

Out-of-the-Box: Artificial Intelligence and professional roles

Luke Tredinnick

2017 promises to be the year that artificial intelligence (AI) moves out of film and fiction and into the workplace. While automation has become commonplace in retail and services, from the ubiquitous uptake of automatic tills to the growing influence intelligent service agents, anxiety about the impact of AI on professional roles has been gradually increasing in popular discourse. Writing in the Guardian in December 2016 theoretical physicist Stephen Hawking warned that “the rise of artificial intelligence is likely to extend this job destruction deep into the middle classes, with only the most caring, creative or supervisory roles remaining” (Hawking, 2016). A report in the Telegraph has warned that “many middle-class professionals will be outsourced to machines within the next few decades” (Knapton, 2016). Automation - once the preserve of manufacturing - is threatening to overtake the professional and managerial classes.

Anxieties about the impact of AI on professional jobs are not new. Half a century ago the father of modern computing Alan Turing predicted “great opposition from the intellectuals who were afraid of being put out of a job” and added that “once the machine thinking method had started, it would not take long to outstrip our feeble powers (1951: 475). At some stage, Turing argued, “we should have expect the machines to take control” (1951: 475). Although we are not yet at that stage, it seems inevitable that AI will have a significant impact on professional jobs over the next few decades, and may even perhaps threaten the very existence of numerous professional fields. This coming social and economic transformation has become known as the fourth industrial revolution; following on from the ages of steam, electricity, and information technology, the age of automation, machine learning, and artificial intelligence is almost upon us.

But what is artificial intelligence, in what kinds of areas is it currently being exploited, and what is coming in the near future? This special edition of *Out-of-the-Box* explores the emerging area of Artificial Intelligence, automation and machine learning in business and professional contexts. It will review the current state of play of intelligent systems, how they are making their way into commercial contexts

Artificial intelligence

Artificial intelligence (AI) is a general term that currently refers to a cluster of technologies and approaches to computing focussed on the ability of computers to make flexible rational decisions in response to often unpredictable environmental conditions. These strands of AI include natural language processing, machine learning, intelligent agents and rational decision making. Representations from popular culture, such as for example in Kubrik's *2001 a Space Odyssey*, or Chanel 4's *Humans*, have often focussed on general machine intelligence: the ability of computers to respond to their environments with the same kind of flexibility and creativity as humans. While this may be a long term goal of AI research, most current practice is focussed on addressing specific problems in specific contexts rather than in simulating general human intelligence. The age of the autonomous thinking machine is not yet here, but that does not mean that AI is not already making a significant impact on our lives, work and careers.

The road to intelligent machines has been a long one. From the very early days of computing in the 1940s the possibility of machine intelligence was recognised and discussed. Turing wrote a number of papers on the top in the 1950s, including the one that describes his test for machine intelligence

(the simulation game) (Turing, 1950). A year later he argued that by “the end of the century it will be possible to programme a machine to answer questions in such a way that it will be extremely difficult to guess whether the answers are being given by a man or a machine” (1951: 484). Some would argue that the latest generation of chatbots have accomplished that goal, but it is clear that the general machine intelligence Turing had in mind is still a distant dream. Indeed many would follow Searle (1984) and Penrose (1989) in suggesting that the computational theory of mind at the heart of strong AI is itself fundamentally naïve and flawed. Nevertheless through sixty years of research much progress has been made which has made us both rethink the difficulty of the task Turing originally defined, and what AI means.

The first significant step on the path to AI was Joseph Weizenbaum’s *Eliza* chatbot originally demonstrated in 1966 (see: Weizenbaum, 1966). *Eliza* famously posed as a psychotherapist which provided a framework for its interaction, and was convincing enough to fool some people into thinking they were conversing with an intelligent machine for a while. While ultimately limited in its application and giving the impression of a trick or novelty, *Eliza* kick-started research into natural language processing and analysis that have become integral to contemporary AI interfaces. There is a direct lineage between *Eliza* and for example intelligent customer service interfaces used by many companies today.

In 1997 IBM’s Deep Blue became the first chess computer to beat a reigning World Champion (Garry Kasparov) in a match using standard tournament conditions (see: Campbell et al, 2002). The rule-bound nature of Chess made this a limited achievement, but the victory of Deep Blue was important in part because Chess had traditionally been a test case for artificial intelligence, as in for example Turing’s 1947 “Lecture on the Automatic Computing Engine”. More recently in 2011 IBM’s Watson competed on the television quiz show *Jeopardy* against former winners, scooping the top prize (see: Gabbott, 2011). Watson’s victory marked a significant moment because of its ability to mine unstructured and structured data sets using parallel semantic parsing to find solutions to general knowledge questions posed in natural language. In 2014 organisers of a competition held at Reading University claimed the Turing Test had been passed for the first time when the chatbot Eugene Goostman convinced one third of the human judges that it was human, although this remains contested.

The achievement of Deep Blue, Watson, Eugene Goostman, and the development of AI in general over recent years has led some commentators to propose the imminence of a technological singularity. This is the projected point in time when machine intelligence will overtake human intelligence leading to runaway technological growth; according to its proponents heralds a significant threat to humanity and upheaval to global socio-economic structures. Vernor Vinge, who was in part responsible for popularising the idea of the technological singularity, argued that “within thirty years, we will have the technological means to create superhuman intelligence. Shortly after, the human era will be ended” (1993). But while many see this as an inevitable eventual outcome of computing technology, the current status of AI suggests no direct route to the technological singularity. Indeed we are nearing the end point of the timescale suggested by Vinge and the singularity seems no nearer. More than anything else the past few decades of research into Artificial Intelligence has taught us how much more complex the task of simulating human intelligence is than was ever originally anticipated, and how much further we are from this ultimate goal than was ever expected.

These developments may not have led to the kind of general purpose thinking machine of the kind Turing once imagined, but they have led to a far better understanding of the complexity of the task of simulating human cognition, and the many distinct parts of that general aspiration. As a field AI

has generally turned away from the aspiration to simulate general human intelligence, and has instead turned to more well-defined and application-rich parts of the problem. These include:

- natural language processing – particularly speech recognition which is a more difficult problem than parsing writing (essentially the ability to function in response to natural language input);
- semantic reasoning (the ability to make decisions on the basis of the meanings of data and statements);
- machine learning (the ability to improve at performing tasks on the basis of iteration), and
- intelligent data processing and retrieval (the ability to apply the to meaningful functions on data and information).

Through the improvements of these processes AI and machine learning have found important applications in a range of commercial contexts where rapid decision-making in response to large but fuzzy data sets is a priority. These include manufacturing, logistics, power generation and supply, and search applications. What is new about the current discussion of AI is not its immanent arrival, but the way in which what have been niche applications are finding more general uses in consumer and professional contexts.

AI and Machine Learning

Two fields in particular have drawn wider cultural attention over the past couple of years. The first is the rise of autonomous vehicles. Much research has been invested in automated vehicles, particularly self-driving cars and drones, in both the consumer and military context. The two leading companies involved in the consumer market this are Tesla and Google who have both been testing autonomous cars on public roads for several years; Google's autonomous cars for example have undertaken almost two million miles of fully-autonomous road testing. Mercedes-Benz, BMW and Audi all also have systems ready to deploy, and the technology is currently waiting for legislation and public opinion to catch-up. The first autonomous taxi services are set to be launched in major cities over the next few years, and driverless taxis and road haulage look like they may become commonplace within the next ten to fifteen years. Consumer automated vehicles may take a little longer.

Autonomous cars are a very good example of the current capabilities of AI and machine learning technologies, and of the kinds of significant changes that AI may bring in its wake. One the one hand driving takes places within controlled and well-understood environments (the public road system) according to well defined rules and conventions (the law). On the other hand autonomous driving is made complex by a number of factors including:

- the volume of data that is generated from on-board sensors;
- the variable quality, patchiness and fuzziness of this data;
- the complexity that arises from many independently driven vehicles occupying the same space and behaving according to their own (human or automated) perceptions of the conditions;
- the variable nature of ambient conditions such as lighting, weather, and road surface.

Many of the situations that arise in driving are neither predictable nor repeatable, but are both unique and context dependent, requiring rapid real-time decision making. In practice therefore autonomous vehicles cannot function by following prescriptive rules-based decision-making approaches - such as the typical algorithmic process - but need to respond to a dynamic, changing

and unpredictable environmental conditions on the basis of often incomplete or partial data. These kinds of tasks requiring fuzzy decisions on the basis of large volumes of nevertheless partial and incomplete data are the kinds of contexts within which Artificial Intelligence applications have found their niche. Through applying probabilistic approaches to such decisions, AI tools can improve their performance at unique tasks over time.

As an AI application autonomous vehicles are a little like computer chess: it looks like a simpler problem than it actually is, and it has taken longer than may have been anticipated for computers to overtake humans in their abilities at the task, nevertheless it is a relatively well-defined problem. Autonomous vehicles may have monopolised a lot of the public attention, but many of us are carrying the other mainstream consumer application of AI around in our pockets without necessarily realising it. And these applications illustrate some of the difficulties of expanding AI applications beyond the kinds of well-defined problems in which AI applications currently excel.

Siri on iOS, and Google Now on Android and Amazon's *Alexa* are amongst a range of products that have been marketed as intelligent personal assistants, able to respond to voice commands, understand the contexts of those commands, and anticipate user needs drawing information from a variety of places including the World Wide Web, email, calendar applications, and text messages. These applications rely on mining messy data drawn from real-world and often natural language sources. While the experience of using these services may for the time being be frustrating and limiting, they have several more important functions. In the first place they mark-out an AI space for the major tech companies. Secondly they provide a vast data source of real-world human-computer interactions that feed-into the improvement and development of Intelligent Agents both within this consumer context and beyond it. Google in particular has a business model that trades free software and services for data allowing it to improve its future software and services and (at the moment) target advertising, and its Google Now service is a key part of that.

In many ways everything that we need to know about AI and Learning Technology is wrapped-up in these personal assistant applications. They are beautifully clever and intuitive in what they do, but maddeningly dumb in their day-to-day use. Google Now can predict where you are likely to be traveling to and tell you how to get there, but currently only if it is a journey that you make regularly enough that you are unlikely to require assistance. It can predict the kind of news and information you may be interested in, but follows Amazon's recommendation engine in over-interpreting the significance of every incidental search that we make and failing to understand the contexts of these searches. And this highlights a fundamental difficulty in slotting AI into human-orientated processes: human behaviour may be largely predictable and routine, but it is the small proportion of unpredictable and irregular events with which we need assistance. That is the bit of our behaviour that is almost impossibly difficult to pre-empt. Machine Learning and AI currently works best in contexts that are well defined and limited, and this helps explain how AI will come to feature in professional and commercial contexts.

None of these systems are intelligent in the ways in which human intelligence would normally be defined and understood. They are not able to make flexible decisions in response to entirely novel or unfamiliar circumstances. Nevertheless within the contexts in which they are applied they are able to make decisions that have real and material consequences. Automated vehicles will be able to thread their way through busy streets avoiding hazards, and will in all likelihood make *better* choices about the most efficient routes reducing the environmental impact of road traffic, and to take and *safer* choices about how to respond to environmental conditions along the way.

Artificial intelligence in the workplace

Over the past year or so there have been numerous reports of AI beginning to make its way into professional and business services organisations; many of these developments echo and expand upon some of the basic AI research touched on above. In December it was reported that the hedge-fund managers Bridgewater Associates were developing software to help automate key parts of the firm's operation, including recruitment and other strategic decisions. The development is being headed by David Ferrucci, who previously led the development of IBM's Watson (Solon, 2016). In May 2016 the law firm Baker & Hostetler announced that they would be employing the ROSS system, developed by ROSS Intelligence (Turner, 2016). ROSS is an intelligent system incorporating machine learning that undertakes legal searches – a key paralegal and law librarianship role. The system has subsequently been deployed by a number of large law firms. It was also developed out of IBM's Watson system, and illustrates the ways in which natural language processing techniques can be used to mine messy data sets (such as law reports) for contextually related information.

While the law sector is currently leading the pack in the use of AI tools, it is not the only sector where AI is beginning to make an impact. Last year also saw KPMG Australia's announcement that it would be teaming up with IBM's Watson to automate some auditing functions (KPMG, 2016). In 2017 it was announced that the British National Health Service was testing an intelligent application that might replace its non-emergency advice service (Murgia, 2017). While news reports about the development of autonomous digital lawyers, doctors or accountants are exaggerated, it is clear that AI is beginning to have a significant impact in those parts of professional and paraprofessional roles that are currently most systematic and routine in nature: search; audit, and elements of due diligence. Nevertheless we should be careful about over-estimating both the scale and the speed of this transformation of professional work.

Machine learning is currently at the very peak of inflated expectations according to Gartner's influential hype cycle (2016); inevitably around the corner is a trough of disillusionment as companies come to realise their aspirations and hopes far exceed the present-day reality. While techniques drawing on AI research are starting to find their way into real-world business services and professional processes, that infiltration is still limited both in scale and in scope. AI is generally being used in contexts which are already highly automated, and depend heavily on digital tools. The fuzzy core of much commercial practice, including the very central role of human relationships in fostering and developing commercial relationships, currently remain impervious to AI. There may therefore be some comfort that AI is currently incapable of simulating the softer skills of professional practice. This, however, may be a false comfort.

Discussion around the impact of automation and artificial intelligence on professional occupations often starts from the assumption that technology will over time reproduce and replace current roles, functions, and tasks, particularly those that are most routine or mundane within professional practice. This is the automated supermarket check-out model of automation: a critical function of the organisation (checking out groceries) is delegated to largely autonomous technology (the automates till) freeing up available staff time to deal with more complex, unusual, or challenging situations and tasks. For all intents and purposes the technology is dropped-into and existing processes offsetting some staff costs. From this emerges the anxiety that our jobs will be replaced by machines.

This is how the technology is largely being deployed in the kinds of experimental contexts that we have seen above. It is also however the wrong way to think about the emerging impact and future opportunities of automation, AI and learning technologies. While the first implementations of

autonomous decision making technology in commercial professional contexts is focussing of reproducing existing functions – such as legal search or financial reporting – it is in altering the commercial contexts within which existing functions take place that AI will eventually prove to be transformational. The automated till is less significant to the transformation of retail than are the intelligent agents utilized in various digital retailing contexts. Amazon’s use of AI techniques in both its recommendation engine and Alexa personal assistant, but also in logistics contexts within its fulfilment centres, has had a more significant impact on how we shop than automation in the high street.

It is not the case that professional roles are likely to be automated either in entirety or in a piecemeal fashion that presents a future challenge for professional work, but rather that emerging intelligent technology is likely to transform the contexts in which professional occupations exist perhaps to the extent that those roles no longer have the same meaning, value or status that they once had. That is to say that lawyers, doctors and accountants may not be replaced by digital agents, but that the demand for services from traditional law and accounting firms, universities, health services, and business services firms is likely to change as technology allows their customers to furnish their needs in new ways.

In his influential book on the influence of social media *Here Comes Everybody*, Clay Shirky made the following observations about the economic and societal function of professional roles:

Most professions exist because there is a scarcity that requires ongoing management: librarians are responsible for organizing books on shelves, newspaper executives are responsible for deciding what goes on the front page... In these cases professionals become gatekeepers, simultaneously providing and controlling access to information, entertainment, communication, or other ephemeral goods... A profession becomes, for its members, a way of understanding the world. Professionals see the world through a lens created by other members of their profession... Sometimes, though, the professional outlook becomes a disadvantage, preventing the very people who have most at stake – the professionals themselves – from understanding major changes to the structure of their profession. In particular, when a profession has been created as the result of some scarcity, as with librarians or television programmers, the professionals are often the last ones to see it when that scarcity goes away. It is easier to understand that you face competition than obsolescence. (2008: 57-59).

In the professions in general the biggest challenge AI may pose in the coming decades is not from replacing professional roles, but from nibbling away at the edges of such roles, and undermining both professional identity and professional bodies of knowledge that Shirky identifies. These changes are however longer-term, and while AI and Machine Learning poses challenges for the future of professional work, it also presents tremendous opportunities for working in new ways and for developing new spheres of information work and new realms of expertise. The challenge for the information profession in the immediate future is to rethink the ways in which it exploits and complements emerging technological tools, particularly through emphasising soft skills including empathy and emotional intelligence, and their role in complementing professional practice. .

References:

Campbell, M. Hoane, A. J. & Hsu, F (2002), Deep Blue, *Artificial Intelligence*, 134 (1-2): 57-83.

Gabbott, A. (2011), IBM Computer Watsons wins Jeopardy Clash, *The Guardian*, Thursday 27th February 2011, available at: <https://www.theguardian.com/technology/2011/feb/17/ibm-computer-watson-wins-jeopardy> [accessed: 16 January 2017]

Gartner (2016), Gartner's 2016 Hype Cycle for Emerging Technologies Identifies Three Key Trends That Organizations Must Track to Gain Competitive Advantage, available at: <http://www.gartner.com/newsroom/id/3412017> [accessed: 17 January 2017]

Hawking, Stephen (2016), This is the most Dangerous Time for our Planet, *The Guardian*, Thursday 1st December 2016, available at: <https://www.theguardian.com/commentisfree/2016/dec/01/stephen-hawking-dangerous-time-planet-inequality> [accessed: 16 January 2017]

Knapton, S. (2016), Robots will take over most jobs within 30 years, experts warn, *The Telegraph*, 13th February 2016, available at: <http://www.telegraph.co.uk/news/science/science-news/12155808/Robots-will-take-over-most-jobs-within-30-years-experts-warn.html> [accessed: 16 January 2017]

KPMG (2016), KPMG invests in game-changing cognitive technologies for professional services, available at: <https://home.kpmg.com/au/en/home/media/press-releases/2016/06/ibm-watson-cognitive-technologies-investment-29-jun-2016.html> [accessed: 16th January 2017]

Murgia, M. (2017), NHS to trial artificial intelligence app in place of 111 helpline, *The Financial Times*, 4th January 2017, available at: <https://www.ft.com/content/aefee3b8-d1d8-11e6-b06b-680c49b4b4c0> [accessed: 16 January 2017]

PENROSE, Roger (1989), *The Emperor's New Mind: Concerning Computers, Minds, and the Laws of Physics*, Oxford University Press, Oxford.

Searle, John (1984), *Minds, Brains and Science: The 1984 Reith Lectures*, Harvard University Press.

Shirky, C. (20018), *Here Comes Everybody: how change happens when people come together*, London: Penguin Books.

Solon, O. (2016), World's largest hedge fund to replace managers with artificial intelligence, *The Guardian*, Thursday 22nd December 2016, available at: <https://www.theguardian.com/technology/2016/dec/22/bridgewater-associates-ai-artificial-intelligence-management> [accessed: 17 January 2017]

Turing, A. (1947), Lecture on the Automatic Computing Engine, available in Copeland, J. (ed.) 2004), *The Essential Turing*, Oxford: Clarendon Press, pp. 378-394.

Turing, A. (1950), Computing Machinery and Intelligence, available in Copeland, J. (ed.) 2004), *The Essential Turing*, Oxford: Clarendon Press, pp. 441-464.

Turing, A. (1951), Intelligent Machinery: A Heretical Theory, available in Copeland, J. (ed.) 2004), *The Essential Turing*, Oxford: Clarendon Press, pp. 472-475.

Turner, K. (2016), Meet 'Ross,' the newly hired legal robot, *The Washington Post*, 16th May 2016, available at: https://www.washingtonpost.com/news/innovations/wp/2016/05/16/meet-ross-the-newly-hired-legal-robot/?utm_term=.24d96f702a7f [accessed: 16 January 2017]

Wizenbaum, J. (1966), ELIZA—a computer program for the study of natural language communication between man and machine, *Communications of the ACM*, 9 (1): 36-45.