

The Caring Personal Agent

SUSAN BULL¹, JIM GREER², GORD MCCALLA²

1 Educational Technology Research Group, Electronic, Electrical and Computer Engineering, University of Birmingham, Edgbaston, Birmingham B15 2TT, U.K.

s.bull@bham.ac.uk, <http://www.eee.bham.ac.uk/bull>.

2 ARIES Laboratory, Department of Computer Science, University of Saskatchewan, Saskatoon, Saskatchewan S7N 5A9, Canada.

{greer, mccalla}@cs.usask.ca, <http://www.cs.usask.ca/homepages/faculty/greer>, <http://www.cs.usask.ca/homepages/faculty/mccalla>.

Abstract. Self (1999) argues that the essence of having a computer-based learning system that “cares” about its learners is that the system model its learners so as to be able to adapt to their needs. In this paper we discuss the notion of personal agents who care for their “owners” by representing the owners’ interests in the learning system. We contextualise this discussion by showing how such personal agents are used in I-Help, a system that promotes caring and sharing by encouraging learners to help one another. In I-Help, personal agents themselves care for their learners by helping them to discover useful information and/or to find “ready, willing, and able” peer learners who can aid them in overcoming problems.

INTRODUCTION

Self (1999) maintains that "a student model is what enables a system to care about a student": i.e. an intelligent tutoring system (ITS) is concerned with the individual student's knowledge, misconceptions, goals and often many other issues, which are represented in a student model. The ITS takes these attributes into account to adapt an interaction to best suit the individual user. In this sense the student model enables a caring approach to computer-assisted learning.

The traditional student model, inferred from diagnosis of a learner’s actions, has evolved much in the last few years. The original notion of the single learner-system interaction of an ITS has changed as more systems are being developed to support collaborative learning. Many systems now facilitate the matching of learners for computer-mediated human-human help sessions, using multiple student models (Bull, 1997; Hoppe, 1995; Mühlenbrock et al., 1998).

Such help environments are increasingly important as class sizes grow, since in large-scale, impersonal learning settings many learners find it difficult to ask for academic help when they experience difficulties. Indeed, it has been suggested that it is often those learners who need help the most who refrain from seeking it (Karabenick & Knapp, 1988). This is not an easy issue to address, since there are a range of reasons for learners not to request help even when they recognise that they need it. These include lack of time (Karabenick, 1990; Moncada & Sanders, 2000); the belief of students that they must simply try harder (Knapp & Karabenick, 1988); not being able to reciprocate (Karabenick, 1990); students' perception of the instructor’s support (or lack thereof) for help-seeking (Galloway & Easterday, 1998; Karabenick & Sharma, 1994; Moncada & Sanders, 1990; Perrine et al., 1995); embarrassment (Shapiro, 1983); and threat to self-esteem (Karabenick & Knapp, 1991; Moncada & Sanders, 2000). Many of these reasons are affective, and would perhaps be alleviated with the guidance of someone who cares. Clearly, it is unrealistic to assign to individual learners a caring person who would be available to help them at any time they require assistance. It is equally unrealistic to expect an ITS to be able to help learners on the whole range of problems they may encounter during a course.

However, it is possible to assign to individual learners a caring computational agent who will locate for them a ready, willing and able person, who can answer their questions or help them solve their problem, at the time they need help.

This involves extending Self's (1999) idea, that it is the student model that allows ITSs to care about their students, by employing personal agents – autonomous programs which make their own decisions in reaction to stimuli from their environment – to care about their owners. This requires using student model information about the agent's owners and about other users, in order to consider their owners' educational, psychological and social needs in the academic setting. Educational, psychological and social knowledge must therefore be represented in the user models of all participants, congruent with the goal of ITS research of making this kind of knowledge precise and explicit (Self, 1999).

Some systems provide learners with an agent that communicates with other agents to find suitable helpers for their owner, focussing on, for example, expertise level (Vivacqua, 1999), interests (Thibodeau et al., 2000) or social networks (Ogata et al., 1999). With large classes there is much scope for basing the selection of helpers on a greater variety of information, and weighting the different attributes according to the preferences of the user. We introduce the I-Help system, which takes this approach.

This paper discusses caring personal agents in the context of I-Help, a system comprising multiple agents (one for each learner) that model educational, psychological, social and availability attributes of their owners. Agents also model attributes of other users with whose agents they have had contact. User model data about any individual is therefore fragmented and distributed across the various I-Help agents, and the traditional single, monolithic student model no longer exists (Bull et al., 2001b; McCalla et al., 2000; Vassileva et al., forthcoming). A user's personal agent communicates with other personal agents to reconstruct relevant user model data according to its owner's preferences (along educational, psychological, social and availability dimensions) for a helper when the owner faces some academic problem. The agent also uses student model information to calculate its owner's likely interest in incoming help requests.

The next section illustrates the general caring approach of I-Help. In the following section the paper focuses on caring personal agents. Then, the general question of what it means for a personal agent to be caring is answered, and some research implications of building such caring agents are discussed. The paper concludes with a brief discussion of future research directions.

PROMOTING CARING AND SHARING IN I-HELP

I-Help is based on the metaphor of a help-desk (Greer et al., 1998), with two components: a help forum open to all course participants (Bull et al., 2001a), and the agent-based one-to-one peer help environment (Vassileva et al., 1999). It is the latter that we focus on here.

In I-Help, the personal agents are implemented using a Java-based agent development framework known as DICE (Deters, 2001). Each personal agent has responsibility for maintaining communication with the user whenever the user is connected through a web-based client program. Personal agents communicate with the matchmaker agent to obtain an ordered list of candidate helpers. They also communicate with other personal agents to negotiate help requests. Agent communication follows a simple protocol derived from the agent communication language KQML. The DICE agent framework for agent development has been demonstrated to have overcome scalability concerns inherent in earlier agent implementations. Thus, using our framework, the prospect of deploying thousands of personal agents is feasible given current server technology.

I-Help has three major educational aims. The first is to provide just-in-time academic help to students. A second aim is to promote collaboration and knowledge sharing amongst students, through the creation of an online community of learners to enrich the learning experience. Thirdly, it is important to enhance learning for all participants, i.e. people being helped (“helpees”) and helpers.

Each of these aims illustrates the notion of caring in an intelligent learning environment, with a focus on the human perspective. First, as illustrated in the introduction, learners do have

problems with their courses, but often do not seek help for a variety of reasons. I-Help provides a resource of people who are willing to offer assistance to those experiencing difficulties. As well as helping those who have little time to interact with others in person (e.g. part-time students with jobs and/or families), it also provides a resource for those who do not request help for affective reasons such as embarrassment, since communications occur via personal agents, and identities can be hidden if participants wish. Those selected as potential helpers for an individual are all “ready, willing and able” helpers – people whose previous I-Help usage patterns suggest that they are often available; people who have stated that they are prepared to help (and are not currently over-burdened with help requests); and people who have sufficient knowledge of the topic of the question, who are sufficiently skilled at providing help, who have a suitable cognitive style to give a good explanation for the particular type of question, etc. (Details of the selection of an appropriate helper are given in the following section.) Using this principle of the ready, willing and able helper greatly increases the likelihood of effecting a caring relationship amongst participants: helpers are dedicated, and are not receiving help requests for which they have too little time or knowledge; helpees are more likely to receive the kind of help they require, and will perceive the positive attitude of their helper. In return they may be more enthusiastic to assist their helper should an occasion arise where they are in a position to do so (students can add others to a 'friends list', indicating to their agent their particular willingness to help that person in the future).

The second aim of I-Help is the promotion of collaboration and knowledge sharing. The above illustrates how a community of learners may be built up. In addition to the one-to-one relationships that grow through I-Help use, learners can form into groups with similar concerns who support each other on specific issues. This might be a problem solving session with one or more learners assisting each other with similar problems, or it might be a longer term study group where each member contributes to mutual understanding of the subject matter in some way. Interaction at the one-to-one and the group level is designed to bring learners closer together, and contribute to the evolution of a community of learners anxious to support each other in their learning – i.e. a caring community.

As suggested in the third aim of I-Help, this kind of community of learners is designed to benefit *everybody*. Most obvious is the case of the helpees, who receive personalized assistance when they have a problem. In addition, the very act of formulating their question may help some learners to solve their problem for themselves, potentially furnishing them with an additional learning strategy they can apply in the future. Furthermore, it is not only helpees who benefit from the interaction. Similar to the case of the helpee who solves their own problem through formulating a question, is the helper who improves understanding of the topic through peer tutoring. This learning effect for helpers and helpees has been seen in the discussion forum component of I-Help (Bull et al., 2001a) and, given the similar aims and context, is likely to also hold true in the agent-based, peer matching component. Thus the caring community is also a *learning* community.

THE CARING AGENTS OF I-HELP

In addition to fostering a caring attitude amongst members of a community of learners, individuals have their own personal agent in I-Help, working on their behalf. Thus, there is also a technical side to the caring approach. To enable helpers to collaborate and support their peers, the personal agent works to locate suitable partners in a manner not possible by the participants themselves. In the same way that it is the student model that allows ITSs to care about users (Self, 1999), it is precisely the many student models held by the many I-Help agents that allow successful human collaborations to be formed: the I-Help student models encompass a range of attributes along educational, psychological, social and availability dimensions. This diversity of characteristics ensures that there is sufficient information available to identify useful pairings for a variety of situations. The personal agent takes account of its owner's needs and preferences in a helper, and negotiates with the agents of other users to locate people who are ready, willing and able to answer the particular question. The caring personal agent, therefore, sets up the

partnerships that allow the human-human interactions to take place. Table 1 illustrates the attributes modelled in I-Help.

Table 1. Attributes modelled in I-Help

educational	psychological	social	availability
knowledge interests banned topics	cognitive style satisfaction	eagerness helpfulness helper preferences banned people preferred helpers	readiness help-load currency

Educational attributes are the learner's level of knowledge of the various course topics; their interest in the course topics; and topics about which they wish not to offer help (perhaps due to lack of knowledge, or perhaps for other reasons). Psychological attributes are the learner's cognitive style on Riding and Cheema's (1991) two dimensions of wholist-analytic and verbal-imagery (see Bull & McCalla, in press); and their satisfaction with peer interactions. The social attributes are an individual's eagerness to participate in I-Help; their helpfulness; their preferences in a helper (the relative importance of each of the attributes in a potential helper: the importance of the helper's knowledge of the topic versus their helpfulness versus their speed of response, etc.); people who they do not want to help; and people they wish to be given preference as potential helpers, if they are otherwise suitable. Availability refers to the learner's readiness – whether they are currently online, or likely to come online soon; their help-load – whether they have reached their chosen maximum of concurrent help sessions; and virtual currency, (I-Help Credit Units: Vassileva et al., 1999), which can be used to regulate the amount of incoming requests. For example, a good helper might charge more in order to reduce the number of help requests received. Currency is used as a points reward system. This motivates some learners to provide help, while others help ignoring such rewards (Greer et al., 2001).

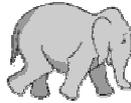
The above attributes are obtained from three sources: I-Help's observations, the learners themselves, and the student's peers (Bull et al., 2001b), and the data resides across the various agents in the system. A personal agent gathers relevant information to identify the best helpers for its owner when required, through interaction with other agents. Relevant information will be different for each agent, depending on its owner's preferences in a helper and other issues, such as the urgency of the help request. The agent also ensures that it only passes on help requests in which its owner is interested.

Self's notion of the student model affording ITSs the ability to care about their users is therefore useful here. However, even ITSs may not appear to the student to be caring. Even if learners appreciate the adaptivity of a system, they might still perceive the experience as impersonal. The personal agent is able to cater for an individual's needs in a more salient manner. Users ask their questions through their personal agent, as it is the agent who can access all the relevant user model data. Interacting with the personal agent in this manner also helps make users aware of the extent of the effort directed towards locating a *suitable* helper for them. Users tell their agents any important information that should be considered – e.g. that they know little about the topic of growing peas, and that someone who would help on this topic should be someone who has been voted a good helper by others in the past. (While this might seem desirable to every help request, it might actually be the case that the learner has a straightforward question, and that speed of response is more important, for example, than depth of knowledge in a topic.)

To clarify the above we introduce five characters and present an example interaction among these characters, facilitated by their agents. First we introduce Muddled McCalla, a man with a problem. We also have several potential helpers. Versatile Vassileva knows much about many topics. Studious Self is also very knowledgeable. In addition we have Banned Bull, who although perhaps able to help, will not be contacted because McCalla has banned her. Finally we have Greer, who is charging a lot for his undoubtedly excellent help. Greedy Greer will therefore only be contacted if Muddled McCalla is either generous or rich. We also have the

personal agents of the above: Forgetful, Sticky, Hacknowledge, Hairy and Stropky (such aliases are important to preserving privacy, an issue discussed in the next section).

Muddled McCalla contacts his agent, Forgetful, to send out a help request.



Forgetful

Forgetful knows that McCalla prefers a helper who is eager and able to respond quickly, because McCalla is usually trying to solve problems at the last minute, and his requests are consequently urgent. He is therefore less interested in the other helper characteristics. The representations Forgetful has concerning McCalla's preferences in a helper (provided by McCalla), are illustrated in Table 2.

Table 2. McCalla's preferences in a helper

<i>Attribute Weighting</i>	helper's helpfulness	helper's eagerness	helper's readiness	helper's cog. style	helper's knowledge	currency
McCalla	3	9	10	1	6	2

Scores are based on a 10-point scale with 10 being highest.

Table 3 illustrates the various attributes of the potential helpers. Recall that data is fragmented across the various agents; thus Table 3 represents 'fully computed' user model excerpts, that mostly do not exist in practice.

Table 3. Helper attributes

<i>Attribute</i>	helpfulness	eagerness	readiness	cog. style	knowledge	currency
Vassileva	7.32	7.48	2.3	I-V	9.1	5
Self	10	5.2	2	A-I	10	2
Bull	6.44	10	10	A-B	8.2	2
Greer	6.08	10	7.6	W-V	8.24	10

Scores are based on a 10-point scale with 10 being highest.

Cognitive style is represented on the continua: wholist-intermediate-analytic; verbal-bimodal-imagery.

As shown in Table 2, McCalla is uninterested in cognitive style, currency and helpfulness of potential helpers, therefore these attributes are given low weightings in Forgetful's consideration of the suitability of an individual as a helper for McCalla. Knowledge is given some weight, but most attention is paid to the potential helpers' eagerness and readiness. Thus Forgetful is seeking a helper who has the characteristics that McCalla considers important at the present time.

Bull, while both eager and online, has been banned by McCalla as a result of unfortunate incidents in the past, and so Hairy, Bull's agent, is ignored by Forgetful.



Hairy

The other main candidate in the areas of eagerness and readiness, as discovered from his agent Stropky, is Greer. While not currently online, he is likely to be back soon. However, Greedy Greer expects high payment for his services.



Stroppy

Vassileva and Self, while reported by their agents Sticky and Hacknowledge to be more helpful and knowledgeable than Greer, score lower on the crucial eagerness and readiness measures.



Sticky



Hacknowledge

McCalla is desperate, and is therefore willing to pay well for an answer. Thus Greer is ranked highest, and suggested by Forgetful to be the most appropriate helper. Forgetful agrees a price with Greer's agent Stroppy. Grateful, Greer accepts the help request and takes the money.



In reality there is much agent-agent interaction taking place when a personal agent is seeking potential helpers for its owner, as illustrated in Figure 1.

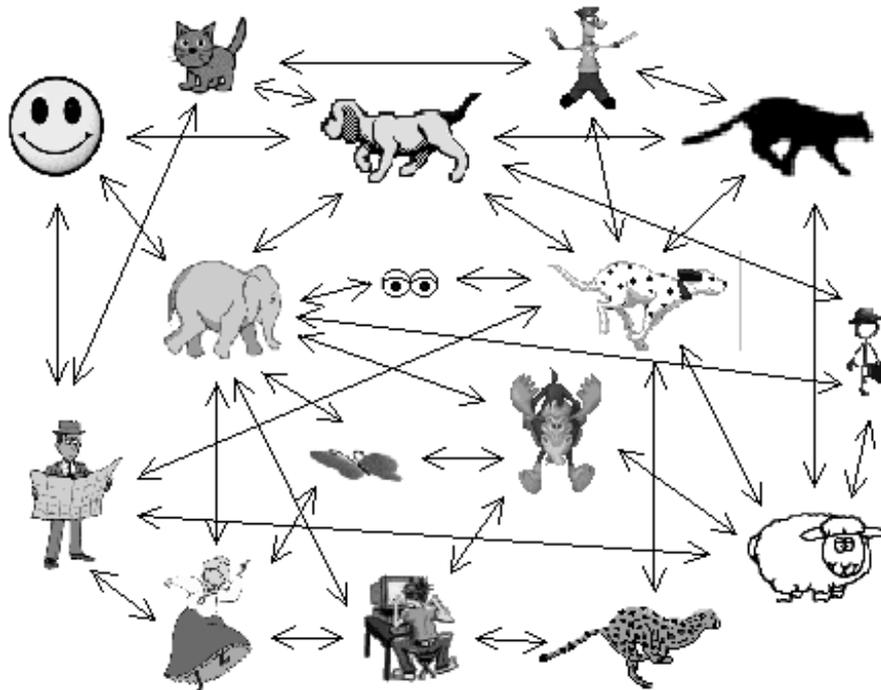


Figure 1. Agent-agent interactions

When a help request is issued, the personal agent of the helpee interacts with the agents of potential helpers, and agents who have interacted with the agents of potential helpers, etc. Thus Forgetful is able to assess the suitability of Greer by direct interaction with Stropky, but can also find further information by interacting with an agent (Hound) who has interacted with Stropky, as Hound's owner might have evaluated Greer, and Hound will therefore carry some viewpoint on him. This should obtain a more balanced view of Greer. (Stropky may not have been giving an accurate description if Greer had falsely claimed his own eagerness or knowledge level to be high – for example, to earn more currency.) Similarly, Forgetful can solicit additional information about Vassileva also by interacting with Hound. In the case of Self, Forgetful can obtain indirect information from Dancer and Moose, who have communicated with Hacknowledge. Since Bull was banned, Hairy is avoided and no agents are sought who have information about Bull.

The situation is, in fact, rather more complex than the above, as Forgetful is trying to find out as much information from as many different sources as possible, to find someone to help Muddled McCalla. Thus, not only will he be talking to Hound to find out indirectly about Stropky's owner, but he will also contact Dancer, who has communicated with Newsreader, who has communicated with Stropky, and so on. In this manner, Forgetful can reconcile all the different viewpoints on Greer, before making his decision. It is at this point that the relevant fragments of the user models of potential helpers are computed and then compared.

The above indicates that the data in Table 3 is only one viewpoint on the helper user models: the manifestation of the models from the viewpoint of Forgetful, based on Forgetful's assessment of the relative reliability of the various agent sources of information on each helper. Table 4 illustrates this, showing various agents' models of Greer, which need to be reconciled by Forgetful.

Table 4. Greer's attributes as modelled by various agents (agents' estimates of Greer's attributes)

<i>Agent</i>	helpfulness	eagerness	readiness	cog. style	knowledge	currency
* I-Help	6.24	10	7.2	W-V	6.84	-
* Stropky	10	10	-	W-V	10	10
Hound	5.3	10	-	-	9.12	-
Newsreader	6	8.2	-	-	5.22	-
Hairy	6	8.2	-	-	6	-

Scores are based on a 10-point scale with 10 being highest.

I-Help's own global assessment of Greer's helpfulness is based partly on peer evaluations subsequent to a help session, and partly on voting on the utility of his postings in the public help forum. His eagerness is measured by his response rate in both the one-on-one component and the public help forum. Readiness indicates whether he is currently online, or whether his login pattern predicts that he will be online again soon. Cognitive style information is procured through a short questionnaire, where he was identified as a wholist-verbaliser. Knowledge level is obtained from self-report and peer evaluations. Currency is not represented by I-Help, since it may fluctuate.

Stropky, Greer's agent, while holding I-Help's data on Greer, also maintains Greer's own contributions to his user model. It can be seen that, in comparison to I-Help's data, he has overestimated himself in some categories.

The remaining agents are those of other users.

Hound has interacted with Stropky in the past, and Hound's owner (du Boulay) has been involved in help sessions with Greer. Hound therefore also has data relating to his (and du Boulay's) opinion of Greer. Du Boulay clearly rates Greer's knowledge and eagerness highly, since in previous interactions Greer demonstrated that he understands the topic well, and he always accepted a help request from du Boulay promptly. However, du Boulay does not consider him particularly helpful.

Newsreader (Brna's agent) has also had contact with Stropky, and so similarly holds representations concerning Greer. Newsreader is less enthusiastic about Greer than is Hound.

Hairy, Bull's agent, has had no direct communication with Stroppey, but has interacted with both Hound and Newsreader, and has from them obtained information on Greer. Hairy has interacted more frequently with Newsreader than with Hound, and has observed the reliability of Newsreader's reports in general. Thus Hairy gives a higher weighting to Newsreader's ratings, since they conflict with Hound's.

Forgetful must consider all the viewpoints he is able to obtain on Greer, taking into account also the source of the information. Forgetful might decide to disregard Hairy's ratings, since Bull was banned. Alternatively, it might be considered that while no direct contact between McCalla and Bull is desirable, Bull's evaluations may nevertheless be useful, and Hairy's data might provide a valuable insight into conflicting representations. This must be decided by Forgetful.

Of course, agent interactions may be even more complex, since Spot has also communicated with both Newsreader and Hound. Moreover, Code Cop has interacted with both Hairy and Spot, and might be able to offer further insight into conflicting representations. Sheepish has interacted with many different agents, and may have second and third hand (and so on) knowledge from a variety of sources.

The above is only a small excerpt of the work of the personal agent. Courses where I-Help is available have between 40 and 350 students, with most courses having around 100 participants. It can be seen that the I-Help personal agent works very hard to reconcile diverse and fragmented user model information, to accommodate the wishes and needs of its owner, and probably gets as close as is possible to a caring relationship, within the current technology.

IMPLICATIONS OF CARING PERSONAL AGENTS

We have shown how I-Help agents do things for their learners such as finding appropriate helpers and relevant information. We have also shown how interesting and subtle the information flows are in the I-Help agent system. In this section we ask the basic question: what does it mean to have a caring agent? We follow this with a discussion of the implications and issues associated with building such an agent.

A caring agent is an agent that "does the right thing for a person". This does not mean that the agent always does what the person wants: people do not always know what is best for them. It does not mean that the agent always does the same thing for the person: different situations can lead to different actions. It does not even mean that the agent has to do something: there are occasions when doing nothing or having the person do something for themselves is the way the agent can be most caring (especially in learning contexts). In short, caring is adapting to the person's needs in a given situation.

The key to such adaptation is that the agent must understand their person as deeply and broadly as possible. Ideally, the agent should keep a model of their person's knowledge, goals, personality attributes, cognitive styles, past activities, etc. That is, as in I-Help, there is much beyond "just knowledge" that should be tracked. In a given situation, these long term characteristics should be integrated with the person's immediate task, time constraints, and emotional state in order to compute the best (i.e. most caring) course of action (or inaction) for the person. In determining the best course of action the agent must have access to another model describing how circumstances relate to actions.

In traditional user modelling much effort has been directed at long term modelling, especially in keeping the models consistent and accurate. Much effort has also gone into instructional planning as a way of relating circumstances to actions. In order to be useful, however, such efforts have had to be applied to highly focused domains, such as educational domains. In an educational domain the learner is willing to submit to the constraints of the subject being learned and is willing to take guidance from the ITS. The ITS, for its part, can take advantage of the restrictions in the educational situation to capture enough knowledge about the subject and to understand enough about the learner to be appropriately adaptive to the learner's needs.

It is becoming obvious, however, that interactive technology today, including learning technology such as our I-Help system, is much more open ended than were closed systems such as the classical ITS. The implications of these open ended environments, as discussed in McCalla (2000), is that both the environments and the modelling in these environments must be fragmented. As in I-Help, models must often be computed “just in time” based on available information, resource constraints, and current needs. The need to maintain consistency and accuracy in a long term model is not as critical as the need for short term modelling that is efficient and effective in the context of the person’s immediate purposes and goals. We discuss the broader characteristics of such “active learner modelling” in McCalla et al. (2000) and Vassileva et al. (forthcoming). Several on-going research projects in our laboratory are further exploring the implications of, and techniques for, active learner modelling.

Regardless of the mix of long term modelling and short term active modelling, it is important for an agent to know as much as possible about the person being modelled, at least if it is to be a caring agent. As people spend more time using their computer, and as new technology such as eye-tracking hardware and new techniques such as data mining are evolved, the potential bandwidth of interaction with the learner is much wider than in the past. Wide bandwidth, however, can never in and of itself guarantee that the modelling has captured the learner’s characteristics with fidelity, especially the learner’s immediate goals, crucial to active learner modelling. Learners must be willing to correct infidelities in the agent’s model, and to tell their agent about their goals, their problems, their time constraints, and even their emotional state. In effect, the learner must care for their agent as much as the agent cares for the learner. In order to do this, the agent’s long term and actively computed short term models must be open to the learner so that infidelities can be observed and gaps in the agent’s knowledge can be filled in by the learner. Not only should the agent display the propositions in its model, but it should also be open about any inferences it has drawn and how it has used (or will use) the information in the model(s). It should reveal which other agents it has interacted with and what, if any, information it has shared with those agents. Thus, openness is not only important pedagogically (as suggested in a number of studies (e.g. Bull & Pain, 1995; Dimitrova et al., 2001; Mitrovic & Martin, 2000), but also is important for effective and caring learner modelling.

For somebody to tell a personal agent about their goals, their problems, their time constraints, and even their emotional states, they must fully trust the agent to use the information appropriately, and especially to guard their privacy by keeping the information confidential. In fact, as Kobsa (2001) discusses, such confidentiality must be guaranteed if user modelling systems are to conform to privacy laws in many countries. Given that, as demonstrated in I-Help, agents frequently interact with other agents and can exchange information with these agents, it is not at all obvious that such privacy can be guaranteed. Even if prohibited from revealing any information about a person to other agents, the actions of an agent (for example its bargaining stances in agent-agent negotiations about the price of help in I-Help) can be observed by other agents and characteristics of the person can be inferred without any direct information exchange. Indeed, as Borking (2000) states: “The security of the data residing within the agent is only one part of the concerns regarding privacy. The arguably more significant concern is the dissemination of information during transactions, and the general conduct of the agent's activities on behalf of the user.”

Allowing learners to create aliases (as in I-Help) will help overcome this problem to some degree, since even if information about a learner is inferred by an outside party, the learner’s identity may still be protected. Another part of the solution to this problem is to allow the learner to tune an agent to achieve a desired level of privacy. The tuning should be along several dimensions. One dimension is the kind of information: some kinds of information can be disclosed while other kinds remain hidden. Another dimension is who is observing the information: certain other people (or other agents) may be allowed to see a lot of information, while others may be prohibited from observing anything.

Such tunability is important since enforcing absolute privacy at all times is not always the most caring policy for a learner. In fact, it is often vital for other people (or agents) to know about the learner so as to be able to help, e.g. to find an appropriate peer helper. Thus, there is a trade-off between total privacy and usefulness. With no privacy, a system such as I-Help would

be maximally useful, but nobody would use it; with total privacy, people wouldn't fear using the system, but it would be of limited usefulness. For most effective caring some middle ground on privacy must be negotiated between the learner and agent. The open modelling issues involved in finding this middle ground are quite interesting, and research in our laboratory is exploring issues of privacy, agent programmability, and the degree to which models of other learners can be open.

A similar trade-off exists in regard to agent autonomy. A personal agent must have at least some autonomy to do what is best for the person, even if the person would not necessarily agree with the course of action recommended. However, if the agent has too much autonomy, it will likely lose the trust of the learner. To achieve optimum caring, there must be a satisfactory middle ground where the agent is given enough autonomy to act usefully, but not so much as to destroy the learner's belief in the agent. It may be possible, as in privacy, to allow the learner to tune the level of autonomy in the agent, for example to allow independent action by the agent in some activities (e.g. negotiating a price for help), but not in others (e.g. not to allow the agent to actually contact the person selected to help without the permission of the learner). This would seem to be very difficult in general, however, since there are so many activities possible and they are not easy to categorize.

One interesting aspect of the caring personal agent is that a person might have more than one of them! That is, it is perfectly feasible to have different personal agents that play different roles. These roles may be specialized around different tasks (e.g. a personal agent to find information and another personal agent to find a peer helper), or different levels of privacy (e.g. a secure agent that can be called in for some situations and an open agent for others), or different groups (e.g. one agent for one class and another agent for a different class), or different perspectives on the world (e.g. one agent might make decisions based on knowledge levels and another on affective states). Perhaps in the future, true caring will be enforced by a horde of personal agents, each of them each looking out for the learner's best interests, albeit from different perspectives.

FUTURE WORK

Self's notion of systems that (who) care focused on personalisation through modelling the user. Systems capable of providing personalised assistance to learners can be thought to "care" about the individual learner. The caring personal agent is one manifestation of this sort of caring system. Our work in the I-Help project, which aims to develop software agents that care about their owners, opens many interesting issues for future research. A number of these future research directions have been identified for new projects in our laboratory, including work on active learner modelling, learner privacy/learner model security, agent programmability, agent trust, agent personas, and social relationships between people and agents.

Interesting future work in active learner modelling involves gaining a better understanding about learner goals, learning purposes and different contexts in which learning can occur. Developing techniques by which to aggregate bits of knowledge having varying reliability from distributed agents is also an interesting challenge. The decision about what information is most relevant to a particular learning purpose in a particular context, where to locate that evidence, and how to assemble it to make a timely decision comprises a complex decision problem.

Privacy and security in multi-agent learner modelling also poses some interesting future challenges. Different individuals have different preferences with respect to their privacy. Opening learner models (either fully or partially) to inspection by other learners or by the agents of other learners requires rules of information disclosure. More interesting is regulating information flow about events that occur in the multi-agent environment. Restricting knowledge about certain events to certain agents can introduce ways to simulate perceptual limitations, locality of information and proximity of one agent to another.

Future work in agent programmability will focus on developing mechanisms by which learners can communicate their wishes to their personal agent and verify that their agent in fact carries out the appointed tasks. Tasks might range from simple notifications of relevant events

or the appearance of important new information in the system to the management of interruptions and schedules. Agents could be programmed to watch for events triggered by other agents, thus monitoring various kinds of activities of other users (subject to privacy-based regulation).

Work with agent personas/personalities is underway in many research laboratories (e.g. Andre et al., 2000; Johnson et al., 2000; Loyall & Bates, 1997). Our future work involves developing agent personalities and personas so that a more natural style of communication can exist between learner and personal agent. Trust between learner and agent is fundamental to the development of high-fidelity learner models. In addition, caring for the agent (by the learner) is also fundamental to the successful operation of environments like I-Help. Agents must be kept informed, kept busy at interesting tasks, and permitted to monitor activities of their person in order that they can do an effective job of modelling the learner. Ideally a mutual respect (mutually caring relationship) may result between agent and learner. While virtual friends may not be a new phenomenon for children, they may be the source of some controversy among older learners.

In all, the caring agent seems to provide an opportunity for adaptive learning environments to become better at adapting to learner needs and better at supporting the learning process. Furthermore, we believe that many new avenues of research will be opened by this work on caring personal agents.

ACKNOWLEDGEMENTS

This work was funded by the Canadian TeleLearning Network of Centres of Excellence. I-Help was developed in the ARIES Lab, Department of Computer Science, University of Saskatchewan. We thank the many people who have been involved in the I-Help project over the years. We also thank the many I-Help agents. Most of the I-Help agent pictures were taken from the following image websites:

<http://www.fg-a.com/aliens.htm>; <http://www.gifs.net/animate/giflist.htm>.

REFERENCES

- Andre, E., Klesen, M., Gebhard, P., Allen, S. & Rist, T. (2000). Exploiting Models of Personality and Emotions to Control the Behavior of Animated Interface Agents, in Jeff Rickel (ed), *Proceedings of the Workshop on Achieving Human-Like Behavior in Interactive Animated Agents*, Fourth International Conference on Autonomous Agents, Barcelona, 3-7.
- Borking, J.J. (2000). *Privacy Incorporated Software Agent (PISA): Proposal for Building a Privacy Guardian for the Electronic Age*, <http://www.datenschutz.de/files/pisa.htm>.
- Bull, S. (1997). A Multiple Student and User Modelling System for Peer Interaction, in R. Schäfer & M. Bauer (eds), *5 GI-Workshop, Adaptivität und Benutzermodellierung in interaktiven Softwaresystemen*, Universität des Saarlandes, 61-71.
- 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*, AACE, 501-508.
- Bull, S., Greer, J., McCalla, G. & Kettel, L. (2001a). Help-Seeking in an Asynchronous Help Forum, in R. Luckin & B. du Boulay (eds), *Proceedings of Workshop on Help Provision and Help Seeking in Interactive Learning Environments*, International Conference on Artificial Intelligence in Education, 9-21.
Online proceedings: <http://www.cogs.susx.ac.uk/users/bend/aied2001/bull.pdf>
- Bull, S., Greer, J., McCalla, G., Kettel, L. & Bowes, J. (2001b). User Modelling in I-Help: What, Why, When and How, in M. Bauer, P.J. Gmytrasiewicz & J. Vassileva (eds), *User Modeling 2001*, Springer-Verlag, Berlin Heidelberg, 117-126.
- Bull, S. & McCalla, G. (in press). Modelling Cognitive Style in a Peer Help Network, *Instructional Science*.
- Deters, R. (2001). Scalability and Information Agents, *Applied Computing Review* 9(3), 13-20.

- Dimitrova, V., Self, J. & Brna, P. (1999) Applying Interactive Open Learner Models to Learning Technical Terminology, in M. Bauer, P.J. Gmytrasiewicz & J. Vassileva (eds), *User Modeling: 8th International Conference*, Springer Verlag, Berlin Heidelberg, 148-157.
- Galloway, L.L. & Easterday, K.E. (1998). Help-Givers in Developmental Studies Mathematics: Whom Do the Students Prefer?, in P.L. Dwinell & J.L. Higbee (eds), *Developmental Education: Meeting Diverse Student Needs*, National Association for Developmental Education, Morrow GA, 21-28.
- Greer, J., McCalla, G., Cooke, J., Collins, J., Kumar, V., Bishop, A. & Vassileva, J. (1998). The Intelligent Helpdesk: Supporting Peer Help in a University Course, in B.P. Goettl, H.M.Halff, C.L. Redfield, & V.J. Shute (eds), *Intelligent Tutoring Systems*, Springer-Verlag, Berlin Heidelberg, 494-503.
- Greer, J., McCalla, G., Vassileva, J., Deters, R., Bull, S. & Kettel, L. (2001). Lessons Learned in Deploying a Multi-Agent Learning Support System: The I-Help Experience, in J.D. Moore, C.L. Redfield & W.L. Johnson (eds), *Artificial Intelligence in Education*, IOS Press, Amsterdam, 410-421.
- Hoppe, H.U. (1995). The Use of Multiple Student Modelling to Parameterize Group Learning, in J. Greer (ed), *Proceedings of World Conference on Artificial Intelligence in Education*, AACE, Charlottesville VA, 234-241.
- 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.
- Karabenick, S.A. (1990). When Students Need Help, *The Journal of Professional Studies*, Winter 1990, 41-54.
- Karabenick, S.A. & Knapp, J.R. (1988). Help Seeking and the Need for Academic Assistance, *Journal of Educational Psychology* 80(3), 406-408.
- Karabenick, S.A. & Knapp, J.R. (1991). Relationship of Academic Help Seeking to the Use of Learning Strategies and Other Instrumental Achievement Behavior in College Students, *Journal of Educational Psychology* 83(2), 221-230.
- Karabenick, S.A. & Sharma, R. (1994). Perceived Teacher Support of Student Questioning in the College Classroom: Its Relation to Student Characteristics and Role in the Classroom Questioning Process, *Journal of Educational Psychology* 86(1), 90-103.
- Knapp, J.R. & Karabenick, S.A. (1988). Incidence of Formal and Informal Academic Help-Seeking in Higher Education, *Journal of College Student Development* 29, 223-227.
- Kobsa, A. (2001). Tailoring Privacy to Users' Needs, in M. Bauer, P.J. Gmytrasiewicz & J. Vassileva (eds), *User Modeling 2001*, Springer-Verlag, Berlin Heidelberg, 303-313.
- Loyall, B. & Bates, J. (1997). Personality-Rich Believable Agents that Use Language, Proceedings of the First International Conference on Autonomous Agents, Marina del Rey, California.
- McCalla, G. (2000). The Fragmentation of Culture, Learning, Teaching and Technology: Implications for the Artificial Intelligence in Education Research Agenda in 2010. *International Journal of Artificial Intelligence in Education* 11(2), 177-196.
- McCalla, G., Vassileva, J., Greer, J. & Bull, S. (2000). Active Learner Modelling, in G. Gauthier, C. Frasson & K. VanLehn (eds), *Intelligent Tutoring Systems*, Springer-Verlag, Berlin Heidelberg, 53-62.
- Mitrovic, A. & Martin, B. (2002). Evaluating the Effects of Open Student Models on Learning, in P. de Bra, P. Brusilovsky & R. Conejo (eds), *Proceedings of Adaptive Hypermedia and Adaptive Web-Based Systems*, Springer-Verlag, Berlin Heidelberg, 296-305.
- 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.
- Moncada, S.M. & Sanders, J.C. (2000). Promoting Academic Help-Seeking in Accounting Courses, *Accounting Instructors' Report* 23, http://college.hmco.com/AIR/spring_2000/sanders.html.
- Ogata, H., Sueda, T., Furugori, N. & Yano, Y. (1999). Augmenting Collaboration Beyond Classrooms through Online Social Networks, in G. Cumming, T. Okamoto & L. Gomez (eds), *Advanced Research in Computers and Communications in Education – Proceedings of ICCE '99*, IOS Press, Amsterdam, 277-284.
- Perrine, R.M., Lisle, J., & Tucker, D.L. (1995). Effects of a Syllabus offer of Help, Student Age, and Class Size on College Students' Willingness to Seek Support from Faculty, *The Journal of Experimental Education* 64(1), 41-52.
- Riding, R. & Cheema, I. (1991). Cognitive Style – an Overview and Integration, *Educational Psychology* 11(3-4), 193-215.

- Self, J. (1999). The Defining Characteristics of Intelligent Tutoring Systems Research: ITSs Care, Precisely. *International Journal of Artificial Intelligence in Education* 10(3-4), 350-364.
- Shapiro, E.G. (1983). Embarrassment and Help-Seeking, in B.M. DePaulo, A. Nadler & J.D. Fisher (eds), *New Directions in Helping Volume 2: Help-Seeking*, Academic Press, New York, London, 143-163.
- Thibodeau, M-A., Belanger, S. & Frasson, C. (2000). WHITE RABBIT – Matchmaking of User Profiles Based on Discussion Analysis Using Intelligent Agents, in G. Gauthier, C. Frasson & K. VanLehn (eds), *Intelligent Tutoring Systems*, Springer-Verlag, Berlin Heidelberg, 113-122.
- VanLehn, K. (1988). Student Modeling, in M.C. Polson & J.J. Richardson (eds), *Intelligent Tutoring Systems*, Lawrence Erlbaum Associates, Hillsdale NJ, 55-78.
- Vassileva, J., Greer, J., McCalla, G., Deters, R., Zapata, D., Mudgal, C. & Grant, S. (1999). A Multi-Agent Design of a Peer-Help Environment, in S. Lajoie & M. Vivet (eds), *Artificial Intelligence in Education*, IOS Press, Amsterdam, 38-45.
- Vassileva, J., McCalla, G., & Greer, J. (forthcoming). Multi-Agent Multi-User Modelling in I-Help, *User Modelling and User Adapted Interaction*.
- Vivacqua, A.S. (1999). *Agents for Expertise Location*, Technical Report SS-99-03, MIT Media Lab.