2. Scientific part

2.1 Summary

When we solve a problem with a colleague or a friend, we often use a piece of paper to draw a diagram or sketch. Such diagrams help us to build a shared representation of the problem and to repair communication breakdowns. Can we design equivalent artefacts to enrich human-computer collaboration? To answer this question, we propose to study the role of social grounding in human-human collaborative problem solving and to integrate functionally equivalent mechanisms into an interactive inference engine for human-computer collaboration. Our research will contribute to the theory of collaborative problem solving and lay the psychological foundations for new generations of interactive computer systems. We will proceed in three stages.

1 Develop or adapt a computer-supported collaborative system.

   The system will be designed for creating the experimental settings necessary for stage 2 and stage 3. The reasons for mediating human-human collaboration are (1) to use the communication channel as an independent variable controlled by the computer system, (2) to coerce (to some extent) human partners to use interaction techniques that can be used in human-computer collaboration.

2. Study the role of social grounding in human-human collaborative problem solving.

   Efficient collaboration requires that each partner understands what the other meant to a criterion sufficient for the current task. Social grounding is the process by which two partners try to reach the mutual belief that the other has reached this understanding criterion (Clark & Brennan, 1991). This goal can be achieved by various linguistic mechanisms, depending on the media used for communication. This project focuses on two issues: (1) how people use external references (drawing a diagram, pointing to an object) during social grounding and (2) the role of grounding mechanisms in problem solving. Three experimental settings are proposed in our research plan.

3. Develop a human-computer collaborative system

   The goal is to develop human-computer interaction techniques that are functionally equivalent to the gestures observed in human-human social grounding. The difficulty is not to support gestures such as drawing or pointing. The challenge is to integrate these gestures within the reasoning process performed by a knowledge-based system. These gestures can be used for changing the problem state representation, guiding the rule selection, determining how to instantiate a variable, ...

This project, like our research team, is multidisciplinary. It is articulated with the **research programme "Learning in Human and Machines" funded by the European Science Foundation** (see section 2.2.1.5). This international and multidisciplinary program includes five task forces. The main contractor of this project, P. Dillenbourg, is responsible for the task force "collaborative learning". The first co-contractor, Prof. Mendelsohn, is also member of the program steering committee. The research plan explains how the experiments that we propose will serve as a basis for this international program.

Keywords: Collaborative problem solving, human-computer interaction, distributed cognition, social grounding, artificial intelligence, knowledge-based systems, groupware.

2.2 Research plan

2.2.1. State of the art

This project involves concepts from different disciplines: psychology, linguistics and computing. These concepts are not simply juxtaposed. They fit into each other and justify an integrative project. To define precisely the scope of this project, we start from a definition of the term "collaboration": **collaboration is the process of building and maintaining a shared conception of a problem** (Roschelle & Behrend, to appear). This definition excludes two categories of systems that are sometimes been labelled as 'collaborative'. We do not consider as collaborative the interaction between a user and a computational tool (e.g. a word processor), because collaboration implies a symmetrical relationship between partners, each of them being allowed to take initiatives, to disagree, and so forth. We do not consider as collaborative systems where the task is split in clearly distinct subtasks and distributed among the human and the machine. Most research on human-computer collaboration falls into this second category: designers seek a task distribution that takes into account the respective strengths and weaknesses of each partner (Hancock, 1992). Our project addresses the more fundamental issue of how two partners solve jointly the same problem, with the same tools and actions, and, to achieve this, how they build a shared understanding of this problem. This does not deny role distribution in collaborative activities. Even when the task cannot be divided into independent subtasks (vertical division), one observes that different subjects spontaneously adopt roles. For instance, in computer supported tasks, the peer member who holds the mouse tends to be "executor", while the other is likely to be the "reflector". This role distribution, which changes over time, corresponds to some horizontal division of cognitive processing (Myake, 1986; O'Malley, 1987; Blaye et al., 1991).

2.2.1.1 Distributed Cognition

The relationship between social interactions and individual cognition issue lies at the very heart of the "social versus individual" debate that concerns the nature of cognition. One can discriminate three theoretical positions with respect to this issue. Piaget and the socio-constructivist approach (Doise and Mugny, 1984) are often presented as the defender of the individual position while Vygotsky and the socio-cultural approach are viewed as the instigator of the social approach. As Butterworth (1982) pointed out, their opposition has been exaggerated. Both authors acknowledge the intertwined social and individual aspects of development, but they attribute the primacy to the individual (for Piaget) or to the social environment (for Vygotsky). This project is inscribed within a third approach, hereafter referred as 'distributed cognition'. In this approach, **the similarities between the individual and the social planes of cognition** receive more attention than their differences. They view cognition as fundamentally 'shared' or 'distributed' over individuals. Defenders of this approach question the very discrimination between what is social and what is individual: "*... research paradigms built on supposedly clear distinctions between what is social and what is cognitive will have an inherent weakness, because the causality of social and cognitive processes is, at the very least, circular and is perhaps even more complex*" (Perret-Clermont et al., 1991, p. 50). The distributed cognition approach is closer to the Vygotskyan position than to the Piagetian view since it considers the group rather than individual as the primary unit of analysis (Resnick, 1991). By its focus on social structure, distributed cognition is deeply intertwined with the 'situated cognition' theory (Lave, 1988). Interestingly, the situated theory pays also attention to the role of physical artefacts in cognition (Suchman, 1987). This is important for understanding the role that diagrams in collaborative problem solving.

2.2.1.2 Collaborative dialogues

Distributed cognition implies new ways for analysing interactions. Within the individual view of cognition, the researchers attempt to discover the mechanisms that relate (individual) utterances in dialogue with (individual) mental states or behaviour. The proposed approach considers communication as a collective activity, emerging from the group, rather than a one-to-one process. This position changes the analysis of some phenomena observed in collaborative learning. For instance, the peer who produces explanations for his partner seems to learn through this activity (Bargh & Schul, 1980). From the 'individual' perspective of cognition, this phenomena is explained by the self-explanation effect (Chi et al., 1989). From a 'distributed' perspective, explanation is not something delivered by the explainer to the explainee, but constructed jointly by both partners (O'Malley, 1987; Baker, 1992; Cawsey, 1993). The methodology focuses on the emergence of joint constructions, namely joint linguistic forms, joint artefacts (drawings, ...) and shared conceptions.

Social grounding refers to the process through which two discussants try to elaborate the mutual belief that the other partner has understood what he meant to a criterion sufficient for current purposes. Clark & Brennan (1991) describe various mechanisms of social grounding: repeating what has been said in another way, pointing to objects ('you mean this one'), left dislocation (e.g. 'Your square, it is too large.'), using words that invite the partner to confirm his understanding,... These grounding mechanisms change according to the medium of communication. For instance, eye contact is generally not available in computer-supported distance collaboration. **This project focuses on how people use external references (a diagram, a picture, ...) during social grounding**. In this project, we use the term "mouse gesture" to refer to the way users draw something, point to an object, circle an area, and so forth.

We try to avoid the issue of natural language processing in order to keep this project within reasonable boundaries. Of course, some verbal dialogue has to be performed, mouse gestures can only be a complement. The link between this project and computational linguistics does not concern the techniques for natural language processing, but the underlying rhetoric structures (arguments, refutations, refinements, illustration, ...) (Baker, 1993; Cawsey, 1993) and belief structures (knowledge gaps, conflicts, awareness, ...) (Cohen & Perrault, 1979). We focus on dialogue models specific to collaborative problem solving. Mouse gestures will be characterized with respect to a more abstract terminology, as classes of communicative acts (Searle, 1969).

2.2.1.3. Multi-Agents Systems

Distributed artificial intelligence (DAI) studies "*how a loosely coupled network of problem solvers can work together to solve problems that are beyond their individual capabilities*" (Durfee et al., p. 85). DAI proposes computational models of collaboration that are not directly useful for this project (human-computer collaboration) for two reasons. Firstly, those models generally include more than two agents and focus on task distribution rather than on building a shared understanding. Secondly, communication between agents is generally considered as free of noise, while we focus on mechanisms for repairing communication breakdowns. Nevertheless, researchers in DAI have studied several variables which are relevant to our work: the balance 'central control versus agent autonomy', the degree of knowledge inconsistency among agents, the heterogeneity of agents, the size of agents (Jennings, 1992; Bird, 1993). Durfee et al. (1989) report interesting observations:

- "some degree of inconsistency among nodes is beneficial to network problem solving " (p. 117)
- "the costs to obtain consistent, complete, and up-to-date information might be prohibitively large in terms of delays due to synchronisation and message transmission" (p. 121)
- "Nodes that are free to generate potentially incorrect and inconsistent tentative results can more completely explore the space of possible solutions." (p 122).

An interesting aspect in DAI is that the size of a computational agent is arbitrary. An agent can be a single neurone, a functional unit (e.g. the 'edge detector' agent), an individual or a society. The agent's size is not a property of the modelled entity. It is function of the observer's distance to the object. This **tuneable granularity** is interesting for our approach. At a low level of granularity, each partner in a collaborative process is considered as a (large) agent. At a higher granularity level, the processing performed by each partner can be decomposed into a set of (smaller) agents. The individual and social planes of cognition can hence be described with the same formalism. This isomorphism supports the theoretical approach we outlined above: to pay primary attention to the similarities between the individual and social planes. Allowing to this individual/social isomorphism, the reflection performed by an agent can be modelled as a conversation among its sub-agents. This corresponds to the idea, shared by Piaget and Vygotsky, that thinking is a discussion we have with ourselves. Moreover, there are interesting similarities between this idea of inner speech (talking to oneself) and the concept of social grounding. Krauss and Fussell (1991) observed that, during social grounding, the expressions used to refer to objects tend to be progressively abbreviated (provided that the partner confirms his or her understanding along the shortening process). The same phenomenon of abbreviation is observed during internalization (Kozulin, 1990; Wertsch, 1979, 1991), i.e. during the process that transforms social speech into inner speech.

2.2.1.4 Human-computer collaboration

Most research on human-computer interaction (HCI) does not correspond to our definition of collaboration. While this project concerns joint cognitive systems, most systems fall into the 'assistant' or 'replacement' relationship. The 'assistant' relationship is based on task distribution: the system helps the human user by performing sub-tasks for which it is more efficient. The user keeps supervisory control (Sheridan, 1991). The 'replacement' relationship is illustrated by the early work on expert systems: the system proposes a solution and explains it on request (by showing the trace of all fired rules). Many efforts have been devoted to the improvement of explanation facilities. However, empirical results show that a "good advice is more than recommending a solution" (Woods & Roth, 1991, p. 19): a solution cannot be completely understood and agreed on if it has not been jointly constructed. Kantowitz and Sorkin (1987) wrote:

> "Instead of thinking about whether a task should be performed by a person or by a machine, we should instead realize that functions are performed by people and machine together. Activities must be shared between people and machines and not just allocated to one or the other".

In the field of expert systems, this "joint construction" can be understood in two non-exclusive ways. The first interpretation leads to participatory design: the future user must participate in building the system (Clancey, 1993). We propose another interpretation: to increase the depth of user-system interaction, i.e. **the extent to which the interaction with the user influences the system's reasoning**. We discriminate three levels of **i**nteraction:

Level 1: The interaction occurs when the reasoning has been completed. The interaction concerns the explanation and does not impact on the system's problem solving process.

Level 2: The user can modify the problem state between two cycles of the inference engine. Thereby, the user influences the solution process.

Level 3: The user interacts with the system about rule selection and rule instantiation. The solution process is hence jointly driven by the user and the system.

In our previous work on MEMOLAB (see section 2.2.2), we achieved a level 2 interaction between a human learner and a rule-based expert. We therefore developed a dedicated inference engine (Dillenbourg et al., 1993). This work led us to believe that we could reach level 3 by designing joint mechanisms of rule selection and instantiation that are inspired by the mechanisms observed in social grounding. We refer to the potential result as **an interactive inference engine**.

2.2.1.5 ESF program "Learning in Human and Machines"

This project is connected to a multidisciplinary research program of the European Science Foundation (ESF) entitled "Learning in Human and Machines". It concerns the analysis, comparison and integration of computational approaches to learning and research on human learning. The human learning perspective comprises in particular psychological research, but also contributions from sociology and educational science. The computational perspective comprises in particular machine learning research, but also contributions from other areas of artificial intelligence, such as research on intelligent tutoring systems and on multi-agent systems. It should be noted that the program does not concern sub-symbolic approaches (which are covered by another ESF program on neurosciences). The program is divided into five task forces. The main contractor of the current project, Pierre Dillenbourg, is the leader of Task Force 5 "Collaborative learning". The first co-contractor, Prof. Mendelsohn, is also member of the program steering committee. This program will run from 1994 to 1997. It is partly funded by the 'Fonds National Suisse de la Recherche Scientifique' (Division 1). The program budget does not cover research projects but only research workshops.

The reason why we refer to the ESF program is that we designed the research plan accordingly to the research methodology adopted by the task force 5 "collaborative learning". In order to make ESF workshops more fruitful, protocols of human-human collaborative problem solving will be shared by participants. They will serve as a concrete basis to compare theoretical positions and computational models. The second stage of our project aims to collect such protocols. The third stage will not only be driven by our own analysis of the protocols. It will receive multiple enlightenment from scientists with various theoretical standings and from various disciplines. Moreover, they will enrich the analysis by providing protocols collected in other settings.

2.2.2 Related research at TECFA

This project nuilds on our previous work on human-computer collaboration (Dillenbourg, to appear). We designed and experimented PEOPLE POWER, a human-computer collaborative system, i.e. a learning environment in which the human learner collaborated with a second learner, simulated by the machine (Dillenbourg and Self, 1991). More recently, we built and experimented MEMOLAB, a learning environment in which the human learner collaborates with a computerized expert (Dillenbourg et al, 1993). These experiments revealed the need for new interaction techniques in human-computer collaboration. They also showed the feasibility of more interactive inference engines.

In PEOPLE POWER, the human learner and the machine learner played with an electoral simulation. They tried to gain seats by moving wards from one constituency to another. The computerized co-learner had a set of naive rules for reasoning about elections. The experiment conducted showed that human learners were able to interact with this rule-based agent, namely to point to a particular rule and to instantiate it with problem data. However, it appeared that this type of discussion 'at the knowledge level' was secondary. The primary task of subjects was to work on the graphical problem representation (a table of votes per party and per ward). Human-computer collaboration would have been more fruitful if it had concerned this problem representation instead of more abstract knowledge. In the design of the next system, MEMOLAB, we paid attention to this issue.

In MEMOLAB, the human learner and the machine expert jointly constructed an experiment on human memory. Collaboration is based on what can be the most easily shared between a person and a machine: the interface. Let us imagine two rule based systems that use a same set of facts. They share the same representation of the problem. Any fact produced by one of them is added to this shared set of facts. Hence, at the next cycle of the inference engine, this new fact may trigger a rule in either of the two rule-bases. Now, let us replace one computerized expert by a human learner. The principle may still apply provided we use an external problem representation instead of the internal one. The shared set of facts is the problem representation as displayed on the interface (see figure 1). All the conditions of the machine's rules refer only to objects displayed on the screen. The actions performed by the rules modify the problem representation.

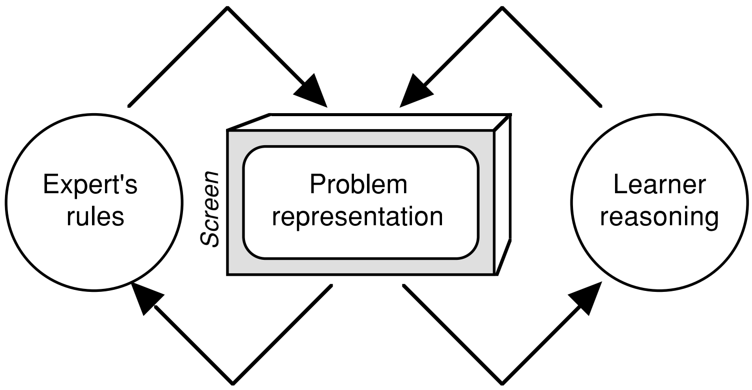

Figure 1: Opportunism in human-machine collaboration in MEMOLAB

In short, the shared representation is visible by both partners and can be modified by both partners. We do not claim that they share the same 'internal' representation. Sharing an external representation does not imply at all that both partners build the same internal representation. The shared concrete representation simply facilitates discussion of the differences between the internal representations and hence supports grounding mechanisms. Some recent experiments with MEMOLAB (Dillenbourg et al., 1993) revealed that such mechanisms may occur between a human and a machine: the learner perceives how the machine understands him (i.e. he makes a diagnosis of the machine diagnosis) and reacts in order to correct eventual misdiagnosis:

- "He supposes that I wanted the subjects to do something during 40 seconds. I wanted the subjects to do nothing."
- "He'll ask me what I wanted to do. And then, since... he'll base himself on wrong things since this is not what I want to do."

These mechanisms can be formalised as nested beliefs (see section 2.2.1.2): "the user believes that the system believes that he believes X" is written "belief (user, belief (system, belief (user, X)))". If the learner notices a misunderstanding, he might start a dialogue with the system to repair the system misdiagnosis. This may sound too heavy for natural dialogues, but we do it in everyday conversations, as illustrated by the fictitious example below:

A "Numbers ending with 4 or 6 are even"
B "778 is also even"
C "I did not say that those ending with an 8 are not even."

When speaker A repairs speaker B's misunderstanding, he does not simply repeat what he said previously. He reinterprets his first utterance from B's point of view in order to repair what B has understood. Interpreting what one said from one's partner's viewpoint corresponds to a learning mechanism referred as 'appropriation' in the socio-cultural theories of learning (Newman, 1989; Rogoff, 1990). These misunderstandings and the subsequent repair mechanisms are necessary to build a shared representation of the problem. Hence, our goal is not to design collaboration techniques that avoid any misunderstanding (if that was possible), but to build techniques that provide the flexibility required for negotiating meanings. In the current implementation of MEMOLAB, rule variables unambiguously refer to screen objects. To support social grounding, mechanisms, the instantiation of variables should not be a internal process, but the result of some interaction with the learner.

2.2.3 Detailed research plan

Stage 1: Adapt / develop a computer-supported collaborative system

Approximate duration : 6 months

The experiments that we want to conduct (see step 2) will concern two people collaborating on remote terminals. It would be easier to observe conversations between two humans sitting side-by-side and using a piece of paper. The reasons for choosing computer-supported collaboration are:

1) The observations in a natural collaborative setting would include a large number of verbal and non-verbal communicative acts that we could not translate in terms of human-computer interactions. With a system for computer-supported collaborative work (CSCW), we restrict the range of interactions available to the users. However, we will not coerce the users to interact only with facilities available in human-computer interaction, otherwise some interesting aspects of human-human conversation might be inhibited.

2) A computerized collaborative setting can be controlled: we can switch some communication channel on or off, and thereby use the communication channel as an independent variable.

The requirements for the system are:

- two users interact on remote terminals;
- they solve a problem presented on the screen;
- they can talk via a microphone;
- they share a sketchpad where they can draw a diagram;
- they share a notepad where they can take notes;
- they see each other's mouse cursor, in such a way that they can point to an object or a zone from the problem representation or from the constructed diagram;
- the communication components (sound canal, sketchpad and notepad) are independent, in such a way that the experimenter can switch off a communication channel.

The sound channel will function in two modes for the purposes of experiments described below: the permanent mode, in which both partners can permanently hear what the other says, and the telephone mode, in which the sound channel is on when one partner pushes some button (and hence has some privacy to think aloud when the channel is OFF).

The core part of this system, the problem solving environment, will later be reused to develop the human-computer collaborative system. This environment may be very simple. It includes an interface allowing the user to solve the problem and the code necessary to respond to the user's actions. The problem to be solved will be selected according to the following criteria:

- the problem solving process must involve several steps and be complex enough to justify a collaboration;
- the problem must be semantically rich, it cannot be a formal problem that inhibits any conceptual negotiation;
- the problem must be solvable by a rule-based system (to be later integrated in the human-computer collaborative system)

This collaborative system will not be developed from scratch, but preferentially by assembling available pieces of software. For the sketchpad for instance, we have tested several products such as WSCROLL or COLLAGE, systems that enable two or more people to draw and type messages on shared sheets, from different machines. We will review the available "groupware" and select the one which is the most appropriate for the sketchpad and the notepad (e.g. COLLAGE does both). Another possibility is to use some tool kit dedicated to the design of collaborative systems (Dewan, 1993; Hill et al, 1993). For the problem solving environment, we have a long experience of developing in Common Lisp. We might however choose another development language that best fits the selected groupware.

Stage 2: Observe grounding mechanisms in human-human collaborative problem solving

Approximate duration : 12 months

Our experiments will address two issues:

1. How do people use external references during social grounding ?
2. What is the role of grounding mechanisms in problem solving ?

We will conduct experiments in three settings. We hope to get a better understanding of grounding mechanisms by varying the communication channels (independent variable). We will compare the protocols collected in setting 1, with those collected in situations 2 and 3. In each setting, pairs of adults will use the computer-supported collaborative software developed in stage 1. Each partner will be located in a different room, on a remote terminal and will interact via the network. The activity of each partner will be videotaped. Computer-mediated communication will be recorded.

In **setting 1**, the communication facilities include the sound channel, the sketchpad and the notepad. This the main setting, in which we aim:

- to identify 'grounding episodes', i.e. protocol segments that correspond to a symbol grounding process;
- to decompose grounding episodes into elementary acts (speech acts, mouse acts), by referring to the categories identified by Clark and Brennan (1991);
- to search if some patterns of elementary acts can be induced over episodes, by referring to the terminology used in dialogue models (Baker, 1993);

In **setting 2**, the sound channel is set to the 'telephone' mode: each partner can think aloud when the microphone is OFF and communicate with the other when the microphone is ON. With think-aloud protocols plus the record of the users' actions in the problem solving environment, we will observe at which stages of the solution process grounding episodes occur. The goal is to understand how the grounding mechanisms participate into the solution process. This second situation will also enable us to analyze protocols from an angle specific to the distributed cognition approach: seeking for similarities between social dialogues and reflective dialogues. Finally, these think-aloud protocols are necessary for designing the rules of the computerized partner to be developed in stage 3.

In **setting 3**, the sound channel is permanently OFF. The users hence communicate via the notepad (written communication). We thereby isolate the role of grounding gestures with respect to other grounding mechanisms. Clark and Brennan (1991) have shown that grounding techniques change with the medium, because media vary with respect to delaying speech, turn taking, making and repairing errors, and so forth. By comparing setting 1 (sound ON) and setting 3 (sound OFF), we will observe how grounding mechanisms adapt to the communication medium. In oral communication, disambiguating sub-dialogues are cheap and fast. Since this is not the case for written messages, the diagrams may get more importance. On the other hand, written communication has several advantages. It leaves a trace on which users can come back to repair misunderstanding. Moreover, it will be the channel for human-computer verbal messages.

The outcome of this second phase will be a high-level description of human-human grounding techniques, as structures of communicative acts with respect to joint problem solving process. This description will be independent from the problem and from the agent (human/machine).

Stage 3: Build and experiment a human-computer collaboration system (HCC)

Approximate duration: 18 months

The goal of this third stage is to develop new interaction techniques that enhance the collaboration between a human user and a knowledge-based system. The user may understand some rules differently from the system because he has not the same frame of reference as the system designer (Clancey, 1991). He or she may not see how a rule fits the problem data. This is why the collaboration between the user and the system requires grounding mechanisms.

A human-computer collaboration system includes three components:

- A problem solving environment (we will reuse the one constructed in stage 1)

- A knowledge-based agent.

  The design of the agent constitutes the main task. The reasoning of this agent must be compatible with (but not necessarily identical to) the way humans solve the same problem. Hence, the design of the rulebase will be inspired by the think aloud protocols collected in phase 2, setting 2. The challenge it to develop an interactive inference engine that satisfies the requirements presented in section 2.2.1.4 (level 2 interaction). Selection and instantiation of rules is jointly performed by the user and the system. Given our experience (see section 2.2.2), such an engine should be object-oriented, i.e. the problem representation is a set of objects. In MEMOLAB, these objects are displayed on the screen. Thereby, the user and the system share a concrete representation of the problem. The joint instantiation of rule variables by the problem data can be done by pointing an object on the screen. This technique raises however a fundamental issue. Screen objects are particular instances, while rule variables concern classes. The generality ratio between a variable and the object by which this variable is instantiated must be negotiated by participants ("Do you mean any piece like that can be moved?"). The precise role of those dialogues will be clarified by the observations conducted in phase 2.

- The interface between the human user and the computerized agent.

  The difficulty is not to build an interface but to integrate interface acts in an inference engine (see previous point). The original piece in this interface will be the sketchpad. This sketchpad will be simpler than the one used in human-human collaboration. Since the system has to understand the drawings, we can not allow the user to draw any kind of diagram, otherwise we will face issues of image interpretation that are beyond the scope of this project. Instead, the human and the machine users will use a limited set of graphical objects (lines, boxes, circles, arrows,...). The set of objects and symbols to be provided will be determined after the observations of the drawings used in human-human collaboration (phase 2, settings 1 and 3). For these issues, we have a contact with Prof. Pun, a specialist in computer vision (Université de Genève).

2.2.4 Research agenda

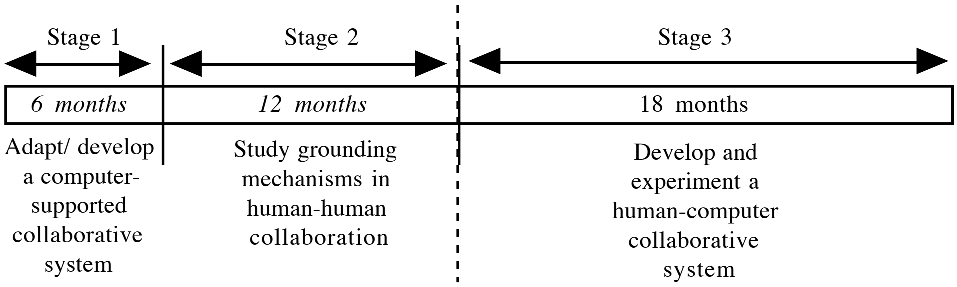

The description of this third stage is rather general since the precise specifications will result from stage 2. If the funding requested for this project was cut down, we could split the project in two parts and submit a new application at the end of stage 2. The third stage could then be more precisely described. The new application could be submitted as a project in computer science.

2.2.5 Importance of work

This project aims to contribute to the understanding of social grounding mechanisms. This understanding is a major step in the validation of theories on distributed cognition. We aim to make these theoretical positions more operational and to translate them in terms of computational mechanisms. Because this fundamental research has a computational side, it will also have an impact at the applied level, namely the development of more efficient human-computer interaction techniques. Such techniques are critical for the future of three software categories: **expert systems, learning environments and CSCW systems.**

Because this project focuses on the role of images and diagrams in joint problem solving, it is also relevant for the development of the **multimedia** technologies. Currently, most multimedia system are weakly interactive: interaction concerns the selection and the display of fixed or animated images, but the system and the user do not interact about the images. Images are considered an add-on. There are few efforts to study how images may more deeply impact on the user's work. This project investigates the role that images could play in collaborative problem solving.

2.2.6 References

Baker, M. (1992) The collaborative construction of explanations. Paper presented at the "2èmes journées Explication du PRC-GDR-IA du CNRS, Sophia-Antipolis.

Baker, M. (1993) Negotiation in Collaborative Problem-Solving Dialogues. Rapport CR-2/93. CNRS, Laboratoire IRPEACS, Equipe Coast, Ecole Normale Supérieure de Lyon.

Bargh, J.A. & Schul, Y. (1980) On the cognitive benefits of teaching. *Journal of Educational Psychology*, 72 (5), 593- 604.

Behrend, S.D. & Roschelle, J. (to appear) The construction of shared knowledge in collaborative problem solving. In C.E. O'Malley (Ed). *Computer Supported Collaborative Learning.* New York: Springer-Verlag.

Bird, S.D. (1993) Toward a taxonomy of multi-agents systems. *International Journal of Man-Machine Studies*, 39, 689-704.

Blaye, A., Light, P., Joiner, R. & Sheldon, S. (1991) Collaboration as a facilitator of planning and problem solving on a computer based task. *British Journal of Psychology*, 9, 471-483.

Butterworth, G. (1982) A brief account of the conflict between the individual & the social in models of cognitive growth. In G. Butterworth & P. Light (Eds) *Social Cognition* (3-16). Brighton, Sussex: Harvester Press.

Cawsey, A. Planning Interactive Explanations. *International Journal of Man-Machine Studies,* 38, 1993, 169-199.

Chi, M.T., 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.

Clancey, W.J. (1991) The frame of reference problem in the design of intelligent machines. In K. Van Lehn (Ed.) *Architectures for Intelligence: The twenty-second Carnegie symposium on cognition* (357-424). Hillsdale: Lawrence Erlbaum.

Clancey, W.J. (1992) Guidon-Manage Revisited: A Socio-Technical Systems Approach. *Journal of Artificial Intelligence in Education,* Vol. 4, 1, 5-34.

Clark, H.H. & Brennan S.E. (1991) Grounding in Communication. In L. Resnick, J. Levine and S. Teasley (Eds).*Perspectives on Socially Shared Cognition* (127-149). Hyattsville, MD: American Psychological Association.

Cohen, P.R. & Perrault, C.R. (1979) Elements of a Plan-Based Theory of Speech Acts. *Cognitive Science*, 3, 177-212.

Dewan, P. (1993) Tools for implementing multi-user interfaces. In Bass and Dewan (Eds) *User Interface Software*, John Wiley.

Dillenbourg, P. (1991) Human-Computer Collaborative Learning. *Doctoral dissertation*. Department of Computing. University of Lancaster, Lancaster LA14YR, UK.

Dillenbourg, P. (to appear) Distributing cognition over brains and machines. In S. Vosniadou, E. De Corte, B. Glaser & H. Mandl (Eds), *International Perspectives on the Psychological Foundations of Technology-Based Learning Environments*. Hamburg: Springer-Verlag.

Dillenbourg, P., Hilario, M., Mendelsohn, P., Schneider D. and Borcic, B. (1993) The Memolab Project. Research Report. TECFA Document. TECFA, University of Geneva.

Doise, W. & Mugny, G. (1984) *The social development of the intellect*. Oxford: Pergamon Press.

Durfee, E.H., Lesser, V.R. & Corkill, D.D. (1989) Cooperative Distributed Problem Solving. In A. Barr, P.R. Cohen & E.A. Feigenbaum (Eds) *The Handbook of Artificial Intelligence*, (Vol. IV, 83-127). Reading, Massachusetts: Addison-Wesley.

Jennings, N. (1992) Joint intentions as a model of multi-agent cooperation. Technical Report 92/18. Department of Electronic Engineering. University of London.

Hill, R.D., Brinck, T., Patterson, J.F., Rohall, S.L. & Wilner, W.T. (1993) The Rendezvous language and architecture: Tools for constructing multi-user interactive systems. *Communications of the ACM*, 36, (1), 62-67.

Kantowitz, B.H. & Sorkin, R.D. (1987) Allocation of functions. In G. Salvendy (ed.), *Handbook of human Factors*. New York: Wiley.

Hancock, P.A. (1992) On the future of hybrid human-machine systems. In J.A. Wise, V.D. Hoipkin, and P Stager (Eds) *Verification and Validation of Complex Systems: Human Factors*. NATO ASI Series F: Computer and Systems Sciences, Vol. 10, 61-85.

Kozulin, A. (1990) *Vygotsky's psychology*. A biography of ideas. Harvester, Hertfordshire.

Krauss, R.M. & Fussell, S.R. (1991) Constructing shared communicative environments. In L. Resnick, J. Levine and S. Teasley (Eds). *Perspectives on Socially Shared Cognition* (172-202). Hyattsville, MD: American Psychological Association.

Lave, J. (1988) Cognition in Practice. Cambridge: Cambridge University Press

Miyake, N. (1986) Constructive Interaction and the Iterative Process of Understanding. *Cognitive Science*, 10, 151-177.

Newman, D. (1989) Is a student model necessary? Apprenticeship as a model for ITS. *Proceedings of the 4th AI & Education Conference* (pp.177-184), May 24-26. Amsterdam, The Netherlands: IOS.

O'malley, C. (1987) Understanding explanation. *Paper presented at the third CeRCLe Workshop 'Teaching Knowledge and Intelligent Tutoring (April)*, Ullswater, UK.

Perret-Clermont, A.-N., Perret J.-F. & Bell N. (1991) The Social Construction of Meaning and Cognitive Activity in Elementary School Children. In L. Resnick, J. Levine and S. Teasley (Eds). Perspectives on Socially Shared Cognition (41- 62). Hyattsville, MD: American Psychological Association.

Resnick, L.B. (1991) Shared cognition: thinking as social practice. In L. Resnick, J. Levine and S. Teasley (Eds). *Perspectives on Socially Shared Cognition* (127-149). Hyattsville, MD: American Psychological Association.

Rogoff, B. (1990) *Apprenticeship in thinking*. New York: Oxford University Press

Searle, J. (1969) *Speech acts: An essay in the philosophy of language*. Cambridge: Cambridge University Press.

Sheridan, T.B. (1991) Task allocation and supervisory control. In M.Helander (Ed) *Handbook of Human-Computer Interaction*, 159-173. Amsterdam: North Holland.

Suchman, L.A. (1987) *Plans and Situated Actions*. The problem of human-machine communication. Cambridge: Cambridge University Press.

Wertsch, J. V. (1979) The regulation of human action and the given-new organization of private speech. In G. Zivin (Ed) *The development of self-regulation through private speech*, 79-98. New York: John Wiley & Sons.

Wertsch, J.V. (1991) A socio-cultural approach to socially shared cognition. In L. Resnick, J. Levine and S. Teasley (Eds). *Perspectives on Socially Shared Cognition* (1 - 20). Hyattsville, MD: American Psychological Association.

Woods, D.D. & Roth, E.M. (1991) Cognitive System Engineering. In M.Helander (Ed) *Handbook of Human-Computer Interaction*, 3-35. Amsterdam: North Holland.