AI in the UK: A Short History

In Wendy Hall & Jérôme Pesenti, *GROWING THE ARTIFICIAL INTELLIGENCE INDUSTRY IN THE UK*
pp.18-20

Kieron O'Hara

University of Southampton
kmo@ecs.soton.ac.uk

The pioneering British computer scientist Alan Turing is widely credited with launching and inspiring much of the development of AI, and the philosophy of AI, with his 1950 paper "Computing Machinery and Intelligence". Turing had already developed the principle of the modern computer in 1936, and played a critical role in breaking ciphers at Bletchley Park during the Second World War. In the 1950 paper, Turing explored what is meant by 'machines' and by 'thinking', and in what became known as the Turing Test, he proposed that if a machine could conduct a conversation in print that could not be distinguished from conversation with a person, it could be said that the machine was "thinking". As his earlier work on the foundations of computing, delivered to the London Mathematical Society, had shown that all digital computers were in effect equivalent (i.e. any computer can simulate the behaviour of all other computers given enough memory and time), this thought experiment expressed a very powerful, elegant and precise concept. The paper is still widely read, discussed, cited and anthologised today.

Early pioneers in AI focused on the development of the necessary tools and techniques to explore Turing's idea. Early approaches focused on symbolic programming (i.e. programs that can manipulate expressions in their own programming languages), as the most promising paradigm available. Many special purpose languages were written as part of that drive, most famously LISP in the US, but also including significant contributions from Britain, such as POP-2 (created by Robin Popplestone and Rod Burstall at the University of Edinburgh) and Edinburgh Prolog (by David H.D. Warren, also at Edinburgh).

In 1952 Christopher Strachey used the Ferranti Mark 1 system at the University of Manchester to write a programme to play draughts, and later programmed the generation of love letters. The performance of AI in increasingly more complex games has been an indicator of progress ever since. Another former Bletchley codebreaker was Donald Michie, later director of the Department of Machine Intelligence and Perception at Edinburgh. His noughts-and-crosses-playing program MENACE was too complex for the computers available to him at the time, and he initially implemented it with 300 matchboxes!

By the 1960s, AI techniques were being applied to far more complex problems with more practical applications. Planning involves developing strategies for problem-solving that produce a series of actions that will approach a goal; example applications include automated reasoning, or planning proofs, as pioneered by Alan Bundy. Understanding natural language was another important strand; for instance Karen Spärck Jones developed ways of retrieving information from documents, and Yorick Wilks' preference semantics was a computational approach to disambiguating word senses, which not only contributed to AI but directly challenged the dominant Chomskyan paradigm in linguistics. Both of these were alumni of the Cambridge Language Research Unit, a legendary crucible of computational linguistics founded by Wittgenstein's student Margaret Masterman. Robotics systems, such as Edinburgh's Freddy I and Freddy II, were able to combine vision, intelligence, versatility and physical engineering to perform tasks such as assembling objects (special-purpose AI languages needed to be developed for robotics). AI systems were influential on the discipline of cognitive psychology, as researchers such as Richard Gregory, Christopher Longuet-Higgins, Philip Johnson-Laird and David Marr realised than many human cognitive processes could be seen as a type of computation, and be modelled as computer programs.

Globally and in the UK, AI has gone through periods of development and periods of relative stagnation (often referred to as "AI winters"). One major such event followed the publication in 1973 of Sir James Lighthill's report on AI, which recommended concentrating AI funding in a smaller group of British universities. Lighthill was sceptical about the ability of AI at the time to scale up to solve the complexity of real-world problems, and indeed the dominant approach of the 1960s, to model complex reasoning as a search through a tree of possible decisions, was vulnerable to the problem of combinatorial explosion. However, in the longer run, the advances in symbolic programming enabled greater understanding of high-level problem-solving intelligence, with especial progress in tools and techniques to simulate or support complex expert reasoning in relatively well-structured domains – ideal for applications in the workplace. So-called Knowledge-Based Systems (KBS) combined AI techniques with other kinds of computing inference and domain-relative expertise to create systems for often quite mundane but important real-world applications.

The unspectacular but practical success of KBSs helped defuse Lighthill's pessimism, and paved the way for a productive expansion of funding with the Alvey Programme. In retrospect, the AI winters that we have seen have been products of excessive hype – overclaiming by boosters leading to an erroneous impression of failure, and consequent undervaluing of the important but unsung successes of the research.

The UK's Alvey Intelligent Knowledge-Based Systems (IKBS) Programme, which ran from 1983 to 1987 was developed in response to progress in other countries, in particular Japan (whose 5^{th} Generation Project rested upon techniques and languages, particularly Edinburgh Prolog, originating from the UK). Alvey influenced development of academic research and research capability but also encouraged industry applications, focusing on the practical problems in which progress had been made, notably natural language processing, interfaces, and KBSs.

These applications gradually coaxed the field of AI away from the idea of producing 'machines that think' (a concept that has always been philosophically problematic), towards the more measurable idea of creating machines that can produce performance that would certainly be ascribed to intelligence if produced by a human (an idea implicit in the Turing Test). Such intelligent performance might be produced by 'brute force' methods that neither mirrored, nor attempted to mirror, human problem-solving. Interestingly, the UK has produced many important philosophers who have helped uncover the concepts behind such distinctions, including, for example, Margaret Boden and Andy Clark.

Post-Alvey, AI funding dipped once more, but the promise of the field was already on an upturn, as new methods of programming, which did not rely on linear combinations of symbolic inferences, became feasible. Whereas symbolic programming is the easiest type of programming for a human to understand, there are also great gains to be had from simulations of natural techniques of inferring information from the perceived environment (e.g. information from the senses), that are often called sub-symbolic, because they do not include direct representations of declarative or propositional knowledge.

One example of inspiration from nature is the genetic algorithm, which encodes a program as a set of ‘genes’, and then modifies them in a way that mimics evolution, looking for a ‘fit’ with an ever-changing environment (pioneering work here includes Richard Forsyth’s BEAGLE system for pattern recognition). Another is neural nets or connectionist systems, in which artificial ‘neurons’ are connected in a system intended to work like the human brain, with the ‘neurons’ stimulating or inhibiting each other. As with symbolic AI, researchers have often rowed back from the aim of copying the human brain to improve performance (for instance, with techniques such as back-propagation, developed by Geoffrey Hinton), but the large-scale neural net SpiNNaker (2005-), led by Steve Furber, remains in the tradition of direct modelling of the brain. Other non-traditional computing approaches with relevance to AI include parallel processing (using multiple processors in parallel to solve problems), multi-agent systems (where many intelligent autonomous agents interact within an environment), and machine learning (algorithms which can learn to find significant structure in data, given training in identifying the interesting patterns).

Other countries and international companies are investing heavily in AI development, but the UK is still regarded as a centre of expertise in research and application of AI, for the present at least. Two of the founders of DeepMind, for instance, met as PhD students in University College, London’s Computational Neuroscience Unit, whose founding director is Hinton. This report recommends that more is done to build on Turing’s legacy to ensure the UK remains among the leaders in AI.