



## Learner Control

JUDY KAY

*Department of Computer Science,, University of Sydney, NSW, 2006, Australia*  
*E-mail: judy@cs.usyd.edu.au*

(Received 16 December 1999; in final form 12 June 2000)

**Abstract.** This paper describes major trends in learner-adapted teaching systems towards greater learner control over the learning process. In the early teaching systems, the goal was to build a clever teacher able to communicate knowledge to the individual student. Recent and emerging work focuses on the learner exploring, designing, constructing, making sense and using adaptive systems as tools. Correspondingly, systems are being built to give the learner greater responsibility and control over all aspects of the learning, and especially over the learner model which is at the core of user-adaptation. A parallel trend is the growing acknowledgement of the importance of the learner's social context. Systems are increasingly being designed for learners working in groups of real or simulated peers. This paper discusses several elements of the shift to greater learner control, with a focus on the implications for learner modelling. The computer may offer the learner a choice of learning tools and companion learners, on-demand learning of various types, control over the elements of the systems and the possibility of controlling the amount of control. Learner control offers promising possibilities for improved learning. At the same time, there are pragmatic issues for achieving the benefits. The paper discusses three of these: the need to evaluate the effectiveness of the emergent learner-controlled systems; problems with learner control; and the need for interoperable and reusable components.

**Key words:** learner model, student model, ITS, CSCL

### 1. Introduction

There is a grand vision for computers as support tools for learning: they promise the possibility of affordable, individualised learning environments. This vision makes user modelling central since it is only possible to individualise a system's behaviour if it has individual models of each learner. So it is natural that the model of the student-user has been an important part of research in teaching systems. We outline elements in the classic teaching system architecture. Then we describe an emerging architectural model. We observe that important changes in the emerging work relate to increased learner control over, and responsibility for, the learning process.

#### 1.1. STUDENT MODELS AND EARLY ARCHITECTURES FOR TEACHING SYSTEMS

The early intelligent teaching systems envisaged a student interacting with their individualised, automated teacher. That teacher was modelled after human tutors in that

it was complex, sophisticated and, very likely, inscrutable. An intelligent teaching system was commonly described in terms of a four-part architecture (Wenger, 1987)

- The student interacted with the teaching system via the *interaction module*. For example, this might support natural language interaction or offer sophisticated graphical interfaces.
- As the student interacted with the system, the *student model* was refined, becoming an increasingly accurate model of the student's knowledge and other relevant preferences and attributes. The student model represented static beliefs about the student and in some cases, was able to simulate the student's reasoning, for example (Paiva and Self, 1995).
- *Domain expertise* enabled the teaching system to perform as an expert. The domain expertise module might well reason in fundamentally different ways from the student. For example, it might use discrete simulations where the student would reason qualitatively.
- The *teaching expertise* module determined the system's teaching actions. This would draw upon the current state of the student model. It might use the domain expertise module. It captured expertise on how to teach the individual student as a good human teacher can do.

From the student modelling perspective, it is important to note that the boundaries between these elements may not always be clear. For example, a system might model misconceptions, bugs and mal-knowledge. These might be considered part of the student model. Alternately, they might be seen as perturbations of the domain knowledge. Similarly, the important task of *diagnosis* requires the system to determine the student's misconceptions. This might be considered part of the student model or the teaching expertise. In general, the four elements of the teaching system architecture are closely coupled and the goal of individualised instruction means that the student modelling will typically link to all three other elements.

The last decade has seen a considerable shift from this architecture. This has been described by many authors, including for example, Andriessen and Sandberg (Andriessen, 1999). There was a similar shift in Soloway's research (Sack et al., 1994). Even so, this four part architecture is still helpful. It identifies the major functional elements of learner-adapted teaching systems. We will refer to it in our discussion of the recent, emerging and future directions of systems.

## 1.2. EMERGING ARCHITECTURAL MODELS

Figure 1 depicts one high level view of the emerging class of learner-adapted systems. At the left, it shows one or more learners working in a learning environment. The system provides up to five types of tools. When the learner has a rich collection of such tools available, one important level of learner-control is the choice of whether to use each tool and how much to use it.

The figure shows the learner(s) interacting with one conceptual computer system,

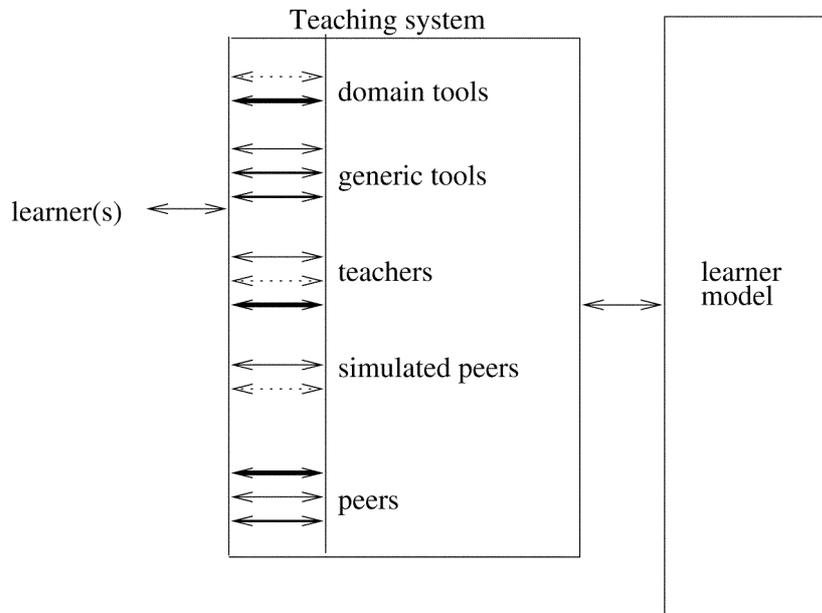


Figure 1. Communities of learners, teachers, peers and tools.

this is the lone bidirectional connection from the learner to the large box depicting the system. This reflects the user's view that the computer system is a single entity.

Within the system, the figure shows many bidirectional links. Each represents the user interacting with one tool or component available at the interface. The figure shows varying thicknesses for these links: these indicate that a particular user interacts with some tools more than others. For example, the dotted top line indicates a domain tool is available at the interface but the learner has not used it. The next bidirectional link is heavy, indicating that the learner makes heavy use of it. The figure indicates that a particular learner will have used different tools to different degrees.

The long box at the right is the learner model. Whereas the early teaching system architecture placed this within a single, tightly coupled collection of system elements, the emerging architecture has this as a separate entity. This single student model is potentially available to many tools. The learner model box is longer than the system box: this is to emphasise that the learner model exists beyond the collection of tools currently available to the learner. The model may be available from a learner model server (Machado, 1999).

We now consider the classes of learning-support tools, their role in user modelling and learner control. At the top are tools which are useful for the learning domain.

For example, these might include a simulation environment which enables the student to explore ideas (Koedinger, 1999) or practice important and time-critical tasks like the management of patients suffering from cardiac arrests (Eliot, 1995). The second group of tools in the figure is the generic tools such as text editors, spreadsheets, calculators and graphing tools. Typically, both these classes of tools are not learner-adapted. However, they have potential as extremely useful sources of user modelling information.

The bottom of Figure 1 shows peers, other students who can communicate with the learner via the system interface. The growing field of CSCL (Computer Supported Cooperative Learning) provides tools which assist learners in co-ordinating their learning activities. For example, a group of students might be working on a large, loosely defined problem. The computer has the potential to provide considerable support for students working in groups: as they divide the task; allocate parts to group members; record their learning achievements; provide assistance to each other; track progress on completed aspects of the problem; reflect on their achievements and difficulties; and manage records and versions of solutions. These tools, like the domain and generic tools, can serve as sources of user modelling information. An important difference is that they involve the learner interacting with peers: a learner may interact with different peers differently. So user modelling information coming from such tools is contextualised in terms of the peer involved, the timing and the task at hand.

The remaining parts of Figure 1 show teachers and simulated peers. We discuss these 'simulated people' as well as the actual peer learners in the following subsections.

### 1.3. FROM STUDENT TO LEARNER, A SHIFT IN LOCUS OF CONTROL

There are many elements in Figure 1. One of the most important, though seemingly subtle, is a shift in terminology. Where the early teaching systems called the user a *student*, the term now favoured is *learner*. This reflects a growing appreciation of the learner's control over the whole learning process. Certainly, the teacher's efforts are irrelevant if the learner pays no attention or they cannot understand the teacher. Constructivists further emphasise the learner's role in actively constructing their own understanding of a learning domain. It seems that we can improve learning effectiveness by giving the learner control over, and responsibility for, their own learning. This is the foundation for teaching strategies such as problem-based learning (Boud, 1991; Curtis, 1995).

There is a general and strong perception that the world is changing quickly. This means that any particular piece of knowledge is likely to be less important than skills for learner-directed life-long learning. So, learners must develop metacognitive skills such as those nurtured by systems like SCI-WISE (White, 1999). These are needed for mastering control activities such as setting personal learning goals, establishing plans for achieving them, assessing progress towards goals and then

revising learning plans. Essentially this involves self-knowledge, planning, reflection and review. There is an important role for adaptive systems to scaffold these activities and to support learning of them. More fundamentally, learner control is of prime importance.

Another example of the shift towards greater learner control is in systems based upon the ZPD, Zone of Proximal Development (Vygotsky, 1978) as, for example, in VIS (Luckin, 1999). The ZPD represents aspects the learner cannot manage alone but which they can do with some assistance, either from a teacher or a peer who is more able. VIS selects learning tasks within the ZPD and provides assistance. This places the teacher in the role of collaborative assistant.

#### 1.4. FROM EXPERTISE TO DOMAIN AND TEACHING BELIEFS

Another terminological shift indicates increasing learner control. Where the classic teaching system architecture dealt with domain and teaching *expertise*, the corresponding parts of newer systems may be referred to as *beliefs*. This reflects a less authoritarian and positivist view, allowing exploration of the possibility of different, even conflicting beliefs.

Part of the reason for this shift is related to the move beyond teaching in domains such as mathematics. This shift does not mean that we equally value the beliefs ' $2 + 2 = 4$ ' and ' $2 + 2 = 5$ '. Rather, consider the case of a domain like law: it is common to model at least two sets of beliefs, associated with two parties in a dispute. In general, there will be at least some different beliefs in these two sets. To teach, or support learning, in law, we may model these different beliefs. Similarly, if we teach history based upon historical sources, it will be common to find that different sources can be used to support different beliefs. Equally, if we teach ethics, we may ask students to reason about a situation from the basis of one of several ethical frameworks. In general, different frameworks will lead to some different beliefs.

So one teacher may have different beliefs from another, and from a learner, and the peer, and so on. Emerging systems, represented by Figure 1, will commonly have different teachers and peers with some different *beliefs* about the domain and how to teach.

## 2. Routes to Learner Control

This section explores several of the elements of learner control in emerging systems. Some of these take the form of the learner's range of choices in learning tools. Another important aspect is the learner's ability to decide when they want to learn. We discuss learner control in terms of access to and control over the detailed beliefs in the learner model as well as the system's domain and teaching beliefs. Finally, we examine the role of mixed-initiative systems (Lester, 1999) which allow the learner the possibility of varying the level of control they wish to exercise.

### 2.1. CONTROLLING CHOICE OF TEACHER, ADVISOR, COACH, . . .

Figure 1 indicates an increasing role for the learner's *community*. The figure shows the possibility of a number of teachers, where the learner can decide which one they will work with. This is akin to variable teaching strategies (Major, 1995) with one important difference. Figure 1 shows the learner controlling the choice of teacher. Where different teachers embody different teaching strategies, this gives the learner control over the choice of teaching strategy. Each teaching tool might need to explain the merits of its teaching strategy. But the Figure 1 model makes the learner responsible for the ultimate choice.

Offering the learner a choice of teachers meshes well with the growing use of teaching agents. For example, work in animated agents (Johnson, 2000) is intended to improve communication by engaging and motivating learners. This acknowledges the learner's control over the learning process, starting with the learner's willingness to pay attention to a teacher or companion at the interface. The agent paradigm may use separate agents to implement the various roles involved in supporting learning. For example, SCI-WISE (White, 1999) assists students in collaborative inquiry and reflective learning. It is implemented as agents to operate in roles such as planner, collaborator and assessor.

### 2.2. CONTROLLING CHOICE OF LEARNING PEERS, COMPANIONS, . . .

The figure also shows simulated peers, often called learning companions. Peers have an important role in collaborating with the learner, creating natural opportunities for the learner to articulate their understanding, reflect on and justify actions. All these activities can improve learning.

There were several ways that simulated peers might support learning: working with the learner as a dyad; learning together from a teacher; the learner teaching the companion (Chan, 1990; Chan, 1996). Such companions can make good use of learner models, as in the case of a companion for Euclidean geometry (Carbanaro, 1995) and a companion which encourages the learner to reflect (Goodman, 1998).

An extension of this approach offered a range of agents, with varying degrees of ability (Hieltala, 1998). Students could elect to work with a peer which worked slowly, made mistakes and was hesitant. Some preferred such a peer to a more competent one; others did not.

From the perspective of user modelling, peer learners call for additional user models, each representing different types of learner. Some of these may be interesting because they require explicit modelling of a particular type of teacher or peer (Self, 2000). For example, one teacher may represent a particular psychological viewpoint. Equally, the peers might be implemented as models of extreme approaches to learning. For example, a rote learning peer would reliably recall facts it had been taught but would not be able to infer new knowledge. The user interacting with this peer would be able to see the implications of that learning preference.

### 2.3. ON-DEMAND LEARNING

On-demand or just-in-time learning is an appealing ideal: just when the learner really needs to know something, they are able to gain the support needed to learn it. The ideal scenario might have the user working with a word processor and realising they need to do a task that they do not know how to do. At that point, they are able to *demand* a teacher (or peer) help them learn about the word processor facilities needed to complete the task.

Learner models are invaluable in this context. First, the system can make use of detailed models of learners in the domain. This may involve extensive empirical work as in the case of Lumiere (Horvitz, 1997) which encapsulates the behaviour of the many users studied. Secondly, the individual learner model can support customisation of the instruction.

One important class of on-demand learning tools supports long term learning in the workplace. For example, the user of a basic tool like a word processor, spreadsheet or database, might use the tool for their work over many years. There is the potential for considerable productivity gains if such users can learn to make more effective use of such tools (Fischer, 2001). Moreover, simple monitoring added to these tools provides a rich source of low grade data for building user models. This approach has been applied to text editors (Linton, 1999; Cook, 1995).

Another important role for on-demand learning arises in conjunction with powerful, exploratory learning environments. For example, simulation tools are an example of the domain tools in Figure 1. These can enable learners to explore ideas that are too costly, dangerous or difficult to try otherwise. While such tools have much potential, they may be more effective when the learner has the opportunity to seek support for their exploration of the learning environment. For example, the Science Learning Spaces (Koedinger, 1999) invites students to define an hypothesis and test it with the simulation and representational tools. In these aspects the learner is in control of the learning. In addition, there are tools which give advice and the student can activate an idea generator.

The long history of help systems fits the model of learning on demand. These may offer natural language interaction, like the early UC system (Wilensky, 1984). Closely related to help systems are the user-adapted online documentation systems (Boyle, 1994; Hook, 1996). These should assist users in learning what they want, more quickly. In terms of learner control, they are only activated by the user at the time they are needed. However, the user may not have control over the adaptation, an issue we consider below.

An interesting variant on a help system is PHelpS (Collins, 1997) which models workers so that it can assist one worker in identifying a peer who can assist them. This is an example of system-facilitated access to actual peers, shown at the bottom of Figure 1.

#### 2.4. CONTROLLING THE LEARNER MODEL

If the learner is expected to take responsibility for their own learning, it seems inconsistent to expect them to tolerate an incomprehensible, inscrutable system which manages their learning. Ultimate control over adaptation requires that the user be able to see aspects of the underlying system control. Since the user model lies at the heart of the individualisation, an open user model is a fundamental part of learner control. We are seeing increasing numbers of systems which give an indication of the current student model as a skillometer (Corbett, 1995; Weber, 1995; Weber, 1999; Koedinger 1999).

The skillometer is a good first step in the direction of learner control. At least, it helps the learner appreciate the current learning goals because they are displayed. It also enables the learner to see how the system models their progress on those goals, relative to standards set by the system. However, the skillometer gives limited access to the learner model. It is tied to a single teaching tool. In this context, it cannot be given much screen space and it is designed to be subordinate to the main teaching activities.

As the learner model becomes an independent entity, separate from any single teaching tool, user control over the model becomes a broader issue. We need tools to support the learner in examining, exploring, correcting and scrutinising the learner model. There are several reasons for making the learner model available.

- Users should have access to and control over personal information.
- Correctness of the model can be assessed by the user. This helps address one of the objections to user-adapted systems on the grounds that they may infer too much from too little information (Henderson, 1995).
- Access to the learner model improves the learner's appreciation of the system's learning goals. This may be useful in dealing with another of the objections to user-adapted systems: it has been argued (Browne, 1990) that if the system and user both attempt to model each other, a situation described as 'hunting' could arise. This means that the user tries to model the system but as they do so, the system changes because it is trying to model the user.
- Programmer accountability is enhanced if the programmer knows the learner will have access to the model.
- There is merit in making complex systems more comprehensible (Fischer, 1991; Fischer, 2000).
- Finally, and importantly, an accessible learner model may aid learning by supporting reflection and planning (Self, 1988; Crawford, 1993; Bull, 1995a; Bull, 1995b; Shute, 1995; Hohl, 1996).

Making the model scrutable (Kay, 1996) is an extreme point in the movement towards access to the learner model. A scrutable learner model enables its owner-learner to see both the model and the processes that contributed to its conclusions about the learner's knowledge and skills.

In building complex learner models, it is common to support the system builder with tools which enable them to scrutinise the model during development (Kobsa, 1995; Paiva, 1995). The challenge is to enable users to really scrutinise their own learner models. The learner models of Figure 1 may be quite large, covering many contexts. If the learner is to have access to, and control of, the learner model, they may need several levels of support.

- *Overview* tools will enable the user to gain a broad-brush understanding of a large model.
- *Detail*-level tools support the user's scrutiny of model.
- *Ontological* support will often be needed. Suppose, for example, that a user explored their learner model for mathematics and found an aspect called 'minuend'. If the user does not know what this means they cannot understand, let alone control, the model. This may be a serious difficulty since learner models often need to represent aspects which are not part of a typical learner's vocabulary.

These tools should also support the user in altering the model. Then, the user could correct it. The user might also explore the impact of arbitrary changes to it. For example, the user might try what-if experiments: alter the learner model to represent a far more knowledgeable user and observe the consequent operation of teachers and peers. In this case, it will be important to support the user in undoing these changes to the model.

Human peers may also have access to the model as well as that of other learners (Bull, 1997a; Bull, 1997b). In the case of users working in a single organisation, OWL (Organisation Wide Learning) approaches can help an individual to choose new learning goals because the system identifies differences between their individual model and models of similar users in the organisation (Linton, 1999).

## 2.5. CONTROLLING DOMAIN AND TEACHING BELIEFS

Shifting from a view of domain *expertise* to domain *beliefs* makes it natural to consider the possibilities for the learner to control those beliefs. One important possibility for this underlies work on *knowledge negotiation* (Moyse, 1992; Baker, 1994) where a system may be willing to negotiate some of its beliefs. This is especially relevant in domains where multiple viewpoints and solutions are usual. It is also appropriate where the teacher may not have complete knowledge.

The last section listed several reasons for enabling user access to the learner model. The last two of these apply in the case of the teacher's and peer's beliefs about the domain and about teaching. It might even be argued that the domain beliefs of the teacher (or simulated peer) are a form of teacher-model (or simulated peer-model). In the cases of simulated peers, learner control of a peer's domain beliefs makes it feasible to explore the implications of such changes.

Because teaching beliefs direct the teaching process, they tend to be very tightly coupled to the other parts of the system (Grandbastien, 1999). They may determine what is needed in the learner model, how that learner model is interpreted and how it affects the teacher's actions. They also interact closely with the domain beliefs: the teacher or peer can only teach things in its beliefs about the domain. It has been argued that the teaching beliefs should be scrutable (Holden, 1999). A primary reason is to give greater learner control. This is a new and little explored area. It is also likely to be very tightly linked to learner modelling issues.

## 2.6. CONTROLLING THE AMOUNT OF CONTROL

Although we have described several ways that the learner can control the teaching system, the learner may not want to maintain control at all times. A helpful hint offered by the system at just the right moment might be greatly appreciated. The Figure 1 architecture may involve teachers and simulated peers which are allowed to take varying levels of initiative. For example, a learner might appreciate a system which takes greater initiative early in the learning and when they have problems (Lester, 1999). Equally, the learning environment might offer a choice of teachers or simulated peers where each exercises a different level of initiative.

Early coaching systems like West (Burton, 1979) emphasised the role of a learner model in assisting a system to decide when to offer advice to the student. The learner model continues to play an important role for managing the way that a system takes the initiative (Haller, 1999).

## 3. Pragmatics

In this section, we briefly review some of the pragmatic issues to be addressed in the near future. These relate to improving our understanding of how best to support learning, how to ensure a desirable level of learner control and the move towards building teaching systems with reusable components that can interoperate.

### 3.1. EVALUATION OF LEARNING EFFECTIVENESS

The emergent model in Figure 1 is motivated by the potential learning benefits. We have much to learn about building these system. Evaluation (Chin, 2001) is critical in assessing whether expected benefits have been achieved. It is also important in tuning approaches and guiding system design to achieve improved learning. A move to rigorous evaluation of the effectiveness of adaptation should be an important future trend.

### 3.2. PROBLEMS WITH LEARNER CONTROL

Given that conventional software is hard enough to build, it may be that the extra demands of learner control will make it too difficult to build practical versions of the tools envisaged in Figure 1. It may also be too limiting if we want to make all parts of the system transparent and scrutable so that the user can control them. Commercial concerns seem to dictate that the implementors devote their energies to building bright, colourful, animated displays which seem attractive and appealing: deeper and subtler issues of learner control may not seem to justify the cost of implementing them.

From the learner's point of view, there are also potential problems. It is likely that greater choice and control will put additional load on the learner: it may even become a distraction from learning! Consider, for example, the considerable size of existing preference files associated with a tool like MS-Word. This may pose serious problems where the learner struggles to understand sophisticated system models.

At quite a different level, learners may not want more control. They may prefer that a teacher takes control. If software support for learner control is never exercised, it is bloatware for that user.

It may seem paradoxical that we have argued the merits of learner control when it is clear that it may introduce potential problems with cognitive overload or may be unimportant to many users much of the time. In fact, there is no paradox. While we have argued the desirability of learner control, we envisage users would only use a small part of their time in scrutinising and taking controlling actions over teaching systems. They may do this as part of the reflection at the end of a learning activity. Since learning is a long term process, learners may only occasionally feel the need to check their learning progress as recorded in the learner model. Indeed, learners may go for very long periods without exerting control. However, when they choose to take control, this should be possible.

Of all the elements involved in providing increased learner control, the learner model is the one where control is essential. This is indicated by its position in Figure 1. It is a persistent object which may be used over long periods of time and across many tools. It holds personal information about the learner. So, in addition to the potential for improved learning outcomes, there are compelling arguments of access to and control over the private information that is the essence of the learner model.

### 3.3. CAN WE TRUST THE USER TO HAVE CONTROL?

If the user has control over their user model, they can fill it with inaccurate information. If they have control over a learning environment, they might sabotage it or simply manage to accidentally reduce its teaching effectiveness. If we ask the user to self-assess their knowledge, they might under- or over-rate themselves. All these are risks of giving learners control.

Future systems will be designed to take account of these possibilities. The case of the user providing inaccurate information can be addressed by ensuring that information is differentiated from other information. For example, the system should be able to determine when a belief is based on information the user volunteered and when it is derived from observing the user doing diagnostic activities like problem solving.

Note that the user always has considerable control, even if it is indirect. For example, in a traditional system like WEST (Burton, 1979) which maintained control, some users purposely stopped playing the main game and began making moves that enabled them to explore the way the system worked. Given the importance of co-operation by the learner if learning is to be achieved, we have no choice but to trust the learner with control.

### 3.4. STANDARDS AND REUSABLE COMPONENTS

Figure 1 highlights the range in types of tools which might be made available to a learner. Although the learning support for a particular domain may involve many of these tools, there may be common elements across them. For example, the various teaching tools and simulated peers will have common subtasks in modelling the domain and teaching beliefs. Equally, there are functional elements which may be common across different domains. Figure 1 emphasises the common learner model: all the learning support tools have the potential to contribute information about the learner; the learner-adapted tools need to access the learner model.

A promising approach to constructing sophisticated learning systems is to build generic tools which can support core activities. Some work in this area involves architectures which permit flexible, interchangeable components (Muhlenbrock, 1998). The central role of the long term learner model needs to be provided by learner modelling shells (Kay, 1995; Paiva, 1995) and servers (Machado, 1999).

Standards should provide an important foundation for interoperability and reuse. This may come from current work (IEEE 1999) to develop standards for various parts of learning tools, especially the relatively solid foundations for basic metadata definitions. The development of standards for learning ontologies and for the learner model will be critical for reuse and interoperability between core tools.

There are also emerging standards for users to manage and control the availability of their user models and associated personal information. For example P3P (Reagle, 1998) supports user control over profile data.

## 4. Conclusions

Research in building intelligent teaching systems has seen a marked shift. Similar shifts appear in user modelling research in HCI (Fischer, 2001) and natural language understanding and generation (Zukerman and Litman, 2001). We have characterised this in terms of the increasing level of learner control in recent and emerging systems.

One might have expected the learner model to become less important for these systems. This is not the case. The learner model continues to have a central role for the continuing work towards individualised learning support tools. The future will see more learning tools in the following broad classes:

- The *domain* and *generic* tools which can serve important roles in the learning environment and, as a side-effect, can contribute evidence about the user, helping to form the beliefs in the learner model;
- *teaching* tools represent a continuation of the history of teaching systems and these teachers can take a variety of roles such as coach, advisor, helper or facilitator;
- Communities of simulated *peers* which may play similar roles to the teachers but will typically have knowledge and skills at a similar level to the learner;
- Interfaces to real peers, be they remote or members of a team working cooperatively.

Learners have at least negative control of any learning situation: they can ‘turn off’ and not pay attention. This paper has described the growing appreciation of the importance of much finer and purposeful learner control at many levels: in choice of teaching tool, simulated and actual peers, in the timing of the learning, in accessing and controlling the learner model as well as parts of the system’s domain and teaching beliefs and in deciding how much control to allow the system in different situations.

The shift from early teaching systems reflects trends in education and computer-based support for learning. Part of this has been due to the growth in our knowledge of how to build systems to support learning. The early systems have provided foundations which make it possible to move beyond tidy domains like mathematics. We now want to support learning in domains where there may be many ways to look at issues, multiple viewpoints, several answers to a problem and a need for less absolutist and positivist approaches to teaching.

There has also been an increasing appreciation of the need to support life-long learning. This is partly due to the view that the world is changing so quickly that many facts learnt during formal education may be superseded by the time students complete that formal education. This means that any set of domain knowledge is likely to be far less valuable than the development of the metacognitive skills which will enable the learner to manage their own long-term learning. We can envisage learner models that span decades of learning development. Learners will need to be able to access and control them. These learner models might be used as a foundation for self-assessment, planning and control of continuing learning.

A similar trend is the increasing role of *flexible* learning. This is yet another dimension of increasing learner control. It creates a need for user-adapted and user-controlled learning support so that the learner can learn ‘at their place, at their pace’.

This paper has described a very rich model of learning support systems, but one

that may be composed of many, small tools which can themselves be built from reusable and interoperable components. We have also noted that the learner model provides the common foundation for adaptation to the individual learner and so, it has a central role to play in supporting learner control.

### Acknowledgements

My thanks to Peter Brusilovsky, Susan Bull and Gerhard Fischer who read the draft and provided insightful comments and advice.

### References

- Andriessen, J. and Sandberg, J.: 1999, Where is education heading and how about AI?, *International Journal of Artificial Intelligence in Education*, **10**, 130–150.
- Baker, M.: 1994, A model for negotiation in teaching–learning dialogues, *International Journal of Artificial Intelligence in Education*, **5**, 2, 199–254.
- Boud, D. and Feletti, G.: 1991, *The challenge of problem based learning*, Kogan Page, London.
- Boyle, C. and Encarnacion, A.O.: 1994, MetaDoc: an adaptive hypertext reading system, *User Models and User Adapted Interaction*, **4**(1), 1–20.
- Browne, D., Totterdell, P. and Norman, M.: 1990, *Adaptive User Interfaces*, Academic Press, San Diego, California.
- Bull, S. and Pain, H.: 1995a, Did I say what I think I said, and do you agree with me?: inspecting and questioning the Student Model, *Proceedings of World Conference on Artificial Intelligence in Education*, AACE, Washington DC, USA, pp. 501–508.
- Bull, S., Brna, P. and Pain, H.: 1995b, Extending the scope of the student model, *User Modeling and User-Adapted Interaction*, **5**(1), 44–65.
- Bull, S. and Brna, P.: 1997a, What does Susan know that Paul doesn't? (and vice-versa): contributing to each other's student model, *International Conference on Artificial Intelligence in Education*, IOS Press, pp. 568–570.
- Bull, S.: 1997b, A multiple student and user modelling system for peer interaction, *Proceedings of ABIS97*, Universität Saarbrücken, pp. 61–71.
- Burton, R.R. and Brown, J.S.: 1979, An investigation of computer coaching for informal learning activities, *International Journal of Man-Machine Studies*, **11**, 5–24.
- Carbonaro, A., Maniezzo, V., Roccietti, M. and Aslomoni, P.: 1995, Modeling the student in Pitagora 2.0, *User Modeling and User-Adapted Interaction*, **4**(4), 233–251.
- Chan, T.W. and Baskin, A.B.: 1990, Learning Companion Systems, *Intelligent Tutoring systems: at the crossroads of Artificial Intelligence and Education*, Ablex.
- Chan, T.W.: 1996, Learning companion systems, social learning systems, and the global social learning club, *International Journal of Artificial Intelligence in Education*, **7**, 125–159.
- Chin, D.: 2001, Empirical evaluation of user models and user-adapted systems, *User Modeling and User-Adapted Interaction* **11**(1–2), 181–194 (this issue).
- Collins, J.A., Greer, J.E., Kumar, V.S. McCalla, G.I., Meagher, P. and Tkatch, R.: 1997, Inspectable user models for just-in-time workplace training, In: A. Jameson, C. Paris and C. Tasso (eds.), *User Modeling Proceedings of the Sixth International Conference UM97*, pp. 327–337, Springer, Chia Laguna, Sardinia, Italy.
- Cook, R., Kay, J., Ryan, G. and Thomas, R.C.: 1995, A toolkit for appraising the long term usability of a text editor, *Software Quality Journal*, Chapman and Hall, London, **4**, 131–154.

- Corbett, A.T. and Anderson, J.: 1995, Knowledge tracing: modeling the acquisition of procedural knowledge, *User Modeling and User-Adapted Interaction*, **4**, 253–278.
- Crawford, K. and Kay, J.: 1993, Metacognitive processes and learning with intelligent educational systems, In: P. Slezak, T. Caelli and R. Clark (eds.), *Perspectives on Cognitive Science*, Ablex, pp. 63–77.
- Curtis, C.A., Carver, A., Howard, R.A. and Lane, W.D.: 1995, A methodology for active, student controlled learning: motivating our weakest students, In: *Proceedings of the 27th SIGCSE Technical Symposium on Computer Science Education*, Philadelphia, Pennsylvania.
- Eliot, C. and Woolf, B.P.: 1995, An adaptive student centred curriculum for an intelligent training system, *User Modeling and User-Adapted Interaction*, **5**(1), 67–86.
- Fischer, G.: 2001, User modeling in Human-Computer Interaction, *User Modeling and User-Adapted Interaction* **11**(1–2), 65–86 (this issue).
- Fischer, G.: 1991, The importance of models in making complex systems comprehensible, In: D. Ackerman and M. Tauber (eds.), *Mental Models and Human Computer Communication: Proceedings of the 8th Interdisciplinary Workshop on Informatics and Psychology*, Schaeding, Austria, Elsevier Science, Amsterdam, pp. 3–36.
- Goodman, B., Soller, A., Linton, F. and Gaimari, R.: 1998, Encouraging student reflection and articulation using a learning companion, *International Journal of Artificial Intelligence in Education*, **9**, 237–255.
- Grandbastien, M.: 1999, Teaching expertise is at the core of ITS research, *International Journal of Artificial Intelligence in Education*, **10**, 335–349.
- Haller, S. and McRoy, S. and Kobsa, A.: 1999, *Computational Models of Mixed-Initiative Interaction*. Dordrecht, Netherlands: Kluwer Academic Publishers.
- Henderson, A. and Kyng, M.: 1995, From customisable systems to intelligent agents, In: R. M. Baeker, J. Grudin, W. Buxton and S. Greenberg (eds), *Readings in Human-Computer Interaction*, pp. 783–792, Morgan Kaufmann.
- Hietala, P. and Niemirepo, T.: 1998, The Competence of Learning Companion Agents, *International Journal of Artificial Intelligence in Education*, **9**, 178–192.
- Hohl, H., Böcker, H-D. and Gunzenhäuser, R.: 1996, Hypadapter: An Adaptive Hypertext System for Exploratory Learning and Programming, *User Modeling and User-Adapted Interaction*, **6**, 12–3.
- Holden, S and Kay, J.: 1999 The scrutable user model and beyond, In: R. Morales (ed.), *AIED, Artificial Intelligence and Education, 99 Workshop W7: Open, interactive, and other overt approaches to learner modelling*, pp. 51–62, <http://www.dai.ed.ac.uk/groups/aiied/Conferences/Ovalm/Papers>
- Hook, K., Karlgren, J. Waern, A., Dahlbeck, N., Jansson, C.G. and Lemaire, B.: 1996, A glass box approach to adaptive hypermedia, *User Modeling and User-Adapted Interaction*, **6**(2–3), 157–184.
- Horvitz, E.: 1997, Agents with beliefs: reflections on bayesian methods for user modeling, *User Modeling, Proceedings of the Sixth International Conference UM97*, Springer-Verlag, New York, pp. 441–442.
- LTSC IEEE: 1999, IEEE Learning Technology Standards Committee (LTSC), In: <http://ltsc.ieee.org/ltsc/>, IEEE.
- Johnson, W.L., Rickel, J.W. and Lester, J.C.: 2000, Animated Pedagogical Agents: Face-to-Face Interaction in Interactive Learning Environments, *International Journal of Artificial Intelligence in Education* **11**, 47–78.
- Kay, J.: 1995, The um toolkit for cooperative user modelling, *User Modeling and User-Adapted Interaction*, **4**(3), 149–196.

- Kobsa, A. and Pohl, W.: 1995, The user modeling shell system BGP-MS, *User Modeling and User-Adapted Interaction*, **4**(2), 59–106.
- Koedinger, K.R., Suthers, D.D. and Forbus, K.D.: 1999, Component-based construction of a science learning space, *International Journal of Artificial Intelligence in Education*, **10**, 292–313.
- Lester, J.C., Stone, B.A. and Stelling, G.D.: 1999, Lifelike pedagogic agents in mixed-initiative problem solving in constructivist learning environments, *User Modeling and User-Adapted Interaction*, **9**(1–2), 1–144.
- Linton, F. Joy, D. and Schaefer, H-P.: 1999, Building user and expert models by long-term observation of application usage, In: J. Kay (ed.), *User Modeling: Proceedings of the Seventh International Conference, UM99*, Springer Wien, pp. 129–138.
- Luckin, R. and du Boulay, B.: Capability, potential and collaborative assistance, In: J. Kay (ed.), *User Modeling: Proceedings of the Seventh International Conference, UM99*, Springer Wien, New York, pp. 139–148.
- Machado, I., Martins, A. and Paiva, A.: 1999, One for all and all for one: a learner modelling server in a multi-agent platform, In: J. Kay (ed.), *User Modeling: Proceedings of the Seventh International Conference, UM99*, Springer Wien, New York.
- Major, N.: 1995, Modelling teaching strategies, *International Journal of Artificial Intelligence in Education*, **6**(2/3), 117–152.
- Moyses, R. and Elsom-Cook, M.: In: *Knowledge Negotiation*, Academic Press, London.
- Muhlenbrock, M., Tewissen, F. and Hoppe, H.U.: A framework system for intelligent support in open distributed learning environments, *International Journal of Artificial Intelligence in Education*, **9**, 256–274.
- Paiva, A. and Self, J.: 1995, TAGUS – a user and learner modeling workbench, *User Modeling and User-Adapted Interaction*, **4**(3), 197–228.
- Reagle, J. and Cranor, L.: 1998, W3C The Platform for Privacy Preferences, In: <http://www.w3.org/TR/NOTE-P3P-CACM/>
- Sack, W., Soloway, E. and Weingard, P.: 1994, Re-writing cartesian student models, In: J. Greer and G. McCalla (eds.), *Student Modelling; The Key to Individualized Knowledge-Based Instruction*, Springer-Verlag, pp. 335–376.
- Self, J.: 2000, The defining characteristics of intelligent tutoring systems research: ITSs care, precisely, *International Journal of Artificial Intelligence in Education* **10**, 350–364.
- Self, J.: 1988, Bypassing the Intractable Problem of Student Modelling: Invited paper, In: *Proceedings of the 1st International Conference on Intelligent Tutoring Systems*, Montreal, pp. 18–24.
- Shute, V.: 1995, Smart: Student Modeling Approach for Responsive Tutoring, *User Modeling and User-Adapted Interaction*, **5**(1), 1–44.
- Vygotsky, L.S.: 1978, *Mind in Society: The Development of Higher Psychological Processes*, Harvard University Press, U.S.
- Weber, G. and Mollenberg, A.: 1995, ELM programming environment: a tutoring system for LISP beginners, In: K. F. Wender and H.-D. Bocker (eds.), Ablex, Norwood, pp. 373–408.
- Weber, G.: 1999, Adaptive learning systems in the World Wide Web, In: J. Kay (ed.), *Proceedings of UM99: Seventh International Conference on User Modeling*, pp. 371–378.
- Wenger, E.: 1987, *Artificial Intelligence and Tutoring Systems – Computational and Cognitive Approaches to the Communication of Knowledge*, Morgan Kaufmann, Los Altos.
- White, B.Y., Shimoda, T.A. and Frederiksen, J.R.: 1999, Enabling students to construct theories of collaborative inquiry and reflective learning: Computer support for metacognitive development, *International Journal of Artificial Intelligence in Education*, **10**, 151–182.

- Wilensky, R., Arens, Y. and Chin, D.: 1984, Talking to UNIX in English: an overview of UC, *Communications of the ACM*, **27**(6), 574–593.
- Zukerman, I. and Litman, D.: 2001, Natural language processing and user modeling: synergies and limitations, *User Modeling and User-Adapted Interaction* **11**(1–2), 129–158 (this issue).

### **Authors' Vita**

**Judy Kay** is a senior lecturer in the department of computer science at Sydney University and co-leader of the User-adapted Systems Research Group. Her primary interest is in architectures and components for scrutable user modelling including representations for user models, interfaces to support user access to user models, and building teaching systems which make use of scrutable student models. Previous research has included work in user-based fairshare scheduling. She has authored over eighty publications, and was recently the Programme Chair for UM'99 User Modeling Conference.