The design of MOO agents: Implications from an empirical CSCW study.

P. Dillenbourg[1], P. Jermann[1], D. Schneider[1], D. Traum[2], C. Buiu[3]

[1]TECFA, School of Education and Psychology, University of Geneva, Route de Drize 9, 1227 Carouge, Switzerland.

{Pierre.Dillenbourg, Patrick.Jermann,Daniel.Schneider}@tecfa.unige.ch

[2]UMIACS, A. V. Williams Building, University of Maryland, USA, traum@cs.umd.edu

[3]Department of Computer Science and Control, University of Bucharest, Romania, cata@lnx.cib.pub.ro

Abstract. We report the results of an empirical study on computer-supported collaborative problem solving. Twenty pairs each participated in a mystery solving game in a text-based virtual environment, enriched with a whiteboard. We analyzed how subjects use different communication tools to build common ground. We present seven observations which inform the design of artificial MOO agents. We propose a new type of artificial agent, observers, which compute statistics regarding interactions and display them to human or artificial tutors or to the collaborators themselves.

1. Introduction

Collaborative learning is not always effective. The design of CSCL environments aims to set up conditions for effectiveness, namely to promote productive interactions. The environment may scaffold the pair interactions with the task and/or the interactions between learners [1]. In both cases, the scaffolding is either achieved through purposely designed interfaces [1, 2, 3] or through the action of human or artificial agent(s). The topic of artificial agents is a large and disparate one, however, we see three main categories of agents:

- *Sub-agents* are autonomous software entities which carry out tasks for the user. In current research, a user generally interacts with a single agent performing a variety of information finding tasks [4, 5] on the Internet. Other examples of tasks accomplished by sub-agents include e-mail filtering [6], room reservation, calendar managing [7] and finding people sharing interests [8].
- *Co-agents* can perform the same actions as the human user they interact with. For instance, the idea of a co-learner was used within various learning paradigms: collaborative learning [9, 10], competition, reciprocal tutoring [11] or learning by teaching [12].
- *Super-agents* provide solutions and monitor the actions of users. Typical examples are tutors in ILEs or critic systems [13]. In CSCL, super-agents monitor the interactions among users. For instance, in Belvedere [14] an advisor may point to specific parts of the learner's argument and proposes ways of extending or revising it. Some argument structures are expected to be more effective in stimulating critical discussions. In COSOFT [15], a super-agent compares its model of on-line students and invites students with specific lacks to collaborate with students who possess the missing skills.

Our work on co-agents revealed that straightforward application of current AI techniques are not appropriate to build collaborative agents [16]. We have hence conducted an empirical study to inform the design of such agents.

2. Empirical study: collaboration in a multi-modal virtual environment.

This study addressed two issues: (1) how two subjects elaborate a shared solution of a problem, (2) how does a whiteboard help to achieve mutual understanding. In the experiments, the two subjects play detectives: they have to find who committed a murder. They have to find the single suspect who has a motive to kill, had access to the murder weapon and had the opportunity to kill the victim when she was alone. They walk in a text-based virtual environment (a MOO) where they meet suspects, ask questions about relations with the victim, regarding what they have done the night of the murder, and so forth. Suspects are MOO robots, programmed to answer relevant questions. When exploring rooms, subjects find various objects which may have a role in the story.

The subjects are located in different physical places, connected by a computer network and using two pieces of software: a shared whiteboard and a standard MOO (tecfamoo.unige.ch - port 7777). The subjects talk to each other on the MOO via two commands: "say..." to communicate with anybody in the same room and "page John..." to communicate with John where ever he is. They each carry a notebook which automatically records the answer to all the questions they ask. The subjects use a MOO client called TKMOOlight which runs Sun, Macintosh, or Windows computers. It includes an elementary whiteboard: multiple users can draw on a shared page, can see and edit the objects drawn by their partner, but they do not see each other's cursor. All actions and interactions in the MOO and in the whiteboard are recorded. We ran the experiments with 20 pairs, the average time to completion was two hours. We present here only some of the main results relevant to the design of agents.

In the following, we refer to four levels of *mutuality of knowledge* [17]: (1) A knows that B can *access* knowledge X; (2) A knows that B has *perceived* X; (3) A knows that B understood X more or less as A *understood* it, and (4) A knows that B *agrees* on X. These four levels cover both the grounding process (achieving mutual understanding) and the negotiation process (reaching agreement).

3. Observations & implications

3.1 Matching knowledge and media persistence

Observations. In our analysis of MOO dialogues or whiteboard drawings, we discriminate four categories of content: *Task* **(**subdivided into *facts* and *inferences*), *management* (who does what, who goes where, etc.), *technical* problems and *meta-communication* (turn taking rules, graphical codes, etc.). The distribution of these content categories differs between MOO dialogues and whiteboard notes, as illustrated by figure 1. This difference can be explained by the persistence of information: utterances regarding task management and meta-comunication, have a short term validity, while task level information (facts and inferences) is more persistent, especially facts (if it is true at time t that Helmut is a colonel, this information will still be true at time t+1). Non-persistent categories represent 45% of MOO dialogues versus only 10% of whiteboard items. Conversely, persistent information represents 90% of whiteboard items versus 55% of dialogue utterances. Our interpretation is that subjects match the content and the medium of communication by comparing the *persistence of information* (how long a piece of information remains true) and *persistence of medium* (how long it can be viewed on a particular medium). Non-persistent items (management, communication) are preferably discussed on a non-persistent medium (the MOO scrolls up every time a new command is performed), while persistent knowledge (facts &

inferences) is displayed on a persistent medium (whiteboard items remain displayed until explicitly erased). The sensitivity to medium persistence is confirmed by another observation: the average delay of acknowledgment is shorter in MOO dialogues (48 seconds) than on the whiteboard (70 seconds). There is no urgency to acknowledge information which will (probably) remain a long time on the whiteboard.

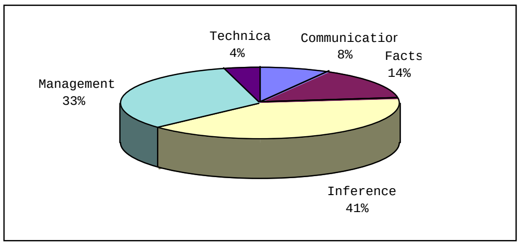

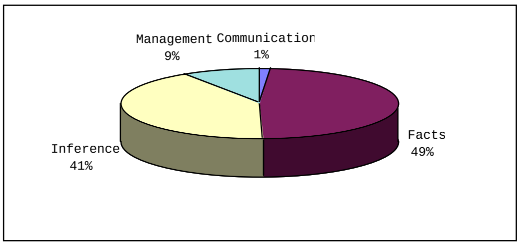

Figure 1: Categories of content in MOO (left) and whiteboard (right) interactions

Implications. An intelligent agent should select a medium of communication whose persistence is adequate to the persistence of knowledge being discussed. Negotiating non-persistent information through a persistent medium leads to display invalid information when its validity has expired. Vice-versa, negotiating persistent information through a non-persistent medium increases the memory load, since the medium does not play its role of group memory. How could an agent match the medium persistence with knowledge persistence? Medium persistence can de determined once and for all. The difficulty consists in reasoning on the knowledge persistence. This issue is not the same as in a truth maintenance system: it requires not only to detect when information is not relevant anymore, but also to anticipate how long it might remain valid. A simple solution would be to associate levels of persistence to knowledge categories: facts and inferences are more persistence than strategies, plans, and MOO positions.

3.2 Reasoning on sharedness.

Observations. We computed the percentage of utterances produced by agent-A which are acknowledged by Agent-B. This *rate of acknowledgment* for facts is 26% while it is 46% for inferences and 43% for task management utterances. The difference of rate between facts and inferences or management reflects the probability of disagreement or misunderstanding. In this task, there is little to disagree about or to misunderstand regarding facts ("A gun in room 5"), while one can of course disagree about inferences ("Heidi has a motive") or strategies. Moreover there is a significant interaction effect on the acknowledgment rate between the knowledge category and the medium ($F(1,18)=6.09$; $p=.001$): Facts are rarely acknowledged on the whiteboard (Fig.2). Our interpretation is the following. Since there is a low probability of misunderstanding or disagreeing about facts, their acknowledgment in MOO conversation (37%) basically means "I read your message". Acknowledgment simply aims to inform one's partner about shared perception (level 2). Conversely, on the whiteboard, mutual visibility is the default assumption, which makes level-2-acknowledgment unnecessary.

Implications. One difficulty we encountered in previous systems was that the artificial agents were behaving too rigidly, e.g. checking agreement at every step of the solution. There are 'social' reasons for not systematically checking agreement, such as demonstrating leadership, and avoiding harassment. These results indicate another reason: the medium may guarantee a level of grounding sufficient for the knowledge / task at hand. Reasoning about mutual visibility is easy to implement for an artificial agent: concerning the whiteboard, it may compare windows; concerning the MOO, simple rules can determine whether two agents see the same thing, given their respective positions and the command being performed. The most difficult aspect is reasoning on the probability that an item of knowledge can be misunderstood or disagreed with.

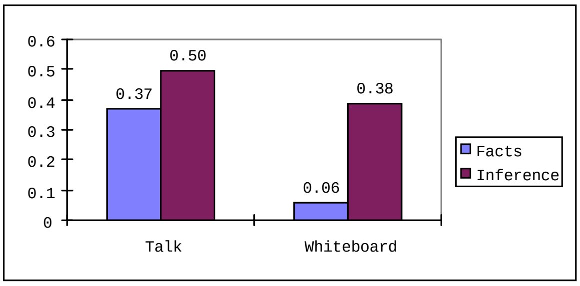

Figure 2: Effect of medium and content of interaction on acknowledgment rate.

3.3 Reasoning on mutual position

Observations. This is a particular case of the previous point. Subjects rarely ask their partners where they are located in the virtual space. They also rarely use the "who" command to check position. Actually, the MOO automatically provides some information about agents' positions. When A pages to B, B receives a notice such as "You sense that A is looking for you in room X". Every time A arrives in or leaves a room where B is located, B is informed of A's arrival/departure. We also showed that when the MOO provides less spatial information, the subjects interact more about space [18]. This means that they do not ignore mutual positions, but the medium provides roughly enough information. We also observed that when subjects are in different rooms, they acknowledge 34% of utterances, versus 50% when they are in the same room. Actually, subjects often go to the same room when they have long discussions. In other words, the MOO spatial metaphor seems to affect the behavior of subjects. Experienced MOO users are more sensitive to space than novices: the average space sensitivity[1] is 75% for novices versus 87% for experienced users ($F(1,9)=4.39$; $p=.05$).

*Implications***.** Since the spatial metaphor appears to be salient even when it is not functionally important, agents should be able to reason about spatial positions. This reasoning is also useful regarding the previous point: they need to reason about which information can be seen by the partner, and visibility in the MOO is itself bound spatially. Through this spatial metaphor, *MOOs provide formal rules to reason about sharedness.*

3.4 Negotiation by action

Observations. The subjects not only acknowledge MOO utterances using other MOO utterances but also with other MOO actions. The first condition for this type of acknowledgment is visibility: Hercule's utterance can be acknowledged by Sherlock's Action-X and only if Hercule sees that Action-X has been performed. In table 1, Sherlock can see that Hercule asks the question (9.6) as Sherlock previously invited him, since they are in the same room. The second condition of acknowledgment by action is that the MOO commands enable the speaker to express dialogue moves that (s)he wants to make. In this experiment, only three dialogue moves could be expressed through MOO actions: simple acknowledgment, straight agreement (one agent suggests an action, the other does it) and disagreement. Therefore, the type of information being acknowledged through action is generally decision about actions, namely spatial moves, asking questions and exchanging objects. Subjects could not use MOO actions for negotiating suspicion, because our

1 The 'space sensitivity' factor is the number of 'say' commands in same room plus the number of 'page' commands in different rooms, divided by the total number of messages.

experimental environment did not include commands conveying this type of information. We could add verbs (e.g., arresting a suspect) or objects (e.g., handcuffs) to indicate degree of suspicion. In other words, the semantics of 'talking by action' are bound by the semantics of the MOO commands created by the designer.

*Implications***.** Negotiation structures should be represented in terms of dialogue moves [20], which could be turned into interactions either as dialogue utterances or as MOO actions. Interpreting an action is no less complex than interpreting natural language, it relies heavily on the context. However, MOO actions may be simpler to interpret because the designer can define commands with clearly different pragmatic values.

9	Bar	S	' ask him what he was doing las night. i am talking to mr saleve
9.4	Bar	S	ask js about last night
9.6	Bar	H	ask giuz about last night

Table 1: Example of talk/action acknowledgment (Column 1: Time, column 2: MOO location , column 3: subject, column 4: actions)

3.5 Dynamically allocate functions to tools

Observations. Different pairs allocate communicative or problem solving functions (store facts, exchange inferences, coordinate action,...) to different tools or even re-allocate a function to another tool on the fly. For instance, the subjects carried a MOO notebook recording the answer to all questions asked to suspects. They could exchange their notebooks or merge the content. Some pairs intensively used these notebooks for sharing facts, while other systematically reported facts on the whiteboard. The former pairs hence had the whiteboard available for sharing inferences, while the latter filled their whiteboard and had hence to discuss inferences through MOO dialogues. The actual [function X tool] matrix varies across pairs. It may also vary within a pair as the collaboration progresses, one function being for instance progressively abandoned because the detectives become familiar with another one.
Implications. From the previous sections, one might infer that the whiteboard and MOO communicative functions could be pre-coded in an agent's design. However, the plasticity we observed, i.e. the ability of the pair to self-organize along different configurations, leads us to think that functions should be dynamically allocated to a medium during interaction.

3.6 Deliberately maintaining the task context

*Observations***.** Our initial hypothesis was that the whiteboard would help to disambiguate MOO dialogues, by supporting deictic gestures or explanatory graphics. We observed almost no deictic gestures because the partner's cursor was not visible on the whiteboard. We observed few explanatory graphics because the experimental task did not require such graphics. The whiteboard was not the place to disambiguate MOO dialogues, but rather appears as the central space of interaction. Information was sometimes negotiated before being put on the whiteboard, as if grounding was a pre-condition to display information in a public space. Conversely, dialogues often aim to ground the whiteboard information items (e.g. "why did you put a cross on...?"). These observations must be explained with respect to the task requirements. The main difficulty of the task was to manage a large amount of collected information. Most whiteboards contain a collection of short text notes which basically support individual and *group memory*. Whiteboards thereby also play a *regulation* role: (1) during the data collection stage, each subject can see on the whiteboard what her partner has done so far; (2) during the data synthesis stage, most pairs use a graphical code for marking discarded suspects, thereby indicating advances towards the solution. The whiteboard reifies the problem state. It is the central space of coordination, probably because *it retains the context* [19]. This context is established *at the task level*: the whiteboard increases mutual knowledge with respect to what has been done and how to do the rest. The whiteboard does not fix the context at the conversational level: the mutual understanding of MOO utterances does not seem to rely on whiteboard information. We even observed cases in which two different contexts are established, i.e. that the subjects participate in parallel into two conversations, one on the whiteboard and the other in MOO dialogues.

*Implications***.** The simplicity of observed graphics is good news for designers (otherwise we should design agents with an ability to interpret complex graphics). However, this observation is bound to the selected task: more complex representations would certainly have been drawn for a problem in physics or in geometry. What can nevertheless be generalized is the role of the whiteboard as a tool for maintaining a shared context. This role is due to knowledge persistence on the whiteboard. Artificial agents should both be able to build and use this shared context.

3.7 Maintaining multiple conversational contexts

Observations. The mechanisms of turn taking in the MOO are very different than in voice conversation. There is no constraint to wait for one's partner's answer before saying more. Moreover, one can more easily refer to utterances earlier than the last one, since they are still visible on the screen. Hence a MOO conversation between two people is not a simple alternation of turns. We computed an index of complexity of turn taking: its value is 0 if, knowing the speaker at turn n we have a probability of 1 for predicting who will speak at n+1. Its value is 1 if knowing the speaker at turn n does not give us any information regarding who will speak at n+1. The average index of complexity on 'say' and 'page' commands is 0.9 (SD = .06), which indicates an almost complete non-systematicity of turn taking! Interestingly, the average index of complexity is the same if we consider pairs with a high acknowledgment rate versus the pairs with a low acknowledgment rate (both 0.9). This indicates that the irregularity of turn taking does not really affect acknowledgment. Moreover, we observed interwoven turns (table 2), simultaneous turns (table 3), and parallel conversations. This phenomena is more important when more than two subjects interact on the MOO.

88.5	r1	H	page sherlock but what about the gun?
88.8	Priv	S	'Hercule which motive jealousy? He would have killed hans no?
89.3	Priv	S	'Hercule he stole it when the colonel was in the bar
90.3	r1	H	page sherlock Giuzeppe wanted to avoid that one discovers that the painting was fake.

Table 2: Example of XYXY turns (translated)

Implications. To participate in parallel conversations, artificial MOO agents need to maintain distinctively different contexts of interactions. If the context was unique, the interwoven turns reported above would lead to complete misunderstanding, which was not the case, or to spend a lot of energy in disambiguating references. Let us note that the multiplicity of contexts raise questions situated cognition theories in which context is often perceived as a whole.

87.1	r3	S	She couldn't have stolen the gun, could she?
87.4	*r3*	*S*	*read giuzeppe from dn1*
87.5	r3	H	I'm just checking something.
87.7	*r3*	*H*	*read Giuzeppe from dn2*
88.2	r3	S	No - Mona was in the restaurant till 9.
88.2	r3	H	No, she left around 9:00. She couldn't have stolen the gun.
88.7	r3	S	So Lisa, Rolf, Claire, Giuzeppe and Jacques are still open.
88.9	r3	S	and Oscar
88.9	r3	H	And Oscar...

Table 3: Simultaneous talk at 88.2 and 88.9 (from Pair 18)

4. Conclusions

Table 4 summarizes the suggestions with respect to the design of MOO agents. Some of them are technically easy to implement (e.g. 3) while others require further research. This list of skills extends the agent reasoning: *an agent does not only reason on the task, it also reasons on the optimal interactions to achieve sharedness*. The boundary between reasoning on the task and reasoning on interaction is shallow: for instance, reasoning on the probability of disagreement concerns both the solution and the interaction. In the slogan "shared solution", there is the word "shared" and the word "solution" and the key is to understand the interplay between these two aspects. We observed that different problem solving strategies lead to different profiles of interaction: for instance, forward chainers often fill the whiteboard with facts during the first hour, the MOO being used during the first hour only for management, while backward chainers very early discuss inferences in the MOO and write them on the whiteboard. However, the relationship we observed between interactions and problem solving is very complex.

1. Match knowledge persistence with medium persistence;
2. Determine to what degree a piece of knowledge is shared before performing any grounding act;
3. Determine what the partner can see according to his position;
4. Translate dialogue moves in terms of MOO actions;
5. Dynamically allocate functions to communicative tools (media);
6. Maintain and use the whiteboard as the shared context;
7. Maintain multiple conversational contexts;

Table 4: Implications: skills required for MOO agents

This experience reveals the importance of fine studies of interactions and leads us to suggest a new type of support for collaborative learning. We mentioned in the introduction that interactions may be either scaffolded by the interface or influenced by an agent. An intermediate approach is to implement *observers*. i.e. agents who collect information regarding users interactions, aggregate observations into high level indicators and display these indicators to the users themselves. We intend to study if pairs could use these indicators to regulate their interactions. Such indicators would also support monitoring by human coaches.

5. Acknowledgments

This study was funded by the Swiss National Science Foundation (grants #11-40711.94 and #7RUPJ048334). We thank the subjects who participated to the experiment. Special thanks to Richard Godard, Patrick Mendelsohn, Philippe Lemay, Andrew Wilson and Richard Joiner.

6. References

[1] C. Pauli & K. Reusser, Supporting collaborative problem solving: supporting collaboration and supporting problem solving. Swiss Workshop on Collaborative and Distributed Systems, Lausanne, 1997. http://tecfa.unige.ch/tecfa/tecfa-activities/loz97contrib.html

[2] R. Ploetzner, U. Hoppe, E. Fehse, C. Nolte & F. Tewissen, Model-based Design of Activity Spaces for collaborative Problem Solving and Learning. Proceedings of Euro AI&ED, Lisbon, October, 1996.

[3] J. Roschelle & S.D. Teasley, The construction of shared knowledge in collaborative problem solving. In: C.E. O'Malley (Ed), Computer-Supported Collaborative Learning, Berlin: Springer-Verlag, pp. 69-197, 1995.

[4] O. Etzioni & D. Weld, A softbot-based Interface to the Internet. Communications of the ACM 37 (7), 1994.

[5] H. Lieberman, Letizia: An Agent That Assists Web Browsing, International Joint Conference on Artificial Intelligence, Montreal, August 1995.

[6] Y. Lashkari, M. Metral, & P. Maes, Collaborative Interface Agents. Proceedings of AAAI '94 Conference, Seattle, Washington, August 1994.

[7] S.R. Bocioneck, Agent systems that negociate and learn. Int. J. Human-Computer Studies, 42, 265-288, 1995.

[8] L. Foner, A Multi-Agent Referral System for Matchmaking, First Int. Conf. on Practical Applications of Intelligent Agents and Multi-Agent Technology, London, UK, April 1996.

[9] P. Dillenbourg & J.A. Self, A computational approach to socially distributed cognition. European Journal of Psychology of Education, 3 (4), 353-372, 1992.

[10] T.-W. Chan & A.B. Baskin "Studying with the prince", The computer as a learning companion. Proc. of the International Conference on Intelligent Tutoring Systems (pp.194-200), Montreal, June 1988.

[11] T.-W. Chan, Learning Companion Systems, social Learning Systems, and the Global Social Learning Club. Jl. of Artificial Intelligence in Education. 7 (2), 125-159, 1996.

[12] S. Ur & K. Van Lehn, STEPS: A simulated, Tutorable Physics Student. Jl. of Artificial Intelligence in Education. 6 (4), 405-437, 1995.

[13] G. Fischer, A.C. Lemke & T. Mastaglio Critics: an emerging approach to knowledge-based human-computer interaction. Int. J. Man-Machine Studies 35, 695-721, 1991.

[14] D. Suthers, A. Weiner, J. Connelly & M. Paolucci, Belvedere: Engaging Students in Critical Discussion of Science and Public Policy Issues. Proceedings of AI-ED 95- Washington, August 1995.

[15] U. Hoppe, The use of multiple student modeling to parameterize group learning. Proceedings of AI-ED 95, Washington, August 1995.

[16] P. Dillenbourg, Some technical implications of distributed cognition on the design of interactive learning environments. J. of Artificial Intelligence in Education, 7 (2), 161-179, 1996.

[17] H.H. Clark, Managing problems in speaking. Speech Communication, 15, 243 – 250, 1994.

[18] P. Dillenbourg, L. Montandon & D. Traum. Where is my partner: coordination in Virtual spaces. Submitted for publication.

[19] S. Whittaker, E. Geelhoed & E. Robinson, Shared workspaces: How do they work and when are they useful? International Journal of Man-Machines Studies, 39, 813-842, 1993.

[20] M.J. Baker. A model for negotiation in teaching-learning dialogues Journal of Artificial *Intelligence in Education*, 5 (2), 199-254, 1994.