Openness and Disclosure in Multi-Agent Learner Models

Julita Vassileva, Jim Greer, and Gordon McCalla

ARIES Lab
Department of Computer Science
University of Saskatchewan
Saskatoon, Saskatchewan, S7N 5A9
CANADA
jiv@cs.usask.ca, greer@cs.usask.ca, mccalla@cs.usask.ca

Abstract

Multi-agent systems raise many new and interesting research issues in AI-based educational collaborative systems. Some of those issues relate to the user model. When many agents inhabit a system, they must cooperate in their agent society, interacting with each other. They must interact with the world outside the computer (i.e. with the learners). They must have adequate knowledge of the users they represent. These new issues also manifest themselves in new issues for open and inspectable learner models. This paper presents some of these issues in the context of the I-Help project.

Introduction

There are a number of reasons for adopting a multi-agent approach to developing AI-based educational systems. The rapid development of new technologies like telecommunications, networking, and mobility leads to emerging environments that can be viewed as "hybrid societies", consisting of real persons and electronic agents. In these environments everyone and everything is connected, so enormous possibilities for sharing resources (computational, applications, human advice etc.) emerge. In order to allow the users to benefit from these possibilities, it is necessary to reduce complexity for them and allow them to concentrate on their primary goals or tasks, as well as support their learning and collaboration. Also, it becomes very important to provide a source of motivation for goodwill and collaboration among users, along with protection mechanisms, ensuring security of communications, privacy of personal data, and equal chances for all to participate. The field of multi-agent systems opens a number of extremely interesting and potentially useful research avenues concerning inter-agent negotiation, persuasion and competition in agent societies. Human societies provide a rich source of paradigms and assumptions for modelling multi-agent systems, e.g. social roles and cultural values, norms and conventions, social movements and institutions, power and dominance distribution.

A multi-agent system involving personal agents and application agents (I-Help)

We have developed peer help environment called I-Help (Greer et al., 1998) which attempts to provide seamless access for students to a variety of distributed help resources

(human resources, e.g. peer help, as well as electronic resources, e.g. threads in discussion forum, FAQ entries, web-resources). The environment is based on a free market economic model. This environment can be viewed as a special case of electronic market because of the following features:

- There are people who possess some goods or resources and those who need or want these goods or resources In our case help (or knowledge) is the resource.
- A person wishing to buy a good must find a seller who offers the good with acceptable quality and at acceptable price. In our case a learner with a specific help request needs to find a competent helper;
- The buyer is willing to pay some amount of money in order to purchase the good and the seller agrees to exchange the good for the payment offered. In our case the learner wishes to achieve the goal of obtaining some help or knowledge, while the helper is willing to give advice or help in exchange for some form of currency. The goal of accumulating currency, which can be exchanged with some other goods (perhaps favourable performance reviews in a workplace environment, or perhaps marks in University environment) creates a motivation for knowledgeable helpers to participate.
- The price of a certain good depends on the offer and the demand for this good on the market. People having exceptional and highly demanded knowledge / expertise, can put higher prices for their advice.
- There is some cost associated with supplying the buyer with the good; helping costs some time for the helper, which could be used for achieving some other goal.

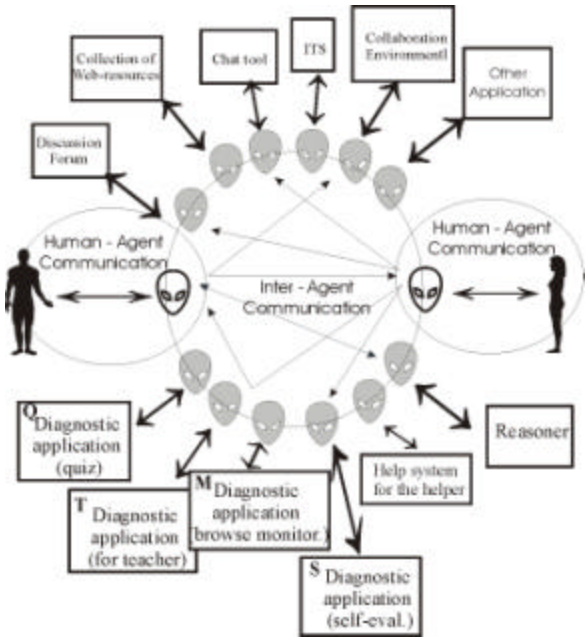

Figure 1: The multi-agent Architecture of I-Help

The adaptation within the multi-agent I-Help system is based on models of human users and models of involved software applications. These are maintained by two classes of agents (see Figure 1): personal agents (of human users) and application agents (of software applications). These agents use a common ontology and communication language. Each agent manages specific resources of the user (or application) it represents,

including for example, the knowledge resources of the user about certain tasks or concepts, or the instructional materials belonging to an application. The agents use their resources to achieve the goals of their users, their own goals, and goals of other agents. Thus all the agents are *autonomous* and *goal-driven*. The agents are *cognitive,* i.e., they can plan the achievement of goals by means of decomposing them into sub-goals and relating them to resources. In their goal pursuit the agents can also use resources borrowed from other agents, i.e. they are *collaborative.* For this they have to negotiate and become involved in persuasion and conflict resolution. Each agent possesses a model of its inter-agent relationships, some of which reflect relationships between human users. Finally, the agents are *mobile* i.e. they can travel from one computer to another, thus optimising resources and bandwidth (Deters, 1999). In this way, we achieve a complex (multi-user, multi-application) adaptive (self-organised) system that supports users in locating and using resources (other users, applications, and information) to achieve their goals. A relatively detailed description of the multi-agent architecture we are using is presented in (Vassileva, 1998; Vassileva et al, 1999).

Each personal agent manages a user model containing information about the *user's goals* (help requests, current goal), about *knowledge resources / competencies* on certain topics or tasks, and about the *relationships* existing between the user and other users. The users communicate with their agents to update their models (actively or passively) and to assert goals for their agent to adopt. The agents communicate with each other and with matchmaker agents to search for appropriate helpers for their users, depending on the topic of the help-request. Once a potential helper is located, the agents then negotiate the price for help. Help is arranged (negotiated) by agents among themselves - in this way the personal agents form an artificial society. It involves various levels of organization, including the negotiation between agents (Mudgal & Vassileva, 1999), an economical model of the society (Kostuik & Vassileva, 1999) and control / policing institutions within the society (Winter, 1999).

Agents' involvement in learner modelling

The personal agents in I-Help maintain models of the goals, resources, and relationships of the users they represent, that is the *user models* in the system. Currently, only two types of goals are represented: the current goal (whatever it might be), and a goal corresponding to a help request. The users' resources are the time available for help (given by the user at login), an amount of currency (fictitious, distributed evenly at the beginning) and the user's knowledge/skill/competence on various topics. The relationships represent real-world relationships between users and their friends, colleagues, or peer helpers or helpees. Users assign an importance factor to current goals, a value to their resources, the cost associated with the consumption of resources, and the importance of each relationship. When a learner needs help (i.e. issues a help-request), that learner's personal agent tries to locate an agent of another user who possesses appropriate resources for achieving the goal of the help request. For example, if a user asks for help about topic X, the personal agent will try to find an agent of a user who has knowledge on topic X and to negotiate with it for a help-session between the users. The personal agent of the potential helper has to check if its user has enough time and no conflicting goals with higher importance, in order to decide whether to adopt the help request and to pass it on to the attention of its user. In this way only important help requests, or requests from people involved in important relationships or offering a high enough payment (which the user is likely to accept) will be passed to the user.

Both the agent of the learner asking for help and the agent of the potential helper make consecutive decisions during the process of negotiation. The helpee agent needs to decide

what amount of currency it is willing to pay for the help; the helper agent needs to evaluate the importance of the help-request, the importance of the relationship, and the amount of payment offered to decide whether it is worth having the user abandon his or her current goal to give help. The representation of the decision problem as an influence diagram enables the agents to analyse the decision situation and find out the most appropriate action that will maximise their utility. Utility is represented as a real valued number that describes how preferable or desirable is the alternative for the agent. Evaluation of influence diagrams is done to choose an action for the decision (i.e. whether to accept an offer, to reject it, or to counter-propose it).

Aspects of openness in the learner models

There are many aspects of the I-Help architecture where openness of a user model are important. First, *the user model is open to its user to be inspected, modified, and initialized*. This type of openness is similar to the one proposed by Paiva (Paiva et al., 1995) and Bull (Bull et al, 1995). The user model contains various types of information about the user, employed for different purposes. It is more important to open some types of data to the user than others. For example, the knowledge resources of a user are important for being selected as a peer helper when some help-request is issued. However, this may not be critical information for the user -- the worst thing that can happen is that the agent of the user is too infrequently approached to provide help, or sometimes called upon to offer help on a topic not known to the user. Since the decision of whether to pass a help request to the user is taken by the personal agent after considering many factors, the danger that the incorrectess of the user model will overload / underload the user is not that big. Still, if the user feels that it is necessary to adjust the cognitive part of the model, the agent can provide an interface for the user to inspect and modify the contents of their competence model, or at least the parts relevant to the current help-request (see Greer et al., submitted).

A second aspect of openness of the user model concerns more dynamic parameters, i.e. *the importance of the user's current goal and the amount of time the user has*. These parameters are dynamic and hard to diagnose automatically. Currently the user directly enters them at login time and adjusts them whenever they need to be updated. In some circumstances, we expect that the personal agent can automatically update these parameters from observations of user behaviour (for example, if the user systematically refuses to help).

A third aspect of openness concerns *the relationships between the user and other users*. We assume that people are generally more inclined to help their friends, or people with whom they have had already useful contacts. That is why one part of the user model is focussed on representing relationships the learner has with other users. These relationships are represented as a set of parameters, denoting the relationship's importance, sign, symmetry and context. The user can enter new relationships manually through an appropriate interface, and the personal agent can extend the set of relationships with "useful contacts" that it has detected (for example, after a successful help-session). The user can view and modify the relationships, delete relationships altogether, or change parameters.

A fourth aspect of openness concerns *the intrinsic importance of the user's own goals* versus the goals that might be adopted from other users *and the importance of certain resources*, like the resource "time", "money" or the importance of relationships with other users. Since every personal agent negotiates (i.e. buys and sells help) on behalf of its user, it has to be "instructed" by the user about how to prefer these parameters, so that the

personal agent can take decisions according to the user's priorities. In this way, the user defines a "character" for its personal agent, for example an "ego-centric" character that gives high priority to its own goals, a "greedy" character for whom currency is important, a "social" character for whom relationship with other users/agents is important, an "altruistic" character that will readily adopt goals of other agents, etc.. This "character" does not necessarily reflect the real "character" of the user. A user can adjust the character of his or her personal agent as desired (e.g. a modest person can hire a very aggressive lawyer). The "character" is defined by a set of variables corresponding to the weights a personal agent will attach to the corresponding resources in decision making during the negotiation process. In a sense, these weights determine the shape of the decision model of every agent. For user convenience they are visualized as sliders and can be adjusted by the user as desired.

A fifth aspect of openness concerns the personal agents when they enter into negotiation about setting up a help session. At the beginning of the negotiation, the *agents for the helpee and the potential helper share important information* that will provide the context for subsequent decisions. This information differs depending on how the personal agents have been instructed by their users (i.e. which information to reveal and which to keep private). For example, the personal agent of the potential helper may be willing to reveal that its user has currently a very important goal and very little time -- this will determine a higher starting price in the negotiation. Users are unlikely to authorize their personal agents to reveal other information, like the importance of the relationship (unless it is a very positive relationship with a friend which both users are aware of and the user has explicitly allowed the personal agent to reveal), the importance of money (greediness), or the state of the user's money resource.

It is possible *that after a successful negotiation the personal agents share more information about their user's resources each providing a view of the other user's knowledge resources, focussed on the topic of the help request*. This knowledge can then be passed on to the user. This capability is particularly useful to allow the helper to get a glimpse of a part of the cognitive model of the helpee. This may help the helper understand better the helpee's problem and tailor an explanation accordingly. This, then, is a sixth aspect of openness of the user models: the user model can be made transparent (to some degree) not just to the user, but to other people. Such transparency may also be useful before the negotiation phase, so that the helpee can see many user models of potential helpers before directing their personal agent to negotiate with one of them. Some of work done in the ARIES Lab (Kumar, 1997; Kumar et al., 1999) addresses this possibility. Such openness can be provided by a broker agent, who receives from the personal agents periodic updates about the knowledge state of their users on different topics. Then the broker agent *can model the competencies of the group of users* and visualize each user as a "knowledge point /vector" in a concept/ topic space. This will create something like a "knowledge map" of the group, (something similar to a shopping mall map) which will help users orient themselves in their "electronic village" (McCalla, submitted).

Since negotiation is a process of interaction, it is natural that during this process, each participant builds a model of its opponent. In our case *the personal agents develop and store models of every other personal agent* that they have encountered and negotiated with. These are models of the other personal agents' decision making process, i.e. the "character" of every other personal agent. Each personal agent tries to improve the model of its opponent from the history of the negotiation process, after the process is over, thus learning about the "characters" of the personal agents it is talking to. There are some

benefits to this (mainly in optimizing negotiation strategies for the personal agents) and also some drawbacks (it can be very expensive, if users change the "characters" of their personal agents frequently; moreover, too much space may be needed if a personal agent keeps a model of every opponent that it has encountered). Since we don't see a direct relevance of this type of openness (personal agents modelling other personal agents) to improving the processes of learning or peer-help, we won't discuss this here any further. However, it provides for potential improvements in the "deals" between personal agents, in terms of improving the personal agents' utility (the quality of help negotiated for a certain price).

Conclusion

We feel our research makes a number of contributions, both in the kind of user modelling we do and in the role of openness in user modelling. First, our approach is a distributed and procedural approach, assigning the responsibility for user modelling to active personal agents representing each user in the system. Each user's personal agent can thus fine tune its representation and reasoning to reflect the particular needs and wishes of its own user. Both cognitive aspects (eg. knowledge of the subject being learned) and social aspects (eg. knowledge of personal relationships and past helpfulness) can be naturally captured and used. The knowledge in the user models can be employed internally by the personal agents in negotiating appropriate peer helpers to handle a help request, as well as externally by the users themselves to better understand a help request.

The user models are thus open to their users. Openness of a user's own user model allows the user to achieve several goals. One of these is the traditional goal (espoused by Kay, 1999, and others) of ensuring the user can check the correctness of the model. It also allows the user to inform the personal agent of changing goals. Perhaps the most innovative capability this provides for a user, however, is that it allows the user to create an alter ego through which that user can interact with other users in the overall distributed system. This alter ego can be changed as needed. All the interactions between a user and their own personal agent about the user model can be contextualized; that is, only the parts of the user model of importance to the user's current goals and problems need to be "in focus".

As well as allowing the user to interact with their own user model, in certain circumstances parts of a user model can be open to other users, again focussed around the particular issues of current relevance to a help request. This can be useful to the helpee in determining which of several possible helpers suggested by the system should be chosen. This can also be useful to the helper in examining the nature of the helpee's problem and formulating the help in terms appropriate to the helpee. Making available a user model to an outsider sets up interesting questions such as which parts of the user model to reveal (in our case largely cognitive parts), how a personal agent should interact differently with a third party than with its own user, and whether an appropriate stance can be maintained on privacy issues and confidentiality of user models. These questions, and other aspects of our agent-based approach to user modelling and intelligent help, are under active investigation in a number of projects.

Acknowledgements

We would like to thank the Canadian TeleLearning Networks of Centres of Excellence project and the Natural Sciences and Engineering Research Council of Canada for funding this research. This paper was written while Gord McCalla was a Visiting Scholar at the Basser Department of Computer Science, University of Sydney, Australia.

References

Bull, S., Brna, P., and Pain, H. (1995) Extending the scope of the student model. User Modeling and User Adapted Interaction 5: 45-65.

Deters, R (1999) Achieving Dependability in a System Consisting of Autonomous Agents in Proceedings of the Workshop on Autonomy Control Software, held in association with the Autonomous Agents 99 Conference, Seattle, 1-5 May 1999.

Greer J.E., McCalla G.I., Cooke J., Collins J., Kumar V.S., Bishop A., & Vassileva J.I. (1998). The Intelligent Helpdesk: Supporting Peer-Help in a University Course. The International Conference on Intelligent Tutoring Systems. San Antonio, TX, USA, 494-503.

Greer, J.E., Zapata-Rivera, J.D, Ong-Scutchings, C. and Cooke, J.E. (submitted) "Visualization of Bayesian Learner Models" Workshop on Open, Interactive, and Other Overt Approaches to Learner Modelling, International Conference on AI in Education, Le Mans, France.

Kay, J. (1999) A Scrutable User Modelling Shell for User-Adapted Interaction, Ph.D. Thesis, Basser Department of Computer Science, University of Sydney, Sydney, Australia.

Kostuik, K. Vassileva, J. (1999) Free Market Control for a Multi-Agent Based Peer Help Environment in Proceedings of the Workshop on Agent based Decision-Support for Managing the Internet-Enabled Supply-Chain, held in association with the Autonomous Agents 99 Conference, Seattle, 1-5 May 1999.

Kumar V. (1997) CSCL Environment for Distributed Collaboration, presented at the Doctoral Consortium of the *Conference of Computer Supported Collaborative Learning CSCL'97*, Toronto, December, 1997.

Kumar, V.S., McCalla, G.I. and Greer. J.E. (1999) "Helping the Peer Helper" Proceedings of the International Conference on AI in Education, Le Mans, France, (to appear).

McCalla, G. (submitted) 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.

Mudgal Ch., Vassileva J. (1999) Negotiation among Agents in a Multi-Agent Environment Supporting Peer-Help: a Decision-Theoretic Approach, in Proceedings of the Workshop on Agent based Decision-Support for Managing the Internet-Enabled Supply-Chain, held in association with the Autonomous Agents 99 Conference, Seattle, 1-5 May 1999.

Paiva, A. Self, J., and Hartley, R. (1995) Externalizing learner models. In J. Greer (Ed.) Artificial Intelligence in Education: Proceedings of AI-ED 95, Charlottesville, VA: AACE, 509-516.

Vassileva J., Greer J., McCalla G., Deters R., Mudgal C. Zapata D., Grant S. (1999) A Multi-Agent Design of a Peer-Help Environment, Proceedings of the AIED'99, Le Mans, France, (to appear).

Vassileva J. (1998) Goal-Based Autonomous Social Agents Supporting Adaptation and Teaching in a Distributed Environment, Proceedings of ITS'98, San Antonio, Texas, LNCS No1452, Springer Verlag: Berlin pp.564-573

Winter M. (1999) The Role of Trust and Security Mechanisms in an Agent-Based Peer Help System, in the Proceedings of the Workshop on Deception, Fraud and Trust in Agent Societies, held in association with the Autonomous Agents 99 Conference, Seattle, 1-5 May 1999.