

User Modeling and User Interfaces

Kathleen R. McKeown

Columbia University
Department of Computer Science
450 Computer Science Building
New York, N.Y. 10027
MCKEOWN@CS.COLUMBIA.EDU

In the past several years, it has become increasingly apparent that there is strong disagreement between researchers within and outside of AI on how to build systems for man-machine communication. Within AI, and specifically within natural language, researchers have noted that human speakers and hearers draw on their knowledge about each other when communicating. This knowledge is used both in understanding, and responding to, a speaker's utterances. *User models* are a means for representing various types of information about speakers and hearers so that systems are able to reason about their users when interpreting input and producing responses. While there is disagreement in the AI community about the form of user model that should be used, there is an implicit assumption that some form of knowledge about users is essential for successful man-machine communication.

In contrast, in the user interface and information retrieval community there has been a reaction against user models. Researchers in this community have argued that with a properly designed interface, users are able to get the information that they need. Modelling a user requires access to and representation of all sorts of ill-defined aspects of human cognition (e.g., how do we know when a user believes that a specific fact is true?). Rather than being so presumptuous as to assume that we can carry out such a task, why not rely on the fact that users, as humans, are smart enough to be able to use interfaces as intended to satisfy their needs?

The purpose of the panel is to bring together these disparate opinions in an open forum. Are we really as far apart as we seem? Are there opinions in the opposing viewpoints that ought to be adopted by the other group?

Just What Counts as a User Model?

This is a good question, since every researcher has a different definition of a user model. In previous heated discussions over the pros and cons of user models, it has become apparent that sometimes what one person thinks is not part of a user model, another counts as part of a user model. There have been at least three different types of user models proposed in the literature:

- A model that represents specific user beliefs, user goals, and possibly plans that the user has

for achieving those goals. Often these beliefs have been inferred from previous dialog with the user. They may or may not have been explicitly stated. In some systems a distinction is made between speaker beliefs, hearer beliefs, and mutual beliefs (those beliefs that both participants recognize both hold). This trend in user modelling was initiated by Allen and Perrault [Allen & Perrault 80] and has been followed up by a variety of researchers [Carberry 83, Litman & Allen 84, Pollack 86, Sidner and Israel 81].

- A model of different types of users that are likely to use the system. These user classes are often called stereotypes. For example, a common distinction is often made between experts and novices in the domain. The system makes assumptions about the sorts of things that a user is likely to know based on the class that s/he falls into. This class of user modelling was initiated by Rich [Rich 79] and also has numerous followers [Chin 86, Wallis and Shortliffe 82].
- A model based on actual observation of the user. This may include actions the system has observed the user carrying out or assertions the user has made. For example, in a help system for software systems, the system may have access to observations about the commands and command sequences that the user has used in the past. Some researchers classify this as a discourse model [Shuster, E. 88], but it has also been used specifically as a user model [Wolz *et al.* 88, Finin 83].

Questions Posed to the Panellists

As moderator, I have asked the panellists to consider several questions in addition to the desirability of user models which, hopefully, will help focus discussion on the central issues. Given that there have been misconceptions in the past regarding what counts as a user model, panelists should be specific about the type of user model they support or oppose.

While it is impossible to restrict discussion to a particular medium for interaction (e.g., menus, NL dialog) given the diversity of backgrounds, I've asked that we avoid discussion on which interaction medium is best for a task. Rather, we need to consider whether user models are desirable regardless of medium.

Since different system tasks may all have different requirements for man-machine interaction, we will restrict ourselves to the same system tasks. In order to allow for the possibility that user modelling is more useful in some domains than others, we will use more than one task. Possibilities include information gathering (e.g., a natural language database system or information retrieval) and a tutorial system.

ACKNOWLEDGMENTS

Research on user modelling and language generation at Columbia is supported by DARPA contract #N00039-84-C-0165, ONR grant N00014-89-J-1782 and NSF grant IRT-84-51438.

REFERENCES

- [Allen & Perrault 80] Allen, J.F. & Perrault, C.R. Analysing Intentions in Utterances. *Artificial Intelligence* 15(1):pages 143-178, January, 1980.
- [Carberry 83] Carberry, S. Tracking User Goal in an Information Seeking Environment. In *Proceedings of AAAI-83*. American Association for Artificial Intelligence, 1983.
- [Chin 86] Chin, D. User Modelling in UC, The UNIX Consultant. In *Proceedings of Computer Human Interaction*, pages 24-28. Boston, Mass., 1986.
- [Finin 83] Finin, T. Providing Help and Advice in Task-Oriented Systems. In *8th International Joint Conference on Artificial Intelligence*, pages 176-8. Karlsruhe, Germany, August, 1983.
- [Litman & Allen 84] Litman, D.J. and Allen, J.F. A Plan Recognition Model for Clarification Subdialogues. In *Coling84*, pages pages 302-311. COLING, Stanford, California, July, 1984.
- [Pollack 86] Pollack, M.E. A Model of Plan Inference that Distinguishes between the Beliefs of Actors and Observers. In *24th Annual Meeting of the Association for Computational Linguistics*, pages pages 207-214. ACL, Columbia University, June, 1986.
- [Rich 79] Rich, E. User modeling via stereotypes. *Cognitive Science* 3(4):pages 329-354, 1979.
- [Shuster, E. 88] Shuster, E. Establishing the Relationship between Discourse Models and User Models. *Computational Linguistics* 14(3):82-85, 1988.
- [Sidner and Israel 81] Sidner, C. and Israel D. Recognizing Intended Meaning and Speakers' Plans. In *Proceedings of the IJCAI*. International Conferences on Artificial Intelligence, August, 1981.
- [Wallis and Shortliffe 82] Wallis, J. and Shortliffe, E. *Explanation Power for Medical Expert Systems: Studies in the Representation of Causal Relationships for Clinical Consultation*. Technical Report STAN-CS-82-923, Stanford Univ. Heurist. Program Proj. Dept. Med. Comput. Sci., 1982.
- [Wolz et al. 88] Wolz, U. *Tutoring that responds to users' questions and provides enrichment*. Technical Report CUCS-410-88, Department of Computer Science, Columbia University, New York, NY, 1988. Also appeared in the conference proceedings of the 4th International Conference on Artificial Intelligence and Education, May 1989, Amsterdam, The Netherlands.