

Coming of age: celebrating a quarter century of user modeling and personalization: Guest editors' introduction

Judy Kay · Gord McCalla

Received: 26 October 2011 / Accepted: 26 October 2011
© Springer Science+Business Media B.V. 2012

When Alfred Kobsa and Wolfgang Wahlster hosted the first workshop on user modeling¹ in 1986, few anticipated how important this field was to become during the next 25 years. That early workshop led to more workshops and then the User Modeling conference, which ultimately merged with the Adaptive Hypermedia conference in 2009 to become the User Modeling, Adaptation, and Personalization conference. In parallel to the evolution of the UMAP conference series, and also under the pioneering leadership of Alfred Kobsa, was the founding of the UMUAI journal, first appearing in 1991, originally published by Kluwer and now by Springer. This special issue of the UMUAI journal celebrates the first 25 years of user modeling research by providing invited (and in the UMUAI tradition also stringently reviewed) papers from leading experts on user modeling, adaptation, and personalization. The articles collectively provide a superb overview of research issues in user modeling that are the likely shape of the next decade of research, as well as an historical perspective of how these issues have evolved since the inception of the field.

User modeling and user adapted interaction has not only grown as a research area, but has become a central applied issue for anybody interested in understanding users as they interact with technology. The foundational applications of learning systems (with *student modeling* to help systems adapt to learner differences) and natural language

¹ See the workshop proceedings (Kobsa and Wahlster 1986) and Sandra Carberry's workshop report (Carberry 1987). An edited collection also followed up this workshop (Kobsa and Wahlster 1989).

J. Kay (✉)
School of Information Technologies, School of IT Building, University of Sydney,
Sydney, NSW, 2006, Australia
e-mail: judy.kay@sydney.edu.au

G. McCalla
Department of Computer Science, University of Saskatchewan, 176 Thorvaldson Building,
110 Science Place, Saskatoon, SK, S7N 5C9, Canada
e-mail: mccalla@cs.usask.ca

dialogue systems (with *user modeling* to track the beliefs of the user so as to enable the system to tailor an appropriate response to them), have been joined by a plethora of applications involving personalization including recommender systems, social computing, targeted on-line advertising, intelligent web search, personalized help systems, adaptive interactive systems, intelligent user interfaces, etc. It is not too much of a stretch to say that personalization is the key issue in an era when there is so much information, and so many people interacting in so many ways, doing so many things, using information and communications technology. Personalization is a critical need to help save us from drowning in this ocean of information and people.

As the user modeling field has grown, the methodologies used have evolved from a few techniques largely drawn from the area of knowledge representation, to techniques drawn from all over artificial intelligence, computer science, and even social science. Ideas have been co-opted from machine learning and data mining, intelligent agents, the semantic web, human-computer interaction, visualization, affective computing, web computing, e-learning, data base systems, information science, social computing, cognitive science, education, anthropology, and so on and on. User modeling researchers have been quick to take advantage of the new insights and algorithms developed in these other areas to enhance the capabilities of their systems for personalization and adaptivity. But, the traffic has not been all in one direction. The user modeling application area is a tough crucible for any technique, full of unpredictable and demanding users, noisy data, constant change, and the need for algorithms that can respond in real time and that are scalable to large user bases. This has, at the very least, provided a rigorous test bed for these techniques and in the best case has generated new versions of existing techniques or even brand new techniques. The papers in this special issue illustrate many of these techniques in the context of many different applications.

Going forward, the future for user modeling is bright indeed. The need for personalization has been widely recognized and almost every major software company has dedicated much effort and many resources to exploring this issue from one angle or another. In some cases, personalization is at the heart of the company's business model. Both the academic and the industrial research community investigating this field are growing rapidly, and are drawn from an increasingly wide range of specializations. Trends such as cloud computing, ubiquitous connectivity, 24/7 technology access, and Web 2.0 participatory applications ensure the need for personalization. Luckily, they possibly also provide both the bandwidth of interaction with users and access to their interaction data that are needed in order to understand user behavior and to appropriately react to user needs.

Several key issues in user modeling have shifted dramatically over the field's existence. The early problem faced by user modeling systems of having too little bandwidth has been replaced by the opposite problem of being overwhelmed by the sheer volume of noisy, contradictory, fine-grained data about users. This new problem, while extremely challenging, is certainly to be preferred to ignorance about user behavior! A second shift has been from exploring toy systems in artificial situations to studying real users "au naturel" with all of their complex real world behavior. This has brought to the field a deep concern with human subject experimental techniques and with doing really sound empirical studies. A third shift has happened from the early concentration

on the user model itself, and its representation, to using the user model, and adapting intelligently to the user in a real world application, essentially providing more balance between the “user modeling” and the “user adapted interaction” aspects of the field. A fourth shift is just getting underway: modeling groups of people, not just individuals. Stimulated by the growing role of social networks in the world, understanding and supporting communities of people is becoming every bit as important as understanding and supporting individuals. Modeling groups of people is not contradictory to modeling individuals: many of the same techniques can be deployed, and understanding individuals in the context of the groups to which they belong helps in understanding both them and their groups. A fifth shift has been in the kinds of things modeled and adapted to. Originally, user models tended to concentrate on the beliefs or knowledge of individuals. Now, all manner of characteristics are modeled, including affective states, learning styles, cognitive or psychological states, physical or social context, etc., and these can be useful in many ways for systems adapting to users. Finally, as the field has become increasingly relevant and increasingly good at what it does, user modelers have realized that the social impact of their work is critical. If user modeling is to reach its full potential, we will have to find ways of assuring people that their privacy will be preserved even as we understand them better. Fortunately, however, the better we can understand a person, the more likely it is that we will understand their privacy needs and how to protect this privacy. Thus, the conflict between user modeling and privacy protection is not insurmountable, and in fact a stream of user modeling research is actively exploring privacy itself, a nice blending of computer and social science.

So, the field of user modeling and adaptivity has adapted and evolved. The one constant is Alfred Kobsa, still at the helm as editor-in-chief of *UMUAI*, a journal that has extremely high standards, great citation statistics and is, in every way, the flagship for the field. It is a pleasure for us as guest editors to present this special issue in honor of the field itself but also of Alfred's 25 years in the trenches, 20 of which he has devoted to *UMUAI*.

Here is a brief synopsis of the papers in the special issue.

Throughout the history of personalization research, a key sub-area has been learner modeling, which deals with aspects of user modeling that are relevant for creating rich learning environments that are tailored to the individual learner. It is then fitting that we begin this special issue with a contribution by Michel C. Desmarais and Ryan S. J. d. Baker who present a carefully crafted overview of learner models that have already left the lab, techniques that have been established for learner modeling, and aspects that are emerging as important directions. The paper provides readers with concrete examples that illustrate key ideas for this review, starting with the nature of skill models, then discussing a selection of the key examples of widely deployed learner models: cognitive tutors, constraint-based tutors, and content sequencing tutors. Desmarais and Baker then reflect on the reasons for the success of these tutors and review techniques for dealing with uncertainty in learner modeling, one of the key technical challenges in skill modeling—and also highly relevant for much other user modeling. The paper then reviews approaches to the important issue of validation of learner modeling research. This contribution concludes by identifying several areas that are the focus of current research energy, some due to the increasing maturity of the field and others being

created by the emergence of new technologies for inexpensive, powerful, pervasive, mobile and ubiquitous technology: affect, emotion, motivation and disengagement; meta-cognition and self-regulation; open learner modeling, group and collaborative learner modeling; and long term learner modeling.

Tanja Mitrovic outlines a longstanding project to take an innovative research idea, Ohlsson's constraint-based student modeling (CBM), from a theoretical proposal into a complete paradigm for developing intelligent tutoring systems of various stripes. Starting with the SQL-Tutor system in 1996 and building on this, Mitrovic and her Intelligent Computer Tutoring Group (ITCG) at the University of Canterbury have created a number of successful CBM-based intelligent tutors and a set of tools that can be used to develop such tutors. The paper shows in some detail how theory, engineering, and empirical evaluation mutually inform one another. It also illustrates how new ideas in the field of intelligent tutoring and user modeling flow into such a project and are illuminated by it. And, it illustrates the value of committing to a particular approach over the long term and exploring the implications of this approach in intricate detail. It seems particularly appropriate for this special issue of UMUAI to be able to provide an inside view of how an important research project that started with a germ of an idea at the beginning of the field of user modeling has grown to maturity over the lifetime of the field.

Closely related to learner modeling and personalized educational systems is personalization in cultural heritage, where visitors to physical and virtual cultural heritage sites can learn and experience aspects that match their preferences, interests, knowledge, context and personal attributes. Liliana Ardissono, Tsvi Kuflik and Daniela Petrelli present an authoritative and very comprehensive view of cultural heritage personalization, past and emerging. They identify the key dimensions of systems that have been created and use these dimensions to provide a new and valuable analysis of the work that has been done. The paper shows the extent of the research, its richness in terms of the settings, technologies employed, presentation styles, forms of adaptation, and user modeling techniques. The paper adds flesh to this with detailed concrete examples of just how key systems have explored these aspects and built our knowledge of them. The final discussion tackles issues related to the enigma that the research appears so promising but has not yet led to wide deployment. The paper discusses the challenges still to be met: the need for standardization; the importance of continuity in the experience, including before and after the actual site visit; social aspects of the visit; exploitation of emerging technologies and techniques; and evaluations. This area is a particularly exciting one for personalization research, both because it appears to be on the cusp of moving into large-scale deployment and because the demands of cultural heritage personalization drive such a rich agenda for user modeling research.

Joseph Konstan and John Riedl explore the evolution of recommender systems, an area of intense research interest (both commercial and academic) for the user modeling community. Their paper shows how, over more than 15 years, the issues in building recommender systems have evolved from a concentration on predictive accuracy of recommendations to concern with a wide range of new issues involving the user experience. This broadening reflects general changes in user modeling research during the same time period, from an early focus on creating and maintaining the user model (the "user modeling" part of the field) to much broader investigations into how systems

can usefully personalize their interactions with the user (the “user adapted interaction” part of the field). Konstan and Riedl explore all manner of issues that arise when taking end use context into account: different kinds of users and different needs of users at different times in the recommender life cycle, trade-offs between individual needs and group and community needs, different notions of value and quality beyond just accuracy, various risks and dangers of recommender systems (and personalization research more generally), and the need for user control and involvement at all stages of the recommendation process. Their paper concludes by arguing that there are three main broad sets of issues for those working on recommender systems: better scalability of the algorithms to tackle the vast amount of data and users and possible end uses available to recommender systems, better exploitation of the immense amount of user-contributed content (especially implicit content through usage), and enhancement to the research infrastructure to allow more systematic exploration of the wide range of issues impacted by recommender system research.

The next paper provides a deep analysis of a particular class of recommenders, which are critiquing-based. These are distinctive in having a foundation that puts the user in the loop: they tightly intertwine user interfaces that explicitly offer the user an opportunity to critique a set of recommendations and the underlying algorithmic recommendation mechanisms. This paper is by Li Chen and Pearl Pu, from the research group that has contributed so much to this class of recommender. The paper presents several examples of three main classes of critiquing-based recommenders: natural language based, system-suggested, and user-initiated critiques. The paper also analyses the relative strengths of each. The authors point to emerging use of this approach in widely used recommenders such as Amazon and MovieLens. They conclude with an analysis of emerging directions in hybrid systems, with adaptation improvements and the move to a broader range of domains.

User modeling research has traditionally been perceived as being concerned with adaptation to individuals, and not having much to do with communities. Georgios Paliouras shatters this perception in his paper about discovering web-based communities. First, he introduces a graph-based notation that clarifies relationships among individuals in different types of web communities. Then, he shows how information about communities can be extracted from the increasingly vast amount of information available about individuals and their interactions on the web. Further, the information about communities thus extracted turns out to be highly useful for helping to achieve many of the traditional roles of individualization explored in user modeling research over the years. The synergy between communities and individuals is getting even tighter in an era of the social web, with everybody being both a producer and consumer of web-based information. This implies an increasingly key role in user modeling research for future community-based modeling and a further broadening of user modeling into groups and communities of people, not just individuals.

Julita Vassileva's paper continues this theme of social applications and the very important issue of encouraging high quality participation in online communities. Here, the role of personalization is to address the differences in what motivates different people and how this changes over time. The paper provides an overview of theoretical foundations for the design of mechanisms to encourage participation. These include classic and behavioral economics, including the introduction of game-like elements,

with the attendant choices of which user actions should drive increases in reputation and status. Another important theoretical foundation comes from theories of motivation, from psychology, organizational science and sociology. Building from these foundations, Vassileva explains how they are combined with mobile and ubiquitous technologies and computer science, to support the tightly inter-related emerging work in persuasion, design of incentive mechanisms and user modeling and adaptation. The paper analyses the ways that personalization techniques can be integrated into the design of incentive mechanisms and the relationship between open user modeling and social visualization, in each case, drawing parallels between user models, purposes and forms of intervention. It concludes with directions for new applications that involve self-directed learning and social activism.

For all of the benefits of personalization, a major downside is the potential of systems that support personalization to seriously violate peoples' privacy. The paper by Eran Toch, Yang Wang, and Lorrie Faith Cranor explores this important issue. The paper first investigates the threat to privacy of three major trends in current technology: social-based personalization, behavioral profiling, and the mobile Web. Then, the research into techniques for reducing these threats is overviewed. The paper concludes by proposing a framework for helping to understand the privacy risks facing user modeling systems. This framework can be used by user modeling researchers to understand the privacy risks of their approach and can provide a starting point for establishing ways of overcoming these risks. It can also help to predict the impact on privacy of new technologies, a key need as personalization technology embeds ever more deeply and broadly into peoples' lives.

References

- Carberry, S.: First International Workshop on User Modeling. *AI Magazine*, vol. 8, no. 3, pp. 72–74 (1987)
- Kobsa, A., Wahlster, W. (eds.): In: *Proceedings of the First International Workshop on User Modeling*, Maria Laach, Germany 30–31 Aug 1986
- Kobsa, A., Wahlster, W. (eds.): *User Models in Dialog Systems*. Springer, New York (1989)

Author Biographies

Judy Kay is Professor of Computer Science at the University of Sydney. She is a principal in the Computer Human Adapted Interaction research group (CHAI), leading research in advanced technologies for human computer interaction, supporting personalisation, pervasive and mobile interaction. She received her Ph.D. in Computer Science from the University of Sydney. Her personalisation research explores ways to ensure that people can maintain control over their user model and its use. She creates new tools for interaction in support of lifelong learning, collaboration and personalised smart services. Dr. Kay's publications are in venues for human–computer interaction, pervasive computing, user modelling and personalisation and artificial intelligence in education and she has served on editorial boards and conference committees in these areas.

Gord McCalla is a Professor in the Department of Computer Science at the University of Saskatchewan and Director of the Laboratory for Advanced Research in Educational Systems (ARIES) there. He received his B.Sc. (Hons) in Mathematical-Statistics and M.Sc. in Computing Science from the University of Alberta, and his Ph.D. in Computer Science from the University of British Columbia. Dr. McCalla has worked in a variety of areas of artificial intelligence, including natural language dialogue, planning, multi-agent systems, artificial intelligence in education (AIED), and user modeling, adaptation, and personalization (UMAP).

His current research focus is on personalization in learning systems, with a specific emphasis on mining fine grained learner behavioural data to find patterns that can help a learning system adapt to individual differences in learners. Dr. McCalla has published widely in both AIED and UMAP, and has a long record of service to both research communities.