Machine Intelligence

A brief history of machine intelligence from a British perspective

Phil Husbands, Department of Informatics, University of Sussex

Major British contributions to the development of intelligent machines can be traced back at least 350 years. The 17th Century philosopher Thomas Hobbes (1588-1679) played an important role in establishing the intellectual climate that would result in the emergence of the modern science of machine intelligence. Although today he is usually remembered as an ethical and political philosopher, Hobbes was one of the most important natural philosophers of his day. His materialist stance emphasised the machine-like qualities of nature, discussing the possible creation of artificial animals: artificial intelligences and artificial life. Hobbes attacked the separation of mind and body advocated by his contemporary Descartes, arguing that human intelligence is the product of physical mechanisms: that mind is a property of suitably organised matter.

Although Hobbe's *Leviathan* included a combinatorial theory of thinking, details of possible mechanisms for intelligence were very sketchy. It was to be some time before much progress was made in this direction: the 18th Century saw the construction of many ingenious mechanical automata, including chess playing Turks and flatulent ducks, but it wasn't until the 19th Century that the next major breakthrough occurred. This was the design of Charles Babbage's programmable Analytical Engine.

The son of a London banker, Babbage (1791-1871) was a brilliant mathematician and engineer who held the same chair at Cambridge University as had Newton. Inspired by Leibniz, whose work was in turn influenced by Hobbes, in 1821 he designed his mechanical Difference Engine for calculating accurate mathematical tables – something of enormous practical importance at the time. However, Babbage's interest in calculating machines ran deeper than the production of mathematical tables. He envisioned such engines as powerful tools for natural philosophy, hoping that their whirring cogs would shed new light on the working of nature. In this spirit, in 1834 he began work on his revolutionary Analytical Engine, a general, programmable machine. The engine was to read instructions from sets of punched cards, adapted from those used by Jacquard looms, and to manipulate partial results in its own internal memory. Rather than being designed to perform just one set of calculations, the machine was intended to be a completely general computing engine – in theory it could be programmed to perform any calculation.

In 1843 Augusta Ada, Countess of Lovelace (1815-1852) translated into English a paper on the Analytical Engine written by the mathematician Luigi Menabrea. Ada was the daughter of Lord Byron, the great poet. Her parents separated almost immediately after her birth and Lady Byron brought up Ada to appreciate mathematics and science, in part because of her own interest in these areas, but also because she hoped it would drive out any Byronic madness her daughter might have inherited. In collaboration with Babbage, Ada added extensive notes to the manuscript which make it clear that they both understood the importance of the general nature of the Engine. Ada wrote of its potential to act as a 'thinking, reasoning machine'. The notes include a detailed description of a method for using the Engine to calculate Bernoulli numbers. This is widely regarded as the first computer program, although there is some controversy over whether the primary author was Lovelace or Babbage.

The Analytical Engine was never completed --- its construction became mired in manufacturing and bureaucratic difficulties which resulted in the British government withdrawing funding -- but in 1991 a team at the Science Museum in London constructed the Difference Engine number 2 according to Babbage's detailed designs. It worked perfectly. In most respects Babbage's remarkable vision of a universal machine anticipated the modern digital computer age by more than a century.

While Babbage was struggling to construct his engines, George Boole (1815-1864), self-educated son of a Lincoln cobbler, was building a formal system of logic which went on to serve as a corner stone of all modern digital technology, but which was also intended to capture the structure of reasoning and thinking. He died after developing a fever following a soaking in a rain storm. His demise was unwittingly aided by his wife who, believing that a cure should mirror the cause, threw buckets of cold water over him as he lay shivering in bed.

All the major players mentioned above were interested in a two edged process we might call the mechanization of mind: a quest to build artificial brains, but also to reveal the underlying mechanisms of natural intelligence. Many key researchers in machine intelligence to the present day regard these as two sides of the same coin, which explains the often close links between this area and the brain sciences.

It wasn't until the mid twentieth century that machine intelligence really took off. When it did, British scientists were at the forefront of many of the key developments. Kenneth Craik (1914 -1945) was an influential figure in this flurry of progress. Craik was a Scottish psychologist of wayward brilliance whose story is made particularly poignant by his tragic and sudden demise at the age of 31 on the last day of the war in Europe. He was killed in a traffic accident while cycling through Cambridge.

After studying Philosophy at Edinburgh University, he began a PhD. in psychology and physiology at Cambridge University in 1936. His love of mechanical devices and his great skill as a designer of scientific apparatus no doubt informed the radical thesis of his classic 1943 book, *The Nature of Explanation*, published in the midst of his war work on factors affecting the efficient operation and servicing of artillery machinery. Noting that 'one of the most fundamental properties of thought is its power of prediction', Craik suggested that the human mind is a kind of machine that constructs small-scale models of reality that it uses to anticipate events. He viewed the proper study of mind as an investigation of classes of mechanisms capable of generating intelligent behaviour both in biological and non-biological machines.

At the same time as Craik was starting to develop his ideas, in another Cambridge college the mathematician Alan Turing (1912-1954) was about to publish a startling paper on one of Hilbert's open problems in mathematics, the *Entscheidungsproblem.* This asked if it was possible to define a formal procedure which could be used to decide whether any given mathematical assertion was provable. Turing's highly original approach to the problem was to define a kind of simple abstract machine. By using such a machine as a very general way of constructing a formal procedure in mathematics, he was able to show that is followed that the answer to the problem was no. The concept of the Turing machine, as it became known, now serves as the foundation of modern theories of computation and computability. In the paper Turing explicitly drew a parallel between the operation of such a machine and human thought processes. Turing also introduced a more general concept which was to have an immense practical impact: the Universal Turing Machine. This machine could interpret and then execute the set of instructions defining *any* given standard Turing machine (each of which corresponded to a particular formal procedure or algorithm). Thus, the Universal Turing Machine embodies the central principle of the computer as we know it today: a single machine that can perform any well-defined task as long as it is given the appropriate set of instructions, or program. A hundred years after Babbage, and by a very different route, Turing had envisaged a completely general super machine. This time the vision was to come to fruition.

The second world war was to prove a major catalyst in the development of practical computers and in further advances in machine intelligence. Turing was recruited into the top secret Government Code and Cypher School, based at Bletchley Park. There Turing and his colleagues Max Newman, Jack Good, Donald Michie and T.H. Flowers made many crucial contributions to code breaking including the development of special purpose electronic code cracking machines which were an important step towards the modern computer.

In Britain there was little explicitly biological research carried out as part of the war effort, so most biologists were drafted into the main thrust of scientific research on communications and radar. To many of them a radar set could be thought of as a kind of artificial sense organ, and they began to see how the theoretical framework associated with it – which focused on how to best extract information from the signal - might be applied to better understanding natural senses such as vision. Some engineers and theoreticians, working on such problems as automatic gun aiming, were influenced by their biologist colleagues and began to see the importance of coordinated sensing and acting in intelligent adaptive behaviour, be it in a machine or in an animal. There was much interdisciplinary discussion about intelligence and electronic brains among these scientists and the intense interest in the subject carried over into peace time. Influenced by Turing and Craik, a group of scientists emerged in the USA who were interested in understanding general principles underlying behaviour in animals and machines. These included mathematicians and engineers (Wiener, von Neumann, Bigelow, Shannon, Pitts) and brain scientists (Lorente de No, Rosenblueth, McCulloch). Weiner named the new enterprise Cybernetics. In Britain a parallel group formed, most of whom became members of the Ratio Club.

The Ratio Club was founded and organized by John Bates, a neurologist at the National Hospital for Nervous Diseases in London. The other twenty carefully selected members were a mixed group of mainly young neurophysiologists, engineers and mathematicians, including W. Ross Ashby, Horace Barlow, Thomas Gold, Jack Good, Donald Mackay, Alan Turing, Grey Walter and Albert Uttley. Many members had a strong interest in developing ‘brain-like’ devices, either as a way of formalizing and exploring theories about biological brains, or as a pioneering effort in creating machine intelligence, or both. Most club meetings occurred between September 1949 and July 1953. During this extremely productive period various members made highly significant contributions to machine intelligence.

At the end of the war Turing used his experience of electronic code breaking machines to put together plans for a practical implementation of his universal machine. In 1948 the world's first stored-program electronic digital computer (the type with which we are familiar today) was developed at Manchester University by a team led by Max Newman and including Turing and Jack Good. At this time Turing wrote highly influential papers, including the 1950 *Computing Machinery and Intelligence*, in which he explored the problems of machine intelligence. He proposed the famous Turing test for identifying intelligence in a machine and sketched ideas for the development of machines that would learn and adapt to their environments. Naturally, he viewed the core of such machines as digital computers.

Meanwhile Grey Walter (1910-1977), a world leader in EEG research, developed a pair of autonomous mobile robots controlled by analogue electronic nervous systems, completing them by late 1949. This was the first explicit use of mobile robots as a tool to study ideas about brain function, a style of research that has become very popular in recent times. Equally importantly, Walter's robots demonstrated an artificial neural mechanism for generating simple adaptive behaviours whereby the machines interacted with each other and returned to a 'hutch' for recharging. Walter was a flamboyant character and something of a media star which helped to propel his 'turtles' to world fame and to be exhibited at the Festival of Britain in 1951.

In the same period, W. Ross Ashby (1903-1972), a psychiatrist who at the time was director of research at Barnwood House Hospital, Gloucester, published a theory of adaptive behaviour in animals and machines which he had been developing over the previous decade. His theories were illustrated with an electromechanical device, The Homeostat, which exhibited a form of learning. Ashby became a very important figure in Cybernetics and Systems theory and his ideas are again influential in present day research in machine intelligence.

At the club's inception Horace Barlow, a great-grandson of Charles Darwin, was a PhD student in Lord Adrian's lab at the department of physiology at Cambridge University. He went on to become an enormously influential neuroscientist, particularly in the field of vision, and was one of the pioneers of using Shannon's Information Theory to understand neural mechanisms, a direct consequence of his involvement in the Ratio Club where much cross-fertilization of ideas and methods took place. Barlow's work has also become very important in machine perception and in the design of learning algorithms.

Albert Uttley (1906-1985) and Donald Mackay (1922-1987) were, among many other things, artificial neural network (ANN) and machine learning pioneers. The development of artificial neural systems, as the basis of artificial brains, has been a major topic in machine intelligence since the seminal 1943 work of McCulloch and Pitts. Several members of the Ratio Club made important contributions to advancing the state-of-the-art in this area. Uttley's work on statistical learning devices was particularly far ahead of its time.

Parallel developments in the USA also focused on biologically inspired brain like devices and included the construction of electronic ANNs, by researchers such as Frank Rosenblatt and Marvin Minsky, that were able to perform simple learning tasks. Oliver Selfridge, grandson of the founder of Selfridge's department store, had left Britain at the age of 14 to study with Weiner at MIT. In the mid 1950s he developed his break-through Pandemonium system which learned to recognize visual patterns, including alpha-numeric characters. The system employed a layered network of processing units which operated in parallel and made use of explicit feature detectors that only responded to certain visual stimuli – a mechanisms that had recently been postulated to exist in biological vision systems by Horace Barlow and which he and other neuroscientists went on to demonstrate.

In 1956 two young American academics, John McCarthy and Marvin Minsky, organized a long workshop at Dartmouth College to develop new directions in what they termed *artificial intelligence*. McCarthy in particular proposed using newly available digital computers to explore Craik's conception of intelligent machines using internal models of external reality, emphasizing the power of symbolic manipulation of such models. At the meeting Newell and Simon, influenced by aspects of Selfridge's work, demonstrated a symbolic reasoning program that was able to solve maths problems. This was the beginning of the rise of symbol manipulating computer programs in the study of machine intelligence. This more abstract, software bound, paradigm came to dominate the field and pulled it away from its biologically inspired origins. For a while the term artificial intelligence, or AI, was exclusively associated with this style of work.

The new AI movement in the USA gained significant financial and industrial support in the 1960s and for a period the field was dominated by American research. However, developments did not cease in Britain. Donald Michie, who had worked with Turing at Bletchley Park, became a geneticist of some distinction after the war but retained a deep interest in machine intelligence. In 1965 he founded what became the department of Machine Intelligence and Perception at Edinburgh University. A wide range of research was conducted there, including important work on ANNs by Christopher Longuet-Higgins and David Willshaw who were among those who kept a more cybernetic style of work alive. ANNs were by now routinely simulated on digital computers, as envisaged by Turing in the late 1940s. Significant centres of research also developed at the Universities of Sussex and Essex.

While Michie was setting up his new department, David Marr (1945-1980) completed a mathematics degree at Cambridge University and began work on a PhD in theoretical Neuroscience. After essentially founding the field of Computational Neuroscience, Marr moved to MIT where he developed an ambitious computational theory of vision that many consider redefined and revitalized the study of human and machine vision. He was diagnosed with acute leukaemia in 1978.

In 1973 the Lighthill report, commissioned by the UK Science Research Council, was highly critical of many areas of basic research into machine intelligence in Britain and funding was slashed. As a result many key British researchers followed Marr to the USA. These included Harry Barrow who, with Marty Tenenbaum, made many important contributions to machine vision and developed a model of visual processing that rivalled Marr's.

Throughout the 1970s research continued in the UK on a reduced scale. During that period Geoffrey Hinton completed a PhD in Longuet-Higgin's and Willshaw's artificial neural networks group. After post-doctoral work at Sussex University in the Cognitive Studies group established by Margret Boden, Aaron Sloman and Gerald Gazdar, Hinton moved to the USA where he was part of the team that developed the back-propagation learning algorithm for multi-layered ANNs, thus considerably advancing the field. He went on to develop other important methods and forms of networks.

A decade later funding reappeared in the UK and research took off again with some major figures returning from the USA. In this period Gerald Gazdar, widely regarded as one of the most important linguists of the twentieth century, was involved in major work in computational linguistics, helping to stimulate important work in that field and in automatic speech and text understanding.

Geoffrey Hinton and other British researchers played a central role in the development of practical pattern recognition and automatic data classification applications for ANNs. These successes were decisive in revitalising the fortunes of this style of research. Since the late 1980s, biologically inspired and sub-symbolic approaches have swept back to take centre stage. These include an emphasis on whole artificial 'creatures' that must adapt to real unforgiving environments. Their brains run on on-board digital computers as Turing foresaw more than fifty years ago. Work in machine intelligence has again become much more closely aligned with research in the biological sciences. Many of the ideas and methods developed by the great British pioneers of the mid twentieth century have once more come to the fore and current UK research is playing an important role in pushing the field forward.

Bibliography

W. Ross Ashby. *Design for a Brain*. London: Chapman and Hall, 1952.
Charles Babbage, *Passages from the Life of a Philosopher*, London: Longman Green, 1864.
H.B. Barlow. Summation and inhibition in the frog's retina. *Journal of Physiology* **119**:69-88, 1953.
H.G. Barrow and J.M. Tenenbaum, Interpreting Line Drawings as Three-Dimensional Surfaces. *Artificial Intelligence*, **17**:75-116, 1981.
George Boole, *An Investigation of the Laws of Thought, on which are founded the mathematical theories of Logic and Probabilities*, London: Macmillan, 1854.
K.J.W. Craik, *The Nature of Explanation*, Cambridge University Press, 1943.
Thomas Hobbes, *Leviathan*, London: Andrew Crooke, 1651.
G. Gazdar, E.H. Klein , G.K. Pullum , & I.A. Sag, *Generalized Phrase Structure Grammar*. Oxford: Blackwell, and Cambridge, Mass.: Harvard University Press, 1985.
Ada Lovelace, Notes on L. Menabrea's 'Sketch of the Analytical Engine invented by Charles Babbage, Esq.', *Taylor's Scientific Memoirs*, vol 3, London: J.E. & R. Taylor, 1843.
D. Marr, *Vision: a computational investigation into the human representation and processing of visual information*. SanFrancisco, CA: W H Freeman, 1982.
N. L. Collins and D. Michie (Eds), *Machine Intelligence 1*, London: Oliver and Boyd, 1967.
D. E. Rumelhart, G.E. Hinton and R.J. Williams, Learning representations by back-propagating errors. *Nature*, **323**:533—536, 1986.
O.G. Selfridge, Pandemonium: A paradigm for learning. In Blake, D., Uttley, A., eds.: *The Mechanisation of Thought Processes*. Volume 10 of National Physical Laboratory Symposia, 511--529. Her Majesty's Stationary Office, London, 1959.

A. Turing, On computable numbers, with an application to the Entscheidungsproblem, Proc. London Math. Soc. (2), **42**: 230-265, 1936.
A. Turing, Computing Machinery and Intelligence*, Mind* **49**:433-460, 1950.
A. M. Uttley. The design of conditional probability computers. *Information and Control*, **2**(1):1-24, April 1959.
G. Walter. An Imitation of Life, *Scientific American*, **182**(5), 42-45, 1950.
Norbert Wiener, *Cybernetics, or control and communication in the Animal and the Machine.* Cambridge, Mass:MIT Press, 1948.
D. Willshaw, O. Buneman and H. Longuet-Higgins, Non-holographic associative memory. *Nature* **222**:960—962, 1969.

Further Reading

M. Boden, *Artificial Intelligence and Natural Man***,** Harvester, 1977.
R. Brooks, *Robot: The Future of Flesh and Machines*, Allen Lane, 2002.
R. Cordeschi, *The Discovery of the Artificial. Behavior, Mind and Machines Before and Beyond Cybernetics*, Kluwer, 2002.
G. Dyson, *Darwin among the Machines*, Addison-Wesley, 1997.
A. Hodges, *Alan Turing: The Enigma of Intelligence*, London: Unwin, 1983.
P. McCorduck *, Machines Who Think: A Personal Inquiry into the History and Prospect of Artificial Intelligence*, San Francisco, California: Freeman, 1979.
P. Husbands, O. Holland and M. Wheeler (Eds.), *The Mechanical Mind in History*, Cambridge, MA: MIT Press.