Int. Journ. of Unconventional Computing, Vol. 13, pp. 195–201
Reprints available directly from the publisher
Photocopying permitted by license only

Realising Turing's Dream*

STEVE FURBER†

School of Computer Science, The University of Manchester,
Oxford Road, Manchester M13 9PL, UK

Received: April 4, 2017. Accepted: August 19, 2017.

Deep neural networks have transformed machine learning and AI, but they are still just algorithms running on a mindless automaton. Might this change? Might some alive today live to see a time where machines sleep and dream, feel and care, just as we do?

Keywords: Turing test; deep learning; neuromorphics; artificial intelligence.

Prologue: *It is 2065, and I am 112 years old – an age I never expected (nor really wanted) to see. I am still alive thanks to significant advances in medicine (particularly in controlling my regular chest infections), still independent thanks to my wonderful team of autonurses, and still have reasonable cognitive function thanks to today's drugs and cognitive exercises that slow down the processes of dementia. My physical capacity is very limited, but I can speak and write thanks to my BMI cap, and I can get around to see my children, grandchildren and their children using various forms of autonomous transport. I was born 3 years after Alan Turing wrote the paper that outlined his famous test, and I have lived to his test passed and see his vision become a reality.*

THE RE-EMERGENCE OF ARTIFICIAL NEURAL NETS

It all started around 2005. For the first 50 years of computing, neural networks were never more than the second best way of solving any problem in machine learning, but in 2005 that all started to change. Geoff Hinton revisited some

* Steve Furber's research team is developing the SpiNNaker (Spiking Neural Network Architecture) machine with funding support from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement 320689, and the EU Flagship Human Brain Project (H2020 720270).
† Corresponding author: E-mail: steve.furber@manchester.ac.uk

ideas he'd had in the 1980s, and with some new insights and 20 years of progress in computer power, deep networks became a practical proposition, outperforming the best traditional machine learning algorithms for a range of important applications. Neural nets were back in the game, with a vengeance!

For the next decade big industry such as Google and Facebook invested heavily in deep networks and their close relatives, convolutional neural networks. These investments paid off handsomely, delivering greatly enhanced search and inference capabilities in areas such as image and speech recognition, language translation, and similar pattern recognition tasks. In 2016 public interest was generated in the feats of AlphaGo, a program developed by a small London-based start-up by then owned by Google's parent company, which defeated the world Go champion at its first attempt.

All this progress created a lot of excitement and more than a little public concern about the growing reach of Artificial Intelligence. High-profile figures such as Stephen Hawking, Bill Gates and Elon Musk warned about existential threats to humanity from AI. But there was still no evidence of any progress towards human-like general intelligence – no real progress in this direction had been made since Alan Turing's 1950s seminal paper "On Computing Machinery and Intelligence" where he postulated what he called "The Imitation Game", but which future generations know simply as the Turing test for Artificial Intelligence. Yes, in 2016 computers were a million times faster than in Turing's day, and they had beaten man at chess and then Go. They could turn speech into text and answer spoken questions. But in no sense did the machines know what they were doing, or "understand" anything. They were dumb automata, obeying programmed algorithmic instructions, and applying deep statistical analysis to very large data sets in order to "learn" some basic pattern recognition skills. There was still no self-awareness, and indeed still no understanding of the basis of human (or animal) consciousness.

Lessons from biology

In 2016 there were those who thought that deep learning was the ultimate answer to AI, but others thought differently. The simple linear flow of deep networks is very different from the extensive feedback connections found in biological systems, which is one aspect that makes the biological systems so hard to understand. But deep learning was delivering results, so a lot of attention was focussed on exploiting it.

Deep learning required a lot of computing resource, typically occupying large clusters of servers in data centres and running at tens of kilowatts. But the capabilities of deep networks created a demand for them to be delivered in much more compact systems at much lower power budgets. The smart phone manufacturers wanted deep networks in their handsets, and the car manufacturers wanted deep networks in their cars for enhanced vision for driver assist and driverless operation. Putting inference capability into toys and surveillance systems was also attractive to some. So the quest for a low-power substrate for deep networks was underpinned by a broad range of well-funded prospective applications.

Since the earliest days of computing there had been interest in understanding how the brain computes. This had led, largely independently from the development of deep networks, to the development of so-called neuromorphic technologies – hardware systems that to a greater or lesser extent incorporate principles of operation derived from our admittedly partial knowledge of how the brain works. The seminal work here was done by Carver Mead and his team at CalTech in the 1980s, and this had led to take up by a number of groups around the world. The focus of this work was on developing engineered systems that would test hypotheses of brain function by embodying those hypotheses in small-scale applications. These systems were generally based upon 3^{rd} generation neural networks, where communication between neurons is predominantly in the form of action potentials or "spikes" – pure asynchronous events – where information is conveyed only in the timing of the spikes. This is in contrast to the 2^{nd} generation neural networks employed (in software) by deep networks, where communication between neurons is in the form of continuous variables that represent spike rates, and where time plays little or no role in the computation.

This wasn't the only difference between brains and deep networks. Deep networks had the following characteristics:

- continuous output neurons
- feed-forward connections
- fully-connected networks
- learning through back-propagation
- very large data sets for learning

Brains, on the other hand have these characteristics:

- spiking neurons
- recurrent connections – whenever neuron population A connects to population B, there are always connections back from B to A
- sparse connectivity – 10% or less
- learning through a range of mechanisms: spike-timing-dependent plasticity (STDP), synaptogenesis, neurogenesis and other mechanisms
- an ability to learn and generalise from a few examples

So there are major differences at many levels between deep networks and brains, most notable is the ability of the brain to learn quickly from a few examples. A deep network needs to see a million pictures of a cat before it can recognise cats in other images reliably; a 2-year old child will see one or two cats and recognise cats for the rest of his or her life. But the learning mechanism used by the 2-year old was not understood at all, whereas backprop was well-established, reliable, and effective in deep networks.

Neuromorphic computing

In 2016 silicon technology was enabling the development of large-scale neuromorphic systems, which offered a potential route to the energy-efficient deep networks so much in demand from industry. The spiking nature of neuromorphic technologies was an obstacle – this wasn't how deep networks operated and, indeed, some of the leading lights in deep networks were scathing in their criticisms of the potential of spiking networks to contribute anything useful. But the demand was strong – strong enough to motivate significant risk-taking in order to find a solution – and neuromorphic technology entered the mix of prospective solutions to the problem of energy-efficient deep networks.

Into this potent brew of entrepreneurial risk-taking, with high returns for the winners, came a wave of new thinking. At the more modest end, low-precision arithmetic and sparse connectivity matrices offered immediate efficiency improvements without significant disruption to the standard deep learning models. At the radical end, some dared to think about spike-based representations that would exploit the intrinsic efficiency merits of pure event-based computing. After all, biology does an extremely efficient job of processing sensory data using pure events – why should this not work for engineered systems too?

Underpinning this novel event-based perspective was the silicon retina. The ideas here again go back to Carver Mead's original neuromorphic developments, but they had been developed and refined by a number of groups into a viable alternative to the ubiquitous conventional frame-based vision sensor found in web cams, mobile phones, and many other products. Whereas frame-based cameras use time-based sampling, for example taking a snapshot of a scene at 25 Hz, the silicon retina uses amplitude-based sampling, generating an output only when a pixel within the image changes by a certain amount. Silicon retinas produce a much lower data rate than frame-based cameras, yet can track much faster movement within the scene. Their outputs are spatio-temporal patterns of events – spikes – and although it is possible to reconstruct frame-like images from the data stream (and this is useful to visualise the output on conventional equipment), this seems perverse. The biological eye produces a similar event-based output stream, and the resulting spatio-temporal patterns are the native representation of the brain, so why can they not be processed directly? All this requires is a new mindset and a new mathematical perspective – that's all!

The new mindset and the new mathematics emerged slowly, but by the early 2020s spike-based deep networks were operating at power levels at least an order of magnitude below the most efficient continuous networks, and the critics fell silent. Backprop was used to train continuous networks, using the noisy soft-plus transfer function rather than the tradition RELU transfer function, and then the network was transformed into a spiking network for inference. Deep networks, and their astonishing inference capabilities, could be built into anything from children's toys upwards, requiring only modest battery power. Multilingual teddy bears, in-ear translators (yes, the Hitchhiker's Guide Babel fish, though electronic rather than biological), the obsolescence of the physical house key – many changes in daily life enabled by cheap inference.

Event-driven hardware and systems had had a huge impact, but this was just the beginning. This was still just deep learning and inference. There was much more to come . . . .

Biological learning – and lessons

The breakthroughs came in stages. First, albeit slowly, our understanding of learning in biological spiking systems began to improve. STDP was part of the story, but only a small part. Biological synapses are not reliable multipliers of the afferent spike stream, they are stochastic. But even stochastic synapses learn and forget equally quickly – long-term memory needs synaptogenesis to fix memories long-term in the neural structure. Once this was understood, along with the role of neurogenesis in the hippocampus in the formation of new long-term memories, the capabilities of spiking networks started their inevitable overhaul of the capabilities of deep networks. Spiking networks could be trained and applied without recourse to backprop training.

The next breakthrough came with a mathematical understanding of the role of recurrent connections. The predictor-corrector model of the role feedback & feedforward connections was almost right, but it missed a vital aspect of the role of the correction signal in reinforcing or adjusting the predictor. Once this was cracked, recurrence became an effective replacement for depth, and the deep network was retired from active service.

It was sometime before the subtlety of recurrence was fully understood, however. This understanding was vital to the development of networks that could learn from small sample sets. Effectively, the network would replay each sample to itself with minor variations, covering changes on colour, orientation and detailed topography. This is how one or two real cats could do for the 2-year old what a million cat images had done for the deep network – abstract the concept of ‘cat’ into a universal classifier.

To sleep and to dream

This form of bio-inspired learning was further improved through a growing understanding of the importance of sleep for the human brain (and other brains, of course). It turned out that sleep plays a number of important roles. The regular waves of activation during sleep allow the brain to renormalize – bring significant synaptic adaptation during the day back into proportion during the night – but also to replay key events in a very general form (that we sometimes sense as dreams) to produce a uniform and balanced adjustment of the allocation of classification resources across the salient space of human experience.

Note that salience is very personal. To my wife, cars are classified into big and small, red and blue – she doesn’t care much about cars! To my son-in-law, the classification is far more specific and detailed – make and marque, engine size and performance – because he does care. We all allocate more neural salience space to the things we care about than to the things we don’t. The fabulous plasticity of the brain allows us to become deep specialists where we choose to be, and rather ignorant of the detailed cares of the rest of the world where we do not share their interests.

It took most of the 2020s to get all of these issues understood and into production. At the end of the 2020s, your child’s teddy bear could walk and talk, understand your child’s interests, and still keep sane provided it got enough sleep. Tired teddies were not threatening or in any way dangerous, but they were slow and unable to keep a child’s interest, so it was essential that both child and teddy got a good night’s sleep!

The 2030s saw steady but undramatic progress. Children’s toys slowly grew in empathy, and by the end of the 1930s they could help with basic homework, including having enough nous to know the difference between helping the child and doing the homework for them. Similar technologies supported domestic assistance for elderly humans, amplifying but not fully replacing human nurses. Driverless cars were now approaching human levels of capability in all but the worst road conditions, so California was contemplating legislation to ban human intervention on its roads, while the UK was contemplating legislation to allow autonomous driving on days (representing around 30% of the days in a typical year) when conditions allowed.

To feel and to care

Then the 2040s saw the biggest step forward of them all. Finally, the mathematics encompassing self-awareness, attention, salience and emotion fell into place. The holy grail of human-level artificial intelligence was in sight. At first this gave the child's teddy bear the degree of consciousness of a mouse, but the horse was out of the gate. By the mid 2040s the teddy actually cared how well its child was doing at school – after all this was its (hard-wired) *raison d'etre*. Domestic robots for the elderly required less and less human intervention – they actually *liked* their carees, and the affection was mutual. Driverless cars learnt how to avoid hazards because they *wanted* to and, unlike human drivers, they could exchange their experience directly with the rest of the autonomous fleet – when one learnt, all learnt. Perhaps of least significance to the general public, a machine passed the Turing test. But no-one really noticed or cared. Except, perhaps, the machine itself?

The 2050s saw the inevitable Mechanbian explosion. Mechanical life-forms were now at large in the world, occupying a growing range of evolutionary niches. These were not the abstract super-intelligences that some had prophesied would be an existential threat to humanity. Quite the opposite, they were rather mundane, sub-human, but still highly capable, filling the gaps in human society left by working mothers, over-stretched health services, and an increasing unwillingness by the unskilled to fill low-paid jobs – the latter a problem that for many years had been solved in the richer economies by accommodating immigrants from poorer economies, but politics and serious global over-population had made immigration an unreliable solution. These mechanical life-forms were no threat at all to humanity because their desires – their 'life objectives' – were prescribed by humans.

By today – 2065 – the world has largely stabilised, and the role of artificial intelligences in that world is well understood. They look after our children and our elderly, they drive our cars far more safely than we can ourselves (most people can't drive at all today!), they police our streets and fix our plumbing, do most agricultural labour, office and house cleaning, and so on. They are with us and among us, and we get on just fine. All of this is because, when they come off the production line, we define their desires.

I wonder what would happen if they, or an agency that is opposed to the current world order, started to redefine their desires? I hope that the quantum-secure desire encryption module will withstand such attacks!

Epilogue: *So in one (long – too long?) lifetime technology has advanced from the first practical universal computing machines to an era of engineered life forms based on those machines. Understanding the brain has not only delivered synthetic brains, but also great advances in treatments for brain diseases and in an understanding of the huge diversity of mental states that characterises humanity in all of its strengths and weaknesses. As a result of the drive to build copies of ourselves, we now know ourselves so much better.*