

The long road from a shared screen to a shared understanding.

Pierre Dillenbourg, David Traum

University of Geneva, TECFA; University of Maryland, UMIACS

Abstract. What is the role of a shared whiteboard in building a shared understanding of the task and its solution? Our hypothesis was that a graphical communication tool would facilitate grounding processes, i.e. the mutual understanding of one or a few utterances and thereby the construction of a shared solution. We conducted an empirical study with 20 pairs solving an enigma in a MUD environment enriched with a whiteboard. The results show that the whiteboard was used less to disambiguate difficult concepts in conversation than as a tool for distributed regulation of the task. The graphical features of the whiteboard were less exploited than the fact that information displayed on the whiteboard was persistent. Subjects selected the communication medium by matching the persistency of display (how long information is displayed) with persistency of information (how long it remains valid). Their grounding behavior also takes into account the probability that some piece of information is misunderstood or disagreed upon.

Keywords: grounding, shared knowledge, virtual environment, whiteboard.

Introduction

The evolution of research on collaborative learning passed through three stages (Dillenbourg, Baker, Blaye & O'Malley, 1995). In the first stage, scholars aimed to prove the effectiveness of collaborative learning per se. Their contradictory results led to a second stage where the goal became to find the conditions that predict the effects of collaborative learning. The conditions are however numerous and interact with each other in a such way that one cannot control a priori these effects. Hence, current efforts to make collaborative learning effective directly target the collaborative process itself, by structuring the collaborative scenario and regulating the interactions. The main challenge, as old as psychology is, is to understand the cognitive processes triggered by social interactions. Social interactions can be described from different perspectives: as explanations (Webb, 1991), as

Cognitive science has for many years been treating social interactions as simple input/output mechanisms for core cognitive processes. This approach is inappropriate for our research field, since a major challenge for explaining the effects of collaborative learning is precisely to understand the deep intertwining between social interactions and problem solving (Dillenbourg 1999). The notion of *shared understanding* is a good candidate concept to use to articulate the connection between social interactions and problem solving. Unfortunately, this concept is not used in the same way by researchers studying problem solving and social interaction. At the linguistic level, 'shared understanding' is concerned with the understanding of a sentence, or even a word in a sentence, while, at the cognitive level, it is concerned with the understanding of a problem and its solution or even with the understanding of a domain. The goal of the research reported here was to articulate the relationships between the cognitive and the linguistic levels, namely to describe how grounding mechanisms contribute to build a shared solution.

There is a large difference of scale between grounding an utterance and sharing a solution through hundreds of interactions. This difference of scale is too large to expect that we would find a direct relationship between patterns of utterances and the final shared solution. Hence, we approach the relation between grounding an utterance and sharing a solution by describing how utterance grounding mechanisms vary according to task criteria.

We explore this issue in the context of a computer-supported collaborative work environment. We expected that a whiteboard would increase both mutual understanding at the utterance level, namely because a schema can disambiguate verbal expressions, and at the task level, by facilitating a shared representation of the solution.

Theoretical framework

Levels of shared understanding

Social grounding is the mechanism by which two participants in a discussion try to elaborate the mutual belief that their partner has understood what they meant to a criterion sufficient for the

current purpose (Clark & Brennan, 1991). Of course, people never understand each other completely. We treat the degree of shared understanding as a discrete variable (Dillenbourg, Traum & Schneider, 1996). We adapted the four communicative functions of Allwood *et al.* (1991), to the task of representing degrees of sharing of information, taking into account the peculiarities of virtual workspaces, namely typed communication and use of a spatial metaphor. If agent A wants to communicate information X to agent B, A have different kinds of information about B's potential sharing of that information, depending in part on the kind of feedback A receives. Each of these aspects can have a negative or positive component:

A can infer that B can (not) access X: For instance, in a virtual space, if A knows that B is in room 7, where information X can be found, then A knows that B can access X.

A can infer that B has (not) perceived X: For instance, if A writes a note on the whiteboard and B moves that note, A can infer that B has read it.

A can infer that B has (mis-) understood X: For instance, in a virtual space, if A says "let's ask *him* a few questions" and B moves to the room where "him" is located, then A can infer that B knows who is referred to as 'him', i.e. that "him" is grounded.

A can infer that B (dis-) agrees. This includes verbal agreement, but also (dis-) agreement by action, as we will illustrate in this paper.

This classification enables us to view grounding and agreement as different levels in a continuum going from complete mutual ignorance to completely shared understanding. We have two reasons to bypass the distinction between misunderstanding and disagreement. First, to be able to disagree requires a certain level of mutual understanding. Second, empirical studies of collaborative learning have shown that 'pure' conflict (p versus ~p) was not a necessary condition for learning – often a slight misunderstanding may be sufficient (Blaye, 1988).

Grounding mechanisms vary according to the task

Two pilots need a higher degree of mutual understanding when they fly a plane than when they talk about politics in a bar. The relation between grounding and the task is encompassed in the so-called 'grounding criterion': "*The contributor and the partners mutually believe that the partners have understood what the contributor meant to a criterion sufficient for the current purpose.*" (Clark & Brennan, 1991). This 'grounding criterion' is a synthetic account of various parameters that are investigated in this study, such as the probability of non-grounding (higher for ambiguous or polemical messages), the cost of non-grounding (how detrimental it is to misunderstand) and persistency of information. Indeed, this study is an investigation of how the grounding criterion varies inside the task.

Grounding mechanisms vary according to the medium

CSCW tools have an impact on the way mutual understanding is achieved. For a review of how medium features impart on grounding, see Clark & Brennan (1991). We here point out three factors which were important in our study.

The degree of sharedness of information.

Environments contribute to grounding by providing peers with the same access to information, i.e. they both see or hear the same things. Copresence decreases the cost of grounding (Clark & Brennan, 1991), This notion is an important principle of CSCW, known as "*What you see is what I see*" (WYSIWIS) (Stefik et al, 1987). It has been mainly

applied to shared workspaces (Terveen et. al., 1991), i.e. places where several users manipulate the same set of objects and see each other's actions.

The degree of mutuality of knowledge.

Sharedness differs from mutuality since we can be mutually aware that we disagree on X. One component of mutuality is visibility, i.e. whether the medium enables peers to see each other. The cost of grounding is lower if visibility is high (Clark & Brennan, 1991). In whiteboards, visibility is reduced to viewing each other's cursors and object moves. In virtual reality environments, visibility is achieved by displaying avatars. When visibility cannot be guaranteed, some CSCW environments provide peers with information about what he partner is looking at through *awareness*, i.e. the *"up-to-the-moment understanding of another person's interaction with the shared space"* (Gutwin & Greenberg, 1998, p. 2).

Persistence of Information.

The third factor is one of a set of features concerning the timing of interactions: synchronous versus asynchronous, turn taking rules, etc. In this contribution, we focus on one of them, **persistence**: objects are persistent (Smith, 1994) if, once created, they stay accessible as long as they are not deliberately erased.

Research setting

The task

Two subjects play detectives solving a mystery: a woman has been killed in a mountain Auberge and they have to find the killer among the (virtual) people present in the Auberge. Subjects navigate a virtual environment, built in a MOO (described in the next section), find various objects, meet suspects and ask them questions. The subjects can ask 3 kinds of question of any suspect: what the suspect knows about the victim, what the suspect did the night before and what the suspect knows about any object. Suspects are programmed robots implemented using the MOO language, and provide pre-defined answers. Subjects are told that they have to find the single suspect who (1) has a motive to kill, (2) had means to kill (i.e., access to the murder weapon) and (3) had the opportunity to kill the victim. The complexity of this task is more the information load (large number of facts to organize) than the intrinsic complexity of the relations to be inferred. This impacts on the way subject use the whiteboard. i.e. as a tool for storing and organizing information (group memory) more than as a tool for disambiguating information. The correct solution was found by 14 out of the 20 pairs. The time for completing the task averaged two hours (123 minutes). It varied between 82 and 182 minutes. The average time of failing pairs was almost the same (113 minutes).

The collaborative environment

These experiments have been run within a standard MOO (tecfamoo.unige.ch), using the TKMOO-lite client on UNIX workstations. The MOO window is split into panes: a pane of 14 X 19 cm, which display about 60 lines of text (any interaction uses several lines) and, just below, a text entry pane which allows the user to enter and edit messages up to 3 lines in length. The experimental environment also includes a MOO-based whiteboard. It supports elementary drawing of boxes, lines, and text objects. Users can move, remove, resize or change the color of the objects created by themselves or their partners (several subjects complained that the Both users see the same area of the whiteboard, there is no scrolling inside the fixed window size. whiteboard was too rudimentary, especially for editing objects). They cannot see each other's cursor. The whiteboard and MOO windows split each users' screen vertically into two equal areas.

The subjects

Twenty pairs of subjects participated in the experiments. We recruited subjects who had no experience of working together. The level of MOO experience was heterogeneous. We compared subjects with intermediate or advanced experience of the MOO with the novices (respectively 16 and 20 subjects). Taken individually, experience has an impact: experienced MOO users sent more messages per minute (mean novices = 0.45, mean experts = 0.68; $F=11.98$; $df=1$; $p=.001$) and were more sensitive to virtual space (mean novices = 75%, mean experts = 87%; $F=4.39$; $df=1$; $p=0.05$). However, there was no difference between novices and experts with respect to the success on task. There was a non-significant difference with respect to task completion time, probably due to shorter latency between messages rather than to problem solving behavior. Most statistics are based on 18 pairs (excluding pairs 1 and 2 who used voice interactions), the statistics involving actions on the whiteboard do not include pair 4 for which the movie of whiteboard interactions was lost. Most quantitative values are presented by pair, since, even if they are sometimes counted individually, most interaction variables make more sense when aggregated by pair.

Variables

We computed many variables in order to investigate grounding processes. We describe only a few of them here, some other variables are also described in the next section.

The mode of communication

This were three modes of communication: Moo dialogues, information put on the whiteboard or virtual actions in the MOO environment, such as moving, asking questions of suspects, etc. An instance of agreement by a MOO action is when A says "Let's go to the kitchen" and B moves to the kitchen. An instance of disagreement in a whiteboard is when A writes a note and B erases it. We observed cross-modal acknowledgment, such whiteboard/talk acknowledgment (A puts a note on the whiteboard and B make comments about it through a Moo dialogue), talk/action acknowledgment (e.g. A says «ask it to Oscar» and B moves to Oscar's room), and so forth. However, 86% of acknowledgment occurred inside MOO dialogues.

The rate of acknowledgment

This is the ratio between the number of interactions that were acknowledged by the partner and the total number of interactions. We parsed the 20 protocols and associated actions by pairs [U1 - U2] when U2 can be interpreted as acknowledging U1. Acknowledgment can be in different forms, ranging from a simple 'hmm' to a clarification question. The average rate of acknowledgment was 41% (For comparison, it was 90% for the pairs with voice communication). Interestingly, the acknowledgment rate does not seem related to surface interaction features. For instance, when split into two groups based on the interaction rate (number of messages per minute), both the high and low groups have the same acknowledgment rate (respectively 41% and 42%). Similarly, the acknowledgment rate was not correlated with the delay of acknowledgment, with acknowledgment symmetry (does A acknowledge B more often than B acknowledges A), nor with the regularity of turn taking.

The content of interaction

The subjects communicated several different types of information. Content types are categorized as shown in Table 1.

Table 1: Categories of content of interactions.

Category	Sub-category	Content and examples

Task knowledge	Facts	Utterances which contain information as it was collected in the MOO by the subjects (e.g. «Rolf was a colleague of the victim »). It often reproduces word by word the answer given by a suspect
	Inferences	An inference utterance involves some interpretation by the subject. (e.g. «Helmut had no motive to kill »).
Management	Strategy	Utterances about how to proceed: how to collect information (which suspects, which rooms, which questions, ...), how to organize data, how to prune the set of possible suspects, who does what in the pair... Utterances regarding spatial positions were generally related to strategy issues and were hence included in this category. . (e.g. «Let's see who could get the gun »).
Meta-communication		Utterances about the interaction itself, such as tuning delay in acknowledgment (e.g. « Sorry I was busy with the whiteboard ») or establishing conversational rules (e.g. « We should use a color coding »).
Technique		Utterances where one subject asks his partner how to perform a particular action in the MOO. . (e.g. «I can't read my notebook »).

Results

Does the whiteboard contribute to grounding utterances?

Our main hypothesis was that whiteboard role in mutual understanding was to support drawing schemata, which are useful for **clarifying an idea** that is difficult to turn into words. Actually, this was rarely the case. The experiment task did not involve the type of misunderstanding which can be easily disambiguated by a schemata. A second hypothesis was that the whiteboard would support grounding by helping partners to **solve references**. An utterance such as "*he lies*" can be disambiguated if it is accompanied by a gesture pointing to "he". In a previous study, we observed that 87% of the numerous gestures performed by two subjects in front of the MEMOLAB environment (Roiron, 1996) were simple deictic gestures. In the present study, deictic gestures were not observed for two reasons. First, the users could not see each other's cursor. Second, even if some gestures were possible (e.g. putting a mark on or moving the object being referred to), it was impossible for the speaker to simultaneously type "he" in the MOO window and move the cursor wherever "he" was located on the whiteboard.

While our hypothesis viewed the whiteboard as a tool for disambiguating dialogues, we observed opposite acknowledgement sequences, we often observed the opposite, i.e. utterances aimed to disambiguate the information displayed on the whiteboard. In short, **the dialogues were instrumental for grounding whiteboard information** rather than the reverse. The preponderance of the whiteboard over dialogues seems related to the persistency of information, as we will see next.

Does the whiteboard contribute to build a shared solution?

Most whiteboards were made of a large collection of text notes. We did not encounter many elaborated graphics. On the 20 pairs, four drew a timeline (e.g. fig. 1), four drew a map and three drew a graph, mainly indicating social relations among suspects. Timelines can be helpful since its is difficult to reason about intervals without visualizing them. However only one complete time line has been produced. The maps did not really help to solve the task since the solution of the enigma did not require any spatial reasoning such as "Hans could not go from room 1 to the bar without being seen by Rolf.", and the basic layout was provided in the experimental instructions. The pairs who drew a map where probably influenced by the spatial metaphor which was very salient in the task. Overall, the graphs were generally abandoned before being completed. In other words, **the expressivity of graphics has been under-exploited by the subjects**. Our interpretation is simply that the task did not require such explanatory schemata,

e.g. to articulate complex causal structures. We did, however, see at least one pair use the schematic power to indicate a chain of logical inference. When one subject asked about the reasons for the latest in a series of inferences displayed on the whiteboard, the other used boxes and arrows to indicate the supporting premises. Thus, even when purely verbal information is expressed, the whiteboard allows a kind of deictic reference as a useful shorthand.




Figure 1: Using the whiteboard to build a shared solution: Uncompleted timeline in Pair 22.

Some graphical features of the whiteboard were used to structure the collection of notes. Three pairs use color codes for indicating the author of the note. Many pairs used the 2-D space to structure the information, either systematically in a two-entry table (2 pairs), more often, by grouping geographically the notes concerning the same suspect.

Delhom (1998) observed five additional pairs solving the same task but with a whiteboard which was not shared (A could not see what B drew). She observed that subjects without a shared whiteboard had a higher acknowledgment rate than those with a shared whiteboard. Conversely, Schwartz (personal communication) observed that subjects with a whiteboard perform more grounding acts than those without a whiteboard. This apparent contradiction differentiates whiteboard functions and confirms the results above. In Schwartz' study, the subjects faced complex routing problems and hence needed the whiteboard to express or repair their contributions, while our task, used by Delhom, required managing a large set of simple factual information, the whiteboard being used for data management, as we explain now.

Does the whiteboard support regulation of the task?

The task has two logical phases, *data acquisition*, in which the subjects search for and gather necessary clues, and *data synthesis*, in which they try to infer the answer from the gathered information. Generally, these phases are temporally distinct, although there can be some overlap. During the first part of the task (data acquisition), most pairs split spatially, one detective visiting the rooms of the upper corridor of the auberge while the other explored the lower corridor. Very often this division of labor did not hold, because one detective was faster than the other and because this division does not specify who will explore the rooms in between corridors (e.g., the bar, and the restaurant). This insufficient specification of the strategy should lead to complementary negotiation of who does what, however, this was not the case. A possible explanation could be that the subjects used MOO commands such as "who" to trace their partner's travels in the MOO. This was not the case either. Indeed, each agent was able to 'trace' his partner by looking at the information (s)he was adding to the whiteboard: If A sees that B puts the information collected from Suspect5 on the whiteboard, A may infer that B is in Suspect5's room. Another explanation of this implicit coordination is related to the specific affordances of the spatial design of the environment. We indeed observed that virtual space

influences collaboration (Dillenbourg, Traum & Montandon, submitted), but we do develop this point here.

During the second part of the task (data synthesis), many pairs started with all the suspects and then sequentially eliminated any suspect who had either no motive to kill, no opportunity to get the weapon, or no opportunity to kill (Fig 2). This process generally takes a concrete form on the whiteboard: the detectives incrementally cross out the notes regarding any suspect who is discarded. The whiteboard provides them with a persistent representation of the set of remaining suspects.

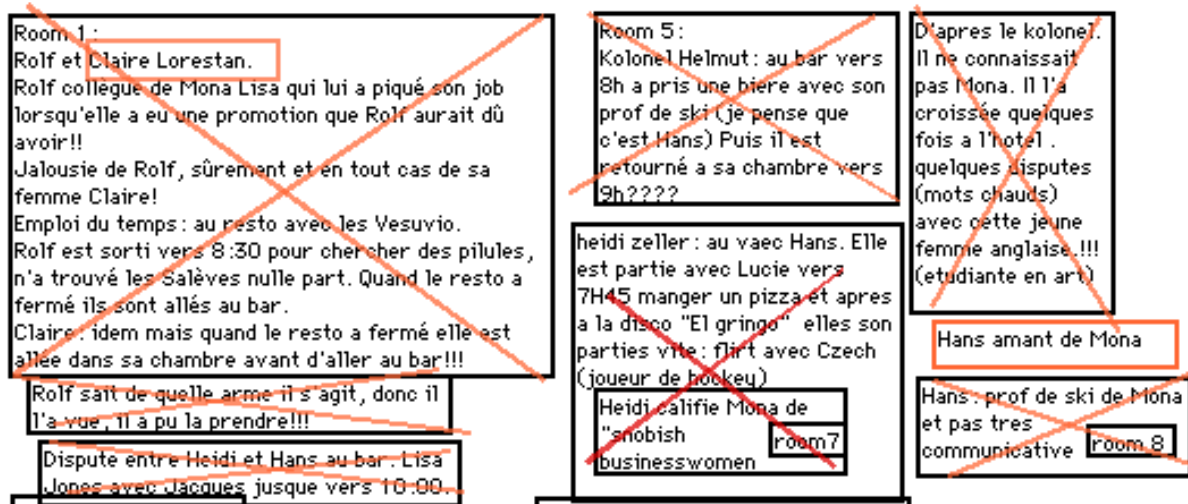


Figure 2: Using the whiteboard for task management, crossing notes = discarding suspects.

In other words, the whiteboard mainly served as shared memory, i.e. accumulating the set of problem data and inferences on which peers agreed or have to agree. The whiteboard reifies the problem state or, in Whittaker, Geelhoed & Robinson (1993) terms, the whiteboard helps to 'retain the context'. Shared memory relies on the persistency of displayed information. The next section confirms this finding by showing that the whiteboard is used as a group memory.

Does grounding predict performance ?

The rate of acknowledgment is a rudimentary appraisal of grounding intensity, but is it related to problem solving? At a first glance, the answer is negative. The acknowledgment rate is not related to global performance measures. When the sample was split into groups consisting of the highest acknowledgment rates versus the lowest rates, we found no difference with respect to correct solution (respectively, 6 correct solutions versus 7) or with respect to time (respectively 120 and 125 minutes, in average). However, low acknowledgers perform significantly more actions in the MOO (ask, move, read, look, etc.) than high acknowledgers (respectively, 237 versus 178 ; $F=5.13$, $df=1$, $p=.05$). This difference seems related to a higher number of redundant actions (respectively 18 versus 6; $F=11$; $df=1$; $p= 0.01$). We discriminate self-redundancy (a subject asked a question that (s)he asked before) from cross-redundancy (a subject repeats a question that her partner asked before). The average self redundancy is almost equal for low and high acknowledgers, respectively 3.4 and 3.2. Hence, the difference between high and low acknowledgers is rather related to cross-redundancy. This result illustrates that redundancy is not just a matter of memory, since memory limitations would also affect self-redundancy. In cross redundancy, we differentiated immediate redundancy (the delay between two redundant questions is less than 5 minutes) from delayed redundancy (over 5 minutes – this threshold has

been selected by inspecting data). Immediate redundancy is not always an indicator of mis-coordination. It may indicate explicit coordination: A, instead of summarizing the information for B, simply invites B to ask the same question again. On average, high acknowledgers ask almost the same number of immediate redundant questions as the low acknowledgers, the mean being respectively 1.20 and 1.40. The difference between the two groups comes from the number of long term redundancies (11.40 for low acknowledgers, 3.40 for high acknowledgers).

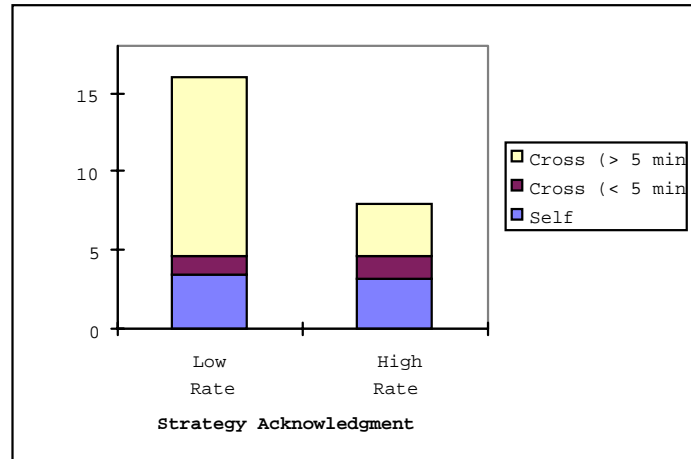


Figure 3: Relationship between the rate of acknowledgment for utterances regarding task management and different indicators of redundancy in problem solving actions.

This comparison reveals a quantitative relationship between specific grounding mechanisms (acknowledging utterances about the strategy) and the efficiency of a distributed problem solving process (measured by the delayed cross-redundancy). This relationship is however not very strong: redundancy may decrease efficiency but increase effectiveness and the cost of acknowledgment might be higher than the cost of redundancy. An interesting conclusion is however that the acknowledgment rate is a sensitive variable (we reuse it in the next sections). The more important conclusion is that the notion of **group memory**, which explains differences in delayed cross-redundancy, is a key to relate grounding utterances and sharing a solution. Actually, the fact that information remains displayed on the whiteboard is only one factor in the selection of a grounding medium, as explained in the next section.

The complex management of shared spaces

In our experiment, the whiteboard was not the only shared space. The MOO was also a shared space. For instance, subjects could disambiguate sentences such as "he lies" not by referring to the whiteboard but also by observing that both subjects are in the same MOO room and that 'he' was the suspect present in the room. *Virtual space creates a micro-context which complements the conversational context.* The previous section revealed that the difference between the MOO dialogues and the whiteboard interaction was less a question of verbal versus graphical interactions than a matter of timing: (1) interactions follow sequentially in the MOO, while there is no temporal order to be respected on the whiteboard, (2) information is persistent on the whiteboard (it remains displayed until it is explicitly erased) while it is semi-persistent in the MOO (past interactions scroll up as new interactions are added to the MOO window). The subjects appear to take these differences into account. As illustrated by figure 4, utterances concerning the strategy, meta-communication and technical issues are rarely conducted via the whiteboard. To understand these data, we need to **discriminate the persistency of display** (how long information is displayed) **from the persistency of validity** (how long information remains

valid). The display persistence on a whiteboard is high, while the validity persistence of strategic knowledge is low. For instance, an utterance such as "I am going to ask questions to Heidi" remains true for a few minutes. If such a piece of information is put on the whiteboard, it might still be displayed while it is not true anymore. If the display persistency is longer than the validity persistency, obsolete information may be displayed (this is becoming a common problem on the world wide web, where sites are moved and removed more often than links updated).

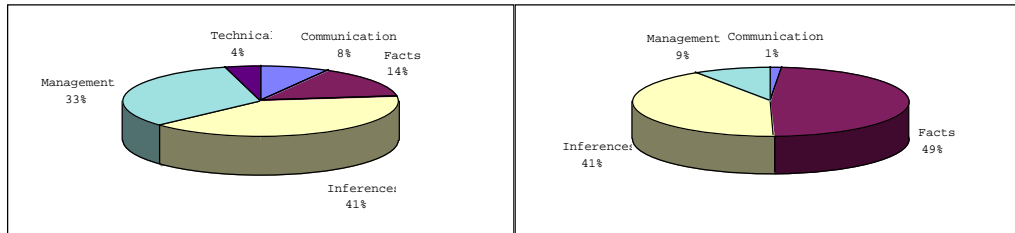


Figure 4: Categories of content in MOO (left) and whiteboard (right) interactions

In summary, subjects seem very effective in matching the display persistence with the validity persistence in such a way that group memory is not 'polluted' with obsolete information. To have a complete picture of how grounding is achieved through different media, we must introduce another dimension, which is the probability of non-grounding.

The average rate of acknowledgment for inferences (46%) is higher than for facts (26%). There is little room for disagreement about facts, and they are also easier to understand, because of having been pre-constructed and tested by the task designers, hence acknowledgment is less necessary. Conversely, inferences are open to and disagreement and, since they are spontaneously constructed, more subject to misunderstanding and thus are more likely to require acknowledgment to be shared.

Acknowledgment is lower in the whiteboard than in Moo dialogues. We set up the size of whiteboard in a way that no scrolling was possible. Hence, partners knew they see the same information as their partners. In Moo dialogues, visibility only occurs when the partners are in the same room and are mutually aware of it.




Figure 5: Interaction on the acknowledgment rate between the mode of communication and the content of communication.

More interestingly, figure 5 shows a significant interaction effect on the acknowledgment rate between the content and the mode of grounding ($F=6.09$; $df=2$; $p=.001$). These data can be interpreted with the four levels of shared knowledge defined in the beginning of this paper. The talk/whiteboard difference occurs on level 2: acknowledgment is more necessary if one cannot be sure that the partner has perceived the message. The facts/inferences differences relate to level 3-

4: acknowledgment is more necessary if one cannot be sure that the partner has understood or agreed. By combining these two factors, table 2 shows why the acknowledgment rate is the lowest for facts in the whiteboard and highest for inferences in MOO dialogues.

Table 2. Explaining variations of acknowledgment rate according to the medium feature (sharedness) and the information features (probability of non-grounding): the number of 'yes' entries in a column is related to the data shown in figure 5.

Is it necessary that B acknowledges information X provided by A?	MOO dialogue (talk)		Whiteboard interactions	
Level of sharedness of information X	Facts	Inferences	Facts	Inferences
4. Agreement	NO: X is presented as true by the environment	YES: B can question both X and the motivations for X	NO: X is presented as true by the environment	YES: B can question both X and the motivations for X
3. Understanding	NO: X was pre-constructed.	YES: B may misunderstand X	NO: X was pre-constructed.	YES: B may misunderstand X
2. Perception	YES: B can be in a different room than A and hence fail to perceive X		NO: A and B see the same information on the whiteboard	
1. Access	YES: B can be in a different room and hence fail to access to information X		NO: A and B have access to the same information on the whiteboard	

Conclusions

In this study, the whiteboard is not used to ground utterances, nor to represent the final solution. Instead, it is used to maintain a shared representation of the state of the problem along the problem solving process. Our whiteboard fulfills its role for this task because of two features: persistency of display and shared visibility. Our observations should be generalized not just to any system labeled 'whiteboard' or 'virtual space' by its designer, but instead to systems including a space of shared and persistent information.,

This group memory is not a static depository of information, but the result of a complex organization of the whole distributed cognitive system in which subjects allocate different functions (grounding information of type X) to different tools (media) according to criteria such as the persistency of display, the persistency of validity of information, the probability of non-grounding due to the medium (level 1 or 2) or due to the information itself (level 3 or 4)

We consider it worthwhile to describe these agents and these tools as a distributed cognitive system because we observed a large variety of ways of distributing the various functions to favorite tools. For instance, if a pair communicates all facts through dialogues, the whiteboard will be more available for inferences; if they exchange all information in the whiteboard, they fill it very soon and hence start to exchange inferences through MOO dialogues, and so forth. 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. This **plasticity**, or system ability to self-organize along different configurations, justifies the descriptions of a pair and the environment as a single cognitive system.

Acknowledgments

This project was funded by the Swiss National Science Foundation (grant #11-40711.94).

References

- Allwood, J., Nivre, J. & Ahlsén, E. (1991). *On the Semantics and Pragmatics of Linguistic Feedback*. Gothenburg Papers in Theoretical Linguistics No. 64. University of Gothenburg, Department of Linguistics, Sweden.
- Blaye, A. (1988) Confrontation socio-cognitive et résolution de problèmes. Doctoral dissertation, Centre de Recherche en Psychologie Cognitive, Université de Provence, 13261 Aix-en-Provence, France.
- Clark, H.H. (1994) Managing problems in speaking. *Speech Communication*, 15:243 – 250.
- Clark, H.H., & Brennan S.E. (1991) Grounding in Communication. In L. Resnick, J. Levine & S. Teasley (Eds.), *Perspectives on Socially Shared Cognition* (127-149). Hyattsville, MD: American Psychological Association.
- Delhom, K. J. (1998) Etude des techniques de collaboration dans la résolution de problème à l'intérieur d'un environnement virtuel. Mémoire de licence non publié, Faculté de Psychologie et des Sciences de l'Education, Université de Genève.
- Webb, N.M. (1991) Task related verbal interaction and mathematics learning in small groups. *Journal for Research in Mathematics Education*, 22 (5), 366-389.
- 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.
- Chi M.T.H., Bassok, M., Lewis, M.W., Reimann, P. & Glaser, R. (1989) Self-Explanations: How Students Study and Use Examples in Learning to Solve Problems. *Cognitive Science*, 13,145-182.
- Dillenbourg, P. , Traum, D. & Schneider, D. (1996) Grounding in multi-modal task-oriented collaboration. In P. Brna, A. Paiva & J. Self (Eds), *Proceedings of the European Conference on Artificial Intelligence in Education*. Lisbon, Portugal, Sept. 20 - Oc. 2, pp. 401-407.
- Dillenbourg, P. (1999) What do you mean by collaborative learning? In P. Dillenbourg (Ed) *Collaborative learning: Cognitive and Computational Approaches*. Oxford: Pergamon.
- Gutwin, C. & Greenberg, S. (1998) The effects of workspace awareness on the usability of real-time

distributed groupware. *Research report 98-632-23*, Department of Computer Science, University of Calgary, Alberta, Canada

Roiron, C. (1996) Expérimentation d'un logiciel éducatif utilisant des techniques d'intelligence artificielle. Rapport de recherche non-publié. TECFA, Faculté de Psychologie et des Sciences de l'Education, Université de Genève.

Stefik, M., Bobrow, D.G., Foster, G., Lanning, S. & Tatart, D. (1987) WYSIWIS Revised: Early Experiences with Multiuser Interfaces. *ACM Transactions on Office Information Systems*, 5(2), 147-167, April.

Terveen L.G., Wroblewski, D.A. & Tighe S.N. (1991) Intelligence Assistance through Collaborative Manipulation, *Proceedings of IJCAI 91*.

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.

Adresses

Pierre Dillenbourg, TECFA, University of Geneva, Route de Drize 9, CH 1227 Carouge Switzerland, Pierre.Dillenbourg@tecfa.unige.ch
David Traum., AV Williams Building, UMIACS, University of Maryland, USA, traum@cs.umd.edu