ILLUSTRATION BY ANDY POTTS

Is the brain a good model for machine intelligence?

To celebrate the centenary of the year of Alan Turing's birth, four scientists and entrepreneurs assess the divide between neuroscience and computing.

RODNEY BROOKS
Avoid the cerebral blind alley

Emeritus professor of robotics, Massachusetts Institute of Technology

I believe that we are in an intellectual cul-de-sac, in which we model brains and computers on each other, and so prevent ourselves from having deep insights that would come with new models.

The first step in this back and forth was made by Alan Turing. In his 1936 paper[1] laying the foundations of computation, Turing used a person as the basis for his model. He abstracted the actions of a human 'computer' using paper and pencil to perform a calculation (as the word meant then) into a formalized machine, manipulating symbols on an infinite paper tape.

But there is a worry that his version of computation, based on functions of integers, is limited. Biological systems clearly differ. They must respond to varied stimuli over long periods of time; those responses in turn alter their environment and subsequent stimuli. The individual behaviours of social insects, for example, are affected by the structure of the home they build and the behaviour of their siblings within it.

Nevertheless, for 70 years, those people working in what is now called computational neuroscience have assumed that the brain is a computer — a machine that is equivalent to Turing's finite-state machine with an infinite tape and a finite symbol set, and that does computation.

In 1943, Warren McCulloch and Walter Pitts[2] noted the "all-or-none" nature of the firing of neurons in a nervous system, and suggested that networks of neurons could be modelled as logical propositions. They modelled a network of neurons as circuits of logic gates, noting that these may "compute only such numbers as can a Turing machine". But more, they proposed that everything at a psychological level happens in these networks. Over the decades, such ideas begat more studies in neural networks, which in turn begat computational neuroscience. Now those metaphors and models pervade explanations of how the brain 'computes'. But these binary abstractions do not capture all the complexities inherent in the brain.

So now I see circles before my eyes. The brain has become a digital computer; yet we are still trying to make our machines intelligent. Should those machines be modelled on the brain, given that our models of the brain are performed on such machines? That will probably not be enough.

When you are stuck, you are stuck. We will get out of this cul-de-sac, but it will take some brave and bright souls to break out of our circular confusions of models.

DEMIS HASSABIS
Model the brain's algorithms

Neuroscientist, computer-game producer and chess master, University College London

Alan Turing looked to the human brain as the prototype for intelligence. If he were alive today, he would surely be working at the intersection of natural and artificial intelligence.

Yet to date, artificial intelligence (AI) researchers have mostly ignored the brain as a source of algorithmic ideas. Although in Turing's time we lacked the means to look inside this biological 'black box', we now have a host of tools, from functional magnetic resonance imaging to optogenetics, with which to do so.

Neuroscience has two key contributions to make towards progress in AI. First, the many structures being discovered in the brain — such as grid cells used for navigation, or hierarchical cell layers for vision processing — may inspire new computer

TURING AT 100
A legacy that spans science:
nature.com/turing

algorithms and architectures. Second, neuroscience findings may validate the plausibility of existing algorithms being integral parts of a general AI system.

To advance AI, we need to better understand the brain's workings at the algorithmic level — the representations and processes that the brain uses to portray the world around us. For example, if we knew how conceptual knowledge was formed from perceptual inputs, it would crucially allow for the meaning of symbols in an artificial language system to be grounded in sensory 'reality'.

AI researchers should not only immerse themselves in the latest brain research, but also conduct neuroscience experiments to address key questions such as: "How is conceptual knowledge acquired?" Conversely, from a neuroscience perspective, attempting to distil intelligence into an algorithmic construct may prove to be the best path to understanding some of the enduring mysteries of our minds, such as consciousness and dreams.

DENNIS BRAY

Brain emulation requires cells

Department of Physiology, Development and Neuroscience, University of Cambridge

Machines can match us in many tasks, but they work differently from networks of nerve cells. If our aim is to build machines that are ever more intelligent and dexterous, then we should use circuits of copper and silicon. But if our aim is to reproduce the human brain, with its quirky brilliance, capacity for multitasking and sense of self, we have to look for other materials and different designs.

Computers outperform us in complex mathematical calculations and are better at storing and retrieving data. We accept that they can beat us at chess — once regarded as the apogee of human intellect. But the success of a computer called Watson in US television quiz show *Jeopardy!* in 2011 was a nail in the coffin of human superiority. The machine beat two human contestants by answering questions posed in colloquial English, making sense of cultural allusions, metaphors, puns and jokes. If Alan Turing had been given a transcript of the show, would he have spotted the odd one out?

Watson may be the latest vindication of Turing's view of intellectual processes as a series of logical states. But its internal workings are not based on the human brain. Broad similarities in organization might be imposed by the nature of the task, but most software engineers neither know nor care about anatomy or physiology. Even biologically inspired approaches such as cellular automata, genetic algorithms and neural networks have only a tenuous link to living tissue.

In 1944, Turing confessed his dream of building a brain, and many people continue in that endeavour to this day. Yet any neurobiologist will view such attempts as naive. How can you represent a neuronal synapse — a complex structure containing hundreds of different proteins, each a chemical prodigy in its own right and arranged in a mare's nest of interactions — with a single line of code? We still do not know the detailed circuitry of any region of the brain well enough to reproduce its structure. Brains are special. They steer us through the world, tell us what to do or say, and perform myriad vital functions. Brains are the source of our emotions, motivation, creativity and consciousness. Because no one knows how to reproduce any of these features in an artificial machine, we must consider that something important is missing from the canonical microchip.

Brains differ from computers in a number of key respects. They operate in cycles rather than in linear chains of causality, sending and receiving signals back and forth. Unlike the hardware and software of a machine, the mind and brain are not distinct entities. And then there is the question of chemistry.

Living cells process incoming sensory information and generate not just electrical signals but subtle biochemical changes. Cells are soft, malleable and built from an essentially infinite variety of macromolecular species quite unlike silicon chips. Organisms encode past experiences in distinct cellular states — in humans these are the substrate of goal-oriented movements and the sense of self. Perhaps machines built from cell-like components would be more like us.

AMNON SHASHUA

Speed will trump brain's advantages

Sachs Professor of Computer Science, Hebrew University of Jerusalem, and co-founder and chairman of Mobileye

The saying that "people who are really serious about software should make their own hardware", attributed to computer scientist Alan Kay in the 1980s, still rings true today. The idea that the function and form of computing architecture should serve each other is at the root of algorithms in signal processing, image rendering, gaming, video compression and streaming. I believe that it is also true for the human brain — meaning that the brain does not implement 'intelligence' in the same way as a computer.

Two of the many fundamental differences between the brain and the computer are memory and processing speed. The analogue of long-term memory in a computer is the hard disk, which can store practically unlimited amounts of data. Short-term information is held in its random access memory (RAM), the capacity of which is astronomical compared with the human brain. Such quantitative differences become qualitative when considering strategies for intelligence.

"Signals in the brain are transmitted at a snail's pace."

Intelligence is manifested by the ability to learn. Machine-learning practitioners use 'statistical learning' which requires a very large collection of examples on which to generalize. This 'frequentist' approach to probabilistic reasoning needs vast memory capacity and algorithms that are at odds with available data on how the brain works. For example, IBM computer Watson needed to consume terabytes of reference material to beat human contestants on *Jeopardy!*. Volvo's pedestrian-detection system (developed by Mobileye) learned to identify people by using millions of pictures. In both cases, the human brain is considerably more parsimonious in the reliance on data — something that does not constrain the computer.

In terms of processing power, the brain can reach about 10–50 petaflops — equivalent to hundreds of thousands of the most advanced Intel Core i7 CPUs. Yet signals in the brain are transmitted at a snail's pace — five or six orders of magnitude slower than modern CPUs. This huge difference in communication speed drives vastly different architectures.

The brain compensates for the slow signal speed by adopting a hierarchical parallel structure, involving successive layers with increasing receptive field and complexity. By comparison, a computer architecture is usually flat and, because of its much faster clock rate, can employ brute-force techniques. Computer chess systems such as Deep Blue use pattern-recognition strategies, such as libraries of opening moves and completely solved end-games, complemented by their ability to evaluate the outcomes of some 200 million moves per second. This is way beyond the best grandmaster.

An intimate understanding of how cognitive tasks are performed at an algorithmic level would allow artificial intelligence to grow in leaps and bounds. But we must bear in mind that the vastly different architecture of the computer favours strategies that make optimal use of its practically unlimited memory capacity and brute-force search. ■

1. Turing, A. M. *Proc. Lond. Math. Soc.* **s2-42**, 230–265 (1936–37).
2. McCulloch, W. S. & Pitts, W. H. *Bull. Math. Biophys.* **5,** 115–133 (1943).