

# PREFACE

## 1 THE HISTORY OF THIS VOLUME

User modeling is now approaching its second decade. Its exact age, however, is difficult to determine. Was it born in 1978, when Allen, Cohen, and Perrault started to publish their seminal papers on dialog processing in natural language systems based on Searle's theory of speech acts and the speaker's intentions behind it? Or was it in 1979, when Elain Rich reported on her GRUNDY system, which contained an explicit representation of what personality traits influence users' preferences of books to read, and built a model of its current user by drawing assumptions as to the degree to which these personality traits apply to the specific user, for being able to give better advice?

There were predecessors to these systems, though, which certainly do not deserve to be neglected. Hayes and Rosner (1976) envisaged a dialog system that would play the role of a curious party guest and try to find out as much as possible about the beliefs of other party guests. In the course of this fictitious dialog, the system built a model of the dialog partner. In 1979, Power published a report on a system dating back to 1974 which simulates a cooperative dialog between two agents trying to achieve some simple real-world task. Their dialog planning is based on a rudimentary model of the beliefs and goals of the other agents.

The idea that, for understanding (the intentions behind) natural language utterances, "it is necessary to have a model of the beliefs of others" probably dates back within AI to papers by Bruce (Bruce and Schmidt 1974, Bruce 1975). Its roots, however, can certainly be traced far back within philosophy, at least to Leibniz and Locke (see e.g., their epigone Christian Freyherr von Wolff, 1712: "Wenn also zwey Personen miteinander reden, und einer den andern verstehen soll; so wird erfordert, 1. daß der, so da redet, bey einem jeden Worte sich etwas gedencken könne: 2. daß der, so ihn reden höret, eben dasjenige sich bey einem jeden Worte gedencken kan, was der andere dencket."').<sup>1</sup>

Now, after 10 years of user modeling, it is certainly acknowledged in artificial intelligence that, in order to be capable of exhibiting *pragmatically correct* dialog behavior, an AI system must include a model of the user containing assumptions about his/her background knowledge as well as his/her goals and plans in consulting the system. Research in the field of user models

investigates how such assumptions can be automatically created, represented, and exploited by the system in the course of interaction with the user.

A dozen major and several more minor user modeling systems have been designed and implemented in the last decade, mostly in the context of natural language dialog systems. The goal of UM86, the first international workshop on user modeling, was to bring together the researchers working on these projects, so that results could be discussed and analyzed, and hopefully general insights be found, that could prove useful for future research. The meeting took place in Maria Laach, a small village some 40 miles south of Bonn, West Germany. Twenty-five prominent researchers had been invited to participate. The pleasant setting of the conference site close to the medieval abbey of Maria Laach and the volcanic Lake Laach fostered a nice atmosphere for intensive discussions and the exchange of ideas until the early hours of the morning.

What were the direct results of the workshop? It was agreed that a user model is "merely" a special case of the more general concept of "agent model" (for further details on all results mentioned here, see Kass and Finin's survey paper in this volume). In certain conversational settings, it will be necessary for a dialog system to maintain both a model of its current user and of other agents being talked about, though in most cases a model of the user alone is sufficient. Several classification criteria for user models have been proposed and discussed which are based on the nature of the task domain and the conversational role that the system is supposed to fill. No consensus could be reached in a lively discussion on the relationship between user models and so-called "discourse models". Opinions ranged from regarding these notions as being distinct both on a conceptual and an implementational level to claiming that user models subsume discourse models at both levels.

It was agreed that two forms of publication should result from the meeting: a number of papers that are directly related to natural language should appear in a special issue of the *CL* journal, and a more general survey book on user modeling should be published by Springer. Moreover, the discussion on the relationship between user models and discourse models which popped up at the workshop should be made available in

Copyright 1988 by the Association for Computational Linguistics. Permission to copy without fee all or part of this material is granted provided that the copies are not made for direct commercial advantage and the *CL* reference and this copyright notice are included on the first page. To copy otherwise, or to republish, requires a fee and/or specific permission.

0362-613X/88/0100●●\$03.00

printed form (see the discussion section in this issue, and its special introduction).

All papers were subjected to an unusually extensive reviewing procedure: Speakers had to distribute tentative versions of their papers long before the workshop, which were then discussed at the meeting. Tape recordings of these discussions were made and distributed to all speakers, and written reviews of each paper were prepared by all participants. Selected speakers were then invited to submit their revised and extended papers to the special *CL* issue. The *CL* reviewers and the editors added additional comments which the authors had to take into account.

## 2 SURVEY OF THE CONTRIBUTIONS TO THIS ISSUE

The volume starts with an excellent survey of the field of user modeling in natural language dialog systems by Robert Kass and Tim Finin. The authors present definitions for the concept of "user model" and the more general concept of "agent model". They make suggestions as to what should be regarded as the content of a user model, and list a number of dimensions along which user models can be classified. Several techniques are discussed as to how the system can acquire assumptions about the user's goals, plans, and beliefs in the course of the dialog. Finally, for assessing the costs and benefits of integrating a user model into an application system to be developed, a number of relevant dimensions are proposed along which interaction with the system can be classified.

The order of the following papers reflects the processing sequence in user modeling systems, starting with papers that focus on the recognition of beliefs and goals of the user, and ending up with papers that focus on the exploitation of these assumptions.

The contribution by Sandra Carberry is concerned with the recognition of users' plans in an ongoing dialog. In her model, a number of domain-independent heuristics are used for inferring candidate goals and plans from user statements (these heuristics are extensions of inference rules proposed by Allen and Perrault 1980, Allen 1983). Focusing strategies are then employed for selecting the candidate goals and plans most related to the current dialog context and for integrating them into a model of the user's actual plan. These strategies are based on recent research on dialog organization and focusing done, for example, by McKeown (1985a,b) and Grosz and Sidner (1986).

Even before all its details have been determined, the assumed user plan can be exploited for generating expectations about dialog contributions of the user. Carberry's model uses this potential for handling two forms of linguistically problematic dialog contributions, namely pragmatically ill-formed sentences and intersentential ellipsis (only the former is described in this contribution, however). Pragmatically ill-formed user sentences, in her definition, are syntactically and se-

matically correct, but do not fit the system's model about the domain. Carberry discusses pragmatic ill-formedness resulting from attribute confusions and missing specifications in a user's question, and proposes repair strategies based on the expectations from the user model.

Up till now, nearly all plan recognition models have assumed that the system's model of the user's plan is perhaps incomplete, but never wrong. In a final section, Carberry discusses how her analysis could be extended to be able to detect clues indicating disparities between the system's model of the user and the user's actual goals and plans, reasoning on the system's model and the system's domain knowledge to form hypotheses as to the source of these disparities, responding to the user and negotiating with him/her to isolate the errors, and appropriately repairing the user model.

The advisor model of Alex Quilici, Michael Dyer, and Margot Flowers presents an explanation-based approach to the problem of recognizing and responding to user misconceptions. Their model assumes that a number of user beliefs and goals have been recognized through analysis of the user's natural language reports about problems in the use of UNIX, and entered into the user model. The modeled beliefs concern applicability conditions, enablements, and effects of UNIX commands.

For each of the three belief types discerned, a number of explanation patterns exist for determining why beliefs of that type might not be held by the system. For example, the system's disbelief in *action Ap causes state Sp* can be explained by the beliefs *Ap causes So* and *Ao causes Sp*, i.e., *Ap* has a different effect and *Sp* is achieved by a different action. When a recognized user belief is not shared by the system, the explanation patterns pertaining to the specific belief type are matched with the system's knowledge about UNIX in order to find an explanation for why the system does not hold this belief. The explanation thus discovered may then be forwarded to the user.

A frequent type of user misconception is the misclassification of objects (i.e., their subsumption under a wrong superconcept in a concept hierarchy). Another type is the belief that objects have certain attributes or attribute values that they do not actually have. Kathleen McCoy's contribution deals with how to behave after such misconceptions have been detected and classified as belonging to one of these categories. For each type of user misconception, she proposes three strategies for generating appropriate responses. The generated response patterns are then transformed into natural language responses using McDonald's (1980) MUMBLE system.

In all of these strategies, the similarity of objects plays an important role. Similarity of objects is defined via the similarity of the attributes of these objects. However, not all attributes are taken into account, but only those which are salient in the current dialog

context. McCoy discerns three types of salience, each stemming from a different source: *static salience* encoded directly in the generalization hierarchy, *dynamic salience* gained by an attribute being explicitly mentioned in a discourse, and *perspective salience* gained by viewing the domain objects through the active perspective.

Perspective salience turns out not to be a simple substructure of the concept hierarchy, but to form an independent structure (namely a set of highlighted attributes), which is orthogonal to the concept hierarchy. Detailed examples from the field of financial advice are given to illustrate the effect of perspective on attribute salience and on one of the above-mentioned response strategies. Suggestions are made as to what influences the choice of a perspective in a given dialog.

As is the case with McCoy, Cécile Paris' model also assumes that inferences have already been drawn about the experience of the user with respect to the domain of discourse. These assumptions are represented by a list of those items in the system's knowledge base which are known to the user, and by information about whether the user understands the basic concepts underlying the domain of discourse.

In Paris' model, these assumptions are exploited for generating user-adapted object descriptions. The author compared a number of encyclopedias for children and adults and found that different types of descriptions are used in each of these two text types, which are obviously aimed at two different types of readership: for readers who can be expected to possess some background knowledge about the domain, objects or concepts are described in terms of their subparts and the properties of these subparts. This pattern of description has been identified in studies on text organization by McKeown (1985a,b), who called it a "constituency schema". Descriptions for "naive" readers, on the other hand, essentially describe the processes associated with the operation of the object. This pattern is called a "process schema" by Paris. The author explains her findings by assuming that readers with background knowledge—in contrast to "naive" readers—can be expected to figure out how parts fit together to form an object capable of performing a function. Providing such information would thus contradict Grice's (1975) maxim of quantity and should therefore be avoided.

Paris' analysis suggests that the user's level of domain knowledge affects not only the amount, but also the kind of information provided in descriptions. The author suggests that, based on a user model, dialog systems should also employ these two distinct strategies for generating descriptions that are well-adapted to the user. Since users are hardly ever either completely naive nor completely expert, she also suggest that a dialog system should combine both strategies, for example in describing complex physical systems: Whenever the user model indicates that the user possesses

local expertise about an involved object that is to be described, the constituency schema can be employed; otherwise the process schema must be used. Paris outlines the points in each description schema at which a switch to the other schema is possible.

The proposed principles have been implemented by the author in TAILOR, a system that—roughly—takes a request for an object description as its input, accesses the user model described above, and generates a conceptual representation of an object description that is adapted to the user's level of expertise. After its concepts have been replaced by lexemes, this description is passed on to a generator that unifies it with a functional grammar to produce English sentences. The operation of the program is illustrated by a detailed example. In a final section, Paris briefly presents heuristics that might be used to acquire the user model preassumed by her.

#### SOME WORDS ABOUT THE AUTHORS

**Sandra Carberry** received a B.A. degree in mathematics from Cornell University and an M.S. degree in computer science from Rice University, and was a Member of Technical Staff at Bell Telephone Laboratories. She is now an assistant professor of computer and information sciences at the University of Delaware, where she received her Ph.D. in 1985. Portions of her paper are based on this dissertation research. Her current interests are robust models of plan recognition and the application of user models to the generation of helpful responses.

**Timothy W. Finin** is a technical director at the UNISYS Paoli Research Center. He received a master's degree in electrical engineering from MIT, and M.Sc. and Ph.D. degrees in computer science from the University of Illinois at Urbana-Champaign. He did research at MIT in the area of robotics and computer vision, and was an assistant professor in the Department of Computer and Information Science at the University of Pennsylvania from 1980 to 1987. His current research interests include knowledge representation, expert systems, and computational linguistics, as well as their applications to intelligent user interfaces.

**Robert Kass** just finished his Ph.D. studies at the University of Pennsylvania. His dissertation concerned the implicit acquisition of user models. He is now working at the Center for Machine Intelligence in Ann Arbor, Michigan.

**Kathleen McCoy** is an assistant professor of computer and information sciences at the University of Delaware. She received her Ph.D. in computer science from the University of Pennsylvania in 1985. Portions of her paper are based on this dissertation research. Her current interests center on contextual effects of previous discourse on representation of user models and generation of subsequent responses.

**Alex Quilici** is currently completing his Ph.D. studies at the University of California, Los Angeles. **Michael G.**

**Dyer and Margot Flowers** did their graduate work at Yale University before joining the faculty at UCLA. Quilici's research concerns automatically detecting and correcting plan-oriented misconceptions that can occur in argumentative dialogs. The current paper summarizes aspects of this research involving the misconceptions of novice computer users. This work is part of ongoing research programs by Dyer and Flowers in argumentation, user modeling, language acquisition, and connectionist models for language comprehension.

**Cécile L. Paris** received her bachelor's degree from the University of California at Berkeley and her Ph.D. in computer science from Columbia University. She is now working at the Information Sciences Institute, continuing her research on natural language generation and user modeling. Her thesis work has focused on how a system should tailor a response depending on how much the user knows about the domain under consideration. Her paper reports major results from this work.

#### ACKNOWLEDGEMENTS

The preparation of the Maria Laach workshop and the production of this publication was supported by a number of individuals or institutions, to which we would like to express our gratitude:

- The Collaborative Research Programme on Artificial Intelligence and Knowledge-Based Systems (SFB 314) of the German Science Foundation (DFG), which funded the workshop;
- Mark Line and Bernd Schäfer, who prepared transcripts and rendered the editors' English intelligible even to native speakers;
- Bernd Nessen and Sokrates Evangelidis for their work as sound engineers; and
- Doris Borchers, Gabriele Jacquinet, Johannes Reinert, and Parinaz Mohammadzadeh for their administrative assistance.

The international user modeling community owes a lot to them.

Alfred Kobsa, Wolfgang Wahlster

#### REFERENCES

- Allen, J.F. and Perrault, C.R. 1978 Participating in Dialogues: Understanding via Plan Deduction. AI-Memo 78-4, University of Toronto, Canada.
- Allen, J.F. and Perrault, C.R. 1980 Analyzing Intention in Utterances. *Artificial Intelligence* 15: 143-178.
- Allen, J.F. 1983 Recognizing Intentions from Natural Language Utterances. In Brady, M., and Berwick, R.C. (eds.) *Computational Models of Discourse*. MIT Press, Cambridge, MA.
- Bruce, B., and Schmidt, C.F. 1974 Episode Understanding and Belief-Guided Parsing. NIH Report, CMB-TR-32, Department of Computer Science, Rutgers University, New Brunswick, NJ.
- Bruce, B. 1975 Belief Systems and Language Understanding. BBN Report 2973 (*Artificial Intelligence* Report 21). Bolt, Beranek, and Newman, Inc., Cambridge, MA.
- Cohen, P.R. 1978 On Knowing What to Say: Planning Speech Acts. TR-118, Department of Computer Science, University of Toronto, Canada.
- Cohen, P.R. and Perrault, C.R. 1979 Elements of a Plan-Based Theory of Speech Acts. *Cognitive Science* 3: 177-212.
- Grice, H.P. 1975 Logic and Conversation. In Cole, P. and Morgan, J. (eds.) *Syntax and Semantics 3: Speech Acts*. Academic Press, New York, NY.
- Grosz, B.J. and Sidner, C.L. 1986 Attention, Intentions, and the Structure of Discourse. *Computational Linguistics* 12: 175-204.
- Hayes, P.J. and Rosner, M.A. 1976 ULLY: a Program for Handling Conversations. In *Proceedings of the 1976 AISB Summer Conference, Edinburgh, Scotland*: 137-147.
- McDonald, D.D. 1980 Natural Language Production as a Process of Decision Making under Constraints. Ph.D. thesis, MIT, Cambridge, MA.
- McKeown, K.R. 1985a Discourse Strategies for Generating Natural-Language Text. *Artificial Intelligence* 27: 1-41.
- McKeown, K.R. 1985b Text Generation: Using Discourse Strategies and Focus Constraints to Generate Natural-Language Text. Cambridge University Press, Cambridge, England.
- Perrault, C.R. and Allen, J.F. 1980 A Plan-Based Analysis of Indirect Speech Acts. *American Journal of Computational Linguistics* 6: 67-182.
- Power, R. 1974 A Computer Model of Conversation. Ph.D. thesis, University of Edinburgh, Scotland.
- Power, R. 1979 The Organisation of Purposeful Dialogues. *Linguistics* 17: 107-152.
- Rich, E. 1979a Building and Exploiting User Models. Ph.D. thesis, Department of Computer Science, Carnegie-Mellon University, Pittsburgh, PA.
- Rich, E. 1979b User Modeling via Stereotypes. *Cognitive Science* 3: 329-345.

#### NOTE

1. Congenial translation: "Whenever two personnes, then, shall speake one to the other, such that each understand the other, so is it of necessity that, in the first instance, he, that personne just speaking, shall, for each and every word, think a thing; and that, in the second instance, he, that personne hearing that which be spoken by the other, shall, for each and every word, think just that thing which was thought by the other."