AI Magazine Volume 11 Number 5 (1990) (© AAAI)

Second International Workshop on User Modeling

Sandra Carberry

The Second International Workshop on User Modeling was held March 30-April 1, 1990 in Honolulu, Hawaii. The general chairperson was Dr. Wolfgang Wahlster of the University of Saarbrucken; the program and local arrangements chairperson was Dr. David Chin of the University of Hawaii at Manoa. The workshop was sponsored by AAAI and the University of Hawaii, with AAAI providing eight travel stipends for students.

An excellent response to the call for papers and participants resulted in 46 high quality submissions, of which 24 were selected for presentation and discussion led by invited commentators. Whereas the first user modeling workshop, held in Maria Laach, West Germany in 1986, focused on user modeling in natural language dialogue systems, the 1990 workshop covered a broader range of topics, including user modeling in tutoring systems and psychological foundations of user modeling. The workshop unofficially began with an evening reception on March 29, during which the participants became acquainted and began informal discussion of their research interests. It continued throughout the next three days, including a beach luau at Paradise Cove.

The emergence of user modeling as an active research area in a number of disciplines (psychology, natural language processing, intelligent tutoring systems, expert system interfaces, information retrieval systems) has resulted in the creation of a new international journal, *User Modeling and User-Adapted Interaction,* published by Kluwer Academic Publishers. The editor-in-chief of the journal is Dr. Alfred Kobsa, University of Saarbrucken; the first volume of the journal will appear in early 1991 and will contain extended versions of selected papers from the workshop. In addition, two special issues of the journal, one devoted to plan recognition and the other to user modeling in response generation, are in progress; the guest editors are respectively David Chin and Ursula Wolz.

Workshop Program

The workshop program included talks on a wide spectrum of topics related to user modeling. Issues addressed by the participants included the implications of experimental studies for user modeling research; active versus passive model acquisition; strategies for inferring models of user beliefs, knowledge, and plans; the development of general user modeling shell systems and toolkits; and techniques for reasoning on a user model to generate appropriate responses.

Wolfgang Wahlster opened the proceedings by summarizing previous activities related to user modeling, describing the goals of the workshop and the organization of the program, and inviting discussion and interaction among the participants. Individual sessions were led by invited commentators; in addition to Wahlster, David Chin, and some of the presenters, these included David Benyon and Peter Johnson of the Open University and the University of London, respectively, in the United Kingdom, Jack Edwards of AI Management and Development Corporation in Canada, M. L. Liebe Harkort of Mouton De Gruyter Publishers in West Germany, Aravind Joshi of the University of Pennsylvania, Michael McTear of the University of Ulster in northern Ireland, and Noguchi Naohiko of the Tokyo Information and Communications Research Lab at Matsushita Electric Industrial C. Ltd. in Japan.

Plan Recognition

A number of researchers have demonstrated the importance of reasoning on a model of a user's plans and goals both in identifying the intended meaning of utterances and in generating appropriate responses. Consequently, implicit incremental acquisition of a model of the user's plans and goals from an ongoing dialogue has become an important topic in user modeling and was the focus of a number of workshop presentations. Previous research has concentrated primarily on plan recognition in relatively ideal dialogues, in which the user's beliefs were assumed to be a subset of the system's beliefs, the system had sufficient knowledge to recognize the motivation for and respond to user queries, and default inferences were unnecessary. Unfortunately, such dialogues are not typical of naturally occurring interactions. Current research efforts have begun to address the problem of developing more robust plan recognition systems.

Doug Appelt, SRI International, discussed his work with Martha Pollack, also of SRI International, on formulating plan recognition as weighted abduction. This work builds on previous research by Pollack and Kurt Konolige in which plan recognition is cast as a process of defeasible reasoning and represented in a formal argumentation system in which potential inferences compete with one another and an argument supporting one inference can be defeated by arguments supporting conflicting inferences. Appelt described weighted abduction as a process in which weights on the premises of inferences rules are used to evaluate competing explanations. He showed how Konolige and Pollack's model of plan recognition could be expressed in terms of weighted abduction and compared the two approaches.

Rhonda Eller, University of Delaware, described work with Sandra Carberry on a meta-rule approach to flexible plan recognition. Her work is a departure from the work of Appelt, Konolige, and Pollack. Eller argued that instead of considering a large number of possible ascriptions of belief and intention at the outset, a plan inference system should be tightly constrained initially and then be relaxed in a principled manner to include more robust inferences until a plausible explanation for the user's observed action is produced. Eller described meta-rules for relaxing the plan inference process and outlined a strategy for controlling the relaxation process.

Peter van Beek described work with Robin Cohen, Fei Song, and Bruce Spencer, all of the University of Waterloo, on a robust model of plan recognition in which the system can recognize aspects of plans that are not encoded in the system's plan library and, if appropriate, assimilate them into the system's knowledge base. He described an algorithm for analyzing linguistic expressions to identify the temporal relations between actions in a user's plan and discussed appropriate responses when these relations are a special case of, or conflict with, constraints in the system's knowledge base. Van Beek then proposed an architecture for assimilating novel plans into the system's plan library and discussed how, if the system learned the novel plan during an interaction, it could use its new knowledge to modify its model of what the user is trying to do.

Most presentations on plan recognition were concerned with passively modeling the user's plans and goals. However, Dekai Wu, University of California at Berkeley, argued that there are situations in which the system should initiate a subdialogue to actively acquire parts of its user model, such as when the system's response to a query depends on which of several possible domain plans the user is pursuing. He presented an approach for evaluating potential user plans based on a multiattribute decision-theoretic notion of expected utility and showed how this approach could be used to suggest appropriate subdialogues for disambiguating the user's intentions.

Default Inferencing

Sandra Carberry, University of Delaware, and Paul van Arragon, University of Waterloo, were both concerned with default inferencing. Carberry presented a model of plan recognition that develops a rational hypothesis about an agent's plan by sanctioning appropriate default inferences while deferring unwarranted decisions until further evidence is available. Her process model uses dynamically constructed preference rules in which alternative possible conclusions are ordered by plausibility. Van Arragon presented a formal tool, Nested Theorist, for handling default reasoning in user modeling. Theorist, a formal framework for reasoning with incomplete knowledge, uses facts and defaults to develop an explanation of a formula; Nested Theorist uses Theorist at multiple levels to model default reasoning, reasoning about another agent's default reasoning, etc. Van Arragon described the capabilities of Nested Theorist and illustrated its application to user modeling.

Modeling User Beliefs

A number of participants addressed the problems of hypothesizing and representing an agent's beliefs. Afzal Ballim, Institute Dalle Molle pour les Etudes Semantiques et Cognitives in Switzerland, described his work with Yorick Wilks, New Mexico State University, on belief ascription. He argued that agents ascribe their own beliefs to other agents unless there is evidence to the contrary and showed how the system's beliefs about another agent's beliefs (nested viewpoints) can be generated on demand using this default ascription rule. Ballim then discussed how lambda expressions, along with evaluation relations specifying how different classes of agents would evaluate each expression, can be used to represent the notion of competency in holding beliefs and thus allow one to model the differing beliefs of classes of stereotypical agents about various topics.

On the other hand, Robert Kass, EDS Center for Machine Intelligence, and Harry Bunt, Institute for Language Technology and Artificial Intelligence in The Netherlands, were both concerned with building more accurate models of an agent's beliefs than is possible by just invoking stereotypes. Kass presented a set of heuristic rules for implicitly acquiring a model of an agent's beliefs during the course of an expert advisement dialogue. His heuristics ascribe beliefs to an agent based on inference rules attributed to that agent, generalizations on the agent's already ascribed beliefs, and principles of rational and cooperative interaction. Kass argued that such a detailed model is necessary for generating responses tailored to individual agents. Bunt presented a formalism for modeling beliefs and intentions that can efficiently handle incomplete information and incremental updating. His formal model consisted of a network of data modules in which each data module took the form of a structured cluster of partial valuation functions representing incomplete information about a particular combination of agents and attitudes.

Deriving User Characteristics from Non-linguistic Behavior

Although most user modeling research has been concerned with deriving user characteristics from an analysis of the user's queries and responses, Martha Crosby and Jan Stelovsky of the University of Hawaii at Manoa demonstrated the utility of unobtrusively monitoring and analyzing other aspects of the user's behavior. They described an experiment in which subjects were asked to read and understand a Pascal program. Their results suggest that duration of eye fixation on different areas of the program suggests the user's level of expertise and his cognitive style. In addition, they conjectured that although individual viewing patterns do not appear to be related to comprehension, they may be indicators of user preferences about the system's response strategy.

Tools for User Modeling

Several of the participants were involved in developing application-independent tools for building and maintaining user models. Alfred Kobsa, University of Saarbrucken, described a user modeling shell system called BGP-MS (Belief, Goal, and Plan Maintenance System). This system provides a rich representation language for conceptual knowledge, a partition hierarchy with inheritance, support for user-defined inferences, and a customizable stereotype management facility for activating and retracting stereotypes. In addition, BGP-MS includes a graphics-based interface for interacting with the system and setting up the knowledge base. Judy Kay, University of Sydney in Australia, advocated a cooperative approach to user modeling in which the user has access to his user model and plays an active role in its construction. She presented the UM system for creating, updating, and applying user models. UM contains a rich set of user modeling tools implemented as rule-based programs; one of these tools, CM, is responsible for modeling the user's conceptual knowledge and includes a

facility for interacting with the user to verify unusual or critical items of information.

Reasoning on a User Model to Produce User-Specific Responses

In recent years a great deal of attention has been given to reasoning on a user model to generate cooperative helpful responses. Margaret Sarner, University of Delaware, discussed research with Sandra Carberry on generating tailored explanations of terms during task-oriented dialogues. Sarner argued that the appropriate content of a definition should guide the selection of a rhetorical strategy rather than the reverse. Her approach first weights the strategic predicates that might comprise a definition and the propositions that might be used to fill them and then selects a rhetorical strategy based on including the information deemed most useful. Less important information is added as necessary to adhere to common rhetorical practices. Mark Maybury, Rome Air Development Center and Cambridge University Computer Laboratory, was also concerned with tailoring definitions to the individual user. He presented a hierarchical text planning system in which rhetorical acts, such as describe, define, and compare, are represented as operators in the system's knowledge base. His operators differentiate between desirable and essential preconditions and take into account the system's beliefs about the user's knowledge, beliefs and desires in constructing a plan for generating a response appropriate to the user.

Ingrid Zukerman, Monash University in Australia, was concerned with augmenting responses to overcome anticipated comprehension problems. She differentiated between connection-related impairments, which occur if a listener is unable to make the intended connections between a lexical item and his knowledge base, and content-related impairments, which occur if the listener's knowledge base is not well enough developed for him to understand the concept being communicated. Zukerman presented a strategy that facilitates better comprehension by reasoning about possible impairments to a user's comprehension of a message and generating appropriate supportive rhetorical devices, including descriptions, instantiations, and similes.

In addition to tailoring responses to the user's domain knowledge and the current situation, responses can take into account preferences and concerns. Ira Haimowitz, MIT, presented a system that explicitly models both the agent participating in the dialogue and the patient or person who is the object of the dialogue and that reasons about the information it intends to convey and its cost/benefit to the patient. Based on this assessment, the system produces empathetic responses that stress favorable information while downplaying unpleasant aspects of the situation.

Reasoning on a User Model to Provide Appropriate Advice

Several of the participants were concerned with providing helpful advice about how to use a system. Ursula Wolz, Columbia University, discussed open-ended situations in which the user must learn enough about a system to perform some task but in which it is impossible to predict a priori what the user already knows and what he will need to know. She contended that such situations lend themselves to the use of an expert consultant that opportunistically attempts to extend the user's expertise. Wolz presented a strategy that reasons on a rich user model to construct responses that not only answer user questions but also satisfy appropriate dynamically constructed pedagogical goals. Alex Quilici, University of California at Los Angeles, extended his previous work on responding to user misconceptions by presenting a strategy that reasons about the relationship between a user's justification for his actions and the user's relevant plan-oriented beliefs. His approach assumes that the user's justifications are the result of reasoning chains constructed via common-sense planning heuristics. Quilici's Correction Machine identifies the user's relevant plan-oriented beliefs by reconstructing these reasoning chains from the user's utterance. It then uses the common-sense planning heuristics to construct a reasoning chain contradicting the user's misconception and uses it as the basis for a response.

Student Modeling in Intelligent Tutoring Systems

Gordon McCalla and Jim Greer of the University of Saskatchewan demonstrated the synergy between research in user modeling and research in intelligent tutoring systems. They described the wide array of problems that a tutoring system faces in modeling the student, emphasized that the nature of tutoring compounds these problems since the student will often have misconceptions and his knowledge will constantly and sometimes dramatically change during the interaction, and described some current strategies that have been developed to address these issues.

Several participants argued that if intelligent tutoring systems are to perform as well as expert human tutors, they must employ multiple instructional approaches that take into account the particular student and situation. Franz Schmalhofer described work with Otto Kuehn, both of the German AI Center at Kaiserslautern, in which a Lisp tutor modeled different learning styles and could invoke one of several different tutoring strategies. Their experiments showed that tailoring the tutoring strategy to the student's style of learning enhanced performance. Robert London of Stanford University and Cimflex Teknowledge Corporation presented a student modeling system called IMAGE that derives multiple partial plans suggesting plausible student behavior. These anticipated plans are used to explain the student's response, compare it against other better alternatives, and provide the GUIDON2 tutoring system with information and advice about appropriate instructional strategies based on cognitive economy.

Human Experiments and their Implications for User Modeling Research

Of particular interest to the participants were several experiments that shed insight on human interaction. Susan Brennan of the Department of Psychology at Stanford University and Hewlett Packard Labs described a study that compared human/computer interaction with human/human interaction. She found that although humans use somewhat different language constructs when interacting with computers from those used when interacting with humans, in both settings humans attempt to collaborate on a successful dialogue and adapt the style of their queries to the kind of responses pro-

duced by the other participant. She argued that this ingrained conversational collaboration should be exploited to design successful natural language interfaces.

Paul McKevitt, New Mexico State University, described a Wizard of Oz experiment in which it was found that particular sequences of speech act types have implications for the structure of the ensuing dialogue and can be correlated with certain aspects of the user, such as his experience in the domain. McKevitt contended that such empirical data, rather than subjective decision-making, should be the basis for constructing user models and argued for the development of automatic techniques for deriving the models.

Summary

The second workshop was as successful as the first, with all agreeing that subsequent workshops should be held more frequently than at four year intervals. Since the general trend has been for researchers in different areas of user modeling to operate in isolation, such workshops are particularly important as a means of increasing cooperation and cross-fertilization of ideas among the subdisciplines. The Third International Workshop on User Modeling is planned for the summer of 1992 in the German Computer Science Meeting Center at Schloss Dagstuhl, near Saarbrucken, Germany. Program co-chairpersons are Dr. Robin Cohen of the University of Waterloo, Bob Kass of the EDS Center for Machine Intelligence, and Cecile Paris of the Information Sciences Institute. Local arrangements co-chairpersons are Elizabeth Andre, Winfried Graf, and Wolfgang Wahlster, all of the German AI Center at the University of Saarbrücken.

About the Author

Sandra Carberry is an associate professor of computer science at the University of Delaware. Her research interests include discourse understanding, user modeling, planning and plan recognition, and intelligent natural language interfaces, and she is the author of a new book entitled *Plan Recognition in Natural Language Dialogue* that is part of the ACL-MIT Press Series in Natural Language Processing.

The First International Workshop on Human and Machine Cognition Pensacola, Florida. Topic: The Frame Problem

Eric Dietrich

In 1859 Charles Darwin published *The Origin of Species* and exposed the complex and fascinating mechanism underlying speciation. Before Darwin's book, we were in fact ignorant not only of the mechanisms underlying speciation, but of the fact that speciation had occurred and was still occurring. Indeed, most supposed that the species that were on the planet at that time were immutable, and that they had existed from the beginning of the Earth. A profound problem—how to tailor an organism to suit a particular environment—was being solved right under our noses. After Darwin's book, and in part because of it, a new field emerged whose cumulative insights over the years have revealed many of the mechanisms involved in evolution.

In 1877 the Italian astronomer Giovanni Schiaparaelli announced the existence of *canali* on Mars: a network of straight and curved lines running across the planet. *Canali,* meaning channels or grooves in Italian, was translated by the English press into "canals," and with that an intense love affair with Mars and its inhabitants began. Of course, the affair had its share of dark misgivings: what if the inhabitants should turn out to be smarter than we with the same bent for conquering and enslaving? At the center of all this was Percival Lowell, the builder of the Lowell Observatory in Flagstaff, Arizona. He championed the seemingly profound problem of the Martian canals and the nature of the beings on Mars. He spent his life trying to unravel these mysteries. The image of Mars as a planet inhabited by an ancient, canal-building civilization became the popular vision, even though many other astronomers could not find the canals Lowell observed regularly and frequently. Today we know that there are no canals on Mars, and that Lowell and others had seen illusions and atmospheric distortions.

In 1969 John McCarthy and Patrick Hayes formulated the *frame problem:*

> ...in proving that one person could get into [a phone] conversation with another, we were obliged to add the hypothesis that if a person has a telephone *he still has it* after looking up a number in the telephone book. If we had a number of actions to be performed in sequence, we would have quite a number of conditions to write down that certain actions *do not change* the values of certain [propositional functions describing situations] (McCarthy and Hayes, p. 487, emphases added).

More generally, the frame problem is the problem of blocking the vast number of inferences about what has not changed as the result of performing some action *A* while allowing the small number of inferences about what has changed as a result of *A*.

For some of us, the "Frame Problem Workshop" (as it was called) was an opportunity to discuss a methodological question which has become important in AI and cognitive science: Is the frame problem profound or a mistake? Which of the above episodes from the history of science, the "evolution" episode or the "Martian canal" episode, does the frame problem most closely resemble? To some researchers, solving the frame problem will unravel the secrets of the higher cognitive processes and intelligence. To others, the frame problem is much ado about nothing; its very generation rests on seeing something which isn't there. Many of us who are interested in this methodological question have been influenced by Jerry Fodor's book *The Modularity of Mind* (1983). Fodor