Robotica (2006) volume 24, pp. 143–144. © 2006 Cambridge University Press
Printed in the United Kingdom

# Book Reviews

doi:10.1017/S0263574705212584

**INTRODUCTION TO MACHINE LEARNING,** by Ethem Alpaydin, MIT Press, 2004, xxx + 415 pp., with index. 160 ref. distributed chapter by chapter. ISBN 0-262-01211-1. (Hardback £32.95).

This is an introductory text aimed at senior undergraduate and first year postgraduate students. Any book dealing with machine learning usually presents a varied collection of the different programming techniques which enable performance criteria to be optimized using either example data or past experience. The range of topics presented here is unusually comprehensive and includes some techniques not usually found at this level.

The introductory chapter gives a very quick tour of some examples of machine learning applications which, curiously, are not really followed up in the rest of the book. However, it does conclude with a useful list of journals and web addresses. Next follows 32 pages dealing with the different methods and problems associated with supervised learning, culminating in a discussion of model generation and the requirements for the design of a supervised machine learning algorithm. In the next two chapters, in order to make the link between mathematical probability (which is summarised in a separate appendix) and software inference methods, Bayesian decision theory is discussed in some detail; to be followed by an excellent discussion of parametric methods. This is then extended to deal with multiple inputs and outputs. This is an area of considerable complexity and I feel that this chapter barely scratches the surface – but, of course, this is an introductory text. The problem of reduction of dimensionality is quite properly discussed although the reason for its importance is not really explained. Clustering is the next topic and this leads on quite naturally to nonparametric methods in general.

The rest of the book deals with a variety of examples of nonparametric techniques: decision trees, linear discrimination, neural networks (looked at in some depth and from a refreshingly original but penetrating point of view). These are followed quite naturally by a discussion of hidden Markov models. All of this is rounded off with a discussion of methods of assessing and comparing classification algorithms and how to combine various models to create multiple learning routines. Lastly, the concept of reinforcement learning is used to tie the various strands together.

Some chapters are better than others, no doubt reflecting the interests of the author but overall the coverage is thorough. There is a satisfactory seven page index. First publication date is 2004 and there are some deficiencies which no doubt subsequent editions will remedy. The prerequisites are a basic knowledge of programming, the elements of probability theory, elementary calculus and an outline knowledge of linear algebra. The examples presented at the ends of each chapter are often a little nebulous – they are not clearly posed and could benefit by being sharpened up and by the inclusion of some tangible mathematical problems. Also, some answers and a selection of hints for solutions is always helpful in an introductory text (by definition, anybody needing to use such a text hasn't yet developed a feel for the subject area and needs the extra moral support that a selection of solutions can provide).

In general, each chapter presents a stand-alone topic and ends with a 'notes' section which sets the work in context and gives pointers to further developments in the field. I thought this was a useful innovation because it enables the student to rapidly gain an overall feel for the subject area. A clear presentation of theory is necessary, and the author succeeds admirably in this, but there is not an adequate range of worked examples to back it up. Instead of being grouped at the end of the book, references appear at the ends of each chapter; this is a sensible arrangement since each is effectively a stand-alone topic. In all there are about 160 up-to-date references. There is a very interesting supporting website which includes 16 chapter-specific downloadable pdf Powerpoint presentations. These are mainly intended for instructors who are using the text for class teaching but any student using the book would find them an invaluable aid. Not only are they commendably clear but also they look as if the author has taken the basic text and improved and extended it. Is this the beginnings of a new edition?

I would recommend this book as a very good main text – particularly if it is used in conjunction with the Powerpoint files. However, it would need to be supplemented with other sources to remedy the shortage of worked examples and the lack of graded problems.

G. F. Page
School of Engineering
Liverpool John Moores University
Byrom Street
Liverpool L3 3AF (UK)

doi:10.1017/S0263574705222580

**DESIGNING EVOLUTIONARY ALGORITHMS FOR DYNAMIC ENVIRONMENTS,** by Ronald W. Morrison, Springer, 2004, xii + 148 pp., with index, ISBN 3-540-21231-0. 66 references. (Hardback £38.50).

Evolutionary Algorithms are a group of optimising search algorithms which grew out of three disparate research areas in the 1960s: genetic algorithms (GAs), evolutionary strategies (ESs) and evolutionary programming (EPs). Later, genetic programming (GP) was added to the list. However, they are all similar, their differences being only in the emphasis attached to ten basic algorithmic steps or to the experimental approach. They have become a popular optimisation and search technique because of their ability to render solutions to difficult problems which were often otherwise intractable. However, they tend not to be very adaptable to sudden changes in the problem search area and usually need the intervention of the operator to keep them on track. Self-tuning is definitely not one of their strengths.

This book is a monograph explaining the research performed by the author in the field of dynamic search algorithms. It starts with an overview of the way in which these search strategies have developed in attempting to deal with dynamic environments and gives pointers to the research issues that are still outstanding. The problems which EAs face when applied to such changing environments is then discussed and an attempt is made to outline the desirable characteristics needed if they are still to perform effectively. There follows a very short chapter dealing with possible solutions suggested from nature and from other branches of engineering. This is followed by a very much deeper discussion of diversity measurement. This begins with a review of basic concepts and discusses strengths and weaknesses, leading to an explanation of an alternative

computational method for binary genotype space and a short section on non-binary populations. This is all leading to a discussion of ways of improving diversity measurement in dynamic populations and a definition of, and the interpretation of, a dispersion index as opposed to a diversity index.

After the scene-setting in the first four chapters the author gets to the central topic of the book: the introduction of a new EA architecture for dynamic problems, one that makes use of *sentinels*. These are a uniformly distributed subset of the population whose members are stationary and do not themselves undergo any of the usual manipulations, e.g.: crossover or mutation. An overview of sentinels is provided, together with a thorough discussion of the main problem of where to place them. After discussing the theory, the experimental approach used and the associated practical problems encountered in the sentinel research is discussed in some depth. The performance measurements decided upon are presented and then the experimental results are analysed and interpreted and a comparison made with the various competing techniques. However, the author seems to run out of steam near the end of the monograph. There is a short chapter outlining some experimental results for population initialization and then the final summary and conclusions, which occupy only four and a half pages.

This final chapter is followed by a useful list of 66 references, mostly from the 1990s up to 2003, and but it is let down by a very skimpy two page index. It is fair to point out that this book is based on research submitted for a PhD thesis and is therefore at the cutting edge of developments in the field. It is a research monogram so there are no worked examples and no graded exercises. Overall, the work is presented in a clear manner and gives a useful introduction to what is likely to be a major area of development in the field of evolutionary algorithms. I would definitely recommend the book to all workers in this field who want a clear but rapid overview but it would need to be supplemented by other, more comprehensive, texts.

G. F. Page
School of Engineering
Liverpool John Moores University
Byrom Street
Liverpool L3 3AF (UK)