

Volumes of human learning on machine learning

Max Bramer

Proceedings of the Fifth International Conference on Machine Learning. Edited by John Laird. *Morgan Kaufmann: 1988.* Pp. 467. Pbk \$24.95. In Britain distributed by Afterhurst, 116 Pentonville Road, London NW1 9JB, £19.95.

EWSL 88: Proceedings of the Third European Working Session on Learning. Edited by Derek Sleeman. *Pitman: 1988.* Pp. 263. Pbk £25.

ARTIFICIAL intelligence (AI) is the branch of computer science concerned with developing systems with sophisticated reasoning capabilities — capabilities which, if possessed by human beings, would be described as intelligent. A decade ago AI was seldom heard of outside the research laboratory or the realms of science fiction. The turning point came in 1981 with the setting up in Japan of the ten-year programme to build 'fifth generation' computers for the 1990s, in which "intelligence will be greatly improved to approach that of a human being".

Although it may have seemed a heretical idea at that time, AI systems (for problem solving, computer vision, natural-language processing and other applications) are becoming part of the mainstream of computing. There are now hundreds (maybe thousands) of so-called expert systems for practical use in medicine, geology, chemistry and, increasingly, finance. Essentially, the construction of an expert system involves the elicitation of knowledge from a subject expert and its representation in a computer program, often in the form of rules. Obtaining such knowledge is often a time-consuming and error-prone activity, however, and increasingly attention has turned to finding ways in which this process may be automated. Such considerations provide one — although by no means the only — motivation for the recent resurgence of interest in machine learning.

Human intelligence

Although the exact nature of human intelligence remains unclear (and is likely to do so for a long time), the ability to learn is undoubtedly a crucial element, and a multi-faceted one. To list only some possibilities, human learning can be by rote, by trial and error, by discovery, by generalizing from a series of example cases or by a detailed analysis of a single specific case. Humans can learn by self-instruction or by being taught. Teaching can take the form of imparting an under-

lying theory and/or showing examples of good (or bad) practice.

Much human learning takes place by generalizing from a series of examples. There are many ways in which this can occur: the examples can be presented in a carefully chosen sequence by a teacher, can be chosen at random or in some situations can be self-selecting (for example, patients attending a doctor's surgery). The learner can examine a collection of examples as a whole to develop a strategy for all future cases, or can begin with a simple strategy which is then refined incrementally from further examples as he or she goes along. These and other aspects (such as inductive reasoning) of the human skill known as learning are reflected in the lines of research currently being pursued in machine learning.

Two forms of machine-learning research for which the link with human behaviour is less apparent are 'connectionist' and 'genetic' learning. Connectionist learning is closely related to work on neural nets, a field which has become a bandwagon in the past year or so. Although some interesting results have emerged, it is difficult to evaluate the work concerned without considering the very strong claims which many of its proponents have made about modelling the physiological workings of the human brain. Such claims are notoriously easy to make, but much more difficult to justify. I myself remain unconvinced by them.

The term 'genetic' learning originates from the observation that natural (biological) selection can be regarded as a form of learning, where a simple yet powerful learning strategy is applied repeatedly to a continuous stream of input, whilst the system moves from a rudimentary state (such as a simple single-cell organism) to increasingly more effective ones (such as *Homo sapiens*). In machine learning, genetic algorithms embody some model of natural selection and genetics as their principal learning mechanism. Such algorithms have successfully been applied to tasks like the evaluation of the worth of a position in a game of checkers. Unfortunately, as with connectionist learning, the systems that result are likely to be inscrutable to people, however well they perform their given tasks.

The books under review are the proceedings of two conferences which took place in 1988. That edited by John Laird contains papers presented at the University of Michigan in June; 49 contributions are included, the overall stress being on theory rather than applications.

The way that conference proceedings are broken down into sections is often a good indicator of the balance of research in a particular field at a particular time. In this case, there are 15 papers on empirical learning, seven on explanation-based learning, one on case-based learning, five

on machine discovery, four on connectionist learning, five on genetic learning, three on integrating explanation-based and empirical learning, and nine on other topics. Of particular interest are a number of contributions on the well-known ID3 algorithm for constructing decision trees from large collections of examples, a technique which has been strongly advocated as a powerful way for developing the decision rules to use in expert systems.

Goals and experiments

The EWSL proceedings, which come from a meeting held in Glasgow in October, contain 21 papers. Unfortunately they have not been grouped into sections and have simply been arranged in alphabetical order of author, thus losing whatever structure they had when presented at the conference. As a result two valuable, general articles are buried in the middle of the text when they could usefully have been placed at the beginning as an introduction. These are Donald Michie's "Machine Learning in the Next Five Years", in which he considers the past, present and future of the subject, and the paper by Dennis Kibler and Pat Langley on "Machine Learning as an Experimental Science", in which the authors take a broad look at the goals that underlie the experimental study of learning systems and the types of experiments necessary to achieve them.

Both books will be required reading for those seriously involved in machine-learning research, and will no doubt be frequently cited in the future. Despite its lack of structure, the EWSL volume is the more accessible and perhaps the more valuable collection. But both suffer from the failings of most such volumes — the papers generally lack the broad context necessary to be meaningful to the general reader, or in many cases the researcher in an adjacent part of the field. At a more mundane level, neither book has a subject index, and there are no cumulative lists of references.

These faults may well be unavoidable in books which are put together quickly, so that they can be handed out at a particular conference. But it is surely not asking too much that they should be remedied before the volumes are sold to a wider technical readership. It may seem harsh to criticize the publishers for a sin which is increasingly widespread in the AI world, and doubtless elsewhere. But the lack of some post-conference editing has undoubtedly resulted in two books which are less valuable (if more timely) than they otherwise might have been. □

Max Bramer is Manager of the Knowledge Engineering Programme, Hewlett-Packard Laboratories, Filton Road, Stoke Gifford, Bristol BS12 6QZ, UK, and Chairman of the British Computer Society's Specialist Group on Expert Systems.