Book reviews

Medical expert systems using personal computers reviewed by Dr Jeremy Wyatt, Knowledge-based Systems Group, IBM (UK) Scientific Centre, Winchester

Understanding computers and cognition: A new foundation for design reviewed by Claire O'Malley, Institute of Educational Technology, The Open University

Approaches to knowledge representation—an introduction reviewed by Professor Alain Bonnet, Ecole Nationale Superieure, des Telecommunications, 46 Rue Barrault, 75013 Paris, France

Artificial intelligence and expert systems reviewed by Jim Doran, Department of Computer Science, University of Essex, Colchester, UK

Plans and situated actions. The problem of human machine communication reviewed by Ivan Leudar, Cognitive Science Program, The University of Manchester, Manchester M13

Medical expert systems using personal computers by M K Chytil and R Engelbrecht (Eds.), 204 pages, Sigma Press (distributed by Wiley & Sons) 1987, £14.95.

This book consists of 18 short papers arranged as three sections. The first is a general review of medical expert systems in 20 pages, covering background material and Internist, MED1, HELP and ONCOCIN in more detail, and including a useful summary table of most medical systems. The second section consists of 10 applications papers, and is followed by a 70 page section on methods and theory, with seven papers. In all, there are 48 authors. The title seems topical, particularly with its promise of systems running on economical hardware. However, some might disagree with the wide definition adopted for "expert system", since nine of the systems described in the applications and methods sections lack both the explicit knowledge representation and the ability to give explanations that many would consider necessary. Those working on limited budgets and seeking medical expert systems running on personal computers might again be disappointed: six of the systems actually run on a mainframe or mini, leaving only six systems which are strictly expert systems running on personal computers.

Nevertheless, this juxtaposition of papers about expert systems against those about statistical decision aids and data bases is an honourable attempt to place medical expert systems in their context, as part of the development of medical informatics. Engelbrecht states, in his introduction, that the book's aim is "not to offer general instructions on how to build your own medical expert system. It should contribute to the demonstration of the developmental stage of these systems in the world ..." To achieve this end, the editors have sought contributions from centres as far apart as Eastern Europe (five papers), Japan (three papers) and Stanford.

Two notable themes run through the work: the integration of expert systems with conventional data processing, and the extraction of rules from data. The papers on multi-modal reasoning in Fluidex (Tomizawa et al., Tokyo), on the implementation of Sphinx on a micro (Giusiano et al., Marseille), on ONCOCIN's temporal reasoning (Kahn et al., Stanford) and on Consul/XT (Kowalski et al., Warsaw) are all of a high standard, with careful descriptions of the background to the work described and reviews of the relevant literature; some of the remaining papers lacked both these components. It was disappointing not to find more in the book on knowledge-rich representations, or reports of experience with using PC-based expert systems or development tools.

In places the production and sub-editing of the text has faltered, leaving unexplained abbreviations, unnumbered references at the end of two chapters, poor English grammar and typographical errors, totalling ten on one page (p. 47). Wittgenstein would be flattered by extensive quotes on pp. 134 and 140, but they are a poor substitute for a discussion about differing methods of know-

ledge representation that appears to be missing in the chapter on Prolog. I was disappointed at the two page index, which omitted such terms as "deep knowledge", "LISP", "evaluation", "induction" or "machine learning"—these last two particularly surprising as this is the subject of three of the chapters. The references are located at the end of each chapter and total 182: 30 refer to papers written in foreign languages, and less than a quarter are since 1984.

To whom will this book appeal? Those most likely to benefit include the computer scientist with experience in conventional medical informatics, any AI researcher wishing to find out about work behind the Iron Curtain, and the doctor with an interest in medical decision aids. It may not be of great benefit to the student or the professional AI worker or researcher, unless they have a specific interest in medical applications, as it fails to cover some of the recent developments in the field; neither will it be very useful to those who do not already have some idea both of the methods used in this area, and of the domain itself.

In summary, though one may not expect from the title that the book contains such a wide variety of material, most of it is of interest to those working in the medical informatics field. The professional knowledge engineer or computer scientist will certainly find two or three chapters of value to them, but will perhaps not wish to add it to their own personal library.

Reviewed by Dr Jeremy Wyatt, Knowledge-based Systems Group, IBM (UK) Scientific Centre, Winchester.

Understanding computers and cognition: A new foundation for design by Terry Winograd and Fernando Flores, 207 pages, Ablex 1986, $24.95.

Although for some cognitive science arose out of dissatisfaction with the standard information-processing paradigm (cf. Norman, 1980), in the last decade it hasn't resulted in a shift so much as an elaboration of the pardigm. Central to this—and to most of the core disciplines which make up the field—is the assumption that thought is the manipulation of symbolic representations in the mind that can be understood as referring to objects and properties in the world. Winograd and Flores' book is meant to be a serious challenge to this so-called "representation (symbol processing) hypothesis". The aim is to create a "new understanding" of how to design computer tools (whether or not they involve AI) that are useful and usable. In doing this, the authors have to address some fundamental issues about human language, thought and action, as well as the nature of technology. This means questioning the very traditions and assumptions on which our current understandings are based.

The first part of the book presents a critique of the "rationalist (cognitivist) tradition", with its emphasis on consciously rationalized thought and action, from the view of hermeneutics and phenomenology—in particular, the ideas developed by Maturana and Heidegger. These writers are by no means the easiest to understand and, although Winograd and Flores provide a reasonably clear exposition of the main ideas, it takes at least a couple of readings of the book to get beyond the rather obscure jargon (e.g. "thrownness", "readiness-to-hand", "structural coupling"). The authors apologize for this, explaining that the ideas are so different to our current way of thinking that they require a language of their own to be expressed.

Although I found the ideas stimulating and their exposition reasonably clear, the book was ultimately disappointing, both in terms of their new theory of design and in terms of the system they described as an example of their new approach to design. Winograd and Flores promised much but didn't deliver. But the authors may be justified in replying "Don't bite my finger, look where it's pointing". They address a number of interesting and important issues in HCI and AI, and I feel the book should be taken as food for thought rather than anything more prescriptive.

Winograd and Flores attack AI approaches to natural language understanding for adhering to a correspondence theory of meaning. However, they caricature these approaches to the extent of creating a straw man. They take the view of Maturana, who argues that language is connotative, not denotative. The descriptions to which we refer when we use language lie in a "consensual domain"—a domain which is constitutively social. The authors argue that our implicit beliefs and