
User Modeling and User-Adapted Interaction **10:** 77–80, 2000.

Preface

SANDRA CARBERRY
Department of Computer Science, University of Delaware, Newark, DE, USA.
Email: carberry@cis.udel.edu; home page://www.cis.udel.edu/~carberry

User modeling research began in the early 1980's with several classical research projects, such as that of Elaine Rich on a personalized system for recommending library books (Rich, 1979). Since then, research in this area has flourished and projects have ranged from systems that present tailored explanations to systems that provide customized tutoring. But until recently, user modeling research has remained in the research laboratory, with the assumption that the techniques developed there could readily be applied in actual systems. Unfortunately, such assumptions are rarely true. In their paper, Billsus and Pazzani note that deployment in the real-world forces developers to address issues of efficiency, computational complexity, and scalability that are generally ignored in a research environment. Similar observations are made by the authors of the other contributions in this special issue. For example, Fink and Kobsa note that characteristics like performance, scalability, extensibility, integration of pre-existing user information, and privacy protection are considered far more important in the case of user modeling servers for e-commerce than qualities such as generality, domain-independence, expressiveness, and strong inferential capabilities that have been regarded as crucial in academic user modeling research.

This special issue on deployed systems contains a number of papers that focus on the problems encountered in converting a research project to a system that is actually used by a large population for real-world tasks under real-world constraints. Many of the papers identify compromises that diverge from the ideal research prototype but are necessary for successful deployment.

For many years, researchers at Carnegie Mellon University have been developing tutoring systems based on a tutoring theory that differentiates between declarative and procedural knowledge. This theory assumes that learning first encodes skills as declarative knowledge and that through practice domain-specific procedural knowledge is acquired. The paper by Corbett, McLaughlin, and Scarpinatto discusses two sets of tutors based on this theory, one a Lisp tutor that is limited to a university environment and the other an algebra tutor that has been extensively deployed in public high school systems. They note several differences resulting from deployment in a real-world classroom as opposed to use in a university environment. For example, their system estimates the probability that a student knows how to use

a rule correctly and equates mastery with a 95% probability. In a self-paced university programming course, the system can insist on mastery before moving on to a new topic, but this is not possible in the classroom since student contact hours are limited and the curriculum requires that the entire course content be covered. On the other hand, human teachers are generally not available for a self-paced university course, so the system must eventually resort to telling a confused student precisely how to work a problem. In the classroom environment, this fall-back response is unnecessary since a teacher is available to interact with students who cannot learn from the system's advice.

Without individualized user-adaptation, sophisticated commercial systems can have a very limited user audience. The research described in the paper by Strachan, Anderson, Sneesby, and Evans was motivated by the goal of making a commercial financial planning system of proven utility accessible to users with varied amounts of computer expertise and financial planning acumen. However, the authors emphasize that practical concerns such as the amount of required system modification, implementation cost, performance overhead, and training and support requirements have a strong impact on the extent to which user modeling can be realistically incorporated into commercial systems. They recommend a minimalist approach that utilizes simplified versions of techniques developed for research systems in order to provide essential user adaptation while adhering to the practical requirements of a system that is to be widely deployed. Their paper describes how this minimalist approach was pursued to add a user-adaptive capability to a complex, powerful, commercially deployed financial planning system. It describes the essential adaptations that were identified through user studies, the resultant architecture that employed minimalist strategies, and the ramifications of design decisions on the system's capability and on the complexity of the user modeling component. The paper presents an empirical evaluation to support the hypothesis that minimalist approaches can enhance user satisfaction while addressing practical considerations important to successful commercialization. Throughout the paper, the authors discuss the problems they encountered in developing an effective user-adaptive system and the issues that must be addressed when developing novel systems as part of a commercial enterprise.

Given the explosion of available information, recommender systems have become a source of much research interest. Billsus and Pazzani discuss how the class of interface and costs of interaction (web-based or wireless) affect design of both the overall system and the user model. For example, to develop a usable wireless version of their news recommender, they had to address the issue of limited bandwidth and expensive data transmission. This necessitated minimizing the amount of retrieved data transmitted between the wireless device and the server. Consequently, explicit acquisition of user ratings was not feasible, and the user modeling system had to rely on implicitly obtaining information about user preferences based on which news stories were accessed and how much of the story was retrieved. In addition to contrasting the web-based and wireless versions of their system for adaptive news access, Billsus

and Pazzani argue for a hybrid user model that accounts for both short-term and long-term user interests, and they provide an empirical evaluation both of their user modeling algorithms and of the algorithms' contribution to a system that addresses user needs.

The article by Linton and Schaefer also presents a recommender system that employs implicit user model acquisition, but in the context of informal workplace education where the system recommends new skills that the user should acquire. Although there are many domains where such a system might be useful, the authors have currently deployed it as part of a word processor on a set of networked desktop computers within a corporate environment. To identify skills that should be recommended, the system computes an individualized expert model reflecting the frequency with which the user would be expected to execute various commands if the user had the pooled knowledge of his peers. Large deviations between these expected frequencies and the user's actual performance suggest skills that the user might profitably learn. The minimalist approach espoused by Strachan et al. is seen in this work, in that the expert model is built using a simple analysis of logs of the user's and his peers' commands over an extended period of time.

Rounding out this special issue is the contribution by Fink and Kobsa that provides an overview of user modeling servers for personalized e-commerce web sites. This article is particularly noteworthy because obtaining detailed information about such personalization servers is very difficult. Consequently, as noted by the authors, the research literature contains little about how user modeling has been employed in systems for e-commerce, despite the burgeoning importance of this area. Fink and Kobsa identify a set of features (such as type of user data, acquisition methods, extensibility, degree of privacy protection, and supported hardware and software platforms) along which different user-adaptive e-commerce servers can be characterized. They then compare and contrast four commercial systems in terms of these features. The article provides an excellent overview of the current state-of-the-art in commercial user modeling servers for e-commerce, including both the positive contributions and the limitations of existing systems.

The articles in this special issue show that user modeling has indeed moved from a strictly research exercise to a beneficial component of systems deployed for real-world tasks. It will be interesting to see how future systems build on these successes and enhance their user-adaptive capabilities while recognizing the need to address the requirements and limitations of commercial deployment.

Acknowledgements

I'd like to thank Judy Kay for all of her efforts in handling the review process for this special issue.

References

Rich, Elaine: 1979, User Modelling via Stereotypes. Cognitive Science, **3**(4), 329–354.

Author's Vitae

Dr. Sandra Carberry is Professor and Chair of the Department of Computer Science at the University of Delaware. She received a B.A. in mathematics from Cornell University, an M.S. in computer science from Rice University, and a PhD in computer science from the University of Delaware. Her current research interests include user modeling, computational linguistics with an emphasis on dialogue systems, and human computer interfaces. Her book ''Plan Recognition in Natural Language Dialogue'' was published by MIT Press in 1990.