Approaches to knowledge representation—an introduction by G A Ringland and D A Duce, (Eds.), 260 pages, Research Studies Press Ltd (Distr J Wiley Ltd) 1988, £29.75

In the book *Approaches to Knowledge Representation—An Introduction*, the authors present a collection of articles which describe the different formalisms used to represent knowledge and reasoning mechanisms in the systems developed to date in AI.

The book is not a course manual for AI students, for it stands far from the down to earth problems of computer implementations of the modes of representation it describes.

The most widely used and universally "recognized" formalisms are described with an unbiased view: logic, semantic networks, frames and production rules.

The first part of the book is devoted to these, showing in particular how notorious weaknesses of the pioneering programs have been progressively reduced by recent research.

A second part gives an account of studies lead by psychologists on the representations of knowledge that human beings seem to employ. As the book rightly points out, the use of a single mode (of representation) seems unlikely; the existence of hybrid representations is suggested by the impressive variety of phenomena involved in the everyday performances displayed by humans. The conjunction of propositional and analogical representations is put forward by several authors. Although not mentioned in the book, recent research on connectionist models seem to go in the same direction.

The third part is devoted to more advanced topics such as representation of time or the functional approach advocated by Levesque and Brachman. The article on conceptual graphs by Sowa seems to be misplaced (in this section) for it is not more "advanced" than semantic networks in general, although the formalism and the mechanisms operating on it have improved precision and therefore are closer to computer implementations.

The sole serious reproach I will make concerns the weak relationship displayed here between the representation of knowledge and its expression in natural language. For example, the difficulty resulting from the polysemy of words is not discussed. The representations in a dictionary of the various meanings of words is never tackled, nor the choice of the right definition of a word depending on the context. An uninformed reader might get an impression (when reading the book) that there is an isomorphism between a string of characters and a meaning.

Notwithstanding this abusive simplication, the book is well written and provides an excellent and up-to-date introduction to the crucial problems of knowledge representation.

Reviewed by Professor Alain Bonnet, Ecole Nationale Superieure, des Telecommunications, 46 Rue Barrault, 75013 Paris, France.

Artificial intelligence and expert systems by Stuart E. Savory (Ed.), (Translation ed. M. W. Rogers), 278 pages, Ellis Horwood Ltd, £29.95

It is commonplace to suggest that the (now waning?) enthusiasm and hype for expert systems has generated a flood of publications of varying degrees of mediocrity. This book, alas, does nothing to counter the impression. Indeed, if a camel is a horse designed by a committee, then this book seems to have been written by a whole heterarchy of committees. The reality is that it is written primarily by a group of researchers attached to Nixdorf Computer AG in West Germany and edited by the group leader Stuart Savory. Originally written in German in 1985, it was translated into English in 1987, and updated by the editor. It aims to present the AI/Expert Systems experience through Nixdorf eyes. Different chapters have different authors and range from general introductions to AI and Expert Systems (including commercial aspects), through detailed descriptions of Prolog and of a particular Emycin type shell (TWAICE) developed at Nixdorf, to aspects of automatic learning and a brief presentation of an associative processor for unification.

As a plan for a book, that sounds good. But unfortunately, not only is there little coordination between the chapters (for example, I counted at least three differing brief histories of AI) but the chapters themselves range from seemingly hasty jottings, through extended attempts at formaliza-

tion (which I often found pretentious and repetitious rather than convincing), to the respectable if pedestrian. The prevailing assumption is that Prolog and rule-based systems are all—I sometimes feel that Bob Kowalski and Stanford have a lot to answer for—and the prevailing data structure deployed in the text is the list: of themes, of subareas, of definitions, of techniques, of types of learning and so on. It is characteristic of the book that it provides a list of supposedly relevant Alvey and Esprit projects (including some completed in 1985!) a list of supposedly relevant addresses (which begins with the now defunct Alvey Directorate at Millbank!) but *no index*.

I do not see who could read this book with profit. Specialists in AI and ES will very quickly turn elsewhere, and the uninformed industrial manager, say, keen to find out what it's all about and/or how Nixdorf fared will be confused and possibly intimidated by a surfeit of disorganized definitions, concepts and technical notation. I shall certainly keep it away from my students.

Reviewed by Jim Doran, Department of Computer Science, University of Essex, Colchester, UK.

Plans and situated actions. The problem of human machine communication by Lucy A Suchman, 203 pages, Cambridge University Press 1987, £22.50 (hardback), £7.95 (paperback)

Designing an advanced knowledge-based system does not simply involve representing knowledge in its domain; it inevitably involves *using* that knowledge in social settings. The problems of knowledge representation and knowledge use are closely related but the relationship is often assumed to be unproblematic and ignored. However, the relationship is far from unproblematic and it is becoming clear that the "disengaged cognitive science" (that which ignores the interweaving of thought processes and actions and the physical embodiment of knowledge structures) is either sterile or gives birth to clockwork ghosts. What is needed is an adequate theory of action. Two focal problems for any such theory are as follows: "What makes an action intelligent?" and "What makes it intelligible?" Lucy Suchman's book is an interesting and timely contribution to the debate about the theory of action and it addresses both of these problems.

What does make a behaviour intelligent and intelligible? Well, intelligent actions are intentional and are interpreted intentionally. In the IT related academic disciplines (Cognitive Science, AI, Cognitive Psychology and Human Computer Interaction) this usually means that an agent of an action has goals, and the action is designed on the basis of the agent's beliefs about the environment, to fulfil these goals effectively, avoiding undesired consequences. This view of intentional action has been systematically developed by John Searle in his book *Speech Acts* and applied in the plan based theory of speech acts by the Toronto group (Allen, Cohen, Perrault et al.) amongst others. There are, however, problems. Actions are represented as operators in plans, designed for *virtual environments*. Only the relevant properties of these are specified and it is assumed that the *actual environments* would pose no unforeseen problems and surprises. Very little weight is given to the on-line relationship between models, actions and the actual environments and the systems deal shakily with uncertainty.

Even these cursory remarks should make it clear that this characterization of intentionality is seriously incomplete. In fact, even Searle distinguishes "intentions before actions" from "intentions in actions". This is to say that human actions have an intrinsic intentionality which is not explicitly calculating and is not abstractable from the environment. This notion has been more thoroughly explored in philosophy (e.g. Merleau-Ponty, Heidegger) than in IT and Suchman owes a debt to European philosophers when she argues that the plan based characterization of intelligent action is fundamentally flawed and impoverished. She suggests that all actions are situated and their characteristics are indexical and that plans systematically filter out the detail that characterizes situated actions, in favour of those aspects of the actions that can be seen to accord with the plan.

Suchman introduces (in Chapter 2) the notion of interactive and self-explanatory artefacts. She very briefly summarizes some research comparing human/human and human/artefact interactions and discusses why users should interpret the actions of artefacts intentionally. According to her this