Evolving artificial intelligence

Evolutionary Robotics: The Biology, Intelligence and Technology of Self-Organizing Machines
by Stefano Nolfi and Dario Floreano, MIT Press, 2000. $50.00/£34.95 (xi + 317 pages)
ISBN 0 262 14070 5

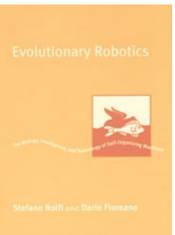

Alan Turing's 1950 paper, *Computing machinery and intelligence*[1], is regarded by many as one of the seminal works in artificial intelligence (AI). It is best known for what came to be called the Turing test – a proposed way of deciding whether or not a machine is intelligent that involves an interrogator asking questions, via a teletype, to players A and B, one of whom is human, the other a machine. If, on balance, the interrogator cannot tell which player is the machine and which human, the machine might be regarded as intelligent. However, tucked away towards the end of Turing's wide ranging discussion of issues arising from the test is a far more interesting proposal. He suggests that worthwhile intelligent machines should be adaptive, should learn and develop, but concedes that designing, building and programming such machines by hand is probably completely infeasible. He goes on to sketch an alternative way of creating machines based on an artificial analog of biological evolution. Each machine would have hereditary material encoding its structure, mutated copies of which would form offspring machines. A selection mechanism would be used to favour better adapted machines – in this case those that learned to behave most intelligently. Turing proposed that the selection mechanism should consist largely of the experimenter's judgment.

It was not until the early 1990s that Turing's long forgotten suggestions became reality. Building on the development of principled evolutionary search algorithms by, among others, John Holland[2], researchers at CNR, Rome, the University of Sussex, UK, and Case Western University, USA, independently demonstrated methodologies and practical techniques to evolve, rather than design, control systems for primitive intelligent machines. However, by this stage all those involved held a view of AI that was different in at least one crucial respect from Turing's original vision. Following Brooks' articulation[3] of a growing dissatisfaction with the disembodied, abstract problem-solving, 'parlour game' notions of intelligence that had come to dominate the field of AI, numerous projects had sprung up with a focus on complete functioning sensorimotor systems, usually simple mobile robots operating in the real world. So-called 'New AI' practitioners regarded the major part of natural intelligence to be closely bound up with adaptive behaviour in the harsh unforgiving environments that most animals inhabit, and the central nervous system was viewed as a fantastically sophisticated control system, not a chess playing computer. Hence the study and development of whole, embodied 'artificial creatures' was, and still is, seen by many as an important way to deepen our understanding of natural intelligence and provide new directions for the engineering of intelligent machines.

It was out of this spirited movement that the field of '*Evolutionary Robotics*', on which this book comments, was born about ten years ago. Initial motivations were similar to Turing's: because the design 'by hand' of intelligent adaptive machines intended for operation in natural environments is extremely difficult, would it be possible, wholly or partly, to automate the process? The vast majority of work in this area has involved populations of artificial genomes (lists of characters and numbers) that encode the structure and other properties of artificial neural networks that are used to control autonomous mobile robots. The robot might be required to carry out a specific task or to exhibit a particular set of behaviours. Other properties of the robot, such as sensor layout or body morphology, might also be under genetic control. The genomes are mutated and interbred, creating new generations of robots along Darwinian lines in which the fittest individuals are most likely to produce offspring. Fitness is measured in terms of how good a robot's behaviour is according to some evaluation criteria; these are usually automatically measured but may, in the manner of 18th century pig breeders, and in accordance with Turing's original proposal, be based on the experimenters' judgment. Work in evolutionary robotics is now carried out in many laboratories around the world and numerous papers have been published.

Stefano Nolfi and Dario Floreano are two of the pioneers who helped to establish this area of research and have continued to make important contributions to it. In this, the first in-depth book on the subject, they provide clear accessible explanations of most of the major techniques used and discuss in detail many of the interesting issues arising and challenges to be met. The book is peppered with the kinds of insights that can only be gained through deep involvement in the area. It gives a thorough and up-to-date picture of the field suitable for interested researchers and postgraduate students. However, the material is presented without over-burdening the reader with technical details, thus making it suitable for advanced undergraduates too. Topics covered include the interaction between evolutionary search and lifetime learning, the competitive co-evolution of multiple robots, and issues in the encoding and decoding of robot genotypes. Throughout, many examples of successfully evolved robot behaviours are described; from simple reactive obstacle avoidance to visually guided navigation and the complex motor coordination needed for walking machines. Mentions of current advanced research, such as the evolution of highly plastic adaptive machines, and the evolution and realization of entire robot body morphologies as well as control systems[4] are also squeezed in.

Personally, I would have liked to have seen more on the varying motivations for this style of research as well as the authors' views on the prospects for the field, including its potential relationships with the biological sciences. However, this is a fine effort at encapsulating a very interesting area that has by no means run out of steam (or rather, silicon) and is poised to make important contributions to the science and technology of adaptive systems, real and artificial.

Phil Husbands
School of Cognitive and Computing Sciences, University of Sussex, Brighton, UK BN1 9QG.
e-mail: philh@cogs.susx.ac.uk

References

1 Turing, A.M. (1950) Computing machinery and intelligence. *Mind* LIX 236, 433–460
2 Holland, J. (1975) *Adaptation in Natural and Artificial Systems*, University of Michigan Press
3 Brooks, R. (1991) Intelligence without representation. *Artif. Intell.* 47, 139–159
4 Lipson, H. and Pollack, J.B. (2000) Automatic design and manufacture of robotic lifeforms. *Nature* 406, 974–978

  PII: S1364-6613(00)01794-0