Brzoska, who translates temporal logic programs with past and metric operators into a constraint logic program over a suitable algebra. This helps to use constraint logic programming as temporal logic programming. The next paper, by Kono, describes a tableau based implementation of an interval temporal logic. Fruhwiirth introduces a family of logic programming languages for representing and reasoning about time by combining both qualitative and quantitative temporal reasoning with time points and periods. The paper lacks both formalisation and any implementation detail. An interesting paper is by Merz, who considers fixed-point temporal logics of linear time and defines a restricted class of safety properties that afford a linear model-construction algorithm.

MetateM is an executable temporal logic programming language developed by Gabbay and others. Concurrent MetateM extends MetateM by the introduction of concurrency and communications for implementing reactive systems. There are two papers in the volume on Concurrent MetateM. Fisher's paper describes a semantics of Concurrent MetateM, and the other paper by Reynold is on first-class Concurrent MetateM.

Koehler and Treinen have proposed yet another interval-based modal temporal logic LLL employing some conventional modal operators, including the Schops operator. The last two papers in the volume have practically nothing to do with temporal logic and little to do with mainstream modal logics. Although the title of the paper by Calvalcanti claims to have solved air-traffic problems with possible worlds, very little evidence is found in the paper itself. The concluding paper discusses the result of a study for representation of Dutch traffic law using deontic modalities.

Overall the papers in the volume are not well polished. The volume has a severe lack of papers describing languages suitable for representing real life problems and their efficient implementation. The area of modal and temporal logics has a lot to offer to the development of computer science. Too much emphasis on theoretical correctness of the logics and not enough attention on how to make them work may severely affect its progress, as in some other fields of artificial intelligence.

Reviewed by S. K. Das, Imperial College, London, UK

Artificial intelligence techniques in Prolog by Yoav Shoham, Morgan Kaufmann, 1994, 327 pp, £38.50, ISBN 1-55860-319-0.

This is another good book from Morgan Kaufmann, a fine textbook, and a suitable companion volume to Matt Ginsberg's *Essentials of Artificial Intelligence* from the same publisher. In this book, Shoham provides a very practical introduction to a large range of well-established methods. The coverage is extensive. Descriptions of the basic technique, such as the various types of search and forward and backward chaining, are accompanied by other methods less often encountered in AI textbooks. These include the belief propagation algorithm for polytrees due to Pearl (1988), a whole selection of reason maintenance systems, and even methods for temporal reasoning (although the latter are not so surprising given the author's research interests).

This breadth of techniques is possible because, unlike, for instance, Bratko's *Prolog Programming for Artificial Intelligence*, this book does not contain much material on Prolog itself. This is not a Prolog textbook; rather it is a study of how various algorithms may be implemented in Prolog. It makes use of the declarative nature of the language to achieve a clear exposition of the algorithms, as well as benefitting from the support that Prolog offers to those writing meta-interpreters and handling definite clause grammars.

The book is suitable for anyone who is interested in the practical application of methods from AI. It could be used as the basis of a very practical course on AI for students who already have a grounding in Prolog, or given the on-line availability of all the code, as a sourcebook of researchers who want to investigate new methods. With the large number of quite advanced exercises it contains, the book could possibly even find a place on an advanced Prolog course.

Acceptance for this latter task would, of course, depend to a large degree on what one thought of the way in which the code is written—the extent to which it provides a suitable example of clarity and efficiency. As far as I am concerned, the authors of the code (who include several of Shoham's research assistants) walk the line between purity of Prolog and efficiency of algorithm reasonably well. While the programs include the liberal use of the cut mechanism to eliminate unnecessary backtracking, they remain suitably easy to follow, and in most cases the comments included in the code are exemplary.

So, to summarise, there is little, apart from a few typographic errors, that I can find wrong with *Artificial Intelligence Techniques in Prolog*. In my opinion, it fully lives up to the author's claim that the material is both fully self-contained and balanced, as well as completely de-mystifying the technique in question. It is a rare example of a textbook that is both well conceived and well executed, and thoroughly deserves the acclaim that I am sure it will receive.

Reviewed by Simon Parsons, Advanced Computation Laboratory, Imperial Cancer Research Fund, PO Box 123, Lincoln's Inn Fields, London WC2A 3PX, UK