

WORKSHOP REPORT

First International Workshop on User Modeling

Sandra Carberry

The First International Workshop on User Modeling in Natural Language Dialogue Systems was held 30-31 August 1986 in Maria Laach, West Germany. Issues addressed by the participants included the appropriate contents of a user model, techniques for constructing user models, strategies for reasoning on user models in both understanding and generating natural language dialogue, and the development of general user-modeling systems. This article includes an overview of the presentations made at the workshop. It is a compilation of the author's impressions and observations and is, therefore, undoubtedly incomplete; and at times might fail to accurately represent the views of the researcher presenting the work.

The First International Workshop on User Modeling in Natural Language Dialogue Systems was held on 30-31 August 1986 in Maria Laach, a small village located about 30 miles south of Bonn, West Germany. The workshop was organized by Dr. Wolfgang Wahlster and Dr. Alfred Kobsa, both of the University of Saarbrücken, and was supported by a grant from the German Science Foundation in its Special Collaborative Program on AI and Knowledge-Based Systems. Twenty-four invited researchers from seven countries participated in the workshop. The program included both long and short talks on current research ideas and projects and lively discussion among the participants; often-times, the participants became so engrossed in the presentations and ensuing discussions that other aspects of the program, including the banquet, had to be delayed. But all agreed the workshop had been an enjoyable experience and extremely worthwhile.

Workshop Program

The workshop program included talks on a wide spectrum of topics related to user modeling in natural language dialogue systems. Issues addressed by the participants included the appropriate contents of a user model, techniques for constructing user models, strategies for reasoning on user models in both understanding and generating natural language dialogue, and the development of general user-modeling systems.

Wolfgang Wahlster began the workshop with a discussion of some basic issues in user modeling, such as the content and function of a user model and user-modeling component, modeling agents versus patients, and dis-

course versus user models. He then described several German projects in which user modeling plays a significant role: Hamburg Application-oriented Natural Language System (HAM-ANS), a transportable natural language system that has been interfaced to a hotel reservation system, an image-sequence analysis system, and a relational database; Expert Translator (XTRA), a natural language interface to an expert system for assisting users in completing tax forms; Visual Translator (VITRA), a natural language interface to a vision system; Sinix Consultant (SC), a natural language help system for SINIX; and Wissenbasierter Beratungsdiallog (WISBER), a natural language financial-consultation system. He then challenged the participants to consider a number of questions during the course of the workshop, including (1) What do the terms user model and user-modeling component mean? (2) What services should a transportable user-modeling component provide, and what architecture is most appropriate? (3) What knowledge representation paradigm is best for user modeling, and how should collective beliefs, such as those of organizations, be dealt with? (4) How can dialogue systems that are transportable to new domains and diverse conversational situations be developed? and (5) What limitations currently exist regarding user modeling, and what major breakthroughs are most likely to be realized in the near future?

The subsequent presentations addressed many of these questions; however, the participants sometimes held opposing views, and it appears too early to form a consensus on these issues.

Modeling the User

Several presentations addressed issues and concerns relevant to all research in user modeling. Alfred Kobsa, University of Saarbrücken, discussed the wide variety of system assumptions about the user's beliefs that are being modeled in proposed and realized systems. He presented a taxonomy of beliefs and goals and proposed using this taxonomy to classify current user-modeling components with respect to the kinds of beliefs modeled by these systems, to evaluate the expressive power of various representation schemes, and to enable the designer of a user-modeling component to consider exactly which beliefs

model construction that is system driven (meaning that the model is derived from a special portion of the session during which users are prompted to characterize themselves) and model construction that is user driven (meaning that the model is inferred during the course of normal interaction). She then showed how the complexity of the system developer's task depended on the particular conversational setting being modeled and the types of changes being made to the model.

Karen Sparck Jones, the Computer Laboratory at the University of Cambridge, emphasized the complexity of the user-modeling task. She distin-

be realistic about what one can hope to accomplish, and not devote extensive system resources to constructing deep models of the user.

A General User-Modeling System

Ethel Schuster, University of Pennsylvania, presented work by Tim Finin, also of the University of Pennsylvania, on a general user-modeling component for constructing and updating long-term user models. The system uses a combination of techniques, including stereotypes (to generalize about members of user classes) and explicit default rules (to make additional inferences likely to be true about a particular user). The architecture for the user-modeling component is domain independent and is intended to be useful for many applications.

...one should adopt a conservative approach to user modeling, be realistic about what one can hope to accomplish, and not devote extensive system resources to constructing deep models of the user.

Plan Recognition and User Modeling

Sandra Carberry, University of Delaware, addressed the problem of dynamically constructing a model of the underlying task-related plan motivating an information seeker's queries. Her modeling framework relates a new utterance to the domain-dependent set of plans and goals to hypothesize a set of candidate-focused plans; focusing heuristics then evaluate the possible relationships between the candidate-focused plans and the existing dialogue context to determine which candidate-focused plan represents the user's actual focus of attention in the dialogue. The user model is expanded to include this plan and any enlarged global context inferred as a result of its inclusion. Carberry then showed how the system could reason on this user model to understand a class of pragmatically ill-formed utterances that current natural language systems are unable to handle.

and goals should be modeled in the particular application. He then described a variety of current systems and categorized their user-modeling components according to his proposed taxonomy.

Katharina Morik, the Technical University of Berlin, addressed the complexity of the user-modeling task. She differentiated conversational settings and types of models according to several criteria. The first criterion concerns what can be changed by actions modeled by the system, such as the domain of discourse itself (that is, changes to the real world), the system's knowledge about the domain, and the system's knowledge about its user's beliefs and goals. The second criterion centers on the means by which changes can be accomplished, such as noncommunicative, nonmental actions; communicative actions; observation; and thinking and inferring. The third criterion concerns who performs the actions, whether the system or the user.

She further distinguished among five types of changes to the user model and differentiated between

guished between the user's role as agent or patient; whether the properties being modeled are objective or subjective, static or dynamic; and whether the user model's function is to enhance the system's effectiveness, efficiency, or acceptability. She contended that the difficulty in obtaining reliable, relevant information limits the user models that can realistically be developed. For example, Sparck Jones defined *decision properties* as properties of the user that influence the system's decision making; although many systems attempt to model such properties of the user, they often are not explicitly communicated but must be inferred indirectly during the course of the dialogue. She showed how a user's input to the system can fail to reliably indicate a decision property or discriminate among several possible properties and suggested that the problem is even worse when attempts are made to identify underlying, rather than superficial, domain goals. Sparck Jones' examples were aimed at the contention that one should adopt a conservative approach to user modeling,

Much of the previous research in plan recognition has adopted a set of restrictive assumptions resembling a near-perfect world in which miscommunication never occurs. Brad Goodman, Bolt Beranek and Newman Laboratories Inc., addressed the problem of plan recognition in situations that are typical of real-world communication. In his previous research, Goodman handled the problem of miscommunication in reference by viewing refer-

ence resolution as a negotiation process during which listeners negotiate with themselves by applying a set of relaxation rules to identify the referent intended by a speaker. He proposed applying the same kind of techniques to the problem of identifying the goals and plans that the user intends to communicate, with clarification dialogues used to resolve any remaining ambiguities or misunderstandings. Carberry was also concerned with removing the restrictive assumptions of current planning systems and proposed a four-phase approach to detecting and correcting errors that might enter into the system's model of the user's plan.

Alexander Quilici, University of California at Los Angeles, presented a system that employs a three-step process to provide intelligent advice to UNIX operating system users. Given a natural language description of a problem the user has encountered, the system understands the problem by modeling the user's intentions; analyzes the user model and similar, previous experiences to classify the user's problem; and then produces helpful advice by applying heuristics associated with the identified problem class. Quilici's talk emphasized construction of the system's model of the user's intentions and classification of user problems.

Elaine Rich, MCC, discussed ongoing research into plan recognition when the interaction with the user includes both a linguistic discourse and a nonlinguistic discourse (such as clicks of a mouse). She described problems that must be addressed in coordinating the understanding of both kinds of discourse, including combining pieces of information whose level of granularity is different.

Modeling the User to Produce User-Specific Understanding

Jaime Carbonell, Carnegie-Mellon University, addressed the problem of robust understanding and discussed several ways that a model of the user can be employed for this purpose. He described the XCALIBUR system and its use of task knowledge to resolve ambiguities. Carbonell then discussed new research being pursued with Jill

Fain; the objective of this current work is to dynamically model a user's idiosyncratic use of language and exploit this user model to adjust to a particular user's language patterns, thereby producing a customized system for parsing and understanding.

Reasoning on a User Model to Produce User-Specific Responses

Aravind Joshi, University of Pennsylvania, described his work with Bonnie Webber and Ralph Weischedel on formalizing the behavior expected of a cooperative expert system. Joshi's revised maxim of quality dictates that it is not sufficient merely to produce truthful answers; a speaker has an obligation to prevent the listener from drawing false conclusions from the speaker's statements. Joshi showed how a cooperative system must modify its response so that the response incorporates further information blocking false inferences the user might otherwise make. He then discussed constraining these extended responses and the role of a user model in this process.

Robin Cohen, University of Waterloo, described work by herself, Peter van Beek, and Marlene Jones directed toward tailoring explanations to meet the user's needs. In the first project, the user's domain-dependent plans are compared with alternative plans generated by the system in order to produce explanations that address the user's higher domain goals. The second project involves incorporating a user model into a system for educational diagnosis. Such a system might be required to supply explanations of its findings to many different kinds of users and, thus, must address the different background knowledge, needs, and desires of each. Cohen described a framework for partitioning the system's knowledge in order to generate explanations that take into account the differences between system and user knowledge.

Kathy McCoy, University of Delaware, presented a strategy for reasoning on a user model to correct user misconceptions about the domain. Her user model reflects the system's beliefs about how the user views the domain, including a notion of object

perspective that causes certain features of the knowledge base to be highlighted as a result of preceding dialogue. McCoy analyzes the structural configurations of the user model to suggest the cause of a misconception. Natural language corrections are generated based on the type of misconception and the user's perspective on the objects involved. Thus, her response strategy is context sensitive and will vary according to the system's model of the user.

Reasoning on a User Model to Address the User's Level of Expertise

Cecile Paris, Columbia University, addressed the problem of tailoring object descriptions to the user's specific domain knowledge. She showed that the user's level of domain knowledge affects not only the amount of information that is normally provided by human question answerers but also the kind of information. She distinguished between two descriptive strategies, one parts oriented (typically used when describing objects to experts) and the other process oriented (typically used when describing objects to those naive about the domain), and showed how they could be mixed to produce a description tailored to the domain knowledge that the system believes the user has.

David Chin, University of California at Berkeley, described a user-modeling component that constructs a model of a user's knowledge about the UNIX operating system. Two sets of stereotypes are used, one representing user expertise and the other the difficulty level of information. Together with metaknowledge about the system's own limited knowledge, the user model is reasoned on to detect user misconceptions and produce responses tailored to the user's knowledge and domain expertise.

Modeling and Dialogue

Ethel Schuster, University of Pennsylvania, discussed the relationship between user models and discourse models. She contended that a discourse model, containing representations of entities, their properties, and relations among entities which the

system believes have been communicated during a discourse, is one component of a user model and that over time some aspects of this temporary discourse model become part of the permanent user model. Schuster's talk focused on the role of this aspect of user modeling in the use of anaphoric expressions.

Marco Columbetti, Milan Polytechnic, discussed his work with Gabriella Airenti and Bruno Bara on the interaction level of communication, in which the participants communicate the relationship of their utterances to the ongoing dialogue. He described several kinds of interpersonal games between system and user, illustrating the role they play in identifying the communicative goal of an utterance. Columbetti contended that a representation of such interpersonal games defining expectations about linguistic goals and actions is an essential part

of a user model that achieves cooperative interaction.

Bill Mann, Information Science Institute, discussed the importance of user modeling in text generation. He surveyed the kinds of knowledge that must be included in a user model if it is to be used to generate effective text; for example, the user model must be sufficiently rich to enable the system to reason about whether a certain fact is already obvious to the user and, therefore, should not be conveyed and to predict the effect a particular utterance will have on an individual user. He concluded that richer user models are necessary for successful user interfaces.

Anthony Jameson, University of Nijmegen, discussed how the impression a speaker wants to create in the mind of the listener affects decisions about what to say. He described a system that models the behavior of an

informant engaged in an evaluation-oriented dialogue; the system pursues several kinds of reasoning, including assessing the impact its utterances are likely to have on the listener's judgments and reasoning about how the listener's likely inferences and possible misconceptions about the speaker's dialogue goals can be used to the speaker's advantage. Overall, the system attempts to respond in a manner that creates the most favorable impression in the mind of the listener.

Summary

Selected papers from the workshop are being expanded and integrated into a book entitled *User Models in Dialogue Systems*, edited by Alfred Kobsa and Wolfgang Wahlster, to appear as part of the Springer Series on Symbolic Computation. Some papers will also appear in a special user modeling issue of *Computational Linguistics Journal*.



Halbrecht Associates, Inc.

Executive Search Consultants

SERVING THE ARTIFICIAL INTELLIGENCE COMMUNITY

Halbrecht Associates, founded in 1957, was the first search firm in the United States to specialize in Information Systems Management, Operations Research, Management Sciences, Telecommunications, Computer Sciences and Decision Support Systems. We are now also in the forefront of the Artificial Intelligence field. Because we have been so extensively involved in this emerging technology, our portfolio of people includes extensive numbers of AI professionals at all levels, from the United States and abroad.

Our client list includes companies coast to coast, large and small, including start-ups with equity opportunities. Accordingly, depending on your needs, we are in a unique position to assist you rapidly with any of your requirements.

Our current assignments range downward from several Chief Scientists, or those who can provide the technical leadership to head up major AI programs, to Senior LISP Programmer. If you would like to be considered to be part of our AI data base, please send us your credentials. This will enable us to give you first notice of appropriate opportunities. All inquiries are in strictest confidence.

Contact: Daryl Furno, Senior Associate

1200 Summer Street • Stamford, CT 06905 • 203-327-5630

For free information, circle no. 193