Lifelong Learner Modeling for Lifelong Personalized Pervasive Learning

Judy Kay

Abstract—Pervasive and ubiquitous computing has the potential to make huge changes in the ways that we will learn throughout our lives. This paper presents a vision for the lifelong user model as a first class citizen, existing independently of any single application and controlled by the learner. The paper argues that this has a key role for a vision of personalized lifelong learning and for augmented cognition that enables learners to supplement their own knowledge with readily accessible digital information based on documents that they have accessed or used. The paper presents work that provides foundations for this vision. First, it outlines technical issues and research into approaches for addressing them. Then, it presents work on the interface between the learner and the lifelong user model, an aspect that is important because the human issues of control and privacy are so central. The final discussion and conclusions outline a roadmap for future research that will underpin this vision of the lifelong user model.

Index Terms—Pervasive computing, lifelong learning, user models, learner models, personalization, open learner models, learner control, scrutability, stereotype user models, reflection, metacognition, mirroring.

1 Introduction

Pervasive and ubiquitous computing collects and moves huge amounts of data that are about us or belong to us. This huge and fast-growing collection of personal data offers great potential benefits. Lifelong learning is one of these, with the possibility that we can learn precisely when and where we need to, with learning resources delivered just for us, taking account of our existing knowledge and preferred ways of learning. There are other closely related benefits in augmented cognition and lifelong memories. We are already seeing the beginnings of this in the form of readily available searches as well as widespread use of computers for personal information management. In the future, we can go further, providing people with ready access to pertinent information about themselves from their own data stores, enabling them to expand their effective cognition and knowledge beyond what they would otherwise remember. The value of augmented cognition tools will be of particular importance for people who suffer cognitive loss, as is common with aging.

The creation of the technology for personalized lifelong learning has been recognized as a Grand Challenge Problem by peak research bodies. The Computing Research Association (CRA) [1] identified the goal to *Provide a Teacher for Every Learner* as one of its five Grand Research Challenges in Computer Science and Engineering. The United Kingdom Computing Research Committee (UKCRC) has nine Current Grand Challenges for Computing [2]. Two of these are about personalized lifelong learning: GC8, *Learning for Life* [3], and GC3, *Memories for Life*. In 2008, the National Academy of Engineering identified 14 Grand Challenge Problems across all of engineering: One of these is *Advance Personalized Learning* [4].

Descriptions of these Grand Challenges include a central role for personalization. For example, they speak of learners being able to "learn at their own pace and in their own style" [1], [3] with the goal to enhance "the effectiveness of learning and the quality of the learning experience by providing a better fit between the needs of learners at a particular time and the learning facilities provided" [3]. They note that technology has the potential to make a real difference for those with special needs, so that all learners can "make the most of their talents, irrespective of their physical and mental disabilities" [1]. Also important is teaching that is "individualized based on learning styles, speeds, and interests to make learning more reliable" [4].

The Grand Challenge visions encompass the full gamut of learning possibilities. For example, they envisage enabling learners to "participate in networked and face-to-face communities of learners composed of peers, teachers, mentors, domain experts, avatars" [1]. They also point to the importance of ensuring that the learner can "receive continuous, customized, and meaningful feedback and assessment" [1] (also quoted in [3]). Closely aligned with explicit teaching is the somewhat related Grand Challenge Problem of augmented cognition, where the goal is that the computer should complement our cognitive abilities, "as a tool to help people store, search, and interpret their *memories*...digital personal information such as e-mails, digital photographs, and Internet telephone calls" [5]. This is inextricably linked to lifelong learning. One very direct link refers to exploiting our existing knowledge by teaching "in the context of the learner's memories and experiences, using language the learner is comfortable with." This Grand Challenge also acknowledges the importance of personalization as it refers to the need "to personalize interfaces to the skills, preferences, and abilities of individual users."

• *The author is with the Human Adaptive Interaction Research Group (CHAI), School of Information Technologies, Room 318, School of IT J12, University of Sydney, Australia. E-mail: judy@it.usyd.edu.au.*

Manuscript received 15 Dec. 2008; accepted 2 Jan. 2009; published online 23 Jan 2009.
For information on obtaining reprints of this article, please send e-mail to: lt@computer.org, and reference IEEECS Log Number TLT-2008-12-0111.
Digital Object Identifier no. 10.1109/TLT.2009.9.

The following scenario illustrates part of a vision for personalized lifelong learning and remembering in terms of a long-term user goal that involves many learning episodes and elements.

Alice decides that she needs to become healthier. She is already aware that she should increase her physical activity and learn more about health. She acquires a pervasive, personalized coaching system as an upgrade to her watch: This unobtrusively monitors her physical activity as well as her heart rate. The system can give advice when and where she requests it, for example, at her workplace desk, at machines in her gym, or on her phone when she is on the train to work. The system augments her memory, enabling her to find information she has seen previously, such as a picture shown to her by the coach as she stretched at the gym last week.

The system has a long-term user model of her fitness and knowledge development. She controls the user model, keeping the full model only on her home machine but releasing selected partial models elsewhere. For example, she enrolls in a course on weight training at a local college. She releases selected parts of her model to the teacher at the start of the course. As the course progresses, her user model is updated. Most mornings, she checks the long-term user model overview. This means she can monitor her progress, celebrating successes or reflecting on slips in progress, seeking advice from the coach and deciding on a plan to improve. At any time, Alice can ask the system to explain why it chose to give particular advice. Every few months, she does this, scrutinizing the details underlying the model and checking the reasoning of the coach to gain a better understanding of its advice. She feels in control of each element in this learning process.

This paper presents a vision built upon several of the elements illustrated in this scenario, which involves a pervasive computing system with a range of novel interfaces. It runs over a very long time, since the goal to become healthier is a lifelong goal. It uses a variety of sources of evidence about Alice, such as evidence from her watch and from her interaction with the desktop coach. Another important feature of the scenario is that Alice has meaningful and comprehensive access to, and control over, her personal information in the pervasive computing environment.

1.1 Lifelong User Models as First Class Citizens

This paper explores the user model as a key to the broad vision for personalized learning and remembering the Grand Challenges [1], [2], [4]. We need to clarify the definition of *user model*, especially as its meaning in human-computer interaction is the system builder's mental model of the user. Typically, this has no explicit representation at all. By contrast, one of the defining aspects of personalization research is the explicit representation of a model of the user. For example, an early definition of user model is the set of beliefs that the machine holds about the user [6]. Importantly, this emphasizes the independent representation of the user model and it makes a distinction between the user model and the *processes* for building, managing, and reasoning about it.

For the lifelong user model, we need to go even further. We distinguish between *evidence* for the user modeling process and the user model *components* modeled. For example, in our scenario, Alice's watch collects data such as accelerometer readings. In this case, accelerometer data is evidence that can be used to reason about the value of the component that models her activity; for example, that she is running or jumping or walking or sitting still. Other sources of evidence come from the coaching system's interactive learning activities and from analysis of the content of documents and Web pages she reads. Alice may also give explicit information, for example, about her goals. Evidence may come from a range of sources. It may be processed or interpreted in the process before it is added to the user model. Evidence is associated with components so that it is possible to reason about the value of those components. Collections of these components will be evaluated in order to drive personalization.

The term *learner model* is used in the Intelligent Tutoring Systems (ITS) research community. Its meaning varies across systems but it often refers to just the learner's knowledge and misconceptions [7]. This is often tightly linked to the ITS representation of the knowledge in the domain of teaching. For example, the overlay model [7] represents the learner's knowledge in terms of the system's own domain knowledge, overlaying this to mark the aspects the learner knows. This paper is concerned with the broader user model, since it subsumes the learner model.

The key to our vision is the recognition that the lifelong user model should be a *first-class citizen* in the sense that it has intrinsic value independent of any one application. We can identify several important roles for the lifelong user model in supporting:

1. personalization in Intelligent Tutoring Systems (ITSs),
2. individual learning because it can support reflection, planning, and control by the learner, as well as their advisors, including teachers and parents,
3. collaborative work based on sharing the user model with partners and team members,
4. augmented cognition via personalized tools for finding, refinding, and remembering, and
5. reuse of the same user model by different applications for all of the above purposes.

The first of these has been the dominant role in ITS learner modeling research. John Self, one of the fathers of the field, observed that the defining characteristic of ITSs is that they know enough about the learner to personalize teaching [8]. There has been much important progress in learner modeling and this paper draws upon it.

The second and third roles are broader, involving learning activities in many contexts, including conventional e-learning environments as well as other tools that people normally use in their work and other activities.

The fourth is even broader, as it treats all the digital artefacts that a person owns or accesses as part of their augmented memory and as potential evidence for their user model. This will help the learner to find artefacts they recall, as well as those they have forgotten about, but which may be useful in their current context. This is particularly important for our aging population and for the even broader class of people who suffer from information overload.

1.2 Widespread, Fragmented User Modeling

A challenge for creating powerful lifelong user models comes from the fragmentation of relevant evidence about the learner. This occurs at several levels. Even if we consider the devices the learner owns, there are several forms of fragmentation. Consider the simplest case, a single desktop computer used only by the learner. Even here, information that could usefully contribute to the individual's user model is typically trapped in forms available only to a single application. Desktop search is partially addressing this. But, there is still much to be done before the information can be made available to a lifelong user model. For example, one maths e-learning tool may have a rich collection of evidence about the learner's knowledge and this is unavailable to another maths e-learning system. In addition, we currently lack tools to exploit potentially valuable information held across several computers and other devices, such as our mobile devices.

Another problem is that our digital footprints are often stored on distant servers, both private as well as publicly available. Consider the case of social networking sites such as Facebook[1] and LinkedIn,[2] where many people choose to share considerable personal information. This is stored on the servers of those sites. Such sites could well provide evidence for the lifelong user model as they provide an API to access the information. A site like Zoominfo[3] builds structured, often extensive, profiles of people. This is an example of the potential for modeling a learner by automatically aggregating information on the Web.

There are even more extensive and rich sources of user modeling evidence in private databases held on various servers. For example, the Learner Management Systems (LMSs) such as Blackboard,[4] Moodle,[5] or Sakai[6] are widely used. Such systems amass data about each learner's activity and assessment results. This is trapped within each class site. It has the potential to be valuable evidence for the learner's lifelong user model.

In summary, there are vast amounts of information that have the potential to contribute valuable evidence to an individual's lifelong user model. At present, it is fragmented across many places, both within the learner's machines and in many public and private databases and systems. The lifelong user model could change this, serving as the repository of user modeling information for learning.

1.3 Overview of Paper

We are currently a long way from being able to build the lifelong user model that can truly support universal, personalized lifelong learning. We still need to make progress in several fields, both technical and broader. At the same time, there has been a huge body of work that will provide foundations for this vision. The remainder of the paper explores some of the challenges to be overcome in achieving that vision. First, it explores the foundation technical research for the *systems issues and representation* of the user model, including a range of issues in storing it and reasoning about it. Next, the paper shifts the focus to the learner and the people around them and their *interaction with the lifelong user model*: We consider the issues of interfaces between learners and their learner models. The final discussion and conclusions link these to a research roadmap for creating the lifelong user model.

1. www.facebook.com.
2. www.linkedin.com.
3. http://www.zoominfo.com.
4. http://www.blackboard.com.
5. http://moodle.org.
6. http://sakaiproject.org/portal.

2 Systems Issues and Representation

To support lifelong user modeling, it will be critical to build an effective set of tools for building, maintaining, and supporting reasoning about user models. This section describes the requirements for the lifelong user model and the drivers for my research to meeting them. Taking first the learner's perspective, the lifelong user model must enable the learner to control:

1. what is allowed into their model,
2. which parts of the model are stored on which devices, and
3. which parts of the model should be shared with particular applications and people.

The first of these involves deciding which evidence and components should be included in the model. Returning to our introductory scenario, Alice must be able to control whether her watch's measurements of her pulse are included in her user model. She may also want to include her model of her exercise partner. This is similar to the way i-Help [9] models both the learner's own attributes and their models of other learners.

The other two requirements concern the control of information flow out of the model. The first is driven by the pervasiveness and ubiquity of computers. So, for example, Alice may decide to keep her complete model on a home computer. She may want some parts on her mobile phone, a different *partial model* on her smart training watch, and yet another on her work computer. This means she can ensure that any device has only the partial models she is comfortable with.

The third aspect is the most common conception of privacy control. For example, Alice should control which partial models are available to her electronic coach or her human training partner.

We now turn to the systems and software level. The lifelong user model must provide the mechanisms for:

1. supporting the controls described above,
2. defining the user model ontologies and defining the components to model and their meaning,
3. reasoning about the user model,
4. distributing the user model across different machines and supporting distributed, flexible access to it, and
5. the special demands of the very long-term nature of the lifelong model.

The remainder of this section discusses each of these and the ways that existing approaches can address these technical demands of lifelong user modeling.

2.1 Supporting Learner Control of Their Model

To enable a user to control what goes into their model, what leaves it, and other privacy concerns, there are challenges at both the systems level and in creating effective user control interfaces. To meet these, the underlying software and its representation provide essential foundations. This view drove the design for a user model representation [10], [11], [12] that supports *scrutability*. It aimed to support the learner in scrutinizing (digging into) a personalized system and its user model to understand what the model means and to see and control the processes that determine the content of the model and reasoning processes associated with it.

A starting point was a theoretical analysis of interaction. This distinguished three classes of information about the learner: that explicitly *given by a person*, for example, when the learner answers a question or the teacher gives a grade; *observations* of the user, such as when their activity is monitored as they edit a document; and information *given by the machine to the learner*, reflecting the way a teacher assumes that the learner knows things they have been told or taught. These forms of evidence have different privacy implications and the representation supports controls based on them.

The representation has two key operations on the user model. *Accretion* is the process of adding evidence to the model, associating it with model components. Learner control of this stage defines what is allowed *into* the model. *Resolution* involves just-in-time evaluation of the meaning of a collection of evidence. Control of this means defining what is released from the model. Our approach gives two levels of such control. One is based on the evidence type and its source. So, for example, Alice may only allow her friend access to information based on evidence *explicitly given by her*. The second control is in the choice of resolution process. Alice can choose which resolvers should be used for particular applications. For the case of location information, we have created many resolvers. One uses ontological reasoning to control the granularity of the location reported [13]. Personalized ontologies give results personalized to the viewer's knowledge [14]. We have created interfaces for users to select which resolver to use for different people and to select the desired parameters [15]. So, for example, people can select a resolver that releases their work location only for the last 6 hours and only within business hours.

The system accepts evidence only from allowed sources. All evidence about the learner is tagged with its source when it is added to their model. All evidence is also timestamped. This too has a role in the privacy controls. We incorporated it in a privacy control interface [15] for sharing location information. Since privacy preferences vary so much among individuals [16], such personalized control is critical for the lifelong user model.

The components of the model are grouped in hierarchical *contexts*, or namespaces. So, for example, we created a context for the learner's knowledge of unix, another for the parts of their model associated with a text editor, and another for the learner's characteristics of machine use, such as typing speed [10]. We created one coaching system for unix and another for the text editor which has a tight integration with unix. The editor coach used components from the editor model, as well as some from the unix model and some generic learner characteristics. For the lifelong learner model, the context provides a mechanism to partition components. Privacy control can operate at the context level.

2.2 Ontologies and Reuse of the User Model

We now turn to the second key mechanism that must be supported, definitions of what is modeled. Ontologies are important for reusing the model across applications. This requires either an agreed upon ontology *understood* by these applications or a mechanism for mapping or harmonizing different ontologies within them. There has been considerable work to establish agreed upon ontologies for some core parts of the model [17] and for elements in Instructional Science and Instructional Design [18]. There is also a clear role for standards, ontology languages [19], pedagogic models [20], and e-learning standards such as PAPI and LIP for learner information [21] and competencies [22], [23]. Such ontologies are clearly very important for the reuse of parts of the user model. This is particularly important for the lifelong user model since the learner develops many skills over many years, making use of many applications. So, for example, if the learner begins using a new teaching system for maths or reading, there is great potential value in that new system being able to build from a model of the model of the learner's existing knowledge, expertise, and goals.

While some parts may be reusable, many have long-term value even if they are used in conjunction with a single application. This is the case for applications that are used over the long term. It is also valuable for cases such as those we describe in the next section, where the learner can benefit from reflecting on their model for purposes such as planning their learning.

Even when reuse is a secondary goal, ontologies are important. They define what will be modeled. As Self [24] pointed out, we should only model what is useful. It is critical that we find ways for teachers to define the learner models that suit their own teaching. This means that a teacher must be able to easily adjust it to the needs of their course. We have explored the challenges of selecting an ontology that meets these needs [25]. An attractive starting point is an online dictionary which we mine to create the ontology [26]. A teacher can augment or refine this ontology simply by writing additional dictionary definitions for local use. For example, we defined the notions of *core* and *advanced* concepts and topics to meet the needs of our teaching [27]. This approach also supports explanations of the ontological reasoning based on the dictionary definitions. Moreover, these serve a useful role as a course glossary. A very important aspect of this approach is that some parts of the ontology are very specific to this course and are useful in that context.

Ontologies affect several other aspects, including inference across granularity levels. This is important because evidence about the learner is typically at one granularity level but the learner may need others. For example, suppose Alice has a high level goal to learn about weight training. A coach may provide fine-grained evidence of her knowledge of triceps training. When she wants to assess her progress at the more general level, an ontology can enable a system to infer this, using the fine-grained evidence.

The ontology can also structure a large user model, making it feasible to build visualization interfaces [28]. We used the same visualization to help teachers tag learning objects [29]. This overcame some of the difficulties that may drive teachers from concept-based approaches to coarser-grained topic-based modeling [30].

2.3 Reasoning about the Learner

For this aspect, one key issue is how much of the reasoning belongs *within* the lifelong user model and how much should be left within applications. User modeling research has explored many approaches to knowledge representation. For example, Bayesian reasoning is demonstrably effective for combining available evidence about the learner. Notable examples are Lumiere's advice based on evidence from user actions [31] and interfaces enabling learners to explore them [32]. Other important work has been based upon rules [33], [34], constraints [35], [36], and conceptual graphs [37]. In all these cases, the user model and associated reasoning has existed within a single application.

The issues change when the learner model, as a first-class citizen, exists outside the application. What knowledge representations are then needed for the lifelong user model? Part of the answer might come from work on user modeling shells [38], designed for generality. These show a trend toward impler representations [38]. This simplicity is also important for supporting scrutability because a simpler system should make it more tractable to build interfaces that explain the user model and modeling processes.

The *stereotype* is a particularly important form of reasoning about users. The term was introduced by Rich [39] and a subsequent pragmatic definition [40] is based on *triggers*, a small set of attributes that are easy to acquire and which can be used to infer many likely attributes. For example, the Unix Consultant [41] showed an elegant use of double stereotypes: A user's answer to a single question about their unix expertise was used to infer a rich default model of their knowledge. Observations of the user were used in the opposite direction; for example, sophisticated usage triggered the inference of high expertise, and then the expert profile.

Such inference may be useful in the lifelong user model, at least to prime it initially for models of new learning contexts. A key aspect is that the stereotype captures statistical knowledge based on common patterns of characteristics among subgroups of users. Orwant [42] introduced a similar notion, the *community*, which differs from stereotypes in that each person can have partial membership of multiple communities. For example, a person may have a 20 percent match with the stereotype for men and an 80 percent match with that for women. So, inferences are based on a weighted combination of the male and female stereotypes. The stereotype is a core, widely used mechanism for reasoning about users, even in nonadaptive systems.

Another important issue for the effectiveness of the lifelong user model will be the interpretation of available evidence to help a learner see whether they have been learning successfully. For example, Alice may judge her progress in the course she is doing on weight training. However, she may also need normalized information to help her assess whether she is actually doing well. So, for example, she may want to compare her performance against that of others who have been successful in weight training on completing the course. Only then can she see that she is already performing at an elite level.

2.4 Distributed Learner Modeling

Our vision of the lifelong user model allows the learner to keep parts of their model on different devices. We have created a software framework that operates seamlessly even when parts of the user's model are stored across different machines [12]. It needs to use discovery mechanisms to determine where the model is, and then it needs to make requests for user model information. We have demonstrated its flexibility for building pervasive applications that can deliver information to the user on any of their devices. Another approach to distributed user modeling sees a person's user model as being distributed across different applications [43]. A very different approach, based on a centralized server for the models of many students, has been explored in the KnowledgeTree project [44].

2.5 Issues for Decades of Modeling

A final set of issues for the lifelong user model relate to time. The modeling system must operate over decades. Although it is technically straightforward to keep large stores of personal digital information, it is less clear that people want to keep all parts of that model indefinitely. Should the model include *use-by* dates, after which the information is deleted? Should there be automatic compaction of collections of old evidence? For example, if there is evidence from the learner correctly solving a large series of arithmetic problems, should this be compacted to a single piece of evidence summarizing the number solved over that period? This would have the merit of simplifying the model as a cost of loss of detail.

There are many interesting issues for what information deserves to be kept within the lifelong user model. For example, our episodic memories are important for learning, as in the ELM-ART system [45]. So, the documents the learner studied or created might be linked to the model. There is also potential value in linking to e-portfolio systems [46].

There are also important issues associated with concept drift [47] as the user changes, an important issue for modeling learners as they learn. We also need to design the user model representation for flexibility of interpretation. This has already been discussed in terms of the case of different applications having different interpretations of the user model evidence. However, the same mechanisms for the important issues are associated with forgetting. This has two aspects. Our user model representation [10], [12] can support modeling of the learner forgetting. It can also apply flexible interpretations, taking account of time.

It will also be important to have effective mechanisms for the learner to ensure that the user model *forgets* parts of its model if that is what the learner wishes. There are also many other intriguing possibilities. For example, just as human memory holds different information in different parts of the mind, so may the user want their distributed user model to move parts of the model to different machines. So, a user may want all model information that

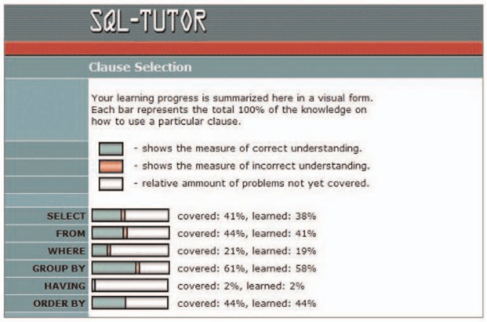

Fig. 1. Example of a skill meter from SQL-Tutor.

Fig. 2. Example of a navigation interface user model.

is more than one year old to be stored only on their home machine.

3 Transparent, Open, and Scrutable User Models

From the technical issues just discussed, this section moves to the user interface challenges of supporting the learner in controlling and using their lifelong user model. This section draws upon the substantial body of research into ways to give the learner a better understanding of their user model within a single application. Much of that work has been motivated by the many learning benefits that might be achieved from enabling learners to see suitable forms of user models. One of these is to empower learners by giving them a better understanding of themselves. This can be the foundation for monitoring progress in learning, for reflection and planning as well as developing metacognitive skills [48]. This requires a meaningful form of relevant parts of the user model, a valuable part of interaction with the lifelong user model.

We also need to draw on research informing the creation of interfaces for users to understand and control personalized systems. Kristina Höök introduced the notion of *glass box* [49], which contrasts with the way that current applications are usually black boxes, hiding details about the user. The glass box enables the user to see into the application, revealing relevant information in a suitable form. This is similar to the idea of *open learner models* (OLMs), which have been described along a range from *transparent* to *scrutable*.

The remainder of this section introduces examples of work that has explored ways to support various levels of transparency and control over user models. The examples have been chosen to illustrate approaches across a range of systems, from personalized teaching through conventional e-learning to general purpose tools used for learning. These examples involve opening the user model within a *single* application: However, they provide lessons for the new challenges of the lifelong learner model.

3.1 Intelligent Tutoring Systems

Learner models have a critical role in Intelligent Tutoring Systems. This makes them an excellent starting point for understanding how to open the lifelong user model to the learner. Many ITSs have very complex representations of the learner. Consider, for example, the Cognitive Tutors [34] and Constraint-Based Tutors [36], which are distinguished in their wide deployment. These have complex learner models with hundreds of rules or constraints. How can such systems usefully share such a complex model with the learner? One answer is the *skill meter* available in both Cognitive and Constraint-Based Tutors. Evaluations indicate that students value a skill meter [34], [35], [36]. In the case of the SQL-tutor, it achieved significant learning benefits [36], particularly for weaker students.

We illustrate this approach with an example of its skill meter, as shown in Fig. 1. This display is available to the learner at any time as they tackle problems posed by SQL-tutor, writing their SQL solutions and studying the results of their solutions. The skill meter shows progress as six bar graphs. In each of these, the leftmost (green) part indicates correct knowledge demonstrated in the tasks done so far. The next (red) part indicates incorrect understanding, and the rightmost (white) part indicates the knowledge for problems to come. This can be used to select the right problems to practice.

The easily understood skill meter is a carefully crafted display based upon a very complex underlying representation. This display is a form that may be readily sent to a lifelong user model. Also, the learner could use this display to decide whether to include it in their model.

Another elegant open learner model provides navigation advice in the ELM-ART tutor [45]. Fig. 2 illustrates this approach with an example from part of the contents page of the SASY-Unix tutor [50]. The top line, *Shell Overview*, is black because the learner has completed that topic. The rest is coded using a traffic light metaphor. The second line is green because the learner knows the prerequisites for this material and is ready to learn it. The remainder is red and not recommended as the learner lacks the prerequisite knowledge.

The SASY framework supports authoring scrutably personalized Web sites such as the page shown in Fig. 3. A link to the user's model is always visible (see *Your Profile* in the figure). Users can then alter their user model if they wish. But, the scrutability goes much further than this. For example, in Fig. 3, the large left pane has teaching material while the one at the right, enlarged in Fig. 4, makes the personalization visible. In this figure, its links indicate to this user that two items were omitted from their page and five were included. Clicking these shows what was omitted/included and the rule driving that. The pane lists just those user model components driving personalization of this page. In the figure, the first is the learner's goal, *You want to get more than a pass grade*. Clicking the *why?* link

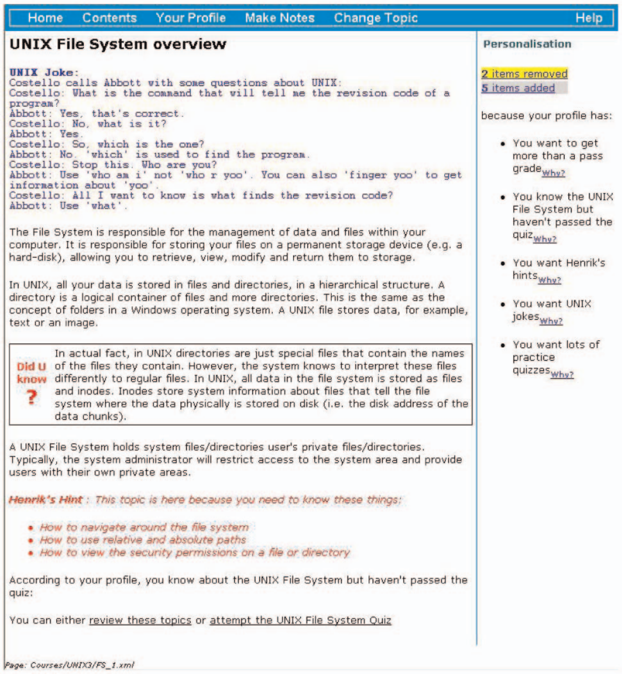

Fig. 3. SASY's interface for scrutinizing personalization.

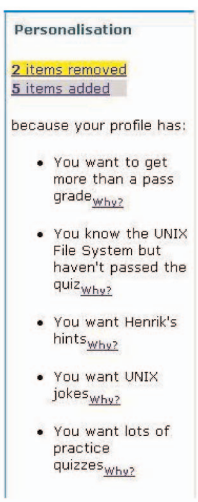

Fig. 4. SASY's scrutiny panel.

shows the evidence for this belief. In field evaluations, learners scrutinized the system extensively, particularly after each quiz [50].

SASY was designed for systems that have small user models, adequate to drive personalization of many applications. Unlike the skill meter discussed above, the user model it shows is very close to the underlying representation.

Of the issues in opening learner models, Bull et al. have explored many, including: negotiated models [51], multiple forms of the learner model [52], presentations for different goals [53], levels of use [54], how students edit their models [55], their use in understanding misconceptions [54], trust gained [56], and sharing the model with peers [57].

For this paper, some important work explored OLMs for young children and their parents. These groups play key roles in lifelong learning. An example of an open model for young children learning basic maths is shown in Fig. 5. The

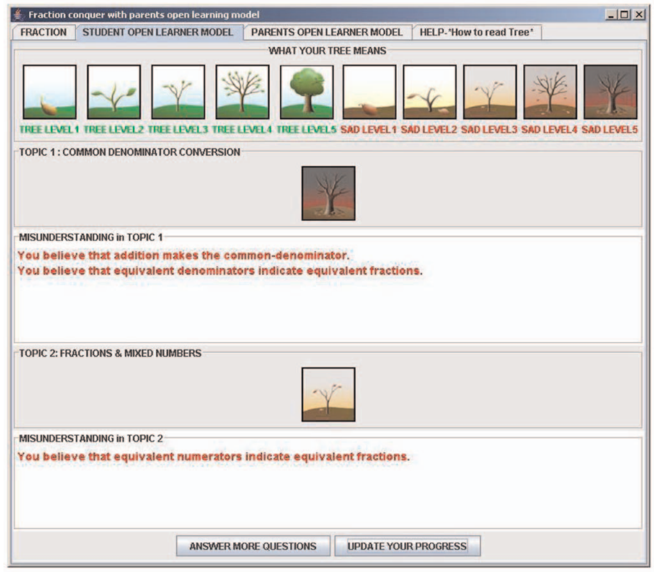

Fig. 5. Overall maths progress display for children.

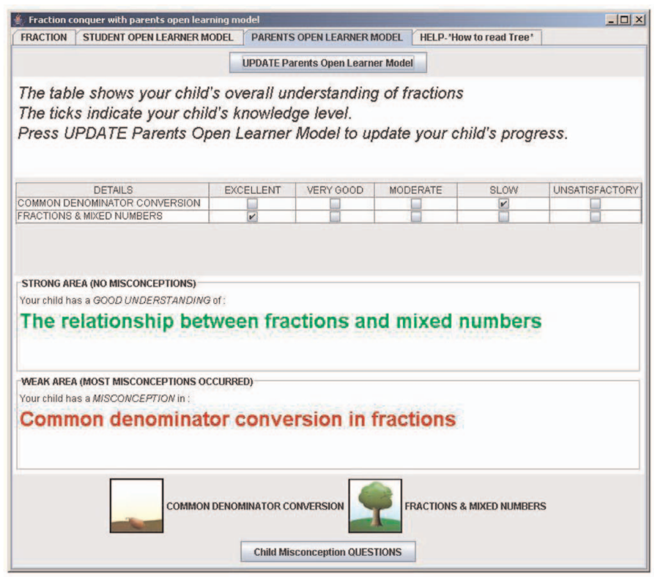

Fig. 6. Parent interface for maths tutor.

tree health shows how well the child has been performing. The interface invites the child to do more problems to improve the tree's health. Once the child has done a problem, they can access this screen, and then click to see the effect on the learner model. The model for parents is illustrated in Fig. 6. It shows information enabling parents to help their children.

Another promising exploration of opening the learner model to parents is the active report [58], as shown in Fig. 7, which displays results from standardized testing. This interface, too, was designed to enable parents to see the details of their child's skills and knowledge and how they can help.

3.2 E-Learning Systems

We now consider the broader classes of widely used e-learning systems. For example, Learning Management Systems (LMSs) collect considerable data over a semester. Typically, this is translated to a final grade which may

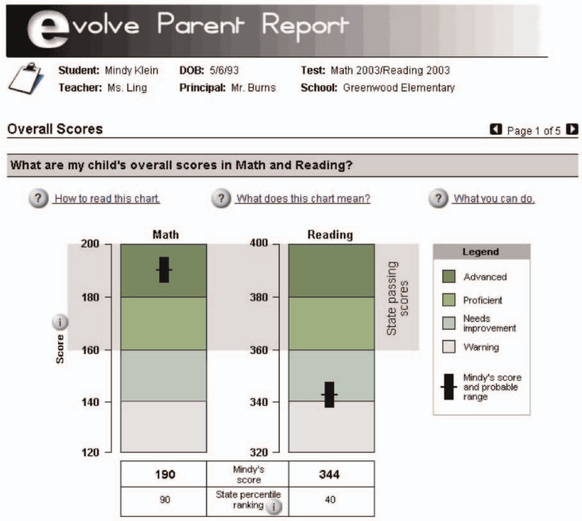

Fig. 7. Example of an active report for parents.

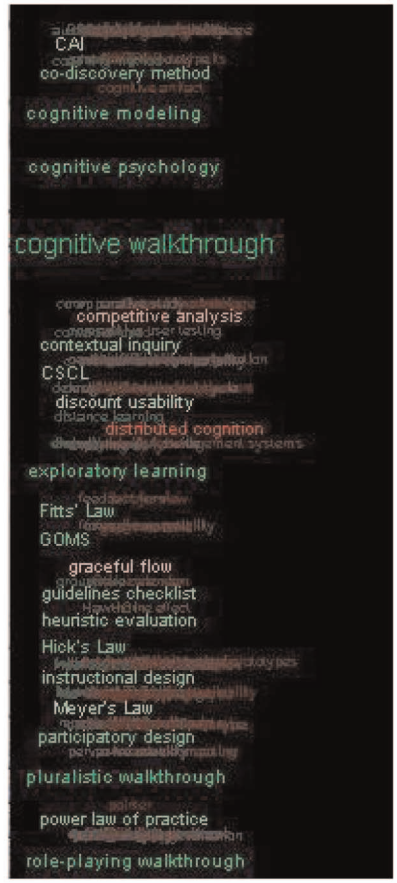

Fig. 8. Example of a visualization interface for large user models.

become part of a transcript. Yet, the detailed data could contribute to a useful and rich model. A major problem is the lack of a *knowledge layer* linking LMS learning data to relevant user model components for the learner's knowledge. We explored ways to address this problem in the context of an LMS used for teaching interface design. We used an automatically extracted ontology as the basic learner model structure [29] and used this for semi-automated tagging of learning objects, such as lecture slides with audio and laboratory assessments. We have also used a different approach in our LMS for a C programming course [27]. In this case, when the teacher adds a new activity, they link it to the relevant course learning goals. This serves three purposes: It captures the teacher's reasons for defining the activity, it shares this with the students improving their understanding of what they are intended to learn, and it supports learner modeling.

For long-term learning, the learner model will be large. We have created and evaluated interfaces of large learner models [59], [28], [60]. The initial design was for a medical course with 600 learning topics. Students had access to a collection of several thousand self-test multiple choice questions. We wanted to create a visualization that would enable a learner to see their overall progress and to see areas of weakness.

The results were the animated visualization interfaces, Viewer for Large User Models (VLUM) [61], [59], [62] and Scrutable Inference Viewer (SIV), an ontology-based version which supported inference across granularity levels [28], [60]. Fig. 8 is an example for a course in human-computer interaction. The learner can gain an overview of their progress by checking the overall color, based on a traffic light metaphor, with known concepts in green (the purer the green, the better known) and misconceptions or unknown ones in red. Yellow indicates the system could not conclude the knowledge level, either due to lack of evidence or to conflicting evidence. The interface permits flexible definitions of the boundaries of knowing; for example, the learner can see their own model compared with the class average [60] or the learner can set the threshold value for a component to become green [28]. Horizontal position indicates certainty, with concepts at the left being more certain as there was more evidence and the rightmost ones having no evidence.

The animation enables the learner to select the focus concept. This is the most visible, having the largest font and spacing around it. In the figure, it is *Cognitive walkthrough* which is bright green as the model indicates that the learner knows it well. The next most visible components are those that the ontology indicates to be most closely related to it, such as *discount usability*. If the learner clicks on any concept, it becomes the focus as the animation alters the display.

Evaluations were conducted with displays for 100, 300, 500, and 700 components visible at once. Users could readily use the model reliably and quickly with up to 500 model components and, even at 700 components, there were only small drops in performance [59]. The lifelong user model will include contexts with many components. User control over the model will require interfaces which, like VLUM and SIV, support scrutiny of large sets of components.

3.3 Conventional Tools

For lifelong learning, it is particularly important to consider the case of learning within normal working contexts. For example, Carroll and Rosson [63] describe the production paradox, where people know they could be more productive if they took the time to learn more about tools like text editors, but they feel too pressured to make the time to do so. If we could provide pointers and tutoring at just the

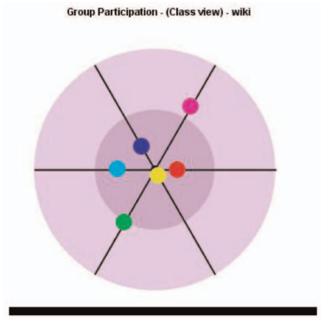

Fig. 9. Visualization of wiki activity.

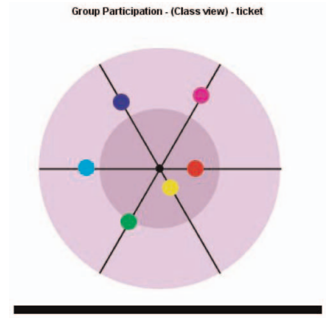

Fig. 10. Planning activity—tickets.

right time, in just the right way, there could be huge benefits to productivity and reduced frustration. This was the inspiration for our building long-term models of users of a text editor [10] and the scrutable model of expertise [64]. Similarly, Linton and Schaefer [65] modeled worker's competence in Microsoft Word. This identified potential learning goals based on the observed knowledge of other workers in the same role. This approach could support an individual's private learning as well as teaching by coworkers identified as knowing those aspects.

We take an example in the important and challenging area of learning group work skills. These are critical for many workplaces and in collaborative learning [66]. Soller [67] identifies three key levels of support for collaboration: mirroring, metacognitive tools, and guiding tools. The simplest, *mirroring*, shows group member activity, requiring learners to reflect on this to decide what to do. *Metacognitive tools* provide supplementary information helping learners judge the mirrored information. *Guiding tools* explicitly inform the learners about how to improve. Long-term group projects involve complexity and subtlety that is beyond currently modeling technology. However, mirroring tools have potential value in summarizing useful information, which learners interpret and which can underpin discussions and guidance by teachers.

Group work is often supported by software such as Trac,[7] a tool for group software development. It has a *wiki* for collaborative document editing. Groups can define tasks to be done in *tickets* and link these to *milestones* which are displayed on a *roadmap*. The *timeline* shows all actions with the most recent first. There is an interface to a *version control system*, such as *svn*. Importantly, all of these are integrated, so that tickets can link to change sets in the version control system and the wiki. While these tools are valuable, it is difficult to see just what each person has done and how team members interact. We concluded that a learner model mirror could help.

We created a set of visualizations guided by the Big-5 Theory for small teams [68]. We have used these over several semesters in a capstone software project course which explicitly teaches students group work skills [69]. Groups of five to seven work for 13 weeks on a large project for a client. A tutor meets each group in weekly labs. The lecturer has role-based meetings. For example, all of the *managers* meet to discuss group progress and problems. We wanted to create learner model mirrors that would facilitate discussions within teams and with teachers.

One of these, inspired by Erickson et al.'s social translucence [70] research, shows activity. Fig. 9 shows an anonymized example for one group's wiki activity. Fig. 10 shows ticket activity and Fig. 11 shows activity on the version control system. Each team member is represented by a dot with consistent color and orientation in all displays. So, for example, the red dot at 3 o'clock is always the same person. Names normally appear at the edge of the circle, but were removed to anonymize these images.

We aggregate an individual's activity and display their dot closer to the center if they have been more active. The darker circle indicates the average level of activity over the whole class. So, the person represented by the 3 o'clock (red) dot was the most active on svn, matching a programmer role in the group. As this dot is right in the center, this person had one of the highest activity levels in the whole class. The yellow dot at 5 o'clock had the highest activity on the wiki and tickets, matching the group manager role. Metacognitive activities of reflection and planning require knowledge of the team member roles. They also must draw on other information; for example, the group may agree that one person should create all tickets.

Activity is important, but interaction between team members is also an indicator of the group dynamics. We created interaction mirrors, like Fig. 12 and Fig. 13. These maintain the same color and clock positions for group members. Interaction is regarded as editing the same wiki page, svn document, or ticket. The heavier the line, the more interaction. So, for example, in Fig. 12, all team members interact but some people interact more. Pathological groups can show isolated individuals. In Fig. 13, the lighter (blue)

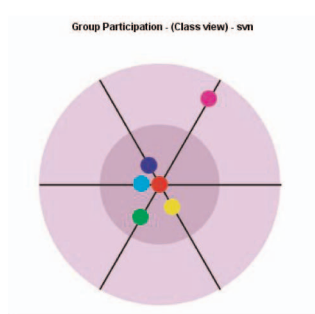

Fig. 11. Activity on svn version control system.

7. http://trac.edgewall.org.

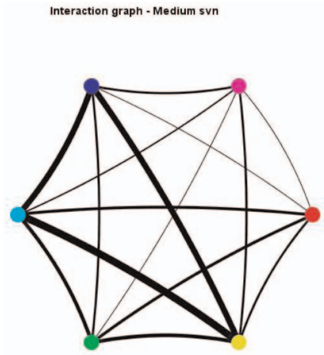

Fig. 12. Interaction on svn version control system.

end of the line indicates the person creating the tickets, often, the manager.

Group members and teachers need to discuss these displays. For example, one person may write poor code with others having to edit it to fix the bugs. Equally, a team member who does not trust others may simply edit their code. Distinguishing these cases requires deep knowledge. The mirrors can be used in reflective classroom activities.

It is also important to gain a sense of the temporal changes in group activity. For this, we created the visualization shown in Fig. 14. This shows a full semester of activity. Time is on the y-axis, starting at the bottom of the figure. Each group member is depicted as one green line and its round *flowers* and line *leaves*. For each day, we display each person's activity: wiki at the left (yellow), svn on the right (orange), ticket creation (dark line), and closing a ticket on task completion (light green line). This visualization makes it easy to see each team member's relative contributions.

2This figure shows a period of inactivity toward the top. This was the semester break. There are clear differences between the activity of different group members. The person fifth from the left does most of their activity on svn, matching a programmer role. There is a very different pattern for the group in Fig. 15, with considerable activity on the wiki but little on svn—an unhealthy profile for a software project. These particular figures were created at the end of the semester and were not available to the groups until then. Subsequently, the visualizations were available throughout the semester and were used in meetings

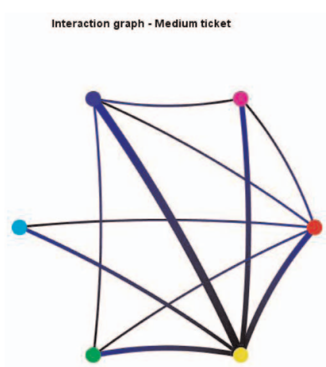

Fig. 13. Interaction on planning using tickets.

Fig. 14. Visualization of activity of each member of a team in a strong group.

between groups and facilitators. Teachers found they could identify group problems very early and help address them. This has largely reduced the most serious group problems. The mirrors enabled individuals to realize whether they are performing at appropriate levels. Surveys indicate that team leaders particularly valued them [71].

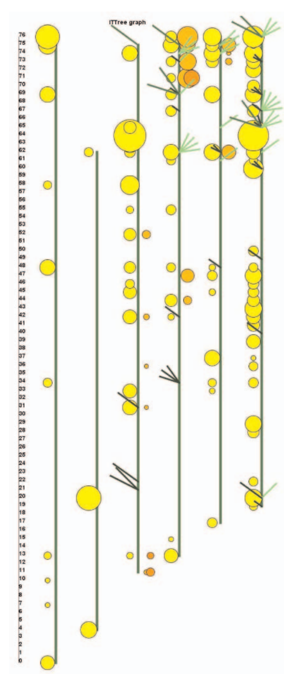

Fig. 15. Visualization for a group with limited svn activity.

Fig. 16. Users interacting at a tabletop interface.

Similar approaches have been explored for e-learning discussion [72]. Mirror tools offer promising solutions for ill-defined and complex learning contexts, such as in group work. They also point to some of the potential benefits of access to a long-term user model.

4 Discussion, Roadmap, and Conclusions

To this point, the paper has focused on two areas that are critical to lifelong user modeling and where there has been substantial work already. First was the systems issues of representation, ontologies, reasoning, distributing the model, and handling long-term issues. Although existing work has primarily involved a single application having its own user model, this paper has discussed how that work provides foundations for the lifelong user model. As a balance to that purely technical focus, the other main area described was some of the explorations of user interfaces that can enable the learner to see and explore their learner model. This section briefly introduces a selection of other important challenges that will be part of a roadmap for lifelong learner modeling.

4.1 Pervasive Computing and New Interaction Media

Pervasive computing is introducing new ways for learners to interact with technology. For example, mobile devices can deliver an open learner model to help direct learning on the device when the learner happens to have time [73]. Taking another example, ambient displays have a potential role for very long-term learning goals associated with health [74]. They can provide aesthetically pleasing and subtle displays of learner behaviors to support changing bad habits and consolidating better ones. Another emerging class of interface is *surface* computing, such as the tabletop interface shown in Fig. 16. It can create new opportunities for collaboration especially if we create mechanisms for effective sharing of documents [75]. Challenges include determining how to make effective use of all the new learning data that might be collected for the lifelong user model as well as creating new interfaces that take advantage of it to personalize teaching and the delivery of information, while ensuring the learner can maintain effective control.

4.2 Educational Data Mining—EDM

EDM is a relatively new field that aims to exploit the data created within software used by the learner. We have already described some valuable uses of that data in the form of the open learner model: It has much more to offer if mined for useful patterns and associations. Mining the individual learner's model has had little exploration but it may be valuable if suitable measures of interesting patterns can be identified [76], [77]. For example, we identified patterns of Trac data associated with effectiveness for the manager role in a team [78]; these can give the individual early warning of potential problems. With tools for mining the data for a whole class, a teacher can identify groups of students with particular constellations of difficulties and address them [79], [80], [81]. Such tools offer the promise that every teacher can become an educational researcher and that the teacher can become aware of the nature of learning difficulties of small groups of students in their class. In larger collections of data, a new form of educational research has become possible. For example, it can give a better understanding of how to provide help in teaching systems [82]. It has also been used to provide high quality, continuous, automated assessment of a learner's achievement in reading without the disruption of tests [83]. Even more broadly, the lifelong user model will be more meaningful when the individual can see their own model with additional information that has been mined from the data of larger cohorts of learners.

4.3 Privacy and Learner Control

As the lifelong user model is a repository of personal information, privacy is a core concern [84]. This paper's vision makes the individual learner the owner of their aggregated data, which is kept on the machines they choose. We need to find ways to give the learner control of this model, and this means we must meet difficult user interface challenges [85]. It requires that the learner can, with modest effort, exert effective control of their model, including sharing it as they wish. For example, as already noted, individual data often has far more meaning when the learner can compare their knowledge or performance against a suitable group of other learners or the defined expectations of a teacher. Such information requires the data of many learners. A reciprocity principle makes it fair to contribute one's own data to get the benefit of knowing about the aggregate data from other learners. In addition, EDM can return greater benefits to the individual in terms of insights derived from mining data from many learners. We need to understand more about the ways that learners are willing to share their model, for example, with anonymity or not [57]. We also need to explore the ways the learner might control the sharing of the model from their mobile device, such as a phone [86]. This must be balanced against learner's concern at the increasing amounts of personal data and potential hazards from mining it [87]. Our vision contrasts with the centralized stores that currently hold and control data for many learners. The case of LMSs has already been discussed. The lifelong user model will draw on many classes of information, such as that at social networking sites; their public APIs mean the learner's user model can include the information [88].

4.4 Augmented Cognition

Our broader vision for augmented cognition includes new ways to support refinding information we have accessed [89], [90], [91], even if we have forgotten about it [92]. There is potential benefit in personalized search [93] based on an implicit user model based on their documents and use of community models to improve search [94]. Rich data from a camera that automatically captures data from normal life [95] might be usefully exploited, especially for people with short-term memory loss. An open question is whether the lifelong user model should include, or link to, our huge stores of documents and other digital artifacts, transforming them from an implicit user model that we cannot scrutinize or control to an explicit one that we can.

4.5 Conclusions

Much of this paper has been devoted to two of the foundations for the lifelong user model: the systems issues and the inextricably linked challenges of creating interfaces to the model. There are many questions yet to be explored. Is the lifelong user model too dangerous because it aggregates so much data about the individual learner? Our vision has been driven by the view that the lifelong user model must have a simple representation, with minimal reasoning, leaving that to applications. Will that give learners enough information to realize the important potential of the lifelong user model for supporting reflection, monitoring, and planning learning? What is the appropriate granularity for the user model ontology and the evidence that is collected about the learner? If is too fine-grained, will it be overwhelming to the point of being useless? If it is too coarse-grained, will we have lost valuable possibilities for data mining that were never envisaged at the time the model was defined? Our vision is based on a lifelong user model where large parts of the model are highly specialized with the reuse and value involving learner interaction with the model, even if applications cannot reuse all of it. Can we create interfaces that make this practical and valuable for learning? It is certain that we are creating huge collections of digital artifacts and that many computers hold large collections of information about us. The lifelong user model aims to harness these and make them really work for the learner, being available as needed, supporting all aspects of the life, and, importantly, ensuring that the learner maintains a sense of control over the model, its use, and their own learning.

Acknowledgments

This paper draws on many discussions with colleagues and students, notably, Bob Kummerfeld for systems aspects and Kalina Yacef for educational data mining. Thanks to Susan Bull, Tanja Mitrovic, and Diego Zapata-Rivera for allowing use of their screenshots. This research was supported by Australian Research Council Projects DP0774532 and DP0877665.

References

[1] CRA, "Grand Research Challenges in Information Systems," http://www.cra.org/reports/gc.systems.pdf, 2002.

[2] UKCRC, "Current Grant Challenges," http://www.cra.org/Activities/grand.challenges, Mar. 2008.

[3] J. Taylor, A. de Roeck, A. Anderson, M. Sharples, T. Boyle, W. Hall, T. Rodden, and D. Scott, "A Grand Challenge for Computing: Learning for Life," *Current Grant Challenges*, http://www.ukcrc.org.uk/grand_challenges/news/learninglife.pdf, 2004.

[4] GCE, *Grand Research Challenges for Engineering*, National Academy of Eng., http://www.engineeringchallenges.org/cms/8996/9127.aspx, 2008.

[5] A. Fitzgibbon and E. Reiter, "Current Grant Challenges: Memories for Life: Managing Information over a Human Lifetime," pp. 13-16, http://www.ukcrc.org.uk/gcresearch.pdf, Mar. 2008.

[6] W. Wahlster and A. Kobsa, *User Models in Dialog Systems*, Univ. des Saarlandes, FB 14, Informatik IV, 1988.

[7] J. Greer and G. McCalla, *Student Modelling: The Key to Individualized Knowledge-Based Instruction*. Springer, 1994.

[8] J. Self, "The Defining Characteristics of Intelligent Tutoring Systems Research: ITSs Care, Precisely," *Int'l J. Artificial Intelligence in Education*, vol. 10, nos. 3-4, pp. 350-364, 1999.

[9] G. McCalla, J. Vassileva, J. Greer, and S. Bull, "Active Learner Modelling," *Intelligent Tutoring Systems*, pp. 53-62, http://www.springerlink.com/content/cbycu121ulv8jg65, Springer, 2000.

[10] J. Kay, "The UM Toolkit for Cooperative User Modelling," *User Modeling and User-Adapted Interaction*, vol. 4, pp. 149-196, 1995.

[11] D.J. Carmichael, J. Kay, and R.J. Kummerfeld, "Consistent Modelling of Users, Devices and Sensors in a Ubiquitous Computing Environment," *User Modeling and User-Adapted Interaction*, vol. 15, nos. 3-4, pp. 197-234, 2005.

[12] M. Assad, D.J. Carmichael, J. Kay, and B. Kummerfeld, "PersonisAD: Distributed, Active, Scrutable Model Framework for Context-Aware Services," *Proc. Fifth Int'l Conf. Pervasive Computing (PERVASIVE '07)*, pp. 55-72, 2007.

[13] W. Niu and J. Kay, "Location Conflict Resolution with an Ontology," *Proc. Sixth Int'l Conf. Pervasive Computing (Pervasive '08)*, pp. 162-179, 2008.

[14] W. Niu and J. Kay, "Pervasive Personalisation of Location Information: Personalised Context Ontology," *Proc. Int'l Conf. Adaptive Hypermedia and Adaptive Web-Based Systems (AH '08)*, pp. 143-152, 2008.

[15] D. Carmichael, "Myplace: Supporting Scrutability and User Control in Location Modelling," PhD dissertation, Univ. of Sydney, 2008.

[16] A. Kobsa, "Tailoring Privacy to User's Needs," *Proc. Int'l Conf. User Modeling*, pp. 303-313, 2001.

[17] D. Heckmann, T. Schwartz, B. Brandherm, M. Schmitz, and M. von Wilamowitz-Moellendorff, "GUMO—The General User Model Ontology," *Proc. Int'l Conf. User Modeling*, p. 428, 2005.

[18] R. Mizoguchi and J. Bourdeau, "Using Ontological Engineering to Overcome Common AI-ED Problems," *Int'l J. Artificial Intelligence in Education*, vol. 11, no. 2, pp. 107-121, 2000.

[19] M. Winter, C. Brooks, and J. Greer, "Towards Best Practices for Semantic Web Student Modelling," *Proc. 12th Int'l Conf. Artificial Intelligence in Education*, pp. 694-701, 2005.

[20] V. Psyche, J. Bourdeau, R. Nkambou, and R. Mizoguchi, "Making Learning Design Standards Work with an Ontology of Educational Theories," *Proc. Int'l Conf. Artifical Intelligence in Education (AIED '05)*, pp. 539-546, 2005.

[21] L. Oubahssi and M. Grandbastien, "From Learner Information Packages to Student Models: Which Continuum?" *Intelligent Tutoring Systems*, pp. 288-297, 2006.

[22] M. Kalz, "Positioning of Learners in Learning Networks with Content, Metadata and Ontologies," *Interactive Learning Environments*, vol. 15, no. 2, pp. 191-200, http://www.informaworld.com/index/781052513.pdf, 2007.

[23] IMS, "IMS Reusable Definition of Competency or Educational Objective (RDCEO)," http://www.imsglobal.org/competencies/index.html, 2008.

[24] J. Self, "Bypassing the Intractable Problem of Student Modelling," *Intelligent Tutoring Systems: At the Crossroads of Artificial Intelligence and Education*, pp. 107-123, 1990.

[25] J. Kay and A. Lum, "Exploiting Readily Available Web Data for Reflective Student Models," *Proc. Int'l Conf. Artificial Intelligence in Education (AIED '05),* pp. 338-345, 2005.

[26] T. Apted and J. Kay, "Mecureo Ontology and Modelling Tools," *Int'l J. Continuing Engineering Education and Lifelong Learning,* special issue on concepts and ontologies in WBES, vol. 14, no. 3, pp. 191-211, 2004.

[27] J. Kay, L. Li, and A. Fekete, "Learner Reflection in Student Self-Assessment," *Proc. Ninth Australasian Computing Education Conf. (ACE '07),* pp. 89-95, http://crpit.com/confpapers/CRPITV66 Kay.pdf, 2007.

[28] T. Apted, J. Kay, A. Lum, and J. Uther, "Visualisation of Ontological Inferences for User Control of Personal Web Agents," *Proc. Seventh Int'l Conf. Information Visualization (IV '03),* E. Banissi, K. Borner, C. Chen, G. Clapworthy, C. Maple, A. Lobben, C. Moore, J. Roberts, A. Ursyn, and J. Zhang, eds., pp. 306-311, 2003.

[29] T. Apted, J. Kay, and A. Lum, "Supporting Metadata Creation with an Ontology Built from an Extensible Dictionary," *Proc. Third Int'l Conf. Adaptive Hypermedia and Adaptive Web-Based Systems (AH '04),* P. de Bra and W. Nejdl, eds., pp. 4-13, 2004.

[30] P. Brusilovsky, S. Sosnovsky, and O. Shcherbinina, "User Modeling in a Distributed E-Learning Architecture," *Proc. Int'l Conf. User Modeling,* p. 387, 2005.

[31] E. Horvitz, J. Breese, D. Heckerman, D. Hovel, and K. Rommelse, "The Lumiere Project: Bayesian User Modeling for Inferring the Goals and Needs of Software Users," *Proc. 14th Conf. Uncertainty in Artificial Intelligence,* pp. 256-265, 1998.

[32] J. Zapata-Rivera and J. Greer, "Inspectable Bayesian Student Modelling Servers in Multi-Agent Tutoring Systems," *Int'l J. Human-Computer Studies,* vol. 61, no. 4, pp. 535-563, 2004.

[33] A. Corbett and J. Anderson, "Knowledge Tracing: Modeling the Acquisition of Procedural Knowledge," *User Modeling and User-Adapted Interaction,* vol. 4, no. 4, pp. 253-278, 1994.

[34] A. Corbett, M. McLaughlin, and K. Scarpinatto, "Modeling Student Knowledge: Cognitive Tutors in High School and College," *User Modeling and User-Adapted Interaction,* vol. 10, no. 2, pp. 81-108, 2000.

[35] A. Mitrovic and B. Martin, "Evaluating the Effects of Open Student Models on Learning," *Adaptive Hypermedia and Adaptive Web-Based Systems,* pp. 296-305, 2002.

[36] A. Mitrovic and B. Martin, "Evaluating the Effect of Open Student Models on Self-Assessment," *Int'l J. Artificial Intelligence in Education,* vol. 17, no. 2, pp. 121-144, 2007.

[37] V. Dimitrova, "STyLE-OLM: Interactive Open Learner Modelling," *Int'l J. Artificial Intelligence in Education,* vol. 13, no. 1, pp. 35-78, http://iospress.metapress.com/content/cena9440xxjqdbf0/, 2003.

[38] A. Kobsa, "Generic User Modeling Systems," *User Modeling and User-Adapted Interaction,* vol. 11, no. 1, pp. 49-63, 2001.

[39] E. Rich, "User Modeling via Stereotypes," *Cognitive Science,* vol. 3, no. 4, pp. 329-354, 1979.

[40] J. Kay, "Vive la Difference! Individualised Interaction with Users," *Proc. Int'l Joint Conf. Artificial Intelligence,* C. Mellish, ed., pp. 978-984, Morgan Kaufman, 1995.

[41] D. Chin, "KNOME: Modeling What the User Knows in UC," *User Models in Dialog Systems,* pp. 74-107, 1989.

[42] J. Orwant, "Heterogeneous Learning in the Doppelgänger User Modeling System," *User Modeling and User-Adapted Interaction,* vol. 4, no. 2, pp. 107-130, 1994.

[43] F. Carmagnola and V. Dimitrova, "An Evidence-Based Approach to Handle Semantic Heterogeneity in Interoperable Distributed User Models," *Adaptive Hypermedia and Adaptive Web-Based Systems,* pp. 73-82, 2008.

[44] P. Brusilovsky, "KnowledgeTree: A Distributed Architecture for Adaptive e-Learning," *Proc. 13th Int'l World Wide Web Conf. Alternate Track Papers & Posters,* pp. 104-113, 2004

[45] G. Weber and P. Brusilovsky, "ELM-ART: An Adaptive Versatile System for Web-Based Instruction," *Int'l J. Artificial Intelligence in Education,* vol. 12, no. 4, pp. 351-384, 2001.

[46] D. Gibson and H. Barrett, "Directions in Electronic Portfolio Development," *Education,* vol. 2, no. 4, pp. 556-573, 2003.

[47] G. Webb, M. Pazzani, and D. Billsus, "Machine Learning for User Modeling," *User Modeling and User-Adapted Interaction,* vol. 11, no. 1, pp. 19-29, 2001.

[48] S. Bull and J. Kay, "Metacognition and Open Learner Models," *Proc. Third Workshop Meta-Cognition and Self-Regulated Learning in Educational Technologies (ITS '08),* 2008.

[49] K. Höök, J. Karlgren, A. Wærn, N. Dahlbäck, C. Jansson, K. Karlgren, and B. Lemaire, "A Glass Box Approach to Adaptive Hypermedia," *User Modeling and User-Adapted Interaction,* vol. 6, no. 2, pp. 157-184, http://www.springerlink.com/content/v54532n417624m86, 1996

[50] M. Czarkowski, "A Scrutable Adaptive Hypertext," PhD dissertation, Univ. of Sydney, http://www.it.usyd.edu.au/research/2006_Marek_Czarkowski_thesis.pdf, 2007.

[51] S. Bull, P. Brna, and H. Pain, "Extending the Scope of the Student Model," *User Modeling and User-Adapted Interaction,* vol. 5, no. 1, pp. 45-65, 1995.

[52] S. Bull and A. Mabbott, "20000 Inspections of a Domain-Independent Open Learner Model with Individual and Comparison Views," *Intelligent Tutoring Systems,* pp. 422-432, http://www.springerlink. com/content/u131742343825504, 2006.

[53] S. Bull, N. Cooke, and A. Mabbott, "Visual Attention in Open Learner Model Presentations: An Eye-Tracking Investigation," *Proc. Int'l Conf. User Modeling,* p. 177, http://www.springerlink.com/content/p2272121j9673136, 2007.

[54] S. Bull, A. Mabbott, P. Gardner, T. Jackson, M. Lancaster, S. Quigley, and P. Childs, "Supporting Interaction Preferences and Recognition of Misconceptions with Independent Open Learner Models," *Adaptive Hypermedia and Adaptive Web-Based Systems,* pp. 62-72, http://www.springerlink.com/content/d46123032j123043, 2008.

[55] S. Bull, X. Dong, M. Britland, and Y. Guo, "Can Students Edit Their Learner Model Appropriately?" *Intelligent Tutoring Systems,* B.P. Woolf, E. Aÿmeur, R. Nkambou, and S. Lajoie, eds., pp. 674-676, http://www.springerlink.com/index/p657115270n70727.pdf, Springer, 2008.

[56] A. Kerly, N. Ahmad, and S. Bull, "Investigating Learner Trust in Open Learner Models Using a 'Wizard of Oz' Approach," *Intelligent Tutoring Systems,* B.P. Woolf, E. Aÿmeur, R. Nkambou, and S. Lajoie, eds., pp. 722-724, http://www.springerlink.com/index/1750n46426511525.pdf, Springer, 2008.

[57] S. Bull, A. Mabbott, and A. AbuIssa, "UMPTEEN: Named and Anonymous Learner Model Access for Instructors and Peers," *Int'l J. Artificial Intelligence in Education,* vol. 17, no. 3, pp. 227-253, 2007.

[58] J.-D. Zapata-Rivera, J.S. Underwood, and M. Bauer, "Advanced Reporting Systems in Assessment Environments," *Proc. Workshop Learner Modelling for Reflection, Int'l Conf. Artificial Intelligence in Education,* pp. 23-31, 2005.

[59] J. Uther, "On the Visualisation of Large User Models in Web Based Systems," PhD dissertation, Univ. of Sydney, http://www.it.usyd.edu.au/research/2001_James_Uther_thesis.pdf, 2002.

[60] A.W.K. Lum, "Light-Weight Ontologies for Scrutable User Modelling," PhD dissertation, http://www.it.usyd.edu.au/research/2007_Andrew_Lum_thesis.pdf, 2007.

[61] J. Uther and J. Kay, "Compact Display of Large User Models," *Proc. Pan-Sydney Area Visual Information Processing Workshop (VIP '98),* VIP Lab, School of Computer Science and Eng., Univ. of New South Wales, pp. 60-64, 1998.

[62] J. Uther and J. Kay, "VIUM, a Web-Based Visualisation of Large User Models," *Proc. Int'l Conf. User Modeling,* P. Brusilovsky, A. Corbett, and F. de Rosis, eds., pp. 198-202, http://www.springerlink.com/content/8666afvdbmj3gueb, Springer, 2003.

[63] J. Carroll and M. Rosson, *Paradox of the Active User.* MIT Press, 1987.

[64] R. Cook and J. Kay, "The Justified User Model: A Viewable, Explained User Model" *Proc. Fourth Int'l Conf. User Modeling (UM '94),* A. Kobsa and D. Litman, eds., pp. 145-150, 1994.

[65] F. Linton and H. Schaefer, "Recommender Systems for Learning: Building User and Expert Models through Long-Term Observation of Application Use," *User Modeling and User-Adapted Interaction,* vol. 10, no. 2, pp. 181-208, 2000.

[66] T. Koschmann, *CSCL: Theory and Practice of an Emerging Paradigm.* Lawrence Erlbaum Assoc. Inc., 1996.

[67] A. Soller, "From Mirroring to Guiding: A Review of State of the Art Technology for Supporting Collaborative Learning," *Int'l J. Artificial Intelligence in Education,* vol. 15, no. 4, pp. 261-290, 2005.

[68] J. Kay, N. Maisonneuve, K. Yacef, and P. Reimann, *The Big Five and Visualisation of Team Work Activity,* http://www.springerlink.com/index/w750g0871vm32474.pdf, Springer, 2006.

[69] J. Kay, P. Reimann, and K. Yacef, "Visualisations for Team Learning: Small Teams Working on Long-Term Projects," *Minds, Mind, and Society. Proc. 6th Int'l Conf. Computer-Supported Collaborative Learning (CSCL '07),* C. Chinn, G. Erkens, and S. Puntambekar, eds., pp. 351-353, 2007.
[70] T. Erickson, W. Huang, C. Danis, and W. Kellogg, "A Social Proxy for Distributed Tasks: Design and Evaluation of a Working Prototype," *Proc. SIGCHI Conf. Human factors in Computing Systems,* pp. 559-566, 2004.
[71] J. Kay, P. Reimann, and K. Yacef, "Mirroring of Group Activity to Support Learning as Participation," *Proc. 13th Int'l Conf. Artificial Intelligence in Education (AIED '07),* R. Luckin, K.R. Koedinger, and J. Greer, eds., vol. 158, pp. 584-586, http://www.booksonline.iospress.nl/Content/View.aspx?piid=6846, 2007.
[72] L. Sun and J. Vassileva, "Social Visualization Encouraging Participation in Online Communities," *Groupware: Design, Implementation, and Use, Proc. CRIWG,* pp. 349-363, 2006.
[73] A. Mabbott and S. Bull, "Student Preferences for Editing, Persuading, and Negotiating the Open Learner Model," *Intelligent Tutoring Systems,* M. Ikeda, K.D. Ashley, and T.-W. Chan, eds., pp. 481-490, http://www.springerlink.com/content/d05131748u40288n, Springer, 2006.
[74] T. Nakajima, V. Lehdonvirta, E. Tokunaga, and H. Kimura, "Reflecting Human Behavior to Motivate Desirable Lifestyle," *Proc. ACM Conf. Designing Interactive Systems,* pp. 405-414, 2008.
[75] A. Collins, T. Apted, and J. Kay, "Tabletop File System Access: Associative and Hierarchical Approaches," *Proc. Second Ann. IEEE Int'l Workshop Horizontal Interactive Human-Computer Systems (TABLETOP '07),* pp. 113-120, 2007.
[76] A. Merceron and K. Yacef, "Interestingness Measures for Association Rules in Educational Data," *Proc. Educational Data Mining Conf.,* pp. 57-66, http://www.educationaldatamining.org/EDM2008, 2008.
[77] A. Merceron and K. Yacef, "Revisiting Interestingness of Strong Symmetric Association Rules in Educational Data," *Proc. Int'l Workshop Applying Data Mining in eLearning (ADML'07),* pp. 3-12, http://ftp.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-305/paper%01.pdf, 2007.
[78] D. Perera, J. Kay, K. Yacef, I. Koprinska, and O. Zaiane, "Clustering and Sequential Pattern Mining of Online Collaborative Learning Data," *IEEE Trans. Knowledge and Data Eng.,* 2008.
[79] A. Merceron and K. Yacef, "A Web-Based Tutoring Tool with Mining Facilities to Improve Learning and Teaching," *Proc. 11th Int'l Conf. Artificial Intelligence in Education,* pp. 201-208, 2003.
[80] A. Merceron and K. Yacef, "TADA-Ed for Educational Data Mining," *Interactive Multimedia Electronic J. Computer-Enhanced Learning,* vol. 7, no. 1, 2005.
[81] K. Yacef, "The Logic-ITA in the Classroom: A Medium Scale Experiment," *Int'l J. Artificial Intelligence in Education,* vol. 15, pp. 41-60, 2005.
[82] C. Heiner, J. Beck, and J. Mostow, "When Do Students Interrupt Help? Effects of Time, Help Type, and Individual Differences," *Proc. Int'l Conf. Artificial Intelligence in Education (AIED '05),* pp. 819-826, 2005.
[83] J. Beck and J. Sison, "Using Knowledge Tracing in a Noisy Environment to Measure Student Reading Proficiencies," *Int'l J. Artificial Intelligence in Education,* vol. 16, no. 2, pp. 129-143, http://iospress.metapress.com/content/9t8nqulr6rdapw44, 2006.
[84] A. Kobsa, "Personalized Hypermedia and Int'l Privacy," *Comm. ACM,* vol. 45, no. 5, pp. 64-67, 2002.
[85] M. Czarkowski and J. Kay, *Giving Learners a Real Sense of Control over Adaptivity, Even If They Are Not Quite Ready For It Yet,* pp. 93-125, http://www.idea-group.com/books/details.asp?id=5248, IDEA Information Science Publishing, 2006.
[86] M. Hitchens, J. Kay, R.J. Kummerfeld, and A. Brar, "Secure Identity Management for Pseudo-Anonymous Service Access," *Proc. Second Int'l Conf. Security in Pervasive Computing (SPC '05),* D. Hutter, ed., pp. 48-55, http://www.spc-conf.org, 2005.
[87] S. Baker, *The Numerati: How They'll Get My Number and Yours.* Jonathan Cape, 2008.
[88] D. Kyriacou, "A Scrutable User Modelling Infrastructure for Enabling Life-Long User Modelling," *Adaptive Hypermedia and Adaptive Web-Based Systems,* pp. 421-425, 2008.
[89] S. Dumais, E. Cutrell, J. Cadiz, G. Jancke, R. Sarin, and D. Robbins, "Stuff I've Seen: A System for Personal Information Retrieval and Re-Use," *Proc. 26th Ann. Int'l ACM SIGIR Conf. Research and Development in Informaion Retrieval,* pp. 72-79, 2003.
[90] D. Barreau and B. Nardi, "Finding and Reminding: File Organization from the Desktop," *ACM SIGCHI Bull.,* vol. 27, no. 3, pp. 39-43, 1995.
[91] S. Fertig, E. Freeman, and D. Gelernter, "'Finding' and 'Reminding' Reconsidered," *ACM SIGCHI Bull.,* vol. 28, no. 1, pp. 66-69, 1996.
[92] D. Elsweiler, I. Ruthven, and C. Jones, "Towards Memory Supporting Personal Information Management Tools," *J. Am. Soc. for Information Science and Technology,* vol. 58, no. 7, pp. 924-946, 2007.
[93] J. Teevan, "Supporting Finding and Re-Finding Through Personalization," PhD dissertation, Massachusetts Inst. of Technology, 2006.
[94] B. Smyth, E. Balfe, J. Freyne, P. Briggs, M. Coyle, and O. Boydell, "Exploiting Query Repetition and Regularity in an Adaptive Community-Based Web Search Engine," *User Modeling and User-Adapted Interaction,* vol. 14, no. 5 pp. 383-423, 2004.
[95] J. Gemmell, G. Bell, and R. Lueder, "MyLifeBits: A Personal Database for Everything," *Comm. ACM,* vol. 49, no. 1, pp. 88-95, 2006.

Judy Kay is a professor of computer science at the School of Information Technologies at the University of Sydney. She is a principal in the CHAI: Computer Human Adapted Interaction Research Group, which conducts both fundamental and applied research in personalization and pervasive human computer interaction. She has more than 200 publications in the areas of personalization, teaching, and learning. She has presented invited keynote addresses at major conferences, such as the 1994 User Modeling Conference (UM), the 1995 International Joint Conference on Artificial Intelligence (IJCAI), the 1997 International Conference on Computers in Education (ICCE); the 2000 Intelligent Tutoring Systems Conference (ITS), the 2006 Adaptive Hypermedia and Adaptive Web-Based Systems Conference (AH), and the 2008 Intelligent Tutoring Systems Conference (ITS). She is an associate editor of the *International Journal of Artificial Intelligence in Education (IJAIED)*, an associate editor of the *IEEE Transactions on Learning Technologies (TLT)*, and an editorial board member of *User Modeling and User-Adapted Interaction: The Journal of Personalization Research (UMUAI)*. She is president-elect of the Artificial Intelligence in Education (AIED) Society and on the steering committees for Intelligent Tutoring Systems (ITS) and User Modeling (UM).