Proceedings of the First EUNITE Symposium, Tenerife, 13–14 December 2001. (Extended abstract of invited presentation.)

User-Adaptive and Other Smart Adaptive Systems: Possible Synergies

Anthony Jameson
German Research Center for Artificial Intelligence (DFKI)
and International University in Germany
Stuhlsatzenhausweg 3, 66123 Saarbrücken, Germany
Phone: (+49)(681) 302-5079, Fax: (+49)(681) 302-5020
Email: jameson@dfki.de, WWW: http://dfki.de/∼jameson/[1]

ABSTRACT: User-adaptive systems are interactive software systems that spontaneously adapt to their individual users–for example, to their interests or their work habits. First, we briefly characterize this type of system and consider its relationship to the broader category of smart adaptive systems. We then give an overview of the computational techniques that are used to realize user-adaptation, ranging from data-based machine learning techniques to theory-based decision-theoretic models. Special attention is given to the question of possible synergies: What methods that have been developed for user-adaptive systems might profitably be transferred to other smart adaptive systems? And what techniques from the latter field deserve increased attention in connection with user-adaptivity?

KEYWORDS: User-adaptive systems, user models, machine learning, Bayesian networks, decision making, transparency, controllability

USER-ADAPTIVE SYSTEMS IN THE EUNITE TOPIC AREAS

One subclass of the class of smart adaptive systems is that of *user-adaptive systems*. A user-adaptive system can be defined as an interactive system that adapts its behavior to individual users on the basis of processes of user model acquisition and application that involve some form of learning, inference, or decision making (cf. [1]).

User-adaptive systems can be found in most of the areas covered by EUNITE.

Telecommunication and Multimedia

One type of user-adaptive system that is currently attracting great interest offers personalized presentation of information (e.g., travel information, news, or product information) on mobile devices such as personal digital assistants and cell phones. Because of the limited communication bandwith and screen size of such devices, there is a great deal to be gained if the form and the content of the information presented is adapted to the interests and usage patterns of individual users (see, e.g., [2], [3]). With the use of multimedia technologies on larger devices, it may be necessary or desirable dynamically to adapt the choice of media to the individual user's available devices or media preferences—and/or to adapt the specific content to the user's interests or knowledge (see, e.g., [4]).

Finance, Trade and Services

Where interaction with (potential) customers is involved, adaptation to individuals often goes under the name of *personalization*. Product recommenders take users' preferences into account while helping them to find suitable products (see, e.g., [5]). One goal is to build up customer loyalty by convincing customers that they can interact more effectively with a system that has built up a useful model of them.

Human, Medical, and Health Care

The tailoring of medical information to individual patients is an attractive idea because of the large differences that exist among patients in terms of (a) their ability to understand various types of information and (b) their medical interests, which are determined to a large extent by their own medical condition (see, e.g., [6], [7]). Medical training offers many possibilities for adapting the form and content of instruction to individual learners' state of knowledge and learning style (see, e.g., [8]). For users with perceptual or physical disabilities, systems that recommend suitable adaptations to a user interface have been developed (see, e.g., [9]).

[1] The presentation slides for this talk are available from http://dfki.de/∼jameson/pdf/eunite01.jameson-slides.pdf.

Transportation

Some user-adaptive systems help travelers to choose their destinations, often using methods similar to those used for recommending products (see, e.g., [10]). Others help them to figure out how to get to their destination, taking into account their transportation preferences (see, e.g., [11]). Tourist information systems may perform both of these functions, as well as providing information on particular locations that the user is visiting. Systems that actually take over part or all of the task of driving an automobile need to adapt not only to the preferences of their passengers but also to the perceived intentions of other drivers on the road (see, e.g., [12]).

Production Industry

Provision of information to operators in industrial plants offers opportunities for fruitful adaptation, partly because of the large amount of potentially relevant information that is typically available and the limited capacity of human operators to process such formation. For example, a system may adapt the displayed information to (a) the expertise or role of individual users and (b) the nature of the situation at hand (e.g., whether time is a critical factor; cf. [13]).

LEARNING, INFERENCE, AND DECISION MAKING IN USER-ADAPTIVE SYSTEMS

Given the great variety of functions that user-adaptive systems serve, it is no wonder that a similar variety of computational techniques are employed for the learning, inference, and decision making that are required for the adaptation. It is useful to distinguish three broad categories of techniques.

Data-Based Inference Methods

With *data-based* methods, the main source of information is data about system usage acquired during interactions with users. Many systems make use of data only about the current user (e.g., which screens this user has visited so far), using one or more machine learning techniques to learn a model of that user. The learning techniques that have been employed include decision trees, probabilistic classifiers, neural networks, case-based reasoning, and specialized text-classification methods (cf., [14]). Less frequently, data from a larger number of users are employed, and the model learned is intended to be applicable to users in general or at least to an entire user group (see, e.g., [15]). A particularly popular type of method is *collaborative filtering* (see, e.g., [16]), which uses data about both the current user and similar other users to make predictions about the current user.

Theory-Based Inference Methods

With *theory-based* methods, the designer specifies some sort of a priori model of users, which may then be adapted to an individual user. *Bayesian networks* allow the designer to formulate a model in terms of variables and the (often causal) relationships among them ([17]). Theory-based models have also sometimes been formalized with the Dempster-Shafer theory of evidence or with Fuzzy Logic, though these formalisms have been used less frequently in this area in recent years ([17]). An approach that can be realized with various formalisms is to define a number of *stereotypes* or user groups, specifying certain typical properties or system actions that are to be associated with each one (see, e.g., [18]).

Decision Making Methods

Adaptation to the individual user involves not only making inferences about the user but also deciding what adaptations are appropriate, given the results of these inferences. Often the adaptation principles are hand-coded by the system designer; but especially where tradeoffs and uncertainty are involved, it may be desirable to use systematic decision-theoretic techniques such as influence diagrams or decision-theoretic planning ([19]).

IDEAS FOR DESIGNERS OF OTHER SMART ADAPTIVE SYSTEMS

Some of the recent research on user-adaptive systems has made contributions to questions that have relevance for other types of smart adaptive system:

Enhancing Transparency

Various ways have been introduced of (a) making the workings of a user-adaptive system clearer to users and/or (b) explaining why the system arrived at particular adaptations (see, e.g., [20]).

Enhancing Controllability

Since spontaneous system adaptation can make the user feel out of control, various approaches have been proposed of allowing the user to control both general aspects of the system and specific aspects of its behavior (see, e.g., [21]).

LESSONS FROM RESEARCH ON OTHER SMART ADAPTIVE SYSTEMS

Use of Particular Computational Techniques

Computational techniques such as neural networks, genetic algorithms, and fuzzy logic are used more frequently with other types of smart adaptive system than with user-adaptive systems. One issue for discussion is whether there are good reasons for this difference or whether some potential advantages of these formalisms are being overlooked by designers of user-adaptive systems.

Adaptation of General Initial Models

One approach that is often employed with smart adaptive systems is to (a) build a system with a default initial model of some sort and (b) provide for that model to adapt in the course of system use. This approach has been used less frequently with user-adaptive systems than one might have expected (cf. [22]). It seems worthwhile to examine the approaches applied with other smart adaptive systems to see whether they can be adapted for application to user-adaptive systems.

REFERENCES

[1] Anthony Jameson. Adaptive interfaces and agents. In Julie A. Jacko and Andrew Sears, editors, *Handbook of Human-Computer Interaction in Interactive Systems*. Erlbaum, Mahwah, NJ, 2002. In press.

[2] Corin R. Anderson, Pedro Domingos, and Daniel S. Weld. Adaptive web navigation for wireless devices. In Bernhard Nebel, editor, *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*. Morgan Kaufmann, San Francisco, CA, 2001.

[3] Daniel Billsus and Michael J. Pazzani. User modeling for adaptive news access. *User Modeling and User-Adapted Interaction*, 10:147–180, 2000.

[4] Elisabeth André, Thomas Rist, and Jochen Müller. Employing AI methods to control the behavior of animated interface agents. *Applied Artificial Intelligence*, pages 415–448, 1999.

[5] J. B. Schafer, Joseph A. Konstan, and John Riedl. E-commerce recommendation applications. *Journal of Data Mining and Knowledge Discovery*, 5(1/2):115–152, 2001.

[6] Graeme Hirst, Chrysanne DiMarco, Eduard Hovy, and Kimberley Parsons. Authoring and generating health-education documents that are tailored to the needs of the individual patient. In Anthony Jameson, Cécile Paris, and Carlo Tasso, editors, *User Modeling: Proceedings of the Sixth International Conference, UM97*, pages 107–118. Springer Wien New York, Vienna, 1997.

[7] Ray Jones, Janne Pearson, Sandra McGregor, Alison J. Cawsey, Ann Barrett, Neil Craig, Jacqueline M. Atkinson, W. H. Gilmour, and Jim McEwen. Randomised trial of personalised computer based information for cancer patients. *British Medical Journal*, 319:1241–1247, 1999.

[8] Roger Azevedo and Sonia L. Faremo. Expertise differences in radiology: Extending the RadTutor to foster medical students' diagnostic skills. In Susanne P. Lajoie and Martial Vivet, editors, *Artificial Intelligence in Education: Open Learning Environments: New Computational Technologies to Support Learning, Exploration, and Collaboration*, pages 501–508. IOI Press, Amsterdam, 1999.

[9] Shari Trewin and Helen Pain. A model of keyboard configuration requirements. In *Proceedings of the Third International ACM Conference on Assistive Technologies*, pages 173–181, 1998.

[10] Mehmet Göker and Cynthia Thompson. Personalized conversational case-based recommendation. In *Proceedings of the 5th European Workshop on Case Based Reasoning, Trento, Italy 2000*, 2000.

[11] S. Rogers, C. Fiechter, and P. Langley. An adaptive interactive agent for route advice. In *Proceedings of the Third International Conference on Autonomous Agents*, Seattle, WA, 1999.

[12] David V. Pynadath and Michael P. Wellman. Accounting for context in plan recognition, with application to traffic monitoring. In Philippe Besnard and Steve Hanks, editors, *Uncertainty in Artificial Intelligence: Proceedings of the Eleventh Conference*, pages 472–481. Morgan Kaufmann, San Francisco, 1995.

[13] Eric Horvitz and Matthew Barry. Display of information for time-critical decision making. In Philippe Besnard and Steve Hanks, editors, *Uncertainty in Artificial Intelligence: Proceedings of the Eleventh Conference*, pages 296–305. Morgan Kaufmann, San Francisco, 1995.

[14] Geoff Webb, Michael J. Pazzani, and Daniel Billsus. Machine learning for user modeling. *User Modeling and User-Adapted Interaction*, 11:19–29, 2001.

[15] Bruce Krulwich. Lifestyle Finder: Intelligent user profiling using large-scale demographic data. *AI Magazine*, 18(2):37–45, 1997.

[16] Jonathan L. Herlocker, Joseph A. Konstan, Al Borchers, and John Riedl. An algorithmic framework for performing collaborative filtering. In *Proceedings of the 1999 Conference on Research and Development in Information Retrieval*, 1999.

[17] Anthony Jameson. Numerical uncertainty management in user and student modeling: An overview of systems and issues. *User Modeling and User-Adapted Interaction*, 5:193–251, 1996.

[18] Alfred Kobsa and Wolfgang Pohl. The user modeling shell system BGP-MS. *User Modeling and User Adapted Interaction*, 4:59–106, 1995.

[19] Anthony Jameson, Barbara Großmann-Hutter, Leonie March, Ralf Rummer, Thorsten Bohnenberger, and Frank Wittig. When actions have consequences: Empirically based decision making for intelligent user interfaces. *Knowledge-Based Systems*, 14:75–92, 2001.

[20] Jonathan L. Herlocker, Joseph A. Konstan, and John Riedl. Explaining collaborative filtering recommendations. In *Proceedings of the 2000 Conference on Computer-Supported Cooperative Work*, 2000.

[21] Alan Wexelblat and Pattie Maes. Issues for software agent UI. Unpublished manuscript, available from http://wex.www.media.mit.edu/people/wex/., 1997.

[22] Frank Wittig and Anthony Jameson. Exploiting qualitative knowledge in the learning of conditional probabilities of Bayesian networks. In Craig Boutilier and Moisés Goldszmidt, editors, *Uncertainty in Artificial Intelligence: Proceedings of the Sixteenth Conference*, pages 644–652. Morgan Kaufmann, San Francisco, 2000.