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

assumptions can never all be made explicit—there is no objective viewpoint from which we examine our beliefs since we always operate in the framework which they provide. It isn't clear whether this is meant to be a strong (logical) claim, or a weaker pragmatic one. They argue further that meaning is fundamentally social and cannot be reduced to the activity of individual subjects. For Maturana, language as a social activity is prior to conscious (reflective) thought. (In this respect his ideas are similar to those of Vygotsky, 1986, who saw cognitive development as proceeding from the interpsychological to the intrapsychological.)

Winograd and Flores argue that few statements can be judged true or false without an appeal to an unstated background, and point to Austin and Searle's work showing that what is unspoken is as much part of the meaning as what is said. However, they criticize work on speech acts for attempting to extend rather than to reject the notion that the meaning of an utterance can be described in terms of conditions stated independently of context. According to Winograd and Flores we need to go beyond the mental to the social to break out of the rationalistic tradition and understand how meaning is shared. Speech acts are jointly constructed and depend on commitments on the part of the listener as well as the speaker. Winograd and Flores argue that language can work without any objective criterion of meaning because all that is required is that there be a sufficient coupling so that breakdowns are infrequent and that there is a standing commitment on the part of the speaker and listener to enter into a dialogue in the face of a breakdown. Few would disagree that language requires a shared understanding (embedded in a background of unstated assumptions), but at some point and in some way language has to refer to the world. This doesn't necessarily imply the naive correspondence theory with which Winograd and Flores attribute AI.

The authors see little future in the fields of natural language understanding and machine translation. They criticize researchers in HCI and KBS for their assumptions that successful interfaces will need to use AI (in particular, natural language) methods. They argue that we should design tools that augment human capabilities, not mimic them: "The key to design lies in understanding the readiness-to-hand of the tools being built, and in anticipating the breakdowns that will occur in their use" (p. 137). In AI the computer operates with a background only to the extent that the background is articulated and embodied in its programs. But the articulation of the unspoken is a never-ending process.

> "A system that provides a limited imitation of human facilities will intrude with apparently irregular and incomprehensible breakdowns. On the other hand, we can create tools that are designed to make the maximal use of human perception and understanding without projecting human capabilities onto the computer." (p. 137)

Winograd and Flores also launch an attack on another aspect of the rationalist tradition, embodied in the planning approach in AI. That is, the notion that behaviour can be described as the execution of precomputed goals and plans.

According to Heidegger, practical (unreflective) understanding is more fundamental than reflective understanding: we do not relate to things primarily through having representations of them—in driving a nail with a hammer, I need not make use of any explicit representation of the hammer; my ability to act comes from my familiarity with hammering, not my knowledge of a hammer.

Maturana and Heidegger both oppose the assumption that cognition is based on the manipulation of mental models or representations of the world: Maturana rejects the separation of organism and environment, and, in particular, he rejects the belief that properties of the environment in which we act are mapped onto or specify the states of structures in the nervous system. He argues that because our interaction is always through the activity of the entire nervous system, the changes are not in the nature of a mapping. They are the results of patterns of activity which, though triggered by changes in the environment, are not representations of it (a view similar to the Gibsonian perspective). Correspondences between structural changes and the pattern of events that caused them are to be seen as historical, not structural. They cannot be explained as a kind of reference relation between neural structures and an external world.

For Maturana, then, learning isn't a process of accumulating representations, and recall doesn't depend on the indefinite retention of a structural invariant that represents an entity (an idea, image

or symbol). It depends on the functional ability of the system to create, when certain recurrent conditions are given, a behaviour that satisfies the recurrent demands or that an observer would class as a re-enacting of a previous one. A cognitive explanation is not a description of the structure of the system and how that structure determines behaviour, but a description of a pattern of interactions by which the structure came to be and the relationship of those changes to effective action. In other words, mental terms don't describe state or structure, but patterns of action—i.e. temporal and historical.

Leaving aside the issue of representation, one important point to draw from this is the implication that, in order to understand the nature of mental processes and structures, we need to look at their genesis and development. The experimental approach common to much of mainstream cognitive science tends to look rather at snapshots of behaviour, isolated from their history and context.

Heidegger makes a similar but more radical critique of the rationalist tradition, questioning the distinction between a conscious, reflective, knowing "subject" and a separate "object". He sees representation as a derivative phenomenon, occurring only where there is a breakdown of our purposeful action. Knowledge lies in the being that situates us in the world—in acting—not in a reflective representation. (There are similarities here with the activity theory of Leont'ev, 1978, who also rejects the notion that actions are preceded and directed by internal goals or plans.)

Winograd and Flores attack the planning approach in knowledge engineering on the grounds that it assumes that the knowledge engineer can articulate an explicit account of the system's interaction with the world—what it is intended to do and what the consequences of its activities will be. They argue that this breaks down when we consider a system with which people interact directly: no simple set of goals and operators can delimit what can and will be done. They argue that we can define the space of possible effects the interaction can have on the system but we cannot model how the system's activity will produce them:

"The most successful designs are not those that try to fully model the domain in which they operate, but those that are 'in alignment' with the fundamental structure of that domain, and that allow for modification and evolution to generate new structural coupling." (p. 53)

Winograd and Flores prefer the term "management" support rather than decision making support. This is because they want to get away from the rationalist notion that decision making and problem solving involve listing all alternative strategies, determining all the consequences that follow and comparatively evaluating these sets of consequences. The view of management as rational problem solving and decision making fails to deal with the question of background:

"Saying that a manager is optimising some value by choosing among alternatives for action is like regarding language understanding as a process of choosing among formal definitions. The hard part is understanding how the alternatives relevant to a given context come into being." (p. 46)

In their view, management is an extended conversation, in which one or more actors may participate, guided by questions concerning how actions should be directed.

Winograd and Flores argue that when we are engaged in successful language activity, the conversation is not "present-at-hand", as something observed. Its structure becomes visible only when there is some kind of breakdown. However, the system they describe as an example of their new approach to design—the coordinator—would seem to force users to pay attention to the very speech acts which are *not* normally present-at-hand (see also Winograd, 1988). In modelling the interface explicitly on a theory of pragmatics of conversation, they appear to be subject to the very same criticisms they level at proponents of natural language interfaces. They add a caveat that this sort of interaction is not suitable for all types of conversation, but fail to justify its use in the more restricted domains to which the system is directed.

The ideas in this book have a couple of interesting implications for work in human-computer interaction. In HCI, the rationalist tradition can be seen in user modelling approaches in intelligent interfaces and models of users in task analyses, which rely on a planning approach and emphasize

the role played by analytical understanding and reasoning in the process of interacting with the system. However, users' behaviour is not the result of purely internal goal-directed activity as it is with AI planners. It involves the interaction between goal-directed activity and the actual physical and functional situation (cf. Suchman, 1987). This results in opportunistic, data-driven behaviour, and the development of partial plans or local tactics as much as more systematic problem solving. While some behaviour might be characterized as plan-driven, there are many other occasions where it is only after users have acted and observed the outcome that they identify that state as being the goal to which the activity was (in retrospect) directed.

Plan-based approaches may assume that users know more than they do in a sense. Models of users in task analyses predict learning times on the basis of information necessary to learn. In other words, these approaches mean looking at the domain representation and assuming the user has to learn and remember everything involved. But as a recent study (Mayes, Draper, McGregor and Oatley, 1988) suggests, we don't necessarily use all the information available to us in learning and interacting with a system. The knowledge represented in a task analysis or elicitation technique may also bear little relationship to the knowledge we actually use to act in the world. For example, local experts in user communities don't necessarily know "more" than anyone else, but they do know how to find the information when it's needed (cf. O'Malley, 1987).

The second implication for HCI is in the design of "direct manipulation" interfaces, which can be seen as one attempt at designing for "thrownness". Hutchins, Hollan and Norman (1985) refer to this notion as "direct engagement"—the feeling that one is operating directly in the world of interest. One example of a direct manipulation interface is the "Alternate Reality Kit" (ARK—Smith, 1987). The ARK interface has many aspects which are literal or direct mappings from what happens in the real world to what happens in the simulated world. However, other aspects deliberately violate this literalism to achieve enhanced functionality. For example, ARK contains "Interactors", which are higher-level tools for which there is no real world correspondence, to represent such things as control over gravity (e.g. by use of a slider for changing its value). This may lead to breakdown or disengagement. But, as Winograd and Flores point out, "Even what we call 'disengagement'—occurs in thrownness: we do not escape our thrownness, but shift our domain of concern." (p. 71) An issue for design is understanding how to make this shift as smooth as possible, to achieve the added functionality of non-literal interface features without losing the experience of "direct engagement" that literal features encourage. Winograd and Flores (or, rather, Heidegger and Maturana) provide some ways of thinking about these issues, but they don't provide a theory for design.

References

Leont'ev, A N, 1978. *Activity, Consciousness, and Personality* (Trans. M J Hall) Englewood Cliffs, NJ: Prentice-Hall.

Mayes, J T, Draper, S W, McGregor, A M and Oatley, K, 1988. "Information flow in a user interface: the effect of experience and context on the recall of MacWrite screens" *Proceedings of the HCI'88 Conference*, Manchester, 5–9, September, Cambridge: Cambridge University Press.

Norman, D A, 1980. "Twelve issues for cognitive science" *Cognitive Science* **4**(1), pp. 1–32.

O'Malley, C E, 1987. "Understanding explanation" *Proceedings of the Third Alvey KBS Workshop on Explanation*, University of Surrey, September 29–October 1, London: IEE.

Payne, S J and Green, T R G, 1986. "Task-action grammars: a model of the mental representation of task languages" *Human-Computer Interaction* **2**, 93–133.

Smith, R B, 1987. "Experiences with the alternative reality kit: an example of the tension between literalism and magic" *Proceedings of the CHI+GI'87 Conference*, Toronto, April 5–9, New York: ACM.

Suchman, L A, 1987. *Plans and Situated Actions*, Cambridge: Cambridge University Press.

Vygotsky, L S, 1986. *Thought and Language* (Trans. A. Kozulin) London: MIT Press.

Winograd, T, 1988. "A language/action perspective on the design of cooperative work *Human-Computer Interaction* **3**(1), 3–30.

Reviewed by Claire O'Malley, Institute of Educational Technology, The Open University.