
The design of MOO agents:
Implications from a study on multi-modal collaborative problem solving.

P. Dillenbourg¹, P. Jermann¹, D. Schneider¹, D. Traum², C. Buiu³

¹ *TECFA, School of Education and Psychology, University of Geneva, Route de Drize 9, 1227 Carouge Switzerland. Phone +41 (22) 705.96.93*

² *UMIACS, A. V. Williams Building 9, University of Maryland, College Park, MD 20742 USA Phone: +1 (301) 405-113*

³ *Department of Computer Science and Control, University of Bucharest, Spl. Independentei 313 77206 Bucharest, Romania*

Address for correspondence: Pierre.Dillenbourg@tecfa.unige.ch

Abstract. This contribution reports the results of an empirical study on computer-supported collaborative problem. Twenty pairs of subjects had to play a mystery solving game in a text-based virtual environment (TecfaMOO). This environment was enriched with a shared whiteboard. We analyzed how subjects use different communication tools to build common grounds. We present seven main observations which could inform the design of artificial MOO agents. They describe the extra-amount of reasoning required to participate into interactions and to build shared understanding. We propose a new type of artificial agent, observers, which do not solve the problem but compute statistics regarding interactions, which can be useful for human or artificial tutors or which could be displayed to the users themselves.

The design of MOO agents: Implications from a study on multi-modal collaborative problem solving.

1. Introduction: Types of Agents

The notion of agent has been present more or less explicitly in Intelligent Learning Environments for some time. Similar types of agents are now investigated in Internet-based systems. We classify several types of internet agents, and present several aspects of interaction in task-based collaboration that a participant must be attuned to. These aspects come from an empirical study of multimodal human collaboration. Attending to these aspects will not only be crucial for agents participating in a collaboration, but for other types of agents. We propose a new type of agent, an **observer**, that calculates statistics regarding collaboration.

Sub-agents. These autonomous software entities which "carry out tasks for the user" are deemed by many to be the next revolution in computing. Since the World Wide Web (WWW) provides an evolving environment of voluminous amounts of unorganized information. Information sources appear and disappear, making information filtering a crucial capability of any Internet-based information gathering system. To date most research on Intelligent Information Agents has dealt with a user interacting with a single agent that has general knowledge and is capable of performing a variety of user delegated information finding tasks (e.g., Etzioni & Weld, 1994; Lieberman, 1995). Other examples of tasks accomplished by such agents are e-mail filtering (Lashkari & al., 1995) room reservation and calendar managing (Bocionek, 1995) or finding people sharing same interests (Foner, 1996). Many of these systems use machine learning techniques to non-intrusively infer the user preferences. Distributed multi-agent systems have been used to overcome limits of machine-learning methods such as: (1) a single general agent need an enormous amount of knowledge to deal effectively with user information requests that cover a variety of tasks, (2) a centralized information agent constitutes a processing bottleneck, the required processing would overwhelm a single agent which would constitute a "single point of failure", (3) a single agent needs considerable reprogramming to deal with the appearance of new agents and information sources in the environment.

Co-agents. In symmetrical systems (Dillenbourg and Baker, 1996), human and artificial agents can perform the same actions. The idea of co-learner, originally introduced as an alternative to learner modelling (Self, 1986), was re-used within various learning paradigms: collaborative learning (Chan & Baskin, 1988; Dillenbourg and Self, 1992), competitive activities, reciprocal tutoring¹ (Chan, 1996), learning by teaching and teacher training (Ur & van Lehn, 1995).

Super-agents. Most intelligent learning environments include agents (coach, tutor, expert, ...) which provide solutions and monitor the actions of users. In the case of multi-user learning environments, super-agents have to monitor the interactions among users. They can be teachers who analyze interactions to detect when they should intervene, judges or referees who intervene only in case of conflict, etc. For instance, the Belvedere system (Suthers & al., 1995) is dedicated to support critical discussions of science issues. Students build a diagrammatic representation with arguments. They can invoke an advisor who points to

¹ In reciprocal tutoring, the learners alternatively play the role of a tutor or a tutee (Palincsar & Brown, 1984).

specific parts of the diagram and proposes ways of extending or revising it. Some argument structures are expected to be more effective in stimulating critical discussions among students. In COSOFT (Hoppe, 1995), 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. Critic systems can also be viewed as super-agents, since they evaluate the user's work, but without a pedagogical perspective: "Critics do not necessarily solve problems for users. The core task of critics is the recognition and communication of deficiencies in a product to the user." (Fisher et al.; 1991).

Observers. We propose a fourth category of agents, observers, i.e. agents who collect information regarding users interactions, aggregate observations into high level indicators and display these indicators to the human coach or to the users themselves. One cannot a priori set up collaborative settings which guarantee effective collaboration, hence a coach must monitor the interactions (Dillenbourg et al, 1995). Observers would support this monitoring process both for human and artificial coaches. These indicators could also be shown to the subjects. We intend to study how a pair could use them to regulate its interactions. We are then close to another approach in which shared external representations are designed to replace intelligent system advice. For instance, Ploetzner et al. (1996) study problem-solving in physics and use a computerized tool which allows subjects to collaboratively build concept maps. These external representations have been designed in order to help students to link qualitative and quantitative knowledge.

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

We report experiments on collaborative problem solving in a virtual environment. Our long term goal is to develop computational models for supporting human-computer collaboration (Dillenbourg, 1996). The specific goal was to study how two human subjects elaborate a shared solution of a problem and how the drawings they make on a shared whiteboard help to support mutual understanding.

In the experiments, the two subjects play detectives in a mystery solving game: Mona-Lisa has been killed and they have to find the killer. They walk in a text-based virtual world (a MOO environment) 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 programmed robots. When exploring rooms, they find various objects which help them to discover the murderer. They are told that they have to find the single suspect who (1) as a motive to kill, (2) had access to the murder weapon and (3) had the opportunity to kill the victim when she was alone.

The subjects are located in different places, connected by a computer network and communicate using two pieces of software: a text-based MOO system and a Whiteboard. The MOO environment is a standard MOO called TECFAMOO². The subjects move in different rooms, talk to each other via two commands: "say...", to communicate with anybody in the same room, and "page John..." to communicate with John where ever he is. The detectives each carry a notebook which automatically records the answer to all the questions that they ask. In this experiment, the subjects use a MOO client called TKMOOLight which runs on UNIX stations. It includes an elementary whiteboard: both users draw on a same page, can see

² Visit [tecfamoo.unige.ch](http://tecfa.unige.ch/tecfamoo.unige.ch) - port 7777 For information:
<http://tecfa.unige.ch/tecfamoo.unige.ch>

and edit the objects drawn by their partner, but they do not see each other's cursor. All actions and interactions in the MOO window and in the whiteboard are recorded.

We ran the experiments with 20 pairs, the average time was two hours. A more complete analysis of observations can be found in Dillenbourg and Traum (1997). We present here only a few particular results which seem interesting for the design of artificial agents. In the following presentation, we refer to four levels of **mutuality of knowledge** (from Clark, 1994): (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 indented levels cover both the grounding process (achieving mutual understanding) and the negotiation process (acceptance or agreement).

3. Observations & implications

3.1 Matching knowledge and media persistence

Observations. Regarding MOO dialogues or whiteboard items, we discriminate four categories of content: **Task** (subdivided into **facts** and **inferences**), **management** (who has collected what, who goes where, who is where, etc.), **technical** problems and meta-**communication** (turn taking rules, graphical codes, etc.). The distributions of these content categories differs between MOO dialogues and whiteboard notes. In dialogues, 33% of interactions concern task management, versus 9% on the whiteboard. Task management, technical problems and meta-communication utterances have a short term validity, while task level information is more persistent, especially facts³: the former categories represent 45% of MOO dialogues (semi-persistent) versus only 10% of whiteboard items (persistent) (Fig.1). Moreover, the average delay of acknowledgment is shorter in MOO dialogues (48 seconds) than in whiteboard acknowledgment (70 seconds).

This difference probably reflects the difference of persistence of these two media: there is no urgency to acknowledge information which will (probably) remain a long time on the whiteboard. We refer to two forms of persistence of knowledge: **persistence of validity**, i.e. how long some piece of information remains true and **persistence of display**, i.e. how long it can be viewed on a particular medium. The latter depends on the medium: MOO information is semi-persistent, it scrolls up every time a new command is performed; whiteboard information is persistent until explicitly erased.

³ If it is true at time t than Helmut is a colonel, this information will still be true at time $t+1$.




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

Implications. An intelligent agent should select a mode of communication adequate to the persistence of knowledge being discussed. Negotiating non-persistent information through a persistent medium leads to the display of invalid information (when its validity has expired - a well-known WWW problem). 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. Selecting the best mode of communication requires reasoning about the match between the persistence of knowledge being discussed and the persistence of the medium through which it is discussed. The latter can be determined once and for all. The difficulty consists in reasoning on the persistence of information validity. The issue is not the same as in a truth maintenance system: it is not only to detect when information is not relevant anymore: it is to anticipate, when information is being discussed, how long it might be 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 between facts and inferences or management issues is the probability of disagreement or misunderstanding. In this task, there is little to disagree about or misunderstand with facts. Moreover there is a significant interaction effect on the acknowledgment rate between the knowledge category and the mode of negotiation ($F=6.09$; $df=2$; $p=.001$): 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 "ok, 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 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 several 'social' reasons for not checking agreement, such as demonstrating leadership, and avoiding harassment. These results indicate another reason: when the medium guarantees a level of grounding sufficient for the interaction. For example, given that mutual visibility (level 2) is guaranteed by the whiteboard and that mutual visibility is enough for facts since facts are not open to misunderstanding or disagreement (levels 3 and 4), facts are rarely explicitly acknowledged on the whiteboard. 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: Interaction effect on the acknowledgment rate between the mode of interaction and the content of interaction (ask knowledge).

3.3 Reasoning on mutual position

Observations. This is a particular case of the former point. We observed that agents 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 information about mutual position⁴. Montandon (1996) showed that, for a task which requires more information about MOO position, when the MOO provides less information, the subjects perform significantly more interactions to get this information. In other words, MOO positions are not ignored, but normally the medium more or less automatically provided enough information, hence making explicit interactions about it very rare. This is confirmed by other observations: (1) When subjects are in different rooms, they acknowledge 34% of utterances, versus 50% when they are in the same room. (2) Subjects often meet in the same room when they have long discussions (making inferences about the facts). (3) Experienced MOO users are more sensitive to space than novices: the average space sensitivity⁵ is 75% for novices versus 87% for experienced users ($F=4.39$; $df=1$; $p=.05$). In other words, the MOO spatial metaphor seems to really affect the behavior of subjects.

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 example 1, Sherlock can see

⁴ 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 leave a room where B is located, B is informed of A's arrival/departure.

⁵ The 'space sensitivity' factor was computed as the number of 'say' commands in same room plus de number of 'page' commands in different rooms, divided by the total number of messages.

that Hercule asks the question (9.6) as Sherlock previously invited him, since they are in the same room.

| | | | |
|-----|-----|---|--|
| 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 |

Example 1: Talk/action acknowledgment (from Pair 16)

The second condition of acknowledgment by action is that the MOO commands enable the speaker to express the dialogue move 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 who was suspect, because our experimental environment did not include commands conveying this type of information. We could include verbs (e.g. putting a suspect in jail) or objects (e.g. putting 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 (Baker, 1989), which could be turned into interactions either as dialogue utterances, or as MOO actions. In everyday life, interpreting an action is not less complex than interpreting natural language, it heavily relies on the context. However, MOO actions may be simpler to interpret because the designer can define commands with clearly different pragmatic values.

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. Let us illustrate that with an example. Our 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 prefer to systematically report 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 from one pair to another. 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 sections 4.1. and 4.2, one could infer that the whiteboard and MOO communicative functions could be directly coded in the design of artificial agents. The plasticity we observed, i.e. the pair ability to self-organize along different configurations leads us to think that functions such as "share facts", "negotiate strategy" should be dynamically allocated to a particular 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 simple deictic gestures or by drawing explanatory graphics. We observed

almost no deictic gestures⁶ and few explanatory graphics. Conversely, information is sometimes negotiated before being put on the whiteboard. Grounding is not achieved through the whiteboard, but appears then as a pre-condition to display information in a public space. Conversely, whiteboard/talk interactions often aim to ground the information put on the whiteboard (e.g. "why did you put a cross on...?"). The subjects do not draw elaborated graphics, probably because the experimental task was not intrinsically spatial. The main difficulty was the management of the large amount of collected information. Most whiteboards contain a collection of text notes (collected facts and inferences). The main function of the whiteboard was to support individual and **group memory**. It thereby also plays 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 rule for marking discarded suspects (generally, crossing notes about these suspects), thereby indicating advances towards the solution. In short, the whiteboard is the central space of coordination, probably because **it retains the context** (Whittaker et al, 1995). 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 context is not established at the conversational level: the mutual understanding of MOO utterances does not seem to rely on whiteboard information. We even observed several 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 able 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 MOO agents should both build in this shared context and use it when reasoning.

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 implicitly or explicitly 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. The average index of complexity⁷ 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 exactly the same if we consider the group of subjects with a high acknowledgment rate versus the group with a low acknowledgment rate (both 0.9 as well). This seems to indicate that the irregularity of turn taking does not really affect acknowledgment. Moreover, we observed simultaneous turns (example 2), interwoven turns (example 3) and even cases where the subjects talk in parallel about one thing in the

⁶ Probably for several reasons: (1) the emitter cannot simultaneously type the message and point on the whiteboard; (2) the receiver cannot look at the same time at the MOO window and at the whiteboard window; (3) the partner's cursor was not visible on the whiteboard.

⁷ This index evaluates the regularity of turn taking. Its value would be zero if, knowing the speaker at turn n we have a probability of 1 for predicting who will speak at $n+1$. Its value would be 1 if knowing the speaker at turn n does not give us any information regarding who will speak at $n+1$.

MOO and another on the whiteboard. This phenomena is more important when more than two agents interact on the MOO.

| | | | |
|------|----|---|---|
| 87.1 | r3 | S | She couldn't have stolen the gun, could she? |
| 87.4 | r3 | S | <i>read giuseppe from dn1</i> |
| 87.5 | r3 | H | I'm just checking something. |
| 87.7 | r3 | H | <i>read Giuseppe from dn2</i> |
| 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, Giuseppe and Jacques are still open. |
| 88.9 | r3 | S | and Oscar |
| 88.9 | r3 | H | And Oscar... |

Example 2: Simultaneous talk at 88.2 and 88.9 (from Pair 18)

| | | | |
|------|------|---|---|
| 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 Giuseppe wanted to avoid that one discovers that the painting was fake. |

Example 3: XYXY turns (from Pair 11, 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 questions situated cognition theories in which context is often perceived as a whole.

4. Conclusions

Table 1 summarizes the suggestions with respect to the design of MOO agents. Some of them are technically easy to implement (e.g. reasoning on MOO positions) while others require further research. This list of skills extends the agent reasoning: **the agent does not only reason to solve the problem, it also reasons in order to interact properly.** They illustrate the additional effort of reasoning involved for building a shared understanding. The boundary between reasoning on the task or on interaction is shallow: for instance, the probability of disagreement or misunderstanding concerns both the task knowledge and the interaction among partners. We attempted to further investigate the relationship between problem solving and interaction. We observed that different problem solving strategies lead to different profiles of interaction: for instance, pure 'forward' chainers often fill the whiteboard with facts during the first hour, then discuss inferences in the MOO, the MOO being used during the first hour only for management. Backward chainers start very early to discuss inferences in the MOO and write them down on the whiteboard (which is still not full). However, the relationship we observed between interactions and problem solving is very complex. In future investigations, we intend to use MOO-Observers, i.e. agents which automatically compute the indicators that we computed by hand when analyzing the protocols such as the rate of acknowledgment, the rate of co-presence in acknowledgment, the delay in acknowledgment, the complexity of dialogue patterns. In order to avoid natural language processing, we could use a semi-structured interface where the users pre-classify their

utterances by using sentence openers and specify to which utterance they acknowledge (Baker & Lund, 1996).

- Match knowledge persistence with medium persistence;
- Determine to what degree a piece of knowledge is shared before any grounding act;
- Determine what the partner can see according to his position;
- Translate dialogue moves in terms of MOO actions
- Dynamically allocate the negotiation of content to communicative tools (media)
- Maintain and use the whiteboard as the shared context
- Maintain multiple conversational contexts

Table 1: Implications for the design of MOO agents

5. Acknowledgments

This project was funded by the Swiss National Science Foundation (grant #11-40711.94). We would like to thank all the subjects who participated to the experiments. Special thanks to Richard Godard, Patrick Mendelsohn, Philippe Lemay, Andrew Wilson and Richard Joiner. The new current project on artificial agents is funded by the Swiss National Science Foundation programme for cooperation with Eastern European Countries (#grant #7RUPJ048334).

6. References

- Baker, M. (1989) Negotiated Tutoring. Doctoral dissertation, Institute of Educational Technology, Open University, Milton Keynes MK76AA, UK.
- Baker, M.J. & Lund, K. (1996) Flexibly structuring the interaction in a CSCL environment. In P. Brna, A. Paiva & J. Self (Eds), Proceedings of the European Conference on Artificial Intelligence in Education. Lisbon, Portugal, Sept. 20 - Oct. 2, pp. 401-407.
- Bocionek, S. R. (1995) - Agent systems that negotiate and learn. *Int. J. Human-Computer Studies*, 42, 265-288.
- Chan, T.-W. & Baskin, A.B. (1988) "Studying with the prince", The computer as a learning companion. *Proceedings of the International Conference on Intelligent Tutoring Systems* (pp.194-200), June 1-3. Montreal, Canada.
- Chan, T.-W. (1996) - Learning Companion Systems, social Learning Systems, and the Global Social Learning Club. *Jl. of Artificial Intelligence in Education*. 7(2), 125-159.
- Clark, H.H.(1994) Managing problems in speaking. *Speech Communication*, 15:243 – 250.
- Dillenbourg, P. (1996) Some technical implications of distributed cognition on the design of interactive learning environments. *J. of Artificial Intelligence in Education*, 7 (2), 161-179.
- Dillenbourg, P. & Baker, M (1996) Negotiation spaces in Human-Computer Collaborative Learning. Proceedings of the *International Conference on Cooperative Systems (COOP'96)*, Juan-Les-Pins (France), June 12-14 1996
- Dillenbourg, P., & Self, J.A. (1992) A computational approach to socially distributed cognition. *European Journal of Psychology of Education*, 3 (4), 353-372.
- Dillenbourg, P., Baker, M., Blaye, A. & O'Malley, C. (1995) The evolution of research on collaborative learning. In E. Spada & P. Reiman (Eds) *Learning in Humans and Machine: Towards an interdisciplinary learning science*. (Pp. 189-211) Oxford: Elsevier.
- Dillenbourg, P., Mendelsohn, P. & Schneider, D. (1994) The distribution of pedagogical roles in a multi-agent learning environment. In R. Lewis and P. Mendelsohn. *Lessons from Learning* (pp.199-216) Amsterdam: North-Holland.
- Dillenbourg, P. & Traum, D. (1996) Grounding in computer-supported collaborative problem solving. GNRS #11-40711.94 research report. Soon on <http://tecfa.unige.ch>.

- Fischer, G., Lemke, A. C. & Mastaglio, T. (1991) - Critics: an emerging approach to knowledge-based human-computer interaction. *Int. J. Man-Machine Studies* 35, 695-721.
- Foner, L. (1996) - A Multi-Agent Referral System for Matchmaking, First Int. Conf. on Practical Applications of Intelligent Agents and Multi-Agent Technology, London, UK, April 1996.
- Hoppe, U. (1995). The use of multiple student modeling to parameterize group learning. Proceedings of AI-ED 95- Washington, DC, August 1995.
- Lashkari, Y., Metral, M. & Maes, P. (1994) - Collaborative Interface Agents. Proceedings of AAAI '94 Conference, Seattle, Washington, August 1994.
- Lieberman, H. (1995) - Letizia: An Agent That Assists Web Browsing, International Joint Conference on Artificial Intelligence, Montreal, August 1995.
- Montandon, L. (1996) Etude des mécanismes de coordination spatiale dans un environnement virtuel de collaboration. Mémoire de Diplôme d'Etudes Supérieures en Sciences et Technologies de l'Apprentissage. Non publié. TECFA, FPSE, Université de Genève.
- Etzioni, O. and Weld, D (1994). A softbot-based Interface to the Internet. *Communications of the ACM* 37(7).
- Palincsar, A.S. & Brown, A.L. (1984) Reciprocal Teaching of Comprehension -Fostering and Comprehension - Monitoring Activities. *Cognition and Instruction*, 1 (2), 117-175.
- Ploetzner, R., Hoppe, H. U., Fehse, E., Nolte, C. & Tewissen, F. (1996) - Model-based Design of Activity Spaces for collaborative Problem Solving and Learning. Proceedings of Euro AI&ED, Lisbon, October, 1996
- Self, J.A. (1986) The application of machine learning to student modelling. *Instructional Science*, 14, 327-338.
- Suthers, D., Weiner, A., Connelly, J. & Paolucci, M. (1995) - Belvedere: Engaging Students in Critical Discussion of Science and Public Policy Issues. Proceedings of AI-ED 95- Washington, DC, August 1995.
- Traum, D. & Dillenbourg, P. (1996) Miscommunication in multi-modal collaboration. In working notes of the AAAI Workshop on Detecting, Repairing, And Preventing Human--Machine Miscommunication, pages 37-46, August 1996.
- Ur, S. and Van Lehn, K. (1995) - STEPS: A simulated, Tutorable Physics Student. *Jl. of Artificial Intelligence in Education*. 6(4), 405-437.
- Whittaker, S., Geelhoed, E. & Robinson, E. (1993) Shared workspaces: How do they work and when are they useful? *International Journal of Man-Machines Studies*, 39, 813-842.