User Modelling in I-Help: What, Why, When and How

Susan Bull, Jim Greer, Gordon McCalla, Lori Kettel, Jeff Bowes

ARIES Laboratory, Department of Computer Science, University of Saskatchewan,
Saskatoon, Saskatchewan, S7N 5A9, Canada.
bull@cs.usask.ca, greer@cs.usask.ca, mccalla@cs.usask.ca

Abstract. This paper describes user modelling in I-Help, a system to facilitate communication amongst learners. There are two I-Help components: Private and Public Discussions. In the Private Discussions learners take part in a one-on-one interaction with a partner (possibly a peer). The Public Discussions are open - everyone in the group has access to all discussion forums relevant to that group. The Public Discussions are most suited to discussion of issues where there might be a variety of valid viewpoints, or different solutions to a problem. It is also useful for straightforward questions and answers that have wide-spread applicability. The Private Discussions are better suited for more intensive interactions involving peer tutoring or in-depth discussions. Because there is only one helper in such situations, I-Help requires a method of selecting an appropriate helper for an individual. We describe the user modelling that takes place in each part of I-Help, in particular to effect this matchmaking for Private Discussions. This modelling takes advantage of a distributed multi-agent architecture, allowing currently relevant user model fragments in various locations to be integrated and computed at the time they are required.

keywords. peer help network, discussion forum, distributed student model

1 Introduction

Discussion forums and one-on-one peer help environments can be useful for providing support outside normal class time. The two environments are best suited to different types of interaction. Discussion forums allow learners to ask and answer questions on a variety of topics, which is especially useful for discussion involving multiple viewpoints or multiple problem solutions, or for straightforward questions requiring a quick response. In contrast, one-on-one peer help networks can assist open-ended intensive interactions between two individuals, such as tutoring. This requires some kind of user modelling in order to locate appropriate helpers (Mühlenbrock et al. [13]; Ogata et al. [14]; Yu et al. [17]). However, despite their utility for different types of interaction, discussion forums and one-on-one peer help networks are rarely integrated. I-Help combines a peer matching component to form pairs for one-on-one help sessions (I-Help Private Discussions), and a discussion forum for small group and large group participation (I-Help Public Discussions).

For peer matching, a method of selecting appropriate partners is a primary purpose for user modelling. Other purposes include assessment and self-assessment, reflection, knowledge management, expertise location. This paper focuses on user

M. Bauer, P.J. Gmytrasiewicz & J. Vassileva (eds), User Modeling 2001: 8th International Conference, Springer-Verlag, Berlin Heidelberg, pp. 117-126.

model data collected during use of I-Help, describing how this is applied to various purposes in I-Help under various circumstances and constraints. Attributes modelled include an individual's knowledge level in the various course topics; their eagerness to participate; their helpfulness; their cognitive style; their learning goals or interests; and their preferences in a learning partner. Sources of model representations are students themselves, peers, and information gathered by the two I-Help components.

I-Help is built on an agent architecture. Agents are also a fundamental part of the user modelling process. Each user has their own personal agent to take care of their interests. In addition to developing a model of their owner, an agent constructs partial models of users of all agents with which it comes into contact. The multi-agent approach results in distributed and fragmented user models which can be computed just-in-time, for the particular purpose at hand, using only the data required for that purpose. The data and user modelling representations and techniques are in place to do many things in I-Help. We demonstrate how our approach takes advantage of some of this user model power to successfully match partners, and suggest how the existing models may be used in the future, to develop the approach even further.

2 User Modelling in I-Help

User models in I-Help are fundamental and central to the system's function, but the models themselves are fragmented and distributed. There is no single monolithic user model describing an individual learner. Rather, learner models are represented, stored, computed, and applied in a distributed fashion and on an as-needed basis. Our overall notion of fragmentation and multi-agent communication in learner modelling has been termed "active learner modelling" (McCalla et al. [11]). In this paper we focus on the types of information used for modelling in I-Help and on techniques for integration of information into meaningful models for use in I-Help.

Various information types are modelled: knowledge; interests; cognitive style; eagerness; helpfulness; interaction preferences; opinions of peers; user actions. These are recorded, observed and inferred in various ways from a range of sources. One source is users themselves. Self-evaluation is employed in a number of user models to take into account a user's views of their competence (Beck et al. [1]; Bull & Pain [5]; Dimitrova et al. [7]; Kay [9]). In I-Help users provide profile information assessing their knowledge on topics of relevance in the I-Help groups to which they belong (topics are pre-assigned by the person in charge of an I-Help group).

Another source of information about users is peer assessment. Although a few examples exist (e.g. Bull et al. [3]), peer evaluations have been less frequently employed in student modelling. In I-Help Private Discussions, feedback on help sessions is sought from helpers and helpees. This updates the knowledge profile data for both helper and helpee. Judgements about content knowledge and interests from these various sources may be inconsistent, but attempts to reconcile this information are not made until some decision forces a summarisation. Indeed, there is no need for reconciliation prior to this point, as the various fragments of user model data will be successively updated as more data is gathered, but this will likely occur at different rates. For example, a person under consideration as a helper should be knowledgeable on the topic. If user-provided profile information on their knowledge level does not agree with observations made by other users, rationalisation of these conflicts is relevant only at that time: incoming data (such as additional peer evaluations) may at any moment increase the weighting of a particular point of view on such a conflict, affording greater credibility for that viewpoint.

As well as being sources of information for modelling, self and peer evaluation, even in this simple form, are useful activities in their own right – to raise awareness of domain issues and learning in general; and perhaps also to contribute to some form of learner assessment. Thus multiple modelling purposes are encompassed.

Browsing patterns and posting activity in the Public Discussions provide another source of information about users. Such activity is analysed to determine topics and concepts of interest through off-line data mining (Bowes [2]). Interest in particular forums, topics, users and individual postings can also be determined by monitoring notification preferences set by a user in the Public Discussions. Learners wishing to be notified by email about specific events in I-Help will likely have an interest in those events. In addition to enabling email notification of new postings of interest, user model representations derived from the Public Discussions are used in matching pairs in the Private Discussions – e.g. a user requesting notification of new postings by a specific author will usually do so because they find this user's postings helpful.

In addition to their knowledge of various topics, users also indicate interaction preferences in the Private Discussions: e.g. characteristics they seek in a helper and the maximum number of concurrent help sessions in which they can participate. They may also ban people with whom they do not wish to interact, and indicate people with whom they will preferentially interact. Topics may be banned if the user does not wish to give help in the area. Thus the learner sets the parameters within which their personal agent should operate when negotiating matches, affording them a degree of control (c.f. Kay [10]). This adds a social dimension to the user modelling.

Other information is also modelled. Data on eagerness to participate is inferred for a user, gathered from the user's I-Help usage patterns. Frequency of login to I-Help, the number of postings read and re-read in the Public Discussions, and time spent responding to help requests in both Public and Private Discussions contribute to the eagerness measure. Similarly, data on helpfulness of users is gathered from many sources. Users may register opinions about their helpfulness. Helpers in the Private Discussions are judged for helpfulness by the helpee. When a potential helper refuses to offer help, some evidence about helpfulness (or lack thereof) is gained. Helpfulness is also judged by the type of participation in Public Discussions – whether they often offer answers, and whether those answers are frequently read and endorsed. The Public Discussions have a voting mechanism where users can vote on the utility of postings. Endorsements of peers may be interpreted as a measure of helpfulness.

The Private Discussions contain a brief questionnaire from which inferences are made to help identify a user's cognitive style, to suggest suitable helpers with reference to different kinds of question. (Users indicate their question type by menu selection when submitting a help request.) The cognitive style categorisation is based on Riding and Cheema's [15] classification of the two dimensions of wholist-analytic and verbal-imagery style (see Bull & Greer [4]).

Table 1 summarises the attributes modelled, and the sources of information for these representations. As stated above, this information is integrated on an as-needed basis, using only the data that is relevant for the particular purpose (for example, knowledge level is relevant for assessment, but eagerness is probably not).

Table 1. User model attributes and information sources

Attributes	Sources of information			
	User	Peers	I-Help Private	I-Help Public
knowledge	•	•		
interests	•			•
eagerness	•		•	•
helpfulness	•	•	•	•
readiness (online)			•	•
cognitive style	•		•	
preferences in helper	•			
preferred people	•			
banned people	•			
banned topics	•			
help-load	•		•	

To enable I-Help to select the most appropriate helper for a particular request in the Private Discussions, it uses a matchmaker agent that communicates with the personal agents of users (Vassileva et al. [16]), which can, in turn, access various kinds of information about their 'owners' and about other users with which their owners have come into contact. The matchmaker performs a quick, coarse-grained survey of users who know about the topic in question. Other factors, such as cognitive style, eagerness and current help-load are also considered. A ranked list of potential helpers is then produced. To rationalise the possible diversity of opinions about a user, the current matchmaker constructs a composite estimate (weighted score). While this estimate is rather crude, preliminary studies have shown it to be sufficient for recruiting satisfactory helpers. A more subtle composite measure might further improve matchmaking. We have constructed a Bayesian modelling agent (service) to take a fragment of a Bayesian network and a set of observations and produce estimates of probabilities of nodes under various assumptions of new evidence (Zapata [18]). This service will be in integrated into an upcoming version of I-Help.

Modelling the helpee also plays a part in the selection of a helper. The preliminary helper list generated by the matchmaker is scanned to identify preferred helpers and banned people. If helpfulness is an important criterion for the helpee, helpers with higher levels of helpfulness are raised in the list. If urgency in finding a helper is an important criterion, then readiness (being on-line) is raised in priority, etc.

Once the helper list is prepared, the personal agent of the helpee is brought into negotiation with the agents near the top of the list. This, too, involves modelling. During negotiation local knowledge the user’s agent has of the potential helper is consulted. Previous help interactions can be recalled; good or bad experiences with the helper that may have been shared by other users’ agents may come into play. Negotiation over availability and willingness to accept the help request is then undertaken. A final list of five potential helpers is prepared, to which the request is sent. The first user to reply becomes the helper.

I-Help user modelling is important to protect helpers from unwanted interruptions and to spread the help-load. Any capable helper would become overwhelmed with requests in an open environment. In I-Help users and agents can set thresholds to regulate the number of requests received. A busy flag can be raised that can, for example, bar anyone but a friend with an urgent request. As agents in I-Help are further developed, inferences made by agents and the richness and independence of the models stored in the agents will grow. Agents will, in effect, each become expert systems able to compute user model information as needed for a particular purpose.

This section has illustrated that there is a substantial quantity of user model information that can be gathered and utilised in I-Help, and a range of uses. The multi-agent, distributed approach enables powerful modelling – much of which is already incorporated in I-Help. It must be emphasised that no effort is made to unify this data into a single user model for each user. The models that must be formulated depend on the purposes to which the models are put (e.g. helper location, information retrieval, promoting reflection, etc.) In fact the information used in the modelling processes may reside in different places – spread across the many agents in I-Help.

3 An Example

To illustrate the issues described in the previous section, we here look at an example drawn from a deployment of I-Help in a first year Computer Science course at the University of Saskatchewan (Spring 2000). We examine in more detail how the user modelling computations were used by the matchmaker agent to rank and select potential helpers in response to a particular help request in the Private Discussions. We then compare the matchmaker's rankings with other evidence of the potential helpers' capabilities drawn from their use of I-Help. The intention is to show how user modelling occurs in the context of a real example. We further explore the context of the help request by providing a snapshot of I-Help usage at the time of the request.

The computations described below are those used in the Spring 2000 version of I-Help, where the following attributes were used by the matchmaker agent in ranking potential helpers: knowledge level, helpfulness, eagerness, and readiness. These attributes are computed for each user and then combined (using a weighted linear equation) into an overall evaluation of each helper. Negotiations between the helpee's agent and the agents for the potential helpers then proceeds, to select the final helper. We consider this process in the context of the following (actual) help request, submitted to the Private Discussions:

> A simple question about looping: I want the program to enter the loop first time without trying the question. I know there is a way. There's the for loop, the while loop and…

```
String currentline;
while (currentline != null) {
        currentline = br.readLine();
```

As described previously, knowledge level information for user modelling is obtained in part from self assessment, and in part from peer evaluations. The self-evaluation of the knowledge level relates to the user's provision of information about their knowledge of the course topic on a scale of 1-10. Peer evaluations are based on previous evaluations of the potential helpers' knowledge of the topic of the current helpee's question. Peers rate knowledge level on a three point scale: low, medium or high. Peers are therefore not expected to rate a user's knowledge as precisely as the user him/herself, since peers only experience a user's understanding of a small area of a topic. However, since individuals may fail to update their own knowledge profile frequently, composite peer evaluations are weighted higher than self evaluations. (In the next version, relative recency of evaluations will also be taken into account.)

Helpfulness is calculated according to answers to questions on previous evaluation forms filled out by previous helpees: "Was the helper helpful?"; "Was the helper easy to understand?" This is used in conjunction with I-Help's observations of frequency and speed of response to questions in the Public and Private Discussions, thus again combining multiple sources of information for a single attribute.

Eagerness is more complex, taking into account a range of factors relating to individuals, in combination with the average and standard deviation of these factors across all users. These issues include logins to the Public and Private Discussions (which, at the time of the first large-scale deployment, were separate); postings made and read in the Public Discussions; information provided by the user to their agent (e.g. the act of updating their knowledge profile).

Readiness refers to whether the user is online. If they are, a weighting is added.

Using the above modelling information and sources, the following ranked list (showing the top eleven) was produced by the matchmaker in response to the above help request, for the purpose of locating a good helper for the question:

Prof1: 6.997 Peer1: 5.612 Prof2: 4.71 TA1: 4.612 banned TA2: 3.91
Peer2: 2.836 Peer3: 2.232 Peer4: 2.18 Peer5: 2.178 Peer6: 2.14 Peer7: 2.126

Table 2. Potential Helper User Models

	Prof 1	Peer 1	Prof 2	TA 1	TA 2	Peer 2	Peer 3	Peer 4	Peer 5	Peer 6	Peer 7
eager	5.45	5.86	4.8	8.36	4.55	6.58	6.36	6.1	5.49	5.9	6.13
helpful	9.69	9	2.5	0	0	0	0	0	0	0	0
self eval	10	10	10	9	10	9	8	8	10	8	7
sys eval	10	4.75	10	10	10	2	2	2	2	2	2
ready	0	0	0	0	0	0.5	0	0	0	0	0

The ranked list includes two professors (Prof) and two teaching assistants (TA), whose score for system_evaluation was set to 10 (maximum) at the start. 5 of the top 6 were considered the most suitable potential helpers by the matchmaker – TA1 was not included as he banned the topic. This ranked list was computed at the time of the help request, integrating all currently available relevant information. Had the request been made three hours later, two lower ranked peers would have made it to the final selection, as peer evaluations of previous help sessions for these users occurred in the meantime, increasing their total score to above those of TA2 and Peer2.

Table 2 shows the breakdown of the scores according to self and system/peer evaluation, helpfulness, eagerness and readiness at the time of the request. All scores are out of 10. The help request in this example happened near the beginning of the course, therefore there are few helpee evaluations to take into account. This explains the low scores for most users for helpfulness and system evaluation.

At this stage negotiation between the helpee's agent and the agents of helpers at the top of the list commenced, with agents referring to many user model fragments at their disposal to act in the best interests of their user: i.e. the helpee's agent will be seeking the best possible helpers for their owner's particular question; agents of potential helpers will give preference to helpees with questions on topics in which they are interested, people they like, etc. Issues such as current help-load for helpers, and preferred helpers for helpees, come into play. In this case the matchmaker's suggestion was upheld, and the shaded individuals in Table 2 received the request. (Peer2 answered it, sending the response within 7 minutes of receipt of the question.)

It is difficult to validate matchmaker choices in I-Help, since so many factors are involved in helper selection. However, there are a number of things that may be considered as indicators. For example, we can look at how many helpees banned their helpers subsequent to the help session in the Private Discussions. In fact, no helpers were banned at any time. We can also look at helpee evaluations of helpers. At the time of data collection, 29 helpee evaluations (and 70 helper evaluations) had been completed. 27 of the 29 helpees found their helper to be helpful; 28 found the helper easy to understand; 14 evaluated the helper's knowledge as high, 14 as medium, and only one as low (this was one of the users who considered their helper unhelpful); 10 helpees added the helper to their preferred helper list (i.e. over one third of helpees requested the same helper to be given high priority in the future). Given the responses to the questions about helpfulness and ease of understanding, it seems that a 'medium' knowledge level is sufficient for help to be considered effective. Thus, from the viewpoint of helpees, the matchmaker appears to nearly always make successful selections even with the more restricted modelling in this earlier version of I-Help.

We can also look at Public Discussions usage statistics. Our purpose here is to show that matchmaker choices are reasonable – at least better than random matching or an open help environment (where the best helpers would be overwhelmed with requests). Focusing on peers (we assume professors and teaching assistants are eager, helpful, and knowledgeable), the two peers contacted as potential helpers for this help request in the Private Discussions were in the 5th and 15th position (of 318 users) in terms of responses to postings in the Public Discussions (24 and 7 replies made; average = 1.82). Prestige measures relating to the Public Discussions can also be used to consider the value of the matchmaker. Prestige of an author is calculated by measuring the average number of times that a user reads the same reply by the author (Bowes [2]). Individuals high on this list have generally been considered good, and they are therefore likely to be good helpers in the Private Discussions. The two peers chosen by the matchmaker were 15th and 16th on the "prestige" list.

Reviewing general activity in the Private Discussions, the two selected helpers were in 1st and 3rd position, suggesting they would probably see the request soon.

Looking briefly at the five peers who almost made the final helper list, two of these were consistently located between the scores of those selected as potential helpers, and one was consistently in first position, above those on the helper list. The remaining two peers that did not make the final list were well above average on the Private activity, Public replies and prestige lists. Other factors were involved in the final listing – for example, the lower ranked helper on the final list (Peer2) was the only one online at the time of the request. Lack of previous peer evaluations has already been explained as affecting the ranking of others. Further factors may have included banning of the topic by some peers, low level of knowledge on the particular topic of the question, the need to spread the help-load, etc.

These results illustrate that the I-Help matchmaker, drawing on information which is available at the time from distributed and fragmented user models, based on a range of data sources, generally makes good decisions about who is an appropriate helper for a particular help request. This applies even with the earlier, simpler modelling process. The current version of I-Help also models cognitive style and a range of choices of the user relating to characteristics important to them, in a helper. Once data from the current version is analysed, we hope to find evidence that the additional user model data further enhances the performance of matchmaking. This is important for planned larger deployments.

4 Conclusions and Future Directions

This paper has shown user modelling to be a crucial process in I-Help. There are many purposes: user modelling is needed in finding and ranking candidate helpers; in negotiation between personal agents over a help request; in protecting users from interruptions; in retrieving relevant information for a user; in providing open models of a user's content knowledge of a domain. There are many attributes: user modelling captures interests, goals, even cognitive styles of users, as well as domain skill levels, eagerness, and helpfulness factors. There are many perspectives on these attributes: user modelling happens from the user's own point of view, from the point of view of other users, from the point of view of application agents. Even though knowledge about users is drawn from a variety of sources with a variety of perspectives for a variety of purposes, thus meaning that user modelling in I-Help is pervasive, it is also essentially fragmented. There is no single, complex, unitary model for each user. Instead, user modelling is a "just-in-time" process, invoked as a by-product of achieving a particular purpose and using information relevant to that purpose. We feel that this perspective on user modelling is scalable and appropriate to the large, distributed, agent-based environments which are increasingly common.

We are extending I-Help, and deploying it ever more widely in University courses, both at the University of Saskatchewan and elsewhere (over 2000 students are currently using our more recent version of I-Help). As I-Help's sophistication and applicability increases in these course deployments, user modelling is becoming, if anything, even more important than in previous versions.

As well as a learning tool, I-Help can be viewed as a framework within which possible future directions are explored and evaluated, with successful projects later becoming integrated into the system. Current investigations include enabling agents to play a role in retrieving information for a use, rather than seeking a human helper, e.g. some FAQ or answer posted in the Public Discussions. To determine whether a resource document or a posting might be relevant to the user, simple keyword queries can be created, latent semantic analysis of the question and available answers can be performed, and/or semantic similarity indices may be used. Any of these methods are complemented through user modelling. For example, the help request itself may be augmented by knowledge of the kinds of browsing occurring prior to the formulation of the help request. If the user had viewed postings on a particular topic, or questions by a particular user, additional query terms could be added to the information retrieval process. Even the knowledge level or degree of eagerness that the user had previously exhibited might give clues as to the kind of information that would be most helpful.

A new project under development for the next version of I-Help is based on anticipating learner interests and offering useful information to users even before help is requested. This resembles a knowledge-based "tip-of-the-day" feature. Whenever information appears in the system, agents should be able to peruse that information and determine the possible relevance for their user. This is akin to information filtering, and once again, is based on the agent modelling its user's goals and interests.

User modelling in I-Help also has an important role in assessment and self-assessment of learners and in supporting reflection on learning. Opening learner models for inspection has been argued to promote reflection (Bull & Pain [5]; Dimitrova et al. [7]; Morales et al. [12]). Some of the current user model fragments of I-Help are open to viewing by the modellee – specifically those parts contributed by the modellee. We are working towards opening other areas: eagerness, helpfulness and a composite score of peer evaluations of knowledge which may be explicitly compared with self-assessments. Research into inspectable interfaces for Bayesian network learner models is also underway (Zapata & Greer [19]).

Another major goal of our investigations is to incorporate a more overt role for groups in I-Help. There are currently many kinds of groups in the system, both explicit and implicit (e.g. the students in a class or project team, friends, banned people, potentially good helpers, etc.). We want to make groups more fundamental, to capture aspects of group behaviour, to stimulate the formation of relevant groups, and to support group interaction. We feel that there is actually a graduated series of functionalities between the one-on-one Private Discussions and the many-on-many Public Discussions. User modelling and group modelling (as in Hoppe [8]) will be important in dealing with the issues raised by a more central role for groups in I-Help.

Even though I-Help has been presented as a learning tool in a University setting, it can be equally applied to a range of corporate and industry activities, e.g. for knowledge management and expertise location, illustrated by the fact that an early prototype of the Private Discussions was piloted in the Prairie Regional Psychiatric Centre of Corrections Canada (Collins et al. [6]); and the Public Discussion forums are now being adapted to serve as a knowledge management tool in a local company. Capturing interests and goals, as well as skill levels, eagerness and helpfulness factors are of great interest in corporate settings. Developing profiles of employees and modelling such nebulous attributes as being a 'team player' or a 'self-starter' become possible with the range of subtle data collected and analysed in I-Help.

References

1. J. Beck, M. Stern and B.P. Woolf. Cooperative Student Models. In B. du Boulay and R. Mizoguchi, editors, *Artificial Intelligence in Education (Proceedings of AIED'97)*, pages 127-13, Amsterdam, 1997. IOS Press.

2. J. Bowes. *Knowledge Management through Data Mining in Discussion Forums*. MSc Thesis, Department of Computer Science, University of Saskatchewan, 2001.
3. S. Bull, P. Brna, S. Critchley, K. Davie and C. Holzherr. The Missing Peer, Artificial Peers and the Enhancement of Human-Human Collaborative Student Modelling. In S. Lajoie and M. Vivet, editors, *Artificial Intelligence in Education (Proceedings of AIED'99)*, pages 269-276, Amsterdam, 1999. IOS Press.
4. S. Bull and J. Greer. Peer Help for Problem-Based Learning. In S. Young, J. Greer, H. Maurer and Y.S. Chee, editors, *Proceedings of ICCE/ICAI'00 (Vol. 2)*, pages 1007-1015, 2000. National Tsing Hua University of Taiwan.
5. S. Bull and H. Pain. 'Did I say what I think I said, and do you agree with me?': Inspecting and Questioning the Student Model. In J. Greer, editor, *Proceedings of World Conference on Artificial Intelligence in Education*, pages 501-508, 1995. AACE.
6. J. Collins, J.E. Greer, V.S. Kumar, G.I. McCalla, P. Meagher and R. Tkatch. Inspectable User Models for Just-In-Time Workplace Training. In A. Jameson, C. Paris and C. Tasso, editors, *User Modeling: Proceedings of the Sixth International Conference*, pages 327-337, Vienna, New York, 1997. Springer Wien New York.
7. V. Dimitrova, J. Self and P. Brna. *Maintaining a Jointly Constructed Student Model*. Technical Report, Computer Based Learning Unit, Leeds University, 2000.
8. H.U. Hoppe. The use of Multiple Student Modelling to Parameterise Group Learning. In J. Greer, editor, *Proceedings of World Conference on Artificial Intelligence in Education*, pages 234-241, 1995. AACE.
9. J. Kay. The UM Toolkit for Cooperative User Modelling. *User Modelling and User-Adapted Interaction*, 4:149-196, 1995.
10. J. Kay. Learner Know Thyself. In Z. Halim, T. Ottman and Z. Razak, editors, *Proceedings of ICCE'97*, pages 18-26, 1997. AACE.
11. G. McCalla, J. Vassileva, J. Greer and S. Bull. Active Learner Modelling. In G. Gauthier, C. Frasson and K. VanLehn, editors, *Intelligent Tutoring Systems (Proceedings of ITS'00)*, pages 53-62, Berlin, Heidelberg, 2000. Springer-Verlag.
12. R. Morales, M. Ramscar, and H. Pain. Modelling the Learner's Awareness and Reflection in a Collaborative Learner Modelling Setting. *Workshop on Current Trends and Applications of AIED, Fourth World Conference on Expert Systems*, Monterrey, 1998.
13. M. Mühlenbrock, F. Tewissen and H.U. Hoppe. A Framework System for Intelligent Support in Open Distributed Learning Environments. *International Journal of Artificial Intelligence in Education*, 9(3-4):256-274, 1998.
14. H. Ogata, T. Sueda, N. Furugori, and Y. Yano. Augmenting Collaboration Beyond Classrooms through Online Social Networks. In G. Cumming, T. Okamoto and L. Gomez, editors, *Advanced Research in Computers and Communications in Education*, pages 277-284, Amsterdam, 1999. IOS Press.
15. R. Riding, and I. Cheema. Cognitive Styles - An Overview and Integration. *Educational Psychology*, 11(3-4):193-215, 1991.
16. J. Vassileva, J. Greer, G. McCalla, R. Deters, D. Zapata, C. Mudgal and S. Grant. A Multi-Agent Design of a Peer-Help Environment. In S.P. Lajoie and M. Vivet, editors, *Artificial Intelligence in Education (Proceedings of AIED'99)*, pages 38-45, Amsterdam, 1999. IOS Press.
17. Y-T. Yu, H-L. Shan, C-Y. Chou, and T-W. Chan. WWW-Based Mentoring System. In Z. Halim, T. Ottman and Z. Razak, editors, *Proceedings of ICCE'97*, pages 381-389, 1997. AACE.
18. J.D. Zapata. *SMODEL Server: Student Modelling in Distributed Multi-Agent Tutoring Systems*. ARIES Lab Technical Report, University of Saskatchewan, 2000.
19. J.D. Zapata and J.E. Greer. Inspecting and Visualizing Distributed Bayesian Student Models. In G. Gauthier, C. Frasson and K. VanLehn, editors, *Intelligent Tutoring Systems (Proceedings of ITS'00)*, pages 544-553, Berlin Heidelberg, 2000. Springer-Verlag.