as they lack the flexibility needed on personalized interface experiences, personalized learning paths, content delivery and the need for a more data-driven approach to learning.

The web has moved on while learning has remained stuck in a fixed world of flat delivery. Almost everything we do online is mediated by AI recommendation engines, with rapid moves into voice recognition and advanced analytics, yet learning is stubbornly static.

Newer technologies that are AI friendly, such as xAPI, allow tracking and data generation well beyond the traditional, fixed SCORM standard. Learning record stores (LRS) open up the possibility of using, not just learning, business data to finally integrate learning into the decision-making that grows the business.

The combination of algorithms, big data and computing power promises to unleash AI into all layers of the learning experience: user interface layer, learning layer and data layer.

In education, if AI can result in accelerating and reducing the cost of learning – even basic skills such as reading, writing and arithmetic – on a global scale, it will make past improvements in pedagogy look like rounding errors.

Since 2014, the Global Learning XPRIZE has offered $15 million for software that will enable children in developing countries to teach themselves basic reading, writing and arithmetic within 15 months (XPRIZE, 2019). It has to be scalable and open source. There have been two winners: the Kitkit School from South Korea and the US, and 'onebillion' from Kenya and the UK.

Conclusion

Little has been written on AI as a technology for learning, little on the pedagogy of AI, and little on its existing and potential applications to help us teach and learn. This book attempts to do just that. We have a unique opportunity. AI has suddenly become a topic of global, public concern. It is the topic of the age. Why now? Why this moment in history? Well, the technology has matured to the point where it can be delivered even to your pocket via mobile devices; the internet means that it can be delivered from the cloud and that tsunamis of data can be used to fuel AI solutions. We have also seen remarkable advances in AI techniques, not just in machine learning but across the board. At the same time the online learning market has matured so that it is ready to adopt AI solutions. At last smart solutions using smart software can be used to produce smart people, through AI-delivered teaching and learning, whether students, employees or citizens.

References

Aeschylus and Vellacott, P (1961) *Prometheus Bound and Other Plays: The Suppliants/Seven Against Thebes/The Persians*, Penguin, London

Asimov, I (1942) Runaround, *Astounding Science Fiction*, **29** (1), pp 94–103

Bloom, BS (1984) The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring, *Educational Researcher*, **13** (6), pp 4–16

Dawkins, R (1996) *The Blind Watchmaker: Why the evidence of evolution reveals a universe without design*, WW Norton & Company, New York

Dennett, DC (2017) *From Bacteria to Bach and Back: The evolution of minds*, W.W. Norton, New York

Descartes, R (2013) *René Descartes: Meditations on first philosophy: With selections from the objections and replies*, Cambridge University Press, Cambridge

Heidegger, M (1954) The question concerning technology, *Technology and Values: Essential Readings*, **99**, p 113

Leibniz, GW (1989) Dissertation on the art of combinations, in *Philosophical Papers and Letters*, pp 73–84, Springer, Dordrecht

McEwan, I (2019) *Machines Like Me*, Jonathan Cape, London

Petrina, S (2004) Sidney Pressey and the automation of education, 1924–1934, *Technology and Culture*, April, Vol 45, pp 305–30 [online] https://www.academia.edu/5065516/Sidney_Pressey_and_the_Automation_of_Education_1924_1934 (archived at https://perma.cc/99AF-35VM)

Russell, S (2019) *Human Compatible: AI and the problem of control*, Penguin Books, London

Sartre, JP (1956/2001) *Being and Nothingness: An essay in phenomenological ontology*, Citadel Press, New York

Searle, JR (1980) Minds, brains, and programs, *Behavioral and Brain Sciences*, **3** (3), pp 417–24

Seldon, A and Abidoye, O (2018) *The Fourth Education Revolution*, Legend Press Ltd, London

Shelley, MW (2009) *Frankenstein, or, The Modern Prometheus, 1818*, Engage Books, AD Classic, San Francisco

XPRIZE (2019) Website, https://www.xprize.org/prizes/global-learning (archived at https://perma.cc/CRK9-JZZP)

Yudkowsky, E (2015) *Rationality: From AI to zombies*, Machine Intelligence Research Institute, Berkeley