

KARL have the feel of “logical modelling” exercises, in which the intention is to produce an elegant and rigorous formal account of some, hitherto informal, methodology. A more pragmatic line is taken in the AIDE system, which is more clearly in the lineage of conventional expert system shells, giving a separation between strategic, domain and case-specific knowledge in a style reminiscent of NEOMYCIN.

Since this book is a description of work in progress, it is unreasonable to expect firm conclusions about the contribution which each of the formal systems might make to KBS development. However, the concluding chapter is effective in drawing together the main lessons from each of the (independent) earlier chapters. Of particular interest are the open problems which are highlighted at the end of the book, since these are shared with the wider software engineering, formal specification and logic programming communities. In summary, these are: integration of system components; describing systems which involve database updates; merging with other software development environments; providing reliable formal mappings between specification languages; coping with (sometimes unexpected) interactions with the external environment. To these we might add the need to demonstrate the benefits of formal frameworks to KBS design problems of realistic size, since the “goodness of fit” between the abstract design of a formal framework and the reality of its use in KBS design is an empirical, as well as theoretical, issue. One of the virtues of this book is that it shows a willingness to address the problem—a trend which, I believe, stems from the growing openness of specialists in the KBS formal methods community. I expect that this book will be helpful in making such techniques more accessible to the wider research community and recommend it as an introduction to formal methods in KBS.

Reviewed by Dave Robertson, Department of Artificial Intelligence, University of Edinburgh, 80 South Bridge, Edinburgh, UK

Genetic algorithms for machine learning edited by John J. Grefenstette, Kluwer Academic, USA, 1993, £72.50, pp 163, ISBN 0-792-39407-0.

This book is concerned with Genetic Algorithms (GA) and their use in machine learning. It is a reprint of the special issue of *Machine Learning* journal (13, No. 2–3). The same five articles in the journal are included in the book, dealing with different aspects of GAs applied to machine learning. The first two articles are about the representation scheme of genetic operators and describe two contrasting approaches. The first approach is to utilize traditional string representation while the second one is to use symbolic rules for problem solving. The third article describes a system called COGIN, showing how to utilize GAs for competition-based learning. In the fourth article, the performance of GA-based reinforcement learning is elaborated and it is compared with a neural network-based learning procedure. The last paper is somewhat different in scope. It is rather a theoretical work, and introduces a tutorial on how to define a good fitness function. It relates the results to machine learning at the end anyway.

I should admit that the book includes very thoughtful ideas, and I have learned something while reading it. It is mainly well organized (apart from two page numbers on one page and a poor index table), and the papers are carefully selected. The results illustrated in the papers are clear enough to understand the ideas presented. They also indicate some future research directions. For example, although the first paper by De Jong et al. describes a learning procedure for a single-class problem, it can direct the reader to think on how to apply the methodology in multi-class problems concerned with real time events in a dynamic environment. The second paper by Janikow provides some ideas to develop justification and validation procedures for learning systems which I believe, is usually ignored in machine learning research. As the authors agree, the third paper by Greene and Smith provides ideas on how to develop a general symbolic induction algorithm for solving continuous-valued classification problems. The fourth paper by Whitley et al. compares reinforcement learning procedures on a single control problem (pole balancing). The results indicate that

further research is required to define challenging control problems and then justify the approach taken.

It should be noted that the book is not meant for the novices to either field. The reader is expected to have some background on the basics of GAs and machine learning. I would prefer to see one more chapter giving the basics of both GAs and machine learning in order to attract new people who are interested and want to study GAs. My second impression was that in most of the papers, GAs are compared with induction based learning procedures. Considering that the machine learning is not all about induction, there should be some ideas on how to integrate other learning procedures with GAs. There is only one example in the book (fourth paper) concerned with neural networks. The results shown are problem dependent. There is a need to justify the advantages and disadvantages of the techniques presented in the paper for different problems. That could be made by including more papers describing different work on GAs and neural networks. I am sure there is some other work on GAs and related machine learning techniques which can be included in the book. If not, it would be nice if the related research issues are highlighted in the book. Inclusion of more papers and chapters in the book would definitely increase the reader interest. As a neural network researcher, I do not see the point of buying the book for only one paper.

Despite all these, it was a pleasure to read the book and I would recommend the book for researchers (postgraduate students or lecturers) in machine learning. Both GA and machine learning researchers will definitely find some hints and ideas within the pages of the book. In particular, if you are a new student and digging for a PhD topic on machine learning and genetic algorithms, I think it is worth buying the book and reading it thoroughly. However, don't forget to learn the basics before reading this book.

Reviewed by Dr Ercan Öztemel, TUBITAK-MRC, Dept. of AI, Gebze 41470 Kocaeli, Turkey
(Also affiliated with SAÜ, Eng. Fac., Kampüs, ADAPAZARI-TURKEY)