Methods and tools for applied artificial intelligence by D. Popovic and V. P. Bhatkar, Marcel Dekker Inc, USA, pp 532, $150.00, ISBN 0-8247-9195-9.

At last, from the glut of specific books on the various aspects of AI, Popovic and Bhatkar's book attempts to bring the whole evolution of the AI system under one umbrella. It fills the gap between the largely theoretical AI methods and working intelligent systems. In a subject embroiled in debate with respect to its ever-increasing boundaries and contents, this book provides a fresh, definitive and timely approach to applying AI methods and tools.

After an introductory chapter in which AI technology to date is briefly surveyed, the main portion of the book details the theoretical foundations and methods found in AI.

Chapter 2 describes knowledge representation formalisms including formal logic, production rules, semantic networks, frames and scripts, model-based techniques and a concise section on representing uncertain knowledge.

Chapter 3 gives a thorough insight to the knowledge acquisition and maintenance phases of an AI system. It covers the standard approaches and existing tools and techniques with insights into the psychological processes involved. A useful chapter for software engineers involved in knowledge acquisition.

The various reasoning methods found in an AI system are detailed in the following chapter, beginning with a basic introduction to application of interference techniques. Reasoning with the formalisms in Chapter 2 are described as well as sections on commonsense, causal, nonmonotonic, plausible, spatial and temporal reasoning. The coverage is wide with simple examples of use, not cluttered with too much depth. It is a useful reference for all the reasoning methods presently available.

Chapter 5 focuses on various machine learning techniques. Although a historical overview is not presented, the milestones and current learning paradigms are explained, including learning by analogy, problem-solving, strategy learning and fault-driven learning. Concept and connectionist learning are covered in greater detail, perhaps emphasising their importance to the area.

Planning towards a solution in AI systems and the scheduling problem for AI systems is the subject of Chapter 6. Typical planning systems such as STRIPS and PLANNER are featured.

The last chapter in this section of the book deals with the implementation of search methods for the purpose of problem-solving. The representation and reduction of a problem using graphical structures is featured, as well as the various search algorithms on the structure.

The final section of the book consists of two chapters which deal with the various tools available for the implementation of AI systems. A taxonomy of the various AI languages and programming paradigms is given in the first chapter. The chapter also addresses the subject of combining AI languages with external devices for creating suitable human–machine interfaces. The natural progression to using parallel processing in AI systems is covered, with an overview of the recent techniques.

The last chapter provides a comprehensive description of the various shells and environments available on the market for creating complex expert systems without programming.

In many ways the book is geared to the material found in the final two chapters, highlighting the necessity to understand the global process involved in building an AI system that is also an efficient, intelligent and appropriate end product.

The sheer size of the area tackled by the book is commendable, though it may be most useful as a reference to applied AI rather than as a textbook. It features many pointers for pursuing aspects of the book and in particular the relevant reference lists given at the end of *each* chapter are extremely useful.

The text is well-presented with important and relevant details often presented in list form and summaries of sections that follow. The typesetting is clear and diagrams are used to depict the processes and examples described. The book presents the interactions in the workings of all parts of the process of AI-modelling without emphasising the theoretical aspects. Possible pitfalls in the process are also described.

Its up-to-date handling of the complete process of creating real-world AI systems is perhaps most relevant to those new to applied AI from related fields, although it assumes some knowledge of computational terminology. The book is suitable for use as a reference guide for all involved in any aspect of AI and its plethora of relevant bibliography will prove useful. However, its plausibility as a text-book for students is questionable especially due to its high price. One consolation for those willing to pay this substantial cost is that its outlook on AI will last well into the next century as it is printed on acid-free paper!

Reviewed by Rashmi Pandya, University of Exeter, Exeter, UK

Neural networks and fuzzy systems: A dynamical systems approach to machine intelligence by Bart Kosko, Prentice Hall, Englewood Cliffs, NJ, 1992, pp 449, £24.96, ISBN 0-13-612334.

The traditional standpoint of artificial intelligence is that real-world concepts can be structurally disassembled into crisp predicates to be treated by first-order logic. Whereas this assumption works quite well in reasoning with simple artificial objects, problems arise with real-world (say, biological) concepts, where uncertainty, imprecision, and noise come into play. Such words as "warmth", "democracy" or "ferocity" resist being described by rigid logical structures. Among the technologies invented to address these issues, Fuzzy-Set Theory (FST) and the philosophy of Artificial Neural Networks (ANN) in particular attract the attention of researchers worldwide.

In FST, each object is associated with a real number $\mu/\varepsilon[0,1]$ understood as the degree of truth of the proposition that the object is a positive instance of some concept. Reasoning is then carried out by applying simple rules for truth propagation along a sequence of fuzzy propositions. Even though fuzziness was invented by Zadeh in 1965, it has raised interest among practitioners only after the recent commercial success of the first fuzzy-logic applications in control systems.

The reason for this delay is probably the common western belief that the only salient grounding for intelligence is in Aristotelian logic. Yet, binary logic suffers from flaws that were pointed out as early as two and a half millennia ago in the form of the familiar Greek paradoxes. In his book *Neural Networks and Fuzzy Systems*, Bart Kosko shows that these paradoxes disappear when the concepts are defined in terms of fuzzy sets. What poses problems is proper definition of truth degrees because their manual setting is somewhat arbitrary and can lead to poor performance in reasoning tasks.

ANNs encode knowledge in terms of synaptic weights between neurons. Importantly, their behaviour is usually robust against small changes in the weights. The major strength of ANN, however, is the ability to *learn* proper weights from examples, often without any prior knowledge about the concept to be acquired. On the other hand, people with an artificial-intelligence background complain that ANNs lack the ability to explain their conclusions in terms of explicit lines of inference.

What the two paradigms share is that dealing with uncertainty and truth grading are intrinsic to them. Bart Kosko deserves credit for being one of the pioneers in the endeavours to combine the learning ability of ANN with the computational simplicity of FST.

The book consists of two major parts. One of them presents, with remarkable erudition, artificial neural networks as gradient systems. The author focuses primarily on the dynamic aspects of neuronal membranes and synapses, and describes the relevant processes by stochastic differential equation. Then, the well-known backpropagation learning algorithm is presented as a stochastic-approximation task. Attention is given to the problems of convergence and stability of feedforward and feedback systems.

The second part examines fuzziness and adaptive fuzzy expert systems. In the author's view, the crucial step for a deeper understanding of the reasoning with *fuzzy adaptive maps*, a concept that unifies the fuzzy-set and neural-network perspectives, is the geometric interpretation of fuzzy sets as points in unit hypercubes. The underlying theory is explained in great detail, and many