

Expertise in Context

Report on the Third International Workshop on Human and Machine Cognition

Robert R. Hoffman and Eric Dietrich

■ The Third International Workshop on Human and Machine Cognition was held in Seaside, Florida, on 13–15 May 1993. Each paper session included presentations on cognitive research, educational research, AI theory and logic, and particular knowledge engineering projects. This mixture encouraged the participants from diverse disciplines to listen and respond to one another. These international workshops are held to allow leading scientists, scholars, and practitioners to discuss current issues and research in particular topics in AI and cognitive science.

The Third International Workshop on Human and Machine Cognition was held in Seaside, Florida, on 13–15 May 1993. These international workshops are held every other year to allow leading scientists, scholars, and practitioners to discuss current issues and research in particular topics in AI and cognitive science. This third workshop was supported by the University of West Florida; the West Florida Regional Medical Center; Taylor and Francis Publishing; John Wiley and Sons Publishers; the American Association for Artificial Intelligence; and the Institute for the Interdisciplinary Study of Human and Machine Cognition, which hosted the workshop.

One aspect that made this workshop special was the setting. Few will forget lounging on a sunny rooftop that overlooked the turquoise waters and white sands and enjoying heated, yet friendly debates on topics that ranged from what Turing meant by his test to whether heads compute and that involved leading scientists such as William Clancey (Institute

for Research on Learning), John McDermott (Yale University), and Pat Hayes (University of Illinois).

Participants at the workshop included AI researchers, cognitive psychologists, social scientists, philosophers, and representatives of many other disciplines interested in expertise. In addition to keynote presentations by Harry Collins (University of Bath, United Kingdom), William Clancey, Neil Agnew (York University, Canada) and Ken Ford (University of West Florida), Vimla Patel (McGill University, Canada), and Michelene Chi (University of Pittsburgh), there were a number of

... discussions broached many of the broad issues in AI and cognitive science, such as the limits of expert system technology.

paper sessions; interestingly enough, however, these were not purely topical. Each paper session included presentations on cognitive research, educational research, AI theory and logic, and particular knowledge engineering projects. This mixture encouraged the participants from diverse disciplines to listen and respond to one another, as did the many informal discussion sessions. To be sure, some of the cognitive psychologists in attendance felt lost

while they listened to presentations that dove into the gory details of abstract logic. Conversely, some philosophy and logic participants were a bit perplexed by the cognitive psychologists' insistence on laying out the gory details of research materials, methods, and results. In general, however, there were few genuine instances of miscommunication, and we suspect that there was general agreement that interdisciplinary meetings of this type are important if one seeks to reveal current issues and significant challenges for further work.

As one would suppose from a workshop subtitled "Human and Machine Cognition," discussions broached many of the broad issues in AI and cognitive science, such as the limits of expert system technology, expertise (or cognition in general) and whether it is computable at all, the relations between AI and cognitive simulation (to what extent are AI models also explanations of cognition), and alternative cognitive models of expertise.

Specific contributions from the perspective of logic and philosophy dealt with the problems of partial entailment, inductive reasoning, fuzzy reasoning, abduction, and abstraction as they are applied to the analysis and description of expertise in particular. Many practical problems were also discussed, such as those involved in prototyping and implementation, the creation of (and need for) intermediate levels of knowledge representation, the avoiding of brittleness, the creation of systems that can be transported into new contexts, uncertain reasoning and reasoning bias, interface design, and fundamental categories of generic problem-solving methods.

The workshop's focus on expertise in context was manifested in two ways. First, the participants seemed to take for granted that diverse domains of expertise will have to be studied if the AI–cognitive science community is to get a real handle, let alone a practical one, on the nature of expertise. Reflecting this presupposition, reports were given on studies of expertise in such domains as medi-

*There can be little doubt
that the concept of
expertise drives home
the recent and
often-touted claim
that there is
more need for
interdisciplinary work.*

cal diagnostic reasoning, manufacturing engineering, athletic coaching, electronics debugging, and satellite image interpretation—as well as our old friends, physics problem solving and chess. The great diversity of domains that have recently been studied empirically is one of the exciting things about modern research on expertise and one of the things that made the workshop so interesting to the participants.

Specific studies focused on the problems of expert identification, training and instructional design, the developmental progression from novice to expert, and representation of perceptual and categorical knowledge, as well as so-called declarative knowledge; empirical studies that reveal sequences of mental operations (for example, forward and backward chaining of inferences); and empirical studies that reveal novice misconceptions that obstruct the process of conceptual change.

The second manifestation of the expertise-in-context theme was the consideration of issues having to do with the social and cultural aspects of expertise and expert system technology. In starting the workshop, Neil Agnew and Ken Ford laid out some of the basic premises of the contextualist orientation: First, expertise is not just a characteristic of individuals; it is also a matter of social selection and attribution (rites of passage; professional, economic, political, or cultural constraints on beliefs and practices; and so on). Hence, expertise is always contextually relative. Second, knowledge acquisition is not only a process of disclosing knowledge but

also a process of constructing it. Third, knowledge engineering is not a process of representing knowledge but of modeling it.

In William Clancey's presentation, he contrasted the situated cognition view with the more traditional view that knowledge is possessed, stored, retrieved, and so on, and that problem solving is merely a process of feature extraction, matching, and so on. To some extent, this line of argument serves to lampoon the view that few people (at least in psychology) might actually have held. However, there are underlying substantive themes. One such theme is that knowledge is not a thing; it is a process. The actual practice of expertise suggests that knowledge is about actions and practices as much as it is about facts. Knowledge is a process residing in local (working group) and global (cultural) contexts as well as in heads (that is, mental representations are themselves dynamic and constantly changing). Hence, knowledge in expert systems is also a process—as in the experience with MYCIN in which the addition of new concepts and rules resulted in the alteration of the meanings of older concepts and rules.

This discussion suggests an alternative approach to knowledge-based systems, but to the cognitive psychologists in attendance—especially those who knew of the recent debate about the situated cognition perspective—this approach was no bombshell. To date, cognitive research on expert knowledge and skill has simply not had to deal with the broader social context to explore basic hypotheses about cognitive mechanisms. Avoidance of social variables is not an artifact of the need for laboratory control and rigor because context can be brought into the lab. However, surely a great deal of potentially fruitful research can and should be conducted on how cognitive factors interact with contextual and social factors.

These themes echoed throughout the workshop, as evidenced by the papers that were presented. Papers discussed empirical analyses of alternative methods of identifying experts in corporations through systematic interviewing and scaling procedures,

analyses of the metaphors that expert system engineers use to conceive the knowledge-acquisition process, issues relating expert judgment to the exercise of authority and decision-making power, and expertise as a social attribution having costs as well as benefits.

There can be little doubt that the concept of expertise drives home the recent and often-touted claim that there is more need for interdisciplinary work. Cognitive psychologists who study expertise can get research ideas from the experiences of practicing knowledge engineers, knowledge engineers working out there in the trenches need the input of experimental psychologists, and so on. There are many lines of cross-fertilization and outright interdependence. It was agreed that the meeting was stimulating and productive and that the debates suggested a number of possible themes for the next workshop.

Robert R. Hoffman received his graduate training in experimental psychology and psycholinguistics at the University of Cincinnati and postdoctoral training at the Center for Research on Human Learning at the University of Minnesota. In his graduate work, he helped pioneer the experimental study of metaphor comprehension. In recent years, his work has focused on applied cognitive science and AI, involving experiments on knowledge elicitation methodology and research on expertise in satellite image interpretation and the design of computer graphic workstations for environmental science. Currently, he is associate professor of psychology at Adelphi University, an honorary fellow of the British Library, and a Fullbright Scholar.

Eric Dietrich is associate professor of philosophy at Binghamton University in New York. He does research on the foundations of AI, analogy making and analogical reasoning, and the semantics of mental representation. He is the founding editor of the *Journal of Experimental and Theoretical AI* and recently edited a book on mental content and the problem of intentionality called *Thinking Computers and Virtual Persons* (Academic Press, 1994).