

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

is because (a) the interactions between the artefacts and the users occur in a common “language” and (b) artefacts are internally opaque to their users and thus require an intentional stance.

In Chapter 3, Suchman reviews plan-based approaches to action. The review is interesting and identifies their common base for critique, but is too brief to do justice to the area covered. (A more thorough review would have made the book a bit more accessible to the uninitiated, and this goes for other chapters too). The main interesting criticism is that the approach does not deal with the open-endedness characteristic of particular contexts and assumes that background assumptions are a part of the actor’s mental state *prior to action*. According to her, if this were the whole truth, agents would not be able to utilize the richness of actual environments as a resource. Suchman suggests that background assumptions are “generated by the activity of accounting for an action when the premise of the action is called into question”. This seems right but not much is said about the process of “accounting for action”.

In Chapter 4 the notion of “situated action” is discussed, including the role of planning. Much of the account is based on ethnomethodology and conversational analysis, and some of it on Mead’s work as well as on Heidegger’s as popularized by Dreyfus. Plans are, according to Suchman, one of the resources people use for practical deliberations about action but they do not normally cause actions in any direct sense. “The source of mutual intelligibility is not a received conceptual scheme, or a set of coercive rules or norms, but common practices”, (pp. 57–58).

Chapter 5 is on communicative resources, and mostly on methods of conversational analysis. The account introduces its basic concepts and will give the uninitiated a flavour of the method rather than working knowledge.

Chapter 6 is an introduction to the issues to be discussed in Chapter 7; Suchman discusses interactions with expert systems, focusing on problems of following and communicating instruction. Unfortunately the chapter does not make much use of the concepts introduced in the previous chapters and as it stands, it doesn’t add much to the book. This is however remedied in Chapter 7, which is an analysis of users’ interactions with an expert help system, using methods of conversational analysis.

The concept of situated action seems right, and so does the criticism of approaches which decontextualize actions and drive them exclusively by reference to representations. But what about explaining “situated actions” in themselves, including how they come to be rational without deliberation? Their enumeration is surely not enough and neither is saying “we just do it” as others have done.

I enjoyed this book very much, even though in places it is too easy, and in places it is too difficult. It is well written and stimulating and it should be read by all working in cognitive science, AI or HCI as an introduction to views of action which takes their context seriously. I think that serious readers will have to admit the force of the criticisms even if they are not persuaded by the offered alternatives.

Reviewed by Ivan Leudar, Cognitive Science Program, The University of Manchester, Manchester M13.