

Preface to Special Issue on Student Modeling

Student modeling researchers constitute one of the oldest and best established subgroups in the user modeling community. A student model is central to the goal of individualised instruction; student modeling has therefore been a focus of activity in intelligent tutoring systems research from its earliest days.

For this journal, which serves the broader user modeling community, it is not at all clear how a student might be fundamentally different from other users. Both have knowledge, beliefs, preferences, goals and plans. Both in student modeling and more general user modeling, these aspects of the user may be significant for effective adaptation of the interface.

The defining aspect of student modeling is that it supports the student user's learning. Let us consider what this means, starting with one definition of learning:

'learning is the transition from *not being able to* to *being able to* as a result of doing something, or as a result of something happening' (Marton and Booth, 1995).

Since this is the goal of a teaching (or learning) system, it suggests that the following should be modeled:

- what the student is able to do at various times
- whatever else is relevant for generating the opportunity for students to participate in an experience that will enable them to learn

The first part of this, determining what the student can do is frequently described as I diagnosis, since it involves developing a model of the learner's knowledge. This is one of the critical tasks of student modeling.

There are two main approaches for representing the student's knowledge. Those that are especially associated with student modeling are cognitively based, relying on theories of human learning and reasoning. Other approaches use pragmatic models that are commonly built from empirical studies. The papers in this special issue reflect both of these.

Commonly the student model goes beyond a single expert performance model. It is the norm that students do not know things that experts in the domain know: it is further the norm that learners know (or believe) things that experts in the domain would not agree with. Most of the papers in this collection operate in terms of a single, expert model. For others, the student model is cast in terms of broader frameworks of the domain. In this collection, for example, this is important for the work reported by Bull, Brna and Pain, where the system models adults learning a new language and the learning is powerfully influenced by the other languages that the students know.

Even where student modeling uses a pragmatic rather than cognitive approach to student modeling, it is characterised by empirical foundations and evaluations for the student model. This means that observations of students are important for all approaches of constructing the model.

Although the student's knowledge is an important part of the student model, it is also common to model some of the other aspects that are relevant for predicting what the student will be able to do after a learning experience that the system might offer. For example, the student's learning preferences, aspirations, current goals and tasks as well as ability may all be important. The papers in this special issue explore various of these.

A relatively recent trend in intelligent teaching systems research is the move to encouraging students to reflect on their own learning. The student model itself can help here and we see this in some of the papers in this issue where the student model is made available to the learner.

The papers in this special issue provide a broad range of views on how to tackle the various problems of student modeling. The articles by Carbonaro, Maniezzo, Rocchetti and Salomoni; Corbett and Anderson; Kashihara, Hirashima and Toyoda will be published in this issue, number 4:4. The articles written by Shute; Bull, Brna and Pain; Eliot and Woolf will be published in issue 5:1. The article by Ragnemalm will be published in issue 5:2.

Carbonaro, Maniezzo, Rocchetti and Salomoni describe the student model in a system that teaches Euclidean geometry. Their system has an expert model for solving problems in Euclidean geometry and a model for the individual student. In addition, in line with a recent and very interesting research direction in student modeling, it also models a companion, a person at a similar level to the student. The system applies a Bayesian model to deal with the problems of noise and uncertainty that are common to broader user modeling.

Corbett and Anderson describe work based on executable cognitive models represented as productions rules. This continues their long term work on mastery learning of programming. The student model uses knowledge tracing, modeling the development of the student's knowledge through the period of the teaching. This enables the system to predict student performance and thereby to individualise the exercises offered.

Kashihara, Hirashima and Toyoda describe a system that individualises explanations to have the 'optimal cognitive load' for the student. Users who can cope with a higher cognitive load are given more substantial and demanding forms of information. The aim is that they learn more efficiently and remain motivated. This system also operates in an interesting and novel way: after giving an explanation

to the student it asks the student to explain how this new information fits into their existing knowledge structure. This supports the system's diagnosis tasks and encourages reflection.

Shute has taken an effective and conventional Computer Aided Instruction system for teaching statistics and added student modeling. This permits a serious assessment of the benefits of adding user modeling. This system is rather unusual in that it applies adaptation at the macro level. This is based on the student's cognitive ability and allows different teaching strategies to be used.

The system also features micro-adaptation, common to most systems, with student knowledge and performance driving the teaching route. This system also explores the use of a visible student model.

Bull, Brna and *Pain* studied actual development of students learning Portuguese, and used the results to develop a model of typical learning sequences. This defines the vocabulary of the modeling process: the essential knowledge and misconceptions. They also modeled the effect of other languages that the student knows, and how this affects transfer. Beyond this, they explicitly model the learning strategies their students employ. They make the student model accessible to the learner to encourage reflection and the development of language awareness.

Eliot and *Woolf* use a simulation-based teaching system to help medical students learn about management of patients suffering cardiac arrests. This system models the importance and difficulty of each topic, as well as the student's apparent comprehension of each aspect. In this system, it is important that students follow established protocols and the system matches student actions against the protocol to generate the simulation. The system alters the actual probabilities of events to ensure students get practice where they need it.

Ragnemalm rounds out this collection of papers by presenting a framework for looking at the various techniques for student modeling. This defines the tasks of student modeling and compares various systems. The paper offers one integrated view of the bewildering range of student modeling approaches and the differing emphasis in the many teaching systems.

It may seem that there are large differences between the student modeling community and the other user modeling communities. Indeed, the role of the user as a learner does ensure some special research issues for student modeling.

However, there is much in common between the problems faced in student modeling work and that in other areas of user modeling. At the most fundamental level, all user modeling has to support two essential tasks: the system's interpretation of the user's actions; and its generation of effective actions.

A primary goal of this journal is to bring together the various strands of user modeling work. This is important because so much of the work that should be shared is contained in largely disjoint literature. One important example of this is the body of research into student modeling from the field of intelligent tutoring systems. This special issue makes this work more accessible to the broader the user modeling community.

References

Marton and Booth: 1995 (to appear). In D. Olson and H. Torrence (eds.): *Handbook of Learning and Human Development*. Oxford: Blackwell.