Helping Learners to Understand Themselves with a Learner Model Open to Students, Peers and Instructors

Susan Bull & Theson Nghiem

Electronic, Electrical and Computer Engineering,
University of Birmingham, Edgbaston, Birmingham, B15 2TT, U.K.
s.bull@bham.ac.uk

Abstract

This paper introduces work in progress on an open student model designed to help learners to better understand their learning. The aim of the system is to investigate issues relevant to open learner models in a large-scale, real use context. This will be achieved initially through the deployment of a simple, domain-independent system. The student model is inspectable by the student it represents, to help focus reflection on their learning. Students may also view the models of their peers, to enable them to compare their own progress against that of others. Furthermore, instructors may view the learner models of their students, to help them support students in their learning. The initial version of the system and the learner model are very simple, to enable early deployment in a variety of contexts. (Subsequent investigations may lead to more detailed recommendations for specific domains.) Planned extensions to enable student-system and peer-peer collaboration with reference to student models (based on existing work on other systems), are also discussed.

Key words: open learner models, learner reflection.

1. Introduction

There has recently been growing interest in opening the learner model to the individual it represents, with several systems demonstrating this approach (e.g. Bull & Pain, 1995; de Buen et al., 1999; Dimitrova et al., 2001; Kay, 1997; Morales et al., 2000; Silva et al., 2001; Specht et al., 1997; Weber et al., 2000). An important reason for rendering the learner model accessible is to help the student to better understand their learning - opening the learner model to the modellee offers a source of information about their relationship with the target domain which is otherwise unavailable, encouraging them to reflect on their beliefs and on the learning process.

Student reflection on their learning has also been encouraged in collaborative learning situations, where pairs of students view their own and each other's learner models, in order to provide a focus for collaborative interaction (Bull & Broady, 1997). Making peer models available more generally could also be beneficial as it would enable students to compare their progress against that of their peers, as proposed by Kay (1997).

In addition, suggestions have been made to open learner models to instructors, allowing tutors to access information about those they teach. This can help instructors to adapt their teaching to the individual or to the group (Grigoriadou et al., 2001; Zapata-Rivera & Greer, 2001); or enable them to suggest suitable peer helpers, or organise learning groups (Mühlenbrock et al., 1998).

Each of the above approaches to open learner modelling (self-access, peer access and instructor access to an individual's student model) can be used to help learners to better understand their learning of a target domain. However, one of the difficulties in effecting learner reflection in this manner on a wider scale, is that typically systems which have inspectable student models as a *focus* for encouraging reflection are quite complex, with the system requiring an understanding of the target domain in order to model in some detail, a student's knowledge,

P. Brna & V. Dimitrova (eds), Proceedings of Workshop on Individual and Group Modelling Methods that Help Learners Understand Themselves, International Conference on Intelligent Tutoring Systems 2002, pp. 5-13.

and infer their misconceptions, in that area. While this results in potentially very effective methods of promoting learner reflection, it also renders the systems relatively expensive to implement, and often restricted to a single or limited set of domains.

To complement such approaches, we suggest employing a very simple student model in a system that can be easily deployed in a variety of course types, with the aim of investigating the potential of open learner modelling in a range of realistic settings, with large numbers of users. While this more straightforward approach to student modelling will not allow the system to adapt its tutoring to specific misconceptions held by specific individuals, it will nevertheless allow investigations of broader questions such as whether students will pay attention to their learner models in a variety of disciplines. Some investigations have suggested that even when they know about the inspectability of their own student model, learners do not necessarily attempt to view it (Barnard & Sandberg, 1996; Kay, 1995). Nevertheless, small-scale studies have suggested that students might indeed use an open student model if it were available, in contexts where they may discuss the contents of their model with the system (Bull & Pain, 1995; Dimitrova et al., 2001). It is worthwhile, therefore, investigating whether the disinterested reactions from some students in the former situations are typical, or whether, as suggested in some other studies, students might find open learner models helpful to their learning at least in some contexts. Furthermore, it would be useful to discover whether instructors in a range of subjects find open learner models a useful way of helping them to recognise student difficulties, enabling them to respond to specific student populations or individuals in appropriate ways. With this as one of the system aims, the limitations in its ability to offer fine-grained adaptive interaction to an individual student, based on the student model, are less crucial. Hence the system is intended more as a practice environment, than a tutoring system.

Another issue that can be investigated with simple open learner models is the means of externalising these to students, peers and instructors. As argued by Morales et al. (2000), presentation modality might be important for the comprehensibility of student models at least in some domains. Similarly, different individuals might work better with different model representations, thus even within a single domain it may be useful to offer alternative or complimentary learner model representations (e.g. Bull, 1998). Furthermore, Zapata-Rivera & Greer (2001) add that the goals of the system may influence the manner in which student model data should be presented.

This paper describes work in progress on a learner model open to the student it represents, their peers and their instructor. A student inspecting their own student model can benefit through reflection on its contents; viewing peer models can enable learners to understand their progress against the context of the advancement of their peers; and allowing teachers to view individuals' learner models can help them to help those specific learners individually, and also adapt their teaching to the group's difficulties, where common problems arise.

The learner model presented here is very simple - in its basic version containing representations only of a student's level of knowledge of various topics. While necessarily restricting the system's ability to adapt to an individual's misconceptions, it does allow it to be straightforwardly deployed more widely (as a practice environment) and, as stated above, the involvement of the teacher (or peers) can compensate for the system's inability to interact with reference to specific misconceptions. Moreover, multiple choice exercises can be created in text files, further allowing the system to be easily deployed in a variety of subjects, by a variety of tutors. The simplicity of both the student model and the method of exercise creation permits investigation of the potential of open student models to help learners to better understand themselves, enabling some of the general questions pertinent to inspectable student models to be investigated in a variety of contexts.

Section 2 introduces the system. Section 3 describes the inspectable learner models from three perspectives: inspection by the student (modellee); inspection of the models of peers; and inspection of learner models by the instructor. In each case, the potential for helping learners to better understand themselves, is discussed. Conclusions are presented in Section 4.

2. The General Learning Environment

The system is suitable for use in domains where multiple choice questions are appropriate as a means of practising the target material. The system presents ten questions at a time, each with a drop-down box where the correct answer should be selected. Figure 1 illustrates the interface for the domain of Japanese particles, where the grammatically correct particle has to be chosen for each of the sentences.

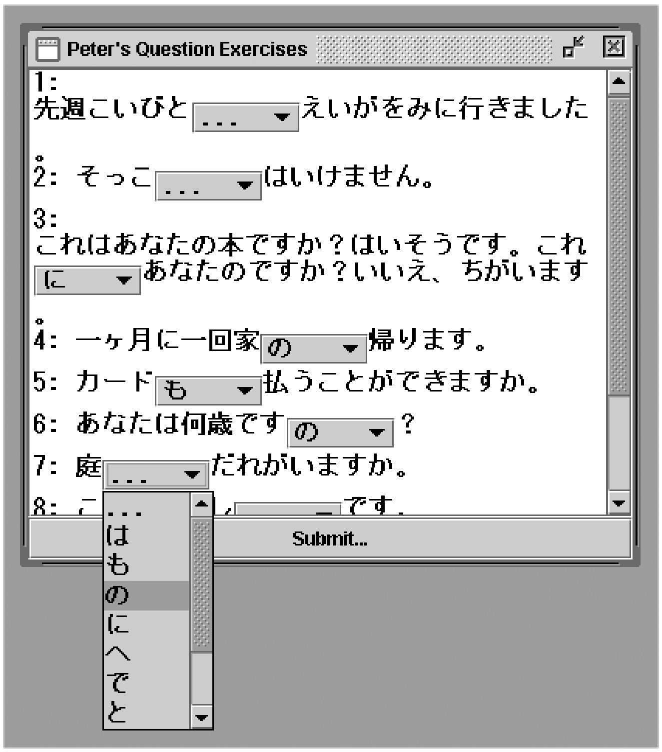

Figure 1: Multiple choice practice questions

Tutors can advise students to use the system to practise material presented during classroom sessions, using textbooks or course notes for reference. Alternatively, instructors can add their own instructional materials, to be accessed from within the system.

Questions for each concept (in this example, Japanese particles) are selected randomly. Candidate concepts are selected according to the student's previous performance. The aim is to provide students with opportunities to practise the areas with which they are experiencing most difficulty, thus each individual will receive exercises targeted in particular at a restricted set of concepts which are represented as least understood in their student model. (During the initial interaction, questions on all concepts are presented.) Standard questions can be loaded for specific domains. New questions can also be added by instructors, thus for particularly complex concepts many practice questions can be made available. These can be divided into various levels of difficulty, to enable the same instantiation of the system to be used by students at a range of levels. It also allows for the same concepts to be practised in different contexts within the same general domain - for example, Japanese particles in business or conversational Japanese. Moreover, it allows the system to be deployed in a range of subjects - not only different languages, but any domain for which appropriate multiple choice practice questions can be created. New questions are simply added to existing text files, or for new concepts (in the same or a new domain), new text files can be created. Thus the system can be easily deployed in a variety of courses, to investigate the research questions.

3. The Open Learner Model

The basic student model is a very simple, numerical model. The total number of questions attempted for each concept is stored, as is the proportion of correct versus incorrect attempts. Greater weighting should be awarded to later answers to enable the system to use this information more accurately to infer where the student's current difficulties lie. (The weighting algorithm for the student model has not yet been implemented.) As described below, these student models can be extended in particular cases, but for the simplest version of the system, the student model is a straightforward representation of knowledge levels in the various areas.

The learner model is open to the individual student, to encourage them to reflect on their knowledge and misconceptions. Students can also access the models of peers, to help them further gauge their progress. Instructors can view the various learner models to enable them to better help individuals; to adapt their teaching for a particular group; or to help them set up optimal peer learning or tutoring groups.

3i. A Learner Model Open to the Student

As described above, the student model contains representations of the student's performance on sets of questions related to different concepts in some domain. Remaining with the example of Japanese particles, Figure 2 illustrates the way in which the model is externalised to the student.

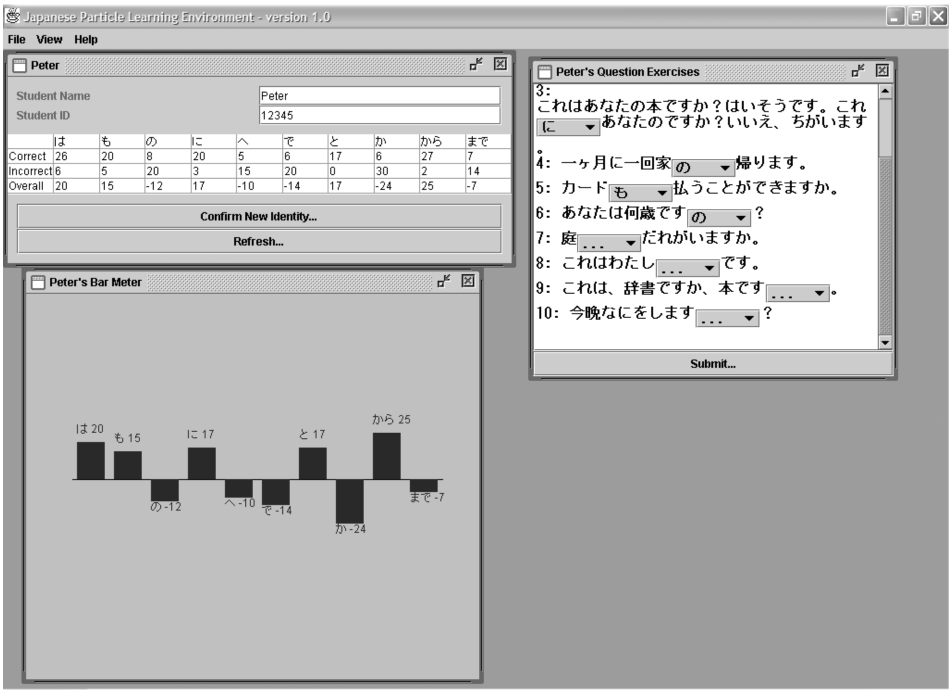

Figure 2: The open learner model

Currently the student model can be presented to the learner in one or both of two forms: tabular and graphical. The graphical representation is similar to the skill or knowledge meters of APT (Corbett & Bhatnagar, 1997), ADI (Specht et al., 1997) and ELM-ART (ELM Research Group, 1998), which display bars where the filled portions represent the student's attained knowledge or skill against the expert level; and OWL (Linton & Schaefer, 2000), which displays in bar form a user's knowledge level against the combined knowledge of other user groups. In contrast to the above, our system illustrates a comparison of correct versus incorrect attempts at questions. While, as pointed out by Linton and Schaefer, the skill meter approach focuses on achievements rather than shortcomings, our system aims to raise learner awareness of their performance in general - i.e. not only the degree to which they have mastered any particular aspect of the domain, but also concepts with which they are having difficulty (in contrast to areas they have not yet attempted).

Other external representations for the student model will be considered at a later date. Nevertheless, as it is knowledge level that is represented, rather than specific concepts and misconceptions, these representations will likewise not be complex.

At present, a score of 1 is awarded for each correct answer for a given concept, and 1 is subtracted for an incorrect response. This can be seen in the tabular representation in Figure 2. (Representations are updated when students complete a set of ten questions.) The overall score, illustrated in both the tabular and the graphical representation of the student model, indicates the student's overall performance. (As stated above, implementation of the mechanism to award greater weighting to later answers is still required. Once implemented, the student model will better reflect the student's actual current knowledge levels.)

The student can use the externalised learner model to see at a glance areas in which they are strongest and weakest, and use these as recommendations of what should be studied further, within or outside the system. If instructors choose to add domain content or explanations of common misconceptions to a particular instantiation of the system, these can be linked to the student model, and examined by the student.

It can be seen that the information shown to students about themselves, is limited. An extension to the system, allowing the learner to negotiate the contents of their student model, is planned. This is based on the MR COLLINS approach of *collaborative student modelling* (Bull & Pain, 1995), where the student can argue with the system if they disagree with the contents of their learner model. Negotiating the content of the learner model is designed in part to enhance learner awareness of their learning. In the case of the new system, such negotiation will be less complex, focusing on the system offering the student a quick test if they claim to know something that they have not yet adequately demonstrated (with the system accepting the student's claim if they can demonstrate their knowledge in the test); and allowing the learner to inform the system if they believe its representations of their knowledge level of any concept is too high (for example, if they had guessed the answer to some of the multiple choice questions, and by chance got them correct; or if they have simply forgotten previously known material). This will enable the student to influence their learner model, and thereby also influence the subsequent selection of practice material.

3ii. A Learner Model Open to Peers

In addition to inspecting their own learner model, a student can compare their performance to that of peers in their group (this occurs anonymously unless students choose to reveal their usernames to others). This enables learners to appreciate how their developing understanding of the target concepts compares to that of other learners, as suggested by Kay (1997). Students can retrieve the student models of good peers, to appreciate the level of knowledge for which they could aim; and also of weaker peers, which in some cases could help learners to realise that they are performing better than they had realised. Figures 3 and 4 show a comparison of the student models of different individuals.

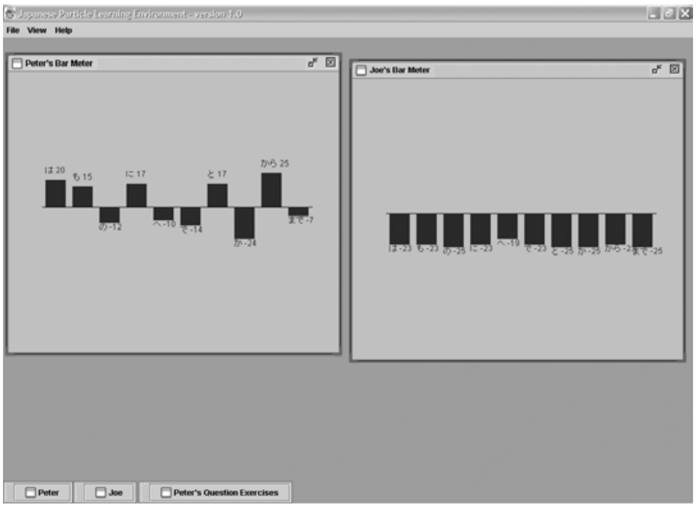

Figure 3: Comparing student models (graphical)

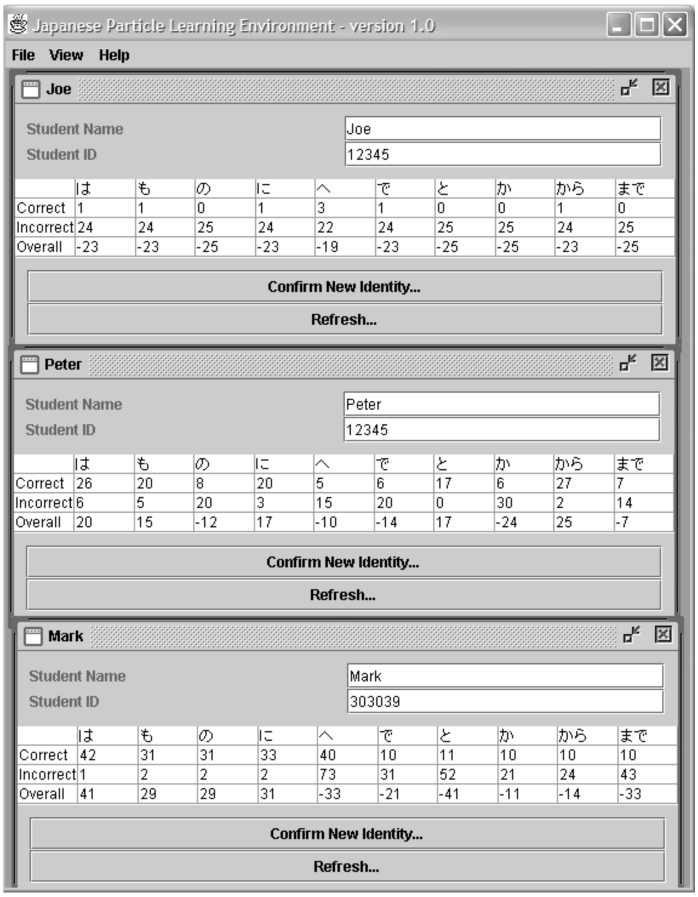

Figure 4: Comparing student models (tabular)

A planned extension to the current system will also allow learners to compare their student model against the 'average' student model of all other learners (or a subset), in their cohort.

The possibility of viewing peer models, in addition to enabling students to gauge their progress against other learners, also allows a student to collaborate with another learner, using their respective student models as a focus for discussion. In 2SM (Bull & Broady, 1997), displaying two student models to their co-present owners was designed to prompt peer tutoring, with each student individually completing an exercise, and then coming together to repeat the same exercise in the presence of their two inspectable learner models. In the new system, a pair (or small group) of learners can also come together, viewing their own and their partner's student models, and learning from collaborative interaction. In contrast to 2SM, each learner will have experienced a different set of questions tailored to their own needs, and the combined exercise will again be different. The pair will commence a fresh session, building a new student model reflecting their joint performance, drawing on the representations in their individual learner models to recognise whose explanation is more likely to be correct, in cases where they disagree about an answer when collaborating about their joint answers. (The system provides correct answers for student consultation after an exercise has been completed. Thus, should the pair agree on an incorrect answer, they are made aware of any inaccurate selection, and can see what the correct choice should have been.) The collaboration and peer tutoring expected to occur with the comparison of learner models (see Bull & Broady, 1997), will focus students' attention more directly on their knowledge and misconceptions. An interesting observation will be how students decide on the representation type to view, if their preferences differ. A planned extension to the current version of the system is to allow it to use the two student models of a collaborating pair of learners, to adapt the joint exercise to best utilise the relative strengths and weaknesses of the pair.

With the extensive sets of questions possible, students can potentially work with a variety of peers on a set of concepts if they prefer to learn with others. They can use the student models of others to help them select suitable collaborative partners or helpers who complement their own abilities. An extension to the system could even enable it to suggest suitable learning partners based on the contents of the various participants' learner models (see e.g. Collins et al., 1997).

3iii. A Learner Model Open to the Tutor

The final use for the open student model is to help tutors to better understand their students. In the same way that students can compare their learner model to the models of others, teachers can do likewise with their students' learner models. This can allow instructors to help individuals with particular problems, or enable them to better target their teaching to the general difficulties of specific groups. The possibility of viewing a combined 'average' student model (see 3i) would facilitate the latter process. This is almost the reverse of the original aim of intelligent tutoring systems - to teach students as a human tutor might, using knowledge about a student contained in their student model. Here the student model is providing information for the human teacher, helping them to adapt their teaching appropriately.

Although not part of the initial planned investigations, teachers could also view models in order to form learning groups (as in Mühlenbrock et al., 1998). Teachers might ultimately be able to provide some additional information to learner models based on, for example, student performance on assignments, and engage in system-mediated discussion with a student, about their learner model (see Bull, 1997).

In the first deployment three instructors will use the system in three different courses: Japanese, physics and interactive systems.

4. Summary

This paper has described a system with an *open learner model* designed to be viewed by the student modelled, their peers, and their tutor. The student model therefore has an even more central role than in the traditional intelligent tutoring system - as a learning resource for the student, to help them reflect on their beliefs in a student-system setting; as a means to prompt collaboration and peer tutoring; and as a source of information for instructors to aid the human teaching process. The system is currently in the early stages of development. While it could now be deployed in its most simple state, a few extensions will increase its utility as a focus for promoting learner reflection on their understanding. These extensions include a simple form of *collaborative student modelling*, and taking into account *pair models* to further support peer collaboration. Most importantly, weightings will be applied to the most recent answers given, in order to ensure that the student model appropriately updates as the student learns.

Even with these extensions, the system will remain straightforward. It is a practice environment rather than a tutoring environment - though the instructor may provide links to course materials. Its primary aim is to investigate some of the questions relevant to the use of inspectable learner models in a range of contexts - both different subject areas, and different students. Such questions include the method by which learners prefer to view (and later interact with) their learner model - the individual's preferred method of accessing the learner model is likely to be more relevant than, for example, the influence of the domain, as it is *knowledge level* (rather than actual knowledge or misconceptions), that is represented. The goal here is to raise learner awareness of where their strengths and weaknesses lie, thus representations similar to those illustrated, could be usefully considered. Other issues include how much attention students pay to their own and to peer models; whether interaction with the learner model occurs with some students more than others; with some domain types more than others; or whether there are differences in the same student's acknowledgement of their learner model in different domains. Also relevant are instructors' use of their students' models.

It is acknowledged that a more complex student model is ultimately likely to be most helpful, in particular in the context of the individual student viewing their own learner model. However, to investigate questions such as the above, it is useful to employ a simple learner model in an easy to deploy system. Of course, more complex student models will have additional requirements for their externalisation - these can be investigated in parallel, or once some of the simpler questions have been addressed. An obvious question is how to show larger models - the current type of display, while clear for small domains, may need to be modified if many interrelated concepts were to be included in a single exercise set. Nevertheless, even with our initial straightforward approach, it is intended that students will be able to benefit from interactions with the system, and that some initial guidelines about the effective externalisation of learner models can be obtained.

References

Barnard, Y.F. & Sandberg, J.A.C. (1996). Self-Explanations, do we get them from our students?, in P. Brna, A. Paiva & J. Self (eds), *Proceedings of European Conference on Artificial Intelligence in Education*, Lisbon, 115-121.

Bull, S. (1997). See Yourself Write: A Simple Student Model to Make Students Think, in A. Jameson, C. Paris & C. Tasso (eds), *User Modeling: Proceedings of the Sixth International Conference*, Springer, Wien New York, 315-326.
Also available from: http://w5.cs.uni-sb.de/~UM97/abstracts/BullS.html

Bull, S. (1998). 'Do It Yourself' Student Models for Collaborative Student Modelling and Peer Interaction, in B.P. Goettl, H.M. Halff, C.L. Redfield & V.J. Shute (eds), *Intelligent Tutoring Systems*, Springer-Verlag, Berlin, Heidelberg, 176-185.

Bull, S. & Broady, E. (1997). Spontaneous Peer Tutoring from Sharing Student Models in B. du Boulay & R. Mizoguchi (eds), *Artificial Intelligence in Education*, IOS Press, Amsterdam, 1997, 143 - 150.
Also available from: http://www.eee.bham.ac.uk/bull/papers/AIED97.htm

Bull, S. & Pain, H. (1995). 'Did I say what I think I said, and do you agree with me?': Inspecting and Questioning the Student Model, in J. Greer (ed), *Proceedings of World Conference on Artificial Intelligence in Education*, Association for the Advancement of Computing in Education (AACE), Charlottesville, VA, 1995, 501-508.
Also available from: http://www.eee.bham.ac.uk/bull/papers/AIED95.htm

Collins, J., Greer, J.E., Kumar, V.S., McCalla, G.I., Meagher, P. & Tkatch, R. (1997). Inspectable User Models for Just-In-Time Workplace Training, in A. Jameson, C. Paris & C. Tasso (eds), *User Modeling: Proceedings of the Sixth International Conference*, Springer Wien New York, 327-337.
Also available from: http://w5.cs.uni-sb.de/~UM97//abstracts/CollinsJA.html

Corbett, A.T. & Bhatnagar, A. (1997). Student Modeling in the ACT Programming Tutor: Adjusting a Procedural Learning Model with Declarative Knowledge, in A. Jameson, C. Paris & C. Tasso (eds), *User Modeling: Proceedings of the Sixth International Conference*, Springer Wien New York, 243-254.
Also available from: http://w5.cs.uni-sb.de/~UM97/abstracts/CorbettA.html

de Buen, P.R., Vadera, S. & Morales, E.F. (1999). A Collaborative Approach to User Modeling within a Multi-Functional Architecture, in J. Kay (ed), *UM99: User Modeling, Proceedings of the Seventh International Conference*, Springer Wien New York, 291-293.
Also available from: http://www.cs.usask.ca/UM99/Proc/short/BuenRodriguez_032414.pdf

Dimitrova, V., Self, J. & Brna, P. (2001). Applying Interactive Open Learner Models to Learning Technical Terminology, in M. Bauer, P.J. Gmytrasiewicz & J. Vassileva (eds), *User Modeling 2001: 8th International Conference*, Springer-Verlag, Berlin Heidelberg, 148-157.

ELM Research Group (1998). *ELM-ART Lisp Course*.
http://www.psychologie.uni-trier.de:8000/elmart

Grigoriadou, M., Papanikolaou, K., kornilakis, H. & Magoulas, G. (2001). INSPIRE: An Intelligent System for Personalized Instruction in a Remote Environment, in P. De Bra, P. Brusilovsky & A. Kobsa (eds), *Pre-Workshop Proceedings: Third Workshop on Adaptive Hypertext and Hypermedia*, User Modeling 2001, 31-40.
Also available from: http://wwwis.win.tue.nl/ah2001/papers/papanikolaou.pdf

Kay, J. (1995). The UM Toolkit for Cooperative User Modelling, *User Modeling and User Adapted Interaction* 4, 149-196.

Kay, J. (1997). Learner Know Thyself: Student Models to Give Learner Control and Responsibility, in Z. Halim, T. Ottomann & Z. Razak (eds), *Proceedings of International Conference on Computers in Education*, Association for the Advancement of Computing in Education (AACE), 17-24.
Also available from: http://www.cs.su.oz.au/~judy/Research_um/index.html

Linton, F. & Schaefer, H-P. (2000). Recommender Systems for Learning: Building User and Expert Models through Long-Term Observation of Application Use, *User Modeling and User-Adapted Interaction* 10, 181-207.

Morales, R., Pain, H. & Conlon, T. (2000). Understandable Learner Models for a Sensorimotor Control Task, in G. Gauthier, C. Frasson & K. VanLehn (eds), *Intelligent Tutoring Systems: Proceedings of the 5th International Conference*, Springer-Verlag, Berlin, Heidelberg, 222-231.

Mühlenbrock, M., Tewissen, F. & Hoppe, H.U. (1998). A Framework System for Intelligent Support in Open Distributed Learning Environments, *International Journal of Artificial Intelligence in Education* 9(3-4), 256-274.

Silva, A., Vale, Z.A. & Ramos, C. (2001). Inspectability and User Controlled Revision on Long Term User Models, in M. Bauer, P.J. Gmytrasiewicz & J. Vassileva (eds), *User Modeling 2001: 8th International Conference*, Springer-Verlag, Berlin Heidelberg, 254-256.

Specht, M., Weber, G. & Schoech, V. (1997). ADI: Ein adaptiver Informations- und Lehragent im WWW, in R. Schaefer & M. Bauer (eds), *Proceedings of ABIS'97: Workshop Adaptivitaet und Benutzermodellierung in interaktiven Softwaresystemen*, Universitaet des Saarlandes, Saarbruecken, 53-60.

Weber, G., Kuhl, H-C. & Weibelzahl, S. (2001). Developing Adaptive Internet Based Courses with the Authoring System NetCoach, in P. De Bra, P. Brusilovsky & A. Kobsa (eds), *Pre-Workshop Proceedings: Third Workshop on Adaptive Hypertext and Hypermedia*, User Modeling 2001, 41-53.
Also available from: http://wwwis.win.tue.nl/ah2001/papers/GWeber-UM01.pdf

Zapata-Rivera, J-D. & Greer, J.E. (2001). Externalising Learner Modelling Representations, *Proceedings of Workshop on External Representations of AIED: Multiple Forms and Multiple Roles*, International Conference on Artificial Intelligence in Education 2001, 71-76.