During the twentieth century, certain mathematicians turned their view inwards to look at foundational aspects of their own subject. In addressing the internal consistency and completeness of mathematics and how theorems related to properties of the real world, they began to look at the relationship between theories expressed as mathematical symbols and the properties of models of those theories. During the 1950s, the phrase "theory of models" was coined (by Tarski) to refer to a field which began to address properties of those models in the abstraction; abstract properties of abstract models of the real world.

Esoteric as it may seem, this does have application to computer science. One can, for example, view a formal specification as a theory; an abstract statement of the properties of an intended system. An implementation then becomes a model of that theory. A process of stepwise refinement whereby a specification is gradually refined into an implementation may be viewed as stepping through a sequence of equivalent models.

Of course, there is also the importance of Herbrand models in studying the semantics of logic programs. Indeed semantic issues in general, for which model theory is an appropriate tool of exploration, are an important part of the agenda of computer science. That is by way of a, rather extended, apology for reviewing this book in this journal. What does it offer?

Comparison must be made with Chang and Keisler (1973). That book begins by setting up a formal language; first order logic. The sentences of the formal language are then used to "say things about the models". Thus far is familiar territory for a software engineer. The bulk of the book then continues with descriptions of methods of constructing models, with example applications of those methods.

Wilfred Hodges' book generally works at a slightly more abstract level. He begins with a discussion of structures; the stuff of which models are made. Concrete examples of such structures might include graphs or groups. Following that, the language for talking about structures is developed. This reflects a slight broadening of outlook over Change and Keisler. The latter focused more on the aspects of formal languages that are of interest to the computer scientist. Hodges' book in contrast is motivated more by a desire to study *all* the classes of structures that are of interest to the mathematician (not just those specifiable in first order logic). As with the earlier book, though, most of the content is focused on model-theoretic methods of construction, although in a more up to date and slightly more comprehensive exposition.

So, as might be expected, that makes the present book more of a mathematician's book than a computer scientist's. Nevertheless, for quality of writing, Hodges' dry and at times almost surrealistic sense of humour, and clarity of presentation, this book makes an extremely valuable source of insight and information. Chapters open with an informal motivating/scene-setting section, and conclude with historical notes and a bibliography. Exercises are provided after each section of each chapter.

Model Theory would perhaps have most interest to those who have mathematics to degree level. But given that, it is a pleasure to read and Hodges' joy at the beauty of, and enthusiasm for, his subject comes across clearly.

References

Chang, CC and Keisler, HJ, 1973. *Model Theory.* North-Holland.

Reviewed by Paul J. Krause, Advanced Computation Laboratory, Imperial Cancer Research Fund, 61 Lincoln's Inn Fields, London WC2A 3PX, UK

Artificial intelligence through search by Chris Thornton and Benedict du Boulay, Intellect, 1992, pp 363, ISBN 1-871516-24-2.

The engineering of knowledge-based systems—theory and practice by Avelino J. Gonzalez and Douglas D. Dankel, Prentice Hall International, 1993, pp 523, ISBN 0-13-334293-X.

Here are two more textbooks aimed at introductory courses on aspects of artificial intelligence, both of them based upon interesting ideas, and both of them ultimately let down by a lack of attention to detail.

The interesting idea behind *Artificial Intelligence Through Search* is, as one might guess from the title, that it is possible to use the idea of automated search as a unifying concept for many of the subfields of artificial intelligence. Thus the book starts with simple state-space search to solve simple problems, and then builds up through solving more complex problems with the help of heuristics, to deal with game tree search. All of this sounds fairly conventional, and it is. What is interesting is that the concept of search is then extended with the idea of searching AND/OR trees which leads naturally to problem reduction and then to planning presented by the authors as a search through possible worlds. Having established this point, it is then possible to present natural language parsing as a search through the possible parsings of a sentence, rule-based expert systems as a search through hypotheses guided by the use of numerical certainty values, and concept learning as a search through a set of descriptions. Finally, the logic programming language Prolog is used to illustrate how search may be used to implement computation.

Thus *Artificial Intelligence Through Search* takes a simple mechanism and builds on it to show that quite complex tasks can be achieved by slight variations of it, and to me this seems like a great way to proceed in an introductory course. However, the book is not without its flaws.

The main problem is that it does not live up to the back cover blurb. This claims that each chapter is broken up into three sections, one of which introduces the technique that the chapter focuses on while the second and third sections implement the idea in POP-11 and Prolog, respectively. This is clearly not true. The real structure of each chapter is the development of the technique by discussion of how one might implement it in POP-11, followed by a brief discussion of how the POP-11 code might be re-implemented in Prolog. In other words, it is not possible to learn about the technique without getting bogged down in the minutae of the POP-11 implementation. Now, while I don't have anything against POP-11 (my only comment on it being that, as far as the examples in the book are concerned, it looks very much like Lisp without brackets), this approach means that, unless one has a good understanding of POP-11, the book becomes rather difficult to read. Furthermore, unless one intends to program in POP-11, a large part of the material that has to be covered in order to understand the technique in question is irrelevant. To me this means that unless the book is intended to be used as a joint introduction to artificial intelligence and POP-11, or to teach students who already know POP-11, it is not very useful. Indeed, if it is to be used to teach impressionable students who know some Prolog it could be positively harmful, since the Prolog is quite awful in places (as the authors admit).

The interesting idea on which *The Engineering of Knowledge-Based Systems* is based is the integration of the relevant theoretical bits of artificial intelligence with the sound software engineering practice needed to implement them. The book starts off with the theory, covering basic search techniques, logic, rule-based representation and reasoning, semantic networks, frames and objects, blackboard systems, handling uncertainty and techniques like qualitative and model-based reasoning. Having dealt with these topics, the focus of the book then turns to more practical issues. These include system lifecycles, feasibility analyses, requirements specification, knowledge acquisition, verification and validation, and even a chapter on legal issues which deals, roughly speaking, with what every system designer should know in case of litigation from an unsatisfied user. The emphasis on practice is such that *The Engineering of Knowledge-Based Systems* even comes with a brace of expert system shells so that inspired readers can immediately start building their own fledgling intelligent agents.

This concentration on practical issues, in my opinion, makes *The Engineering of Knowledge-Based Systems* stand out from other introductory texts on knowledge-based systems. Especially useful is the comforting sense that the authors are writing from their own experience, and actually have some field-tested ideas to pass on. This is especially noticeable in the chapters on knowledge acquisition which, while they could be dismissed as "just common sense" at least appear to be common sense derived from real situations. Once again, however, I feel that there are a number of problems with the book.

The most important of these is that it is really a text on the theory and practice of building rule-based systems rather than the broader area of knowledge-based systems in general. This is clear from the index, where a third of the nine chapters on theory are largely concerned with rules while model-based and qualitative reasoning, two of the mainstays of work on "deep knowledge"[1] are relegated to subsections of the chapter on "advanced techniques". The problem is also clear from the fact that alternative methods of knowledge representation, such as frames, are treated rather briefly, and that the chapter on verification and validation only gives examples concerning rules. Thus it looks a lot as though *The Engineering of Knowledge-Based Systems* started life concentrating on rule-based systems and was expanded rather late in the day by the addition of a couple of chapters which gloss over additional topics. This is a shame because, with a little more effort to make the brief descriptions of areas such as objects and qualitative reasoning, the book could have been very good indeed.

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

Recent advances in qualitative physics edited by Boi Faltings and Peter Struss, MIT Press, 1992, £35.95, pp 449, ISBN 0-262-06142-2.

Qualitative physics, for those who have not come across the subject before, is a subfield of artificial intelligence which attempts to model physical systems. The systems that are considered range in complexity from the sublimely simple, such as a mass oscillating at the end of a spring, to the ridiculously complex, such as the cooling plant of a pressurized water reactor. Thus the domain of the system provides the physics.[2] The qualitativeness comes from the approach used to analyse the systems. A classical analysis would be by the numerical simulation or analytic solution of a set of differential equations, with a simulation being necessary for more complex systems because typically such systems involve differential equations that are beyond analytical solution.[3] Now to either perform a simulation of or obtain results from an analytic solution for a given system, it is necessary to have a good deal of precise numerical information about that system. Without such data, classical engineering techniques are useless. However, the lack of this kind of information does not stop trained engineers performing "back of an envelope" type calculations to assess the approximate behaviour of a system without going into numerical detail, obtaining results such as: "closing this valve will cause the flow-rate through that pipe to increase, and so the pressure in that third chamber will fall". It is precisely this kind of qualitative, that is as opposed to quantitative, reasoning that qualitative physics tries to deliver.

Now, as with most books that include the words "recent advances" in their title, *Recent Advances in Qualitative Physics* is a collection of papers. According to the editors' foreword, it is a selection of the papers presented at the 1989 and 1990 International Workshop on Qualitative Physics, and as such might simply be dismissed as yet another collection of papers whose only connection is that they were presented at the same conference, and whose only merit is that they boost the authors' publications list since the same, or very similar, paper is already available in the proceedings of the meeting. Happily, neither of these criticisms can be applied to this volume.

To begin with, most of the papers have not appeared elsewhere, or at least had not at the time that the collection was put together. This is largely because the proceedings of the workshop are not published, the organizers preferring to keep the meeting small and exclusive (participation is by invitation only), and thus productive for the participants rather than large and open and thus less likely to induce interaction between attendees. Since the potential audience for new work on

[1]That is, knowledge which encodes information about the structure of the domain as opposed to the shallow knowledge of rule-based systems which only deals with how to solve the problem at hand.

[2]Of course, as with all sweeping generalizations there are many, varied, exceptions from systems that reason about the human body to those that consider the propagation of uncertain information in probabilistic networks.

[3]For those unversed in the arcane mysteries of applied integral calculus it should be noted that analytical solutions may only be obtained for the very simplest differential equations—it is simply not known how to solve most of them.