

NEGOTIATION SPACES IN HUMAN-COMPUTER COLLABORATIVE LEARNING

PIERRE DILLENBOURG

TECFA., FPSE, Université de Genève,
9, Route de Drize, CH 1227 Carouge, Suisse
Phone + 41 22 705 96 93, Fax + 41 22 342.89.24
pdillen@divsun.unige.ch WWW: tecfa.unige.ch

MICHAEL BAKER

C.N.R.S.- GRIC-Equipe COAST ENS de Lyon (F)
Allée d'Italie, 46, 69364 Lyon cedex 07, France
Phone +33 72 72 85 38, fax + 33 72 80 80
mbaker@ens-lyon.fr

Abstract. This paper compares the negotiation processes in different learning environments: systems where an artificial agent collaborate with the human learner, and systems where the computer supports collaboration between two human users. We argue that, in learning context, collaboration implies symmetry between agents at the design level and variable asymmetry at the interaction level. Negotiation is described as a collection of different spaces defined with seven dimensions: mode, object, symmetry, complexity, flexibility, systematicity and directness. We observed that human-human negotiation jumps between spaces, switching easily between modes of negotiation, connecting the various objects of negotiation while the 'disease' of human-computer collaborative systems was to be fixed within one negotiation space.

Résumé. Cet article compare les formes de négociation utilisées dans différents environnements d'apprentissage, au sein desquels soit l'apprenant collabore avec un agent artificiel, soit deux apprenants collaborent via un collecticiel. Dans un contexte éducatif, nous défendons l'idée d'une symétrie entre agents sur le plan de leur conception et d'une asymétrie variable sur le plan de l'exécution. Nous analysons la collaboration dans ces systèmes en termes d'espaces de négociation définis au moyen de 7 dimensions: le mode, l'objet, la symétrie, la complexité, la flexibilité, le caractère plus ou moins systématique des agents et la possibilité de communication indirecte. Nous avons observé que lorsque deux utilisateurs négocient, ils passent fréquemment d'un espace de négociation à un autre, alors que la négociation avec un agent artificiel reste souvent bloquée au sein d'un même espace.

Keywords. collaboration, negotiation, apprentissage.

This paper presents a comparative study of negotiation in several Human-Computer Collaborative Learning Systems (HCCLS) and Computer-Supported Collaborative Learning Systems (CSCLS) that we have implemented. To make this comparison, we use two key concepts, with a view to establishing design principles for future HCCLS/CSCLS : variable asymmetry and negotiation spaces. We start with some definition of collaboration, justifying the role of negotiation and then the notion of variable asymmetry. We introduce the notion of negotiation spaces and instantiate these concepts in five systems in order to illustrate our concepts in different problem-solving/learning domains, and with respect to different types of agents.

1. Collaboration implies negotiation

In Distributed Artificial Intelligence research [see e.g. Bond & Gasser 88], negotiation is almost universally viewed as a process by which *conflicts* (with respect to resource allocations) may be resolved. Our own view is more closely related to language sciences research [e.g. Edmondson 81 ; Moeschler 85; Clark & Schaefer 89; Roulet 92] and to some areas of agent theory [Galliers 89].

We do not view the existence of 'conflict' - whether openly declared and recognised or not - as essential to the definition of negotiation. All that is basically required is that the interacting agents possess the *mutual goal* of achieving *agreement*, with respect to some set of *negotia*, or objects of negotiation. Usually, several dimensions of *negotia* will be negotiated simultaneously. The initial state for negotiation is thus an *absence* of such agreement, that may or may not include conflict. In task-oriented interactions, negotiation can occur on three main levels : (1) communication (meaning, signification of utterances, words, ...), (2) task (problem-solving strategies, methods, solutions, ...) and (3) management of the interaction on previous levels 1 and 2 (coordination , feedback on perception, understanding, attitudes) [Allwood et al 91; Bunt 89]. Any interaction, conversation, dialogue involves negotiation, at least on the first level.

The second main defining characteristic of negotiation - especially in task-oriented interactions - is that specific *strategies* exist for achieving agreement in the interaction [Baker 94]: *mutual refinement* (each agent successively refines the contribution of the other), *argumentation*, [Sycara, 88; Baker, *to appear*] (agents attempt to verbally resolve mutually recognised conflicts) and '*stand pat*', (one agent successively elaborates a proposal, receiving elicitation, positive and negative feedback, ... from the other). Each strategy is associated with specific characteristic *communicative acts* [see Baker 94 for more details]. For example, the mutual refinement strategy is associated basically with initiative acts [Moeschler 85] such as OFFERS, and reactive acts that provide positive or negative feedback, such as ACCEPTANCE, REJECTION and RATIFICATION.

We view negotiation as a distinctive feature of collaboration, while it may not occur in cooperation. Actually, many researchers use 'collaboration' and 'cooperation' as almost synonymous terms. They use "collaborative learning" in any situation in which agents interact about some task designed to promote learning. We view collaboration as something much more specific. If, for example, they have divided the task up completely, i.e. if the agents produce separate solutions in parallel with little interaction between them, we would not say that they are 'collaborating' but rather that they cooperate. Roschelle and Teasley [95] have expressed the cooperation/collaboration distinction as follows: "*Collaboration is a coordinated, synchronous activity that is the result of a continued attempt to construct and maintain a shared conception of a problem* ». Cooperative work can be accomplished by the division of labour among participants, as an activity where each person is responsible for a portion of the problem solving. Collaboration is a specific form of synchronous¹ cooperation in interaction where *negotiation* takes place

¹ Actually, the role of synchronicity in collaboration is complex. For instance, MOO communication is generally treated as

simultaneously on all three of the above levels, i.e. the agents are coordinating their problem-solving interaction, developing shared meanings, and co-constructing problem solutions. Note that such co-construction does not exclude the existence of conflict, since, as many researchers have argued [e.g. Mevarech & Light 92; Baker to appear], the *constructive* resolution of conflicts may be a key factor in successful collaboration.

2. Negotiation requires symmetrical interaction possibilities

The relationship between the user and a system is generally asymmetrical : the user and the system do not perform the same actions, they do not have the same role in decision making. For instance, in the 'assistant' or 'slave' metaphors, the balance of control is on the user's side, while most expert systems reason independently from the user, referring to her only for requesting data or generating explanations. Some knowledge-based systems functioning in a critique mode [Fischer et al. 91] are more symmetrical. In fact, when talking about HCCLS, the word '*collaboration*' implies that the balance of control will be more *equilibrated*, or *symmetrical*. Such 'symmetry' does not imply 'identity' between human and artificial agents, i.e. that the user and the system will actually perform the same sub-tasks, nor of course that they have equal competencies. In fact, a major concern in design is to set up *a suitable distribution of roles* that optimises the specific skills of each agent [Woods & Roth 88]. For instance, computational agents process systematically large amount of data, at the syntactical level, whilst human agents intervene on subtasks which require common sense, at the semantic level. The 'partners' can thus be viewed as 'complementary'. Dalal and Kasper [94] use the term '*cognitive coupling*' to refer to the relationship between the cognitive characteristics of the user and the corresponding cognitive characteristics of the system

Although such complementarity is a rational choice, we do not believe that it is useful to impose a *single, fixed distribution of roles* in HCCLS. Instead, we have investigated the opposite direction: *human-computer symmetry*. This symmetry implies to give the user and the system the same range of possible (task-level and communicative) actions and symmetrical rights to negotiate decisions. However, for a given task-interaction, the agents' negotiation behaviour and 'power' will, of course, be largely determined by the differences in their respective knowledge concerning the sub-task under discussion. 'Symmetry' of the interaction, thus understood, therefore concerns principally the actions that agent-X can *potentially* perform rather than those it actually performs: both Agent-A and Agent-B can perform sub-task X, but, very often, if A does it, B does not, and

synchroneous while messages take about the same time to travel through the net as e-mail messages, generally viewed as an asynchronous tool. We probably haveto consider the syncronicity from the users subjective point of view. Our hypothesis is that this subjective syncronicity implies that the partners mutually represent the cognitive processes performed by othe other.

vice-versa. In other words, **the possibility of variable interaction asymmetry** (when the system is running) **implies symmetry at the system design level**. Our argument is based on 4 points.

- In an asymmetrical mode, one agent has always the final word, there is no space for real negotiation. In the symmetrical mode, there is no pre-defined 'winner', conflicts have to be solved through negotiation. Precisely, the cognitive processes triggered to make a statement explicit, to justify an argument or to refute the partner's point probably explain why collaborative learning is sometimes more efficient than learning alone. It has been shown that, when one partner dominates the other to the point that (s)he has not to justify his/her decisions, the benefits of collaborative learning are lower [Rogoff 91].

- The main functions of interactive learning environments (explanation, tutoring, diagnosis) have traditionally been implemented as one-way mechanisms, with the system retaining complete control. Recent work tend however to treat them as two-way processes [Dillenbourg, to appear]: an explanation is not simply built by one agent and delivered to the other, but jointly constructed [Baker 92] or negotiated [Baker et al 94]; the quality of diagnosis depends to a large extent on the collaboration between the diagnoser and the subject being diagnosed. Even in tutoring the learner plays an active role which helps the teacher to focus his interventions [Douglas 91, Fox 91].

- In empirical research on human-human collaborative learning, it has been observed that even when division of labour occurs spontaneously, it is not fixed, boundaries change over time [Miyake 86, Dillenbourg et al 95b].

- The progressive transfer of sub-tasks from the machine agent towards its human partner is the core mechanism in the apprenticeship approach (scaffolding/fading), which now inspires the design of many interactive learning environments[Collins et al 89].

3. Negotiation Spaces

In order for negotiation to take place, there must be some 'degree of latitude' [Adler et al 88] available to the agents - otherwise there would be nothing that is negotiable. This defines the global space of negotiation within which the two agents attempt to build a shared understanding of the problem and its solution. Actually, the negotiation space is not naturally flat. We observed [Dillenbourg et al. *submitted*] that human partners do not negotiate a 'unique' shared representation of the problem (as Roshelle & Teasley's definition above mentioned might suggest), but actually develop several shared representations, i.e. they move across a mosaic of different negotiation spaces. These spaces differ by the nature of information being negotiated (e.g. some aspects require explicit agreement more than others) and by the mechanisms of negotiation (e.g. the media being used). However, in human-computer collaboration, more precisely in the HCCLS presented below, we have often designed a unique negotiation space. We will return to this difference in our conclusions.

We believe that the word 'collaborative learning' is too general and therefore attempt to describe the negotiation space(s) more precisely. At the present state of our research we view the following seven dimensions as particularly relevant. We present first the two main dimensions *what* can be

negotiated in a given space (object) and *how* it is negotiated (mode). The five subsequent dimensions describe more specific parameters of these first two general dimensions.

3.1 Dimension 1 : Mode of negotiation

The first axis is the mode of negotiation. Two agents can negotiate by sending messages to each other (hereafter referred to as the **discussion mode**) or by performing actions on the task (hereafter the **action mode**), for instance if one agent undoes the last action of his partner, thus expressing disagreement. In practice, these two modes usually correspond to two interfaces, respectively the agent-agent interface and the task-agent interface (see section 4). Cumming and Self [89] introduced a similar distinction ("task level" versus "discussion level" in Intelligent Tutoring Systems) to emphasise that learning results not only from problem solving activities (as in 'learning by doing'), i.e. at the task level, but also from the reflection upon these activities, at the discussion level.

We stress here that negotiation occurs in both modes : two agents can for instance disagree by uttering statements, the propositions of which are mutually recognised as contradictory, but also by activating opposed commands. These two modes correspond to different interaction styles, namely direct manipulation versus conversational interfaces. The former reduce the 'referential distance' between expressions and object being referred to, but, because of that, offers low possibilities regarding abstraction (e.g. referring to unseen objects) [Frolhich 93]. Terveen [93] introduced the notion of *collaborative manipulation* in which the user and the system collaborate in a *shared workspace*, involving the representation and manipulation of objects.

In the action mode, the negotiation space is the subset of interface commands available to both agents. There can not be action-mode negotiation on parts of the interface which are not available to both agents. In some of our systems, since one agent is human while the other one is computational, the commands can be concretely different, but functionally equivalent². Hence, the negotiation space is defined by mapping function between the commands respectively available to the user and to the system. Finding a good mapping is a difficult part of the design process (see MEMOLAB experiments below). Finally, we will see (See dimension 7) that the effective negotiation space, can be extended by indirect use of some actions.

² However, in Memolab, some subjects complained that they did not see the expert making actions as if they were themselves doing this actions, e.g. by moving the cursor to a menu, pulling it down, selecting a command. We partially implemented such a functionality but did not test whether this would modify or improve the human-machine negotiation.

3.2 Dimension 2 : Object of negotiation

The second axis is the object of negotiation, i.e. *what* is being negotiated. For example, two agents can negotiate what to do next (negotiating **action**), negotiate the knowledge underlying their decisions (negotiating **knowledge**), negotiate how to represent this knowledge (negotiating **representations**), negotiate their mode of interaction (e.g. 'was that a question or a claim?', negotiating turn taking), (negotiating **interaction**). This dimension crosses the previous one (e.g. negotiating the next action can be done through discussion or simply by doing it). What can be considered as an object of negotiation within a specific system depends on a number of factors, principally : *the nature of the task domain* and *the agents' respective degrees of knowledge* with respect to it. Thus for domains where there exists a single 'correct' solution method, and only one of the agents knows it, then this can not (sincerely) be an object of negotiation. However, the conceptual point of view within which the domain is approached (e.g. functional, procedural, ...) may still be negotiable. In more 'open' domains, where there are several possible solution methods, or where 'plausible' or 'uncertain' reasoning is required, the negotiation space on this level may be wider and more symmetrical (see below).

3.3 Dimension 3 : Degree of symmetry

Although we are principally interested in highly symmetrical collaborative systems, the degree of symmetry is in fact a continuous variable: Designers have to decide whether they provide to artificial and human agents a more or less equivalent a range of interface actions and interaction possibilities (i.e. both in action mode and discussion mode - dimension 1). The particular choice along this continuous axis may also vary along dimension 2 (what is negotiable). It also depends on the rights and obligations inherent in agents' social-institutional status, even when transferred to the human-machine case. For example, in a 'traditional' teacher-student interaction, the teacher may conventionally have the right to make the dialogue move "NEGATIVE-EVALUATION", with respect to a previous move of the student, whereas the student may not have the social right to make such a move. In practice, however, the degree of symmetry is also influenced by technical constraints. For instance, although in certain cases both 'teacher' and 'student' should have the social right to make an EXPLANATION move, often this is excluded for the student simply because the system will lack sufficient natural language understanding competence.

3.4 Dimension 4 : Degree of complexity

The degree of complexity in a specific negotiation space refers to the complexity of the interaction that is supported. This will be mainly a function of the types of objects of negotiation that are supported (dimension 2) and will correspond to a certain degree of symmetry. We can, however,

identify a minimal degree of complexity that must be supported in order for the system to be described as a negotiation. Three dialogue moves must at least be supported, to be realised in either or both modes (dimension 1) : OFFER (different negotiation objects proposed as candidates for mutual acceptance), ACCEPTANCE and REJECTION.

A surprisingly rich interaction can be minimally supported with even these three moves, within each of the three negotiation strategies (mutual refinement, argumentation, 'stand pat'). For example, mutual refinement can be supported by sequences of successive offers from both agents, 'punctuated' by acceptance or rejection. In this case, many indirect effects of the acts will be produced (see below) - an OFFER of "we should move ward2 from North to South" by Agent_2, that follows an OFFER of "we should move ward2" by Agent_1 communicates implicit acceptance of Agent_1's offer (since it subsumes it and elaborates on it). Similarly, the argumentation strategy can be implemented in a rudimentary way since an OFFER that follows a REJECT may be contextually interpreted as a DEFENSE of the previously rejected offer, and so on. However, a richer set of dialogue moves will be required to support more extended (and effective) negotiations, particularly those that work on the meaning of offered negotiation objects (e.g. various QUESTION forms). Finally, it should be noted that the complexity may be more or less symmetrical.

3.5 Dimension 5 : degree of flexibility

The degree of **flexibility** concerns basically the degree of *liberty* accorded to each agent to realise or not the different dialogue moves with which they are provided at a given stage of the interaction. Thus turn taking can be constrained or not ; the system can force the two agents to agree on each step before to perform the next one, or leave them marking agreement or disagreement only when they consider it is relevant ; the topic shifts can be constrained or not ; the system can force the two agents to come to a decision at one point before moving on to another, ... and so on. The degree of flexibility is often determined by technological constraints/limitations. However, in HCCLS, ideally the degree of flexibility should be determined on pedagogical grounds. For example, a designer may decide to "script" the interaction in particular ways (e.g. forcing the student to explain each domain-level proposal) because such 'inflexibility' is viewed as potentially promoting reflection and learning.

3.6 Dimension 6 : degree of systematicity

An agent is systematic if, all relevant attitudes and information are communicated on all occasions when they are so relevant. For example, an agent is systematic if it/(s)he communicates (directly or indirectly) disagreement with a proposal when the agent does in fact disagree, and never

communicates disagreement when this is not the case. Similarly, an agent is systematic if it/(s)he communicates information that would support or 'undermine' a proposal when the agent possesses such information, and it is relevant to communicate it. The degree of systematicity is thus a fundamental manifestation of a form of cooperative behaviour, and relates to basic maxims of cooperation such as sincerity, mendacity and helpfulness.

3.7 Dimension 7 : degree of directness

The degree of directness concerns both the functional (user-system) and the intentional (user-author) interactions. In functional interaction, when an agent says for instance to the *other* "*The delay should be 30 seconds*", this utterance can be classified as an offer. However, it also constitutes an indirect (and implicit) invitation to start negotiating the delay duration. This possibility is simply a function of the *pragmatic interpretation* of utterances in a given *context*. In this case, the interface possibilities are used directly *and* indirectly simultaneously.

In intentional interaction, directness refers to the extent to which the agents use the interaction possibilities explicitly provided by a system in a manner that corresponds to the intention of the system designer. For instance, in an early design of the C-CHENE system, we provided a button that make a 'beeping' sound, where we intended that the students should use it for maintaining mutual perception, i.e. attracting the attention of the other, verifying that the other was attending. However, the students did not directly use the button in this way: they discussed and agreed on an 'indirect' way of using this button for coordinating interface actions ("use the beep to tell me when you want to draw something on the interface"). This is 'indirectness' from the system designer's point of view. We must therefore distinguish the **designer-negotiation-space** from the **user-negotiation-space**, where the latter may often surpass the former. Whilst the space of such a user negotiation space may be *predictable* by the designer, it is clear that all such possible indirect uses can not be so predicted. In other words, the degree of directness is not a design parameter, but rather something that one can only measure when analysing actual interactions.

4. Characterizing negotiation spaces in HCCLS and CSCLS.

The core of the human-computer collaborative systems (HCCLS)(see figure 1) is a triangular relationship between two agents, the human user and one (or more) machine agent, plus the task (generally a computer-based activity). This triangle theoretically includes three interfaces: (1) between the human agent and the task, (2) between the artificial agents and the task and (3) between the two types of agents. We consider hereafter the two agent-task interfaces (2 and 3) as forming a single interface, because our symmetry postulate leads us to seek for the optimal overlapping between these two interfaces.

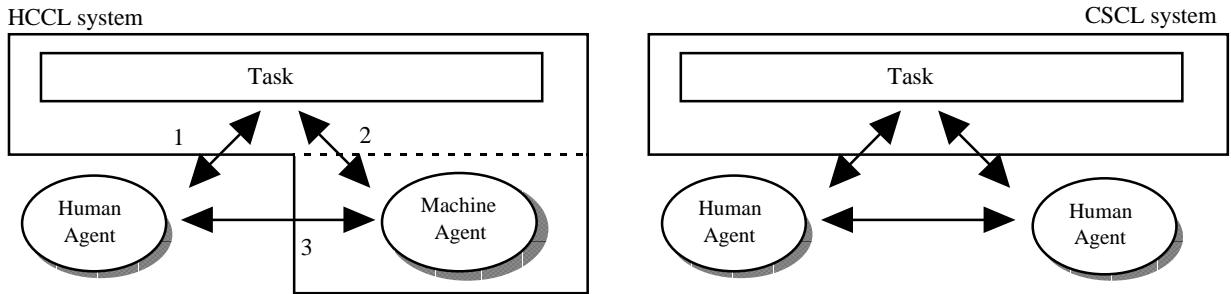


Figure 1: Basic architecture in collaborative learning systems: human-computer (left) or human-human (right)

The first systems to be described below (People Power and Memolab) do not include an explicit model of human-machine collaboration in the sense that ‘model’ can take in AI. The implicit model of negotiation is encrypted in the interface (e.g. in the difference between what each agent can do) and in the design of the computational agents (rule-based systems). The KANT system includes an explicit model of negotiation. The CSCL systems (BOOTNAP and C-CHENE) embody a model for negotiation within the mechanisms that the system provides for supporting collaboration and negotiation between human agents.

4.1 People Power

The task. People Power [Dillenbourg & Self 92] is a microworld in which two agents design an electoral system and run an electoral simulation. They play a game known as gerrymandering. Their country is divided into electoral constituencies, and constituencies into wards. The number of votes per party and per ward is fixed. Their goal is to modify the distribution of wards per constituency in such a way that, with the same number of votes, their party can gain one or more seats in the parliament.

The agents. Both agents are learners, the machine agent simulating a co-learner. It is implemented as a small rule-based system, including rules such as "if a party gains votes, it gains seats". Such a rule is not completely wrong, but still over-general. By discussing the issues with the human learner and by analysing elections, the co-learner progressively finds out when this rule applies (rule specialisation). The inference engine was designed to account for learning mechanisms specific to collaborative problem solving.³

Evaluation. The subjects complained about the high systematicity and low flexibility, i.e. systematic agreement or disagreement quickly turns out to be boring. Human conversation heavily

³ PEOPLE POWER is written in Procyon Lisp for Macintosh. The co-learner's inference engine is implemented as a dialogue between an agent and itself, i.e. as a monologue. Learning is modelled by the fact that, during its monologue, the co-learner replays - modus modendi - the dialogue patterns that it induced during its interaction with the real learner. For a better description of learning algorithms, see Dillenbourg [1992].

relies on assumptions such as default agreement (one infers agreement as long as there is no explicit disagreement). The cost of systematicity was increased by the mode (discussion): the human agent had to read and enter explanations, which takes more time than the action mode. Human conversation smoothly uses non-verbal cues for marking agreement or disagreement (facial expressions, gestures, ...) and does not have to make negotiation always explicitly (indirectness). Another problem was the fact that negotiation occurred almost exclusively in the discussion mode and about knowledge. When we observed the subjects reasoning about the possible changes in the electoral systems, they seemed to move mentally the columns of votes (displayed in a table of votes party X ward). Sometimes this reasoning was supported by hand gestures in front of the screen, but it was not supported by the system. This table-based reasoning was functionally equivalent to the rule-based reasoning, but grounded in concrete problems, whilst arguments canned in rules express general knowledge. Moving negotiation into the action mode means enabling *learners to dialogue with table commands*: two contradictory column moves express a more precise disagreement than two abstract (rule-based) statements on electoral mechanisms. This concern for moving negotiation in the action mode was central to the next system.

Features	Description
<i>Mode:</i> Discussion.	In the action mode, the space is reduced to a single command: choosing which ward has to be moved where. Hence, all negotiation occurs through discussion, even though the discussion interface is rudimentary. When an agent disagrees with the other's action, it may ask him to explain his decision. An explanation is a proof that moving ward X to constituency Y will lead to gain a seat; a refutation is a proof of the opposite. When the machine learner explains, it provides the trace of the rules used in the proof. The human learner introduces explanation by selecting an argument and instantiating its variable with the problem data (parties, wards, etc...).
<i>Object:</i> Knowledge	Agents discuss knowledge about electoral systems, expressed for the machine agent as a set of rules. For the human agent, the list of arguments available for negotiation was the list of rules used by the machine learner. He was hence forced to reason within the conceptual framework of her machine partner. There was no possibility to negotiate a concept (e.g. the constituency concept) independently from its use in a rule, nor to negotiate strategy or interaction.
<i>Symmetry:</i> High	This system is symmetrical. Any step of the explanation can be accepted or refuted by any agent. Any refutation can itself be refuted and so on. However, at the end, the human learner always had the possibility to take over and to impose his decision (finally, he is supposed to learn in this system). Hence the system is not 100% symmetrical.
<i>Complexity:</i> Minimal	The set of dialogue moves is binary: agreement or disagreement. A partner can agree or disagree on the other' statement as a whole, but he cannot express partial disagreement, non-agreement (feeling doubts about an implication without being able to prove the opposite), misunderstanding (on a concept used in the rule), and so forth.,...
<i>Flexibility:</i> Low	Strict turn taking is imposed. Any step in an explanation had to be agreed or refuted by the partner. In a refutation, each step has to be either agreed or refuted on its turn, and so forth.
<i>Systematicity:</i> Tunable	When the co-learner disagreed with the human learner, he only expressed his disagreement in N% of the cases, N being tunable by the experimenter. In human-computer experiments, we set systematicity to 100%, but we varied it computer-computer experiments.

<i>Directness:</i> High	The dialogue interface was so constrained that the users could not really mean more than what was intended by the designer.
-------------------------	---

Table 1: Description of the negotiation space in PEOPLE POWER

4.2 MEMOLAB

The task. The goal of MEMOLAB [Dillenbourg et al. 94] is for psychology students to acquire the basic skills in the methodology of experimentation. A goal is assigned to agents, for instance “study the interaction between the list length effect and the semantic similarity effect”. The learners build an experiment to study this effect. An experiment involves different groups of subjects each encoding different list of words. An experiment is described by assembling graphical objects on a workbench. Each object associate a group of subjects, a task (e.g. read, listen, count down, free or serial recall,...) and a material (a list of words that the fictitious subjects have to study). The system simulates the experiment. The learner can visualize the results and perform an analysis of variance.

The agents. The human agent interacts with a rule-based system which plays the role of an expert. The inference engine was modified in order to make the expert sensible to the user actions on the task interface (the graphical problem representation). The specificity of design also concerns knowledge engineering. When a rulebased system reasons alone, the set of possible states is defined by the combinations of its own rules (and the problem data). In case of collaboration with a human, the set of possible problem states covers the combinations of all commands of both agents. Hence, the rulebase must include any intermediate step that the expert would normally skip but that the user could reach. Hence, the designer has to decompose rule into smaller chunks. Moreover, the learner action may transform the problem into a state which is not covered by the expert because it does not belong to any sensible solution path. We hence had to add sub-optimal rules and repair rules to the system which enable him to cope with these peculiar problem states [Dillenbourg et al. 95a].

Features	Description
<i>Mode:</i> Action.	Designing an experiment requires several dozens of commands. Some of these commands are available to both agents: these are the commands to create, delete or modify an event. Modifying an event means changing its group of subjects, its material or its task. Negotiation does not occur below that level. For instance, if the system disagrees on the features of a material used in an event, it can create another material, with different features and replace the one it disliked. But, the agent can not intervene when the human agent is creating this material (selecting the words, analysing its properties, typing the material names). In the discussion mode, there is no real negotiation: the expert provides

⁴ Such a facility would imply some kind of ‘shared-truth maintainance system’, i.e. an algorithm which does the following: (1) in case of diagremment on decision-X, all decision sub-sequent to decision-X have to be abandoned (If you disagree to go to

<i>Object:</i> Action	an explanation on request, but the learners cannot express disagreement ⁴ . Agents negotiate the actions necessary to build the experiment. The knowledge itself is not directly negotiated. Only the explanations provided by the expert refer to some theoretical constructs from the methodology of experimentation.
<i>Symmetry:</i> Variable	The degree of symmetry is low in the discussion mode, but high in the action mode , both agents can perform the same actions. The distribution of task in actual problem solving is better described by a variable asymmetry: the machine expert performs more at the beginning and fades progressively as the learner becomes more competent.
<i>Complexity:</i> High	The set of dialogue moves was rich although implicit. If an agent replaces in an event material-X by material-Y, its ‘dialogue move’ (in the large sense) is not made explicit as is often the case in the discussion mode, it is implicitly conveyed by the relationship between material-X and material-Y. For instance, if these two materials are equivalent on all properties but word frequency, but that material-X includes average words while material-Y includes only frequent words, replacing material-X by material-Y is equivalent to a dialogue operator such as ‘refinement’ which could be verbalised as “what you have done is not true in general, but only with frequent words”.
<i>Flexibility:</i> Variable	Turn taking was controlled by another agent. The tutor decided whether the expert or the learner has to perform the next step. However, any significant step in problem solving has to be agreed or refuted by the machine expert. This constraint was mainly due to our diagnosis technique: in the model-tracing approach [Anderson et al, 90], after any user action, the expert verifies whether any of its activatable rules would use the same command with equivalent arguments. This was too constraining. When we re-designed the system, we let the human agent deciding about turn taking. When it has completed a number of steps, the learner can ask the system to continue in order to know what the machine expert thinks. After his feedback, the learner can decide who does the next step. When the machine expert makes a step, it stops and gives the learner the possibility to take over or to ask for an explanation.
<i>Systematicity:</i> Moderate	The degree of systematicity is moderate . The expert itself is completely <i>egocentric</i> : it tries to reach its goal, and does whatever he can do in that direction. It does not refrain disagreement. However, when it disagrees, the tutor (another machine agent) may decide to prevent him interrupting the learner.
<i>Directness:</i> Low	Negotiation is mainly indirect from the user side, since most user actions can implicitly initiate a negotiation of the action she just performed. Negotiation is more direct from the machine agent side, since it explicitly expresses agreement or disagreement to the user.

Table 2: Description of the negotiation space in MEMOLAB

Evaluation. The experiments revealed several problems connected with the action mode. The learner action is interpreted every time he has performed a significant command. A difficult design question is precisely what should be treated as a significant. For instance, when an agent creates a list of words, the actual commands for building this list are irrelevant. What matters here are the intrinsic features of this list (e.g. its degree of semantic homogeneity) and which experimental group will study this list. The issue is therefore : at which granularity do we compare the actions of two agents: any interface commands (any click, any character,... even inside a word?) or only at major interface commands. In Memolab, the mapping concerns about 20 commands: for the expert these are the commands which appear as conclusions of one of its rules. The answer is in the question: the mapping must be done at the level where the designer wants the agents to negotiate. If the mapped commands are too elementary, there is nothing to negotiate because individual

Bruxelles, you also disagree to go to Bruxelles by train); (2) in case of agreement with decision-Y, all decisions upstream decision-X are implicitly agreed (if you agree to go to Bruxelles by train, you agree to go to Bruxelles).

commands do not carry enough meaning. Conversely, if mapped commands are too large, two divergent commands may imply various types of disagreement In MEMOLAB, some meaningful operations took the learner several steps, and when diagnosis occurred in between the steps, the subsequent interactions led to misunderstanding between agents.

The second main problem was that the expert was not flexible enough, because its problem solving strategy was not sufficiently opportunistic (an in the action mode, negotiation is not separated from task-level reasoning). We re-designed the expert and made it more opportunistic: there were fewer cases where the expert disagree not because the user decision was unacceptable, but because it had itself previously taken another decision. Although the collaboration was easier after redesign, subjects complained about the lack of negotiation regarding decisions taken by the expert. The problem was due to the absence of action-mode negotiation of the expert decisions. The most 'strategical' part of the expert's reasoning (choosing the independent variables, choosing the modalities for each independent variable) were not negotiable since they were not represented on the screen (in action mode, only displayed objects can be negotiated). However, these 'hidden' decisions inevitably came into focus when agents disagreed. We should either include these decisions in the problem space (i.e. extending the action mode negotiation) or develop an interface for discussing these decisions (i.e. extending the discussion mode), i.e. attempt to fill the gap between the action mode and the discussion mode. Negotiation appeared to be too much centred on action and not enough on underlying knowledge. Once again, a key design issue is deciding which subset of the task commands is available to both agents.

We also experimented Memolab with two human agents playing only with the task interface (two people in front of a single machine). It appears that the type of negotiation should not be modelled at the rule level, but "below the rule level, i.e. as a negotiation of the scope of the variables in that rule [Dillenbourg, to appear]. We should develop inference engines in which the instantiation of variables is not performed internally, but negotiated with the user.

4.3 KANT

The task. "KANT" (Kritical Argument Negotiated Tutoring system) was an ITS designed for teaching analysis of musical phrase structures in tonal melodies [Baker 89c, 92b]. The domain model was based on a (chart) parser that was able to take tonal melodies as input (via a MIDI keyboard) and produce a set of possible parses for the position of phrase and sub-phrase boundaries as output [Baker 89a, 89b]. Drawing on Lerhahl and Jackendoff's "Generative Theory of Tonal Music [83], we showed that multiple knowledge sources interact simultaneously to determine musical structures - for example, rhythmic differences, micro-pauses, intervallic structures, contrasts in dynamics, and, most importantly, harmonic structures. The knowledge to be

taught embodied in the system therefore consisted of sets of *beliefs* concerning musical structures, that were justified in terms of these different knowledge sources. The system also contained a semantic network, with "canned text" explanations for the concepts embodied in beliefs and their justifications, as well as for controlling dialogue focus shifts.

The agents. The two agents involved were 'system = teacher' and 'human user = novice'. Given that the domain to be taught consisted of justified belief sets, a non-directive or *negotiative* tutorial interaction style was therefore deemed to be appropriate.

The negotiation model. We provided the same set of dialogue moves to both the system and the user. The moves were divided into two sets - one for critical argument, and a second for negotiation of interaction control. The first set was represented hierarchically. At the highest level, the system might decide to DISCUSS a particular concept or instance of it. This tree then divided into a 'branch' of moves for making CLAIMs (CONCRETE_CLAIM = specific instance of a musical structure / ABSTRACT_CLAIM = concept embodied in a claim) and justifying them (SUPPORT_CLAIM), and a second 'branch' for CHALLENGING claims (giving reasons for and against). The moves were hierarchical in that the preconditions determining whether they were relevant or not at a particular stage of the interaction, increased in specificity with descent of the tree.

The system's decision as to what move to make was determined according to whether parameters (s - speaker, n - negotiator, c- concept, and inst - instance) referring to models of both agents' beliefs and to the dialogue history were satisfied or not. Whilst the 'speaker' parameter corresponds to the agent who actually performs a dialogue move, this is not necessarily identical to the *role* of 'negotiator'. For example, when 'negotiator = system', the system could negotiate that : 'speaker = student', 'concept = nil', 'inst = (student choice)', i.e. that in the next turn the student should propose an instance of a phrase boundary to be discussed. Similarly, the student could negotiate that the system should generate an explanation of a concept, or the system could negotiate that the system generate such an explanation.

The following example briefly illustrates the type of interaction managed by KANT. The sequence is taken from the middle of an interaction, where it is now the student's turn to negotiate, and the system presents successively the choices available from the common set of dialogue moves.

```
<...>
1. Are you sure you want to continue the discussion? (y or n)
2. y
3. Your turn to negotiate. Please type who you want the SPEAKER to be (student or system)
4. student
5. Please type in the concept you want to talk about (type SHOW for a list of possibilities)
6. phrase_boundary
7. Please choose a dialogue move : CHALLENGE or CLAIM
8. claim
9. Please choose a dialogue move : CONCRETE_CLAIM or NONE
10. concrete_claim
```

```

11. Please choose a dialogue move : MAKE_INSTANCE CLAIM or SUPPORT_INSTANCE CLAIM
12. make_instance_claim
13. Ok, I will cooperate. Please type the instance
14. n7-n8
15. Please choose a dialogue move : SUPPORT_INSTANCE CLAIM or NO_SUPPORT
16. support_instance_claim
17. Please type in the justification(s) for the phrase_boundary instance n7-n8
18. INTERVAL CONTRAST IMPERFECT CADENCE
19. Ok

```

Table 3. An excerpt from a dialogue with Kant. The student's interventions are shown in *italic*

Features	Description
<i>Mode</i> : Discussion.	The action mode could also have been implemented to allow marking of musical structures on a graphical interface
<i>Object</i> :: Multiple.	Action (reasons for and against domain beliefs) and communicative level (choice of speaker, negotiator role, topic, dialogue move); the meaning of domain concepts is not a possible object of negotiation ;
<i>Symmetry</i> : High	Both agents have the same dialogue moves, but with the notable exception of ABSTRACT CLAIM (explanation) which is restricted to the computational agent (see below) ; the argumentation can be <i>cooperative</i> since the system can find reasons to support the students' claims.
<i>Complexity</i> : Low	Complexity is relatively low in terms of the number of possible dialogue moves, but symmetry increases complexity since a given dialogue move could be instantiated by either speaker in a number of different ways
<i>Flexibility</i> : High	The use of any appropriately instantiated dialogue move can be negotiated, with the exception of strict alternance of negotiator roles.
<i>Systematicity</i> : High	Agreement or disagreement is always communicated.
<i>Directness</i> : High	Since the system enforces a restricted set of dialogue moves ; although the user could attempt indirect communicative action, the system has no means for understanding it.

Table 4: Description of the negotiation space in KANT

Evaluation. There were two major limitations to KANT : (1) verbosity and exhaustive explicitness of negotiation - although the successive negotiation of parameters and moves could be condensed into a single intervention, such fine-grained mechanisms are nevertheless necessary since rejected offers need to be 'unpacked' in order to determine precisely what is not agreed (speaker?, topic?, dialogue move ?); (2) the absence of natural language understanding capabilities diminishes symmetry (the student cannot participate in explanation) and excludes negotiation of meaning of utterances and domain concepts.

4.4 C-CHENE

The task. C-CHENE (Collaborative CHENE = CHaîne ENErgetique = Energy Chain) is a system for facilitating collaborative learning of modelling energy in physics in a network. The students' task is to construct "energy chains", or qualitative models for simple experimental situations (e.g. a battery connected to a bulb by two wires) in the form of a diagram for reservoirs, transfers and transformations of energy [Bental et al. 95] The students are provided with a graphical interface for

constructing energy chains (see upper part of Figure 4) and with an interface that allows them to communicate during the performance of such tasks.

The agents Two *human students* communicate and act via the system. We are currently working to design a computational agent that will be able to infer students' beliefs from their problem-solving and communicative actions, and to generate guidance with respect to the domain and the collaborative activity itself [Lund, Baker & Baron *in press*]. C-CHENE is clearly not designed to be a 'passive' means of collaboration - its communication interface was specifically designed to constrain the students' communication towards forms of collaboration that may promote learning.

The negotiation model. The negotiation model in C-CHENE is embedded in the design of the communication interface (lower part of Figure 2). The lower part of the communication area contains a set of buttons to be used by both students for performing different communicative acts, and the upper part the ongoing interaction history displayed for the students. The interaction history is an important resource in collaborative dialogue since it provides a common objective reference to previous activity (unlike oral dialogues) that encourage reflection and effective collaboration [Collins & Brown, 1988; Katz & Lesgold 1993].

The first rationale for designing the new button-based interface was to ease the students' typing load, thus freeing up time for more problem-solving task related discussion. The second was to encourage the students to engage in certain pedagogically preferred communicative activities (e.g. using the "Because" button to give reasons and explanations for intermediary solutions). This hypothesis was confirmed by analysis of transcripts of six pairs of students using the new interface. The third was to avoid some natural language interpretation problems (e.g. illocutionary force recognition), thus facilitating dialogue and belief modelling.

The set of CA buttons provided was designed on the basis of analysis of a corpus of 'chat-box' interactions with C-CHENE, and existing models for information dialogues [Moeschler, 85; Bunt, 89, 95] or collaborative problem-solving interactions [Baker, 94].

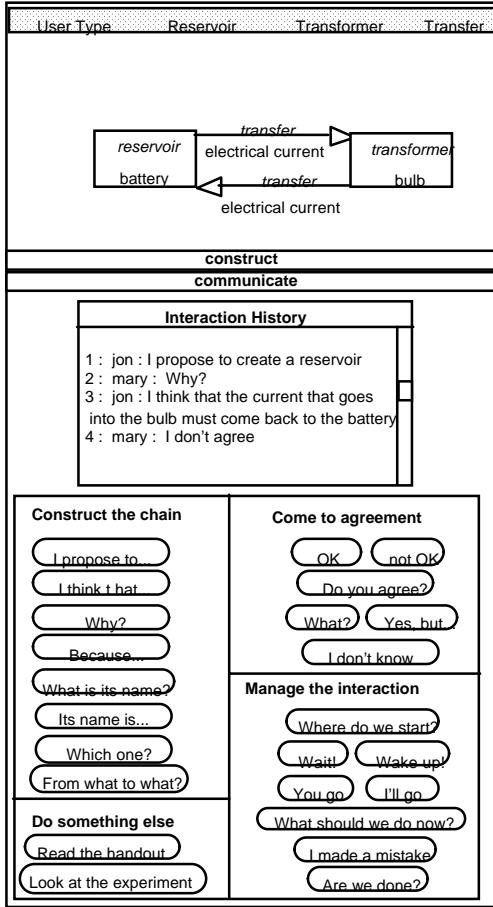


Figure 2: C-CHENE interface

The left hand column of buttons (e.g. "Construct the Chain") contains communicative acts that implement negotiation whose objects are at the *task level*. The right hand column corresponds to CAs that are used for negotiation at the level of *communicative interaction*. (*reactive acts* for maintaining feedback on *agreement* - OK, NOT OK - and perception/understanding - WHAT ? - and aspects of dialogue *management* , e.g. opening or closing the interaction, negotiating who will perform a particular graphical construction action, ...). The following extract has been taken from an interaction between two students who used C-CHENE to create an energy chain for an experiment where a battery was linked to a bulb by two wires. The text enclosed with "<SomeText>" corresponds to the system's automatic trace of the CA button that was clicked, other text having been directly typed by the students. The students successively negotiate opening the dialogue, who will create what, and the value to be assigned to the reservoir.

Time	Line	Loc.	Dialogue	Action
95	22	John	<Where do we begin ?>	
121	23	Mary	<I propose to ...><create a reservoir>	
144	24	John	<Which one ?>	
179	25	John	<Hello there !>	
200	26	Mary	<I don't know>	
213	27	John	<I think that ...> the first reservoir is called battery	
262	28	John	<Do you agree ?>	
287	29	Mary	<Agreed>	
312	30	John		reservoir created = reservoir1

334	31	John		name reservoir1 = battery
-----	----	------	--	---------------------------

Table 5 : An excerpt of dialogue between two human learners with C-CHENE (we translated)

Features	Description
<i>Mode</i> : Discussion and Action.	The students negotiate by drawing on the graphical interface, then discussing what has been done
<i>Object</i> : Multiple.	They negotiate the task (solutions and explanations), the meaning of utterances and different aspects of interaction control
<i>Symmetry</i> : High	Both students have the same set of interaction possibilities. Specific interactions may however become asymmetrical since the students spontaneously adopt <i>roles</i> (e.g. "you draw the diagram, I'll criticise")
<i>Complexity</i> : Medium	The interface complexity is not high, but sufficient for discussions relating to this task. Moreover, complexity is increased given different possible propositional contents for communicative acts
<i>Flexibility</i> : High	Such flexibility can pose coordination problems in this context which may indicate the necessity for imposing appropriate constraints
<i>Systematicity</i> : Undetermined	This is not imposed here: students may be insincere, "agreeing" simply in order to allow continuation of the interaction).
<i>Directness</i> : High	Although the set of possible communicative acts is fixed, the system is not completely direct, since students can communicate indirect meanings in virtue of the sequencing of acts (e.g. an offer of a proposition following a request for reasons is indirectly assumed to communicate a reason)

Table 6: Description of the negotiation space in C-CHENE

Evaluation. Two important results that emerged from case studies of pairs of students using C-CHENE : (1) **students actually are able to auto-classify their communicative action** in a manner required by the interface, (2) the 'passive constraint' on the form of the students' interaction, imposed by the interface, did succeed in encouraging them to engage in more task-related discussion, and *explanation* than in the 'free' chat-box interface.

4.5 BOOTNAP

We conducted experiments on computer-supported collaborative problem solving, using a MOO environment (tecfamoo.unige.ch - port 7777) and a whiteboard system (BeingThere™). The goal is to study social grounding, i.e. the mechanisms by which two humans verify that they understood what the other meant, and if it is not the case, repair misunderstanding [Clark & Brennan 91]. These mechanisms are central to negotiation. In human-human conversation, grounding relies on various techniques, including gestures and drawings. We study which schemata are drawn by two agents who have to solve a problem together and which role do these schemata play in social grounding or negotiation. The results of this work should help us to design more powerful collaboration interface between a human user and a knowledge-based system.

The task. Two agents play a CLUEDO-like game: somebody has been killed and they have to find the killer. They walk in text-based virtual world (a MOO environment) where they meet suspects, ask questions about relations with the victim, regarding what they have done before the murder,

and so forth. The two detectives explore rooms and find various objects which help them to find the murderer. They are invited to draw any representation useful to solve the problem. Both agents are in different rooms, but draw a common schema through a whiteboard. The two detectives are provided with a map of their virtual environment (an auberge), so that the schema focuses on the inquiry solution itself instead of on a (trivial) spatial representation of their environment. Subjects are familiarised with the MOO and the whiteboard through a training task.

The agents. The long-term goal of this project is to improve human-computer collaboration techniques. However, these experiments study human-human collaboration, not necessarily to imitate it, but to come out with functionally equivalent mechanisms. The CSCL setting does not include audio and video communication in order to reduce the bandwidth to something close to currently available interfaces for human-computer collaboration.

Features	Description
<i>Mode:</i> Discussion and Action.	The students negotiate by typed discussion in the MOO, but also by gestures in the whiteboard (e.g. crossing a note put by the partner) and by MOO action (e.g. a partner suggests 'let's go to room 1', the other does answer by moves to another room thereby expressing disagreement)
<i>Object:</i> All.	They negotiate which action to perform on the MOO, the graphical representations (e.g. the colour codes used in the MOO), the knowledge itself (e.g. the motives of various suspects), the problem solving strategy (e.g. "Let's find out who could get the gun between 7 and 9"). They do not negotiate turn taking probably because MOO conversations escape from the turn taking constraints of normal conversation
<i>Symmetry:</i> High	Both agents can for instance modify or delete any drawing made by the other. The symmetry changes however over time, some agent may lead for sometime and then role shift (the MOO includes 'lead' and 'follow' commands).
<i>Complexity:</i> High	The negotiation spaces are connected to each other, especially along the object axis. For instance, negotiating in which colour to put the cleared suspects (object : representation) forces agents to negotiate who is suspect (object: knowledge) (see dialogue in table 8). Vice-versa, negotiating task sub-goals (object: knowledge) is often made through negotiating representation. ("You know what we should do, it's a small table per suspect with who can have motives, who see her last, well a short synthesis").
<i>Flexibility:</i> High	Partners dynamically escape from some discussion episode if any of them bring a new interesting information in the conversation.
<i>Systematicity:</i> Undetermined	We cannot evaluate systematicity (as in C-CHENE), but globally it seems subordinated to the need to reach the solution.
<i>Directness:</i> High	We observed a transfer of negotiative functions according to the artefacts we provide them [Dillenbourg et al submitted]: for instance, one pair did not use the whiteboard at all but exchange factual data though a MOO notebook, using a 'compare notebook' command. We removed this command and the following pairs intensively used the whiteboard to report to their partner any interesting data. This can be related to recent theories on distributed cognition [Hutchins 95].

Table 7: Description of the negotiation space in BOOTNAP

Hercule: Why did you put a second arrow?

Sherlock: Because it is those who may have killed

Hercule: Yes, but why Giuseppe? He has no reason to kill...

Sherlock: No, but I did not know, I tried to see (*She moves the arrow to another name*)

Table 8: Connection between the negotiation spaces for different objects: representation -> knowledge

5. Conclusions

In this paper we have proposed the concepts of variable symmetry and negotiation spaces as part of considerations to be made in designing HCCLS. The concepts appear to be useful in the retrospective assessment of existing systems, and as a consequence, the design of future systems.

Three main conclusions have emerged. The first is **that human-human negotiation jumps between spaces** (C-CHENE and Bootnap), switching easily between modes of negotiation, connecting the various objects of negotiation. The 'disease' of People Power, Memolab and KANT was to be fixed within one negotiation space, the (mode: discussion X object: knowledge) space for People Power and KANT, and the (mode: action X object: action) for Memolab. Frohlich [93] suggested exploiting the complementarity of conversational (discussion mode) and direct manipulation (action mode) interfaces. We view as a main challenge, both at a technical and conceptual levels, to design agents able to conduct negotiation in both modes, action and negotiation, in a closely connected way.

The second conclusion is that complete symmetry is not a universally desirable goal (nor is it even, perhaps, a possible one). On one hand, technological limitations will, in the foreseeable future, mean that some asymmetry must exist in the human-machine interaction. On the other, we may decide that for specific combinations of agents, and for specific tasks, a certain degree of asymmetry is necessary and preferable, in order to best exploit the potential of each. Our claim is that symmetry must be considered as variable, i.e., that symmetry regarding what agents could do (design symmetry) leads to various forms asymmetry at different stages of actual interaction (interaction variable asymmetry). Designer may consider the system (i.e. the task, the communicative artefacts, and the computational agents) plus the human users as a single cognitive system whose various functions are distributed over the different components . This distribution varies over time.

Finally, designers of collaborative systems need to retain a degree of humility in their intentions and ambitions : however much they attempt to define or constrain the negotiation space, with its constituent dimensions, human users may always attempt to adapt the designer space to their own aims in unforseeable (indirect) ways. 'Flexibility within constraints' thus appears to be a reasonable design approach. The ways the whole cognitive system distributes its negotiative functions among agents and artefacts is not controlled by the designer.

6. Acknowledgements

The systems described here have been designed and implemented in collaboration with other members of our respective research teams. We therefore gratefully acknowledge David Traum,

Daniel Schneider, Patrick Mendelsohn, Boris Borcic, Melanie Hilario, John Self, Richard Godard, Andrée Tiberghien and Kris Lund. We would also like to thank the subjects and students, as well as their teachers, for participating in experimentation of the systems.

7. References

- Adler, M.R., Alvah, B.D., Wehmayer, R. & Worrest, W. (1988). Conflict-resolution Strategies for Nonhierarchical Distributed Agents. In *Distributed Artificial Intelligence* : Volume II, 139-161. London : Pitman Publishing.
- 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.
- Anderson, J.R., Boyle, C.F., Corbett, A.T. and Lewis, M.W. (1990). Cognitive modelling and intelligent tutoring, *Artificial Intelligence*, 42, 7-49.
- Baker, M.J. (1989a). An artificial intelligence approach to musical grouping analysis. *Contemporary Music Review*, 3(1), 43-68.
- Baker, M.J. (1989b). A computational approach to modeling musical grouping structure. *Contemporary Music Review*, 4(1), 311-325.
- Baker, M.J. (1989c). *Negotiated Tutoring : An Approach to Interaction in Intelligent Tutoring Systems*. PhD thesis (unpublished), Institute of Educational Technology, The Open University, Milton Keynes, Great Britain.
- Baker, M.J. (1992a). Le role de la collaboration dans la construction d'explications. *Actes des Deuxiemes journées "Explication" du PRC-GDR-IA du CNRS*, pp 25-42, Sophia-Antipolis, juin 1992.
- Baker, M.J. (1992b). Modelling Negotiation in Intelligent Teaching Dialogues. In Elsom-Cook, M. & Moyse, R. (eds.) *Knowledge Negotiation*, pp. 199-240. London : Paul Chapman Publishing.
- Baker, M.J. (1994). A model for negotiation in teaching-learning dialogues *Journal of Artificial Intelligence in Education*, 5 (2), 199-254.
- Baker, M.J. (to appear). Argumentation et co-construction de connaissances. To appear in *Interaction et Cognitions*, Presses Universitaires de Nancy.
- Baker, M.J., Dessalles, J-L., Joab, M., Raccah, P-Y., Safar, B. & Schlienger, D. (1994). Analyse d'explications negociees pour la specification d'un systeme d'aide a un systeme a base de connaissances. Dans les Actes du 4eme colloque ERGO-IA'94, pp. 37 - 47, Biarritz, octobre 1994.
- Bental, D., Tiberghien, A., Baker, M.J. & Megalakaki, O. (1995). Analyse et modelisation de l'apprentissage des notions de l'energie dans l'environnement CHENE. Dans les Actes des Quatriemes Journees EIAO de Cachan, ENS de Cachan 22-24 mars 1995, publiees dans Guin, D., Nicaud, J-F. & Py, D. (eds.) Environnements Interactifs d'Apprentissage avec Ordinateur, Tome 2., pp. 137 - 148. Paris : Eyrolles.
- Bond, A.H. & Gasser, L. (1988). *Readings in Distributed Artificial Intelligence*. San Mateo, Calif. : Morgan Kaufmann.
- Bunt, H.C. (1989). Information dialogues as communicative action in relation to partner modelling and information processing. In *The Structure of Multimodal Dialogue*, (eds.) Taylor, M.M., Néel, F. & Bouwhuis, D.G. , pp. 47-74. North-Holland : Elsevier Sciences Publishers.
- Bunt, H.C. (1995). Dialogue Control Functions and Interaction Design. In R.J. Beun, M. Baker, M. Reiner (Eds.), *Dialogue and Instruction, Modeling Interaction in Intelligent Tutoring Systems*. Proceedings of the NATO Advanced Research Workshop on Natural Dialogue and Interactive Student Modeling, (pp. 197-214). Berlin, Germany: Springer.
- Clark, H.H. & Schaeffer, E.F. (1989). Contributing to Discourse. *Cognitive Science* **13**, 259-294.
- 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.
- Collins, A. & Brown, J.S. (1988). The computer as a tool for learning through reflection. In H. Mandl & A. Lesgold (eds.) *Learning Issues for Intelligent Tutoring Systems*, pp. 1-18, New York : Springer Verlag.
- Collins, A., Brown J.S. & Newman, S. (1989) Cognitive apprenticeship: teaching the craft of reading, writing and mathematics. In L.B. Resnick (Ed.), *Cognition and Instruction: Issues and Agendas*. Hillsdale, N.J.: Lawrence Erlbaum Associates.

- Cumming G. and Self J. (1989) *Collaborative Intelligent Educational System*. Proceedings of the 4th AI and Education Conference. IOS. Amsterdam. pp. 73-80.
- Dalal P.N. & Kasper, G.M. (1994) The design of joint cognitive systems: the effect of cognitive coupling on performance. *International Journal of Human-Computer Studies*, 40, pp. 677-702.
- Dillenbourg, P. (To appear) Some technical implications of the distributed cognition approach on the design of interactive learning environments. To appear in *Journal of Artificial Intelligence in Education*.
- 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., 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. , Borcic, B, Mendelsohn, P. & Schneider, D. (1995a) "Influence du niveau d'interactivité sur la conception d'un système expert. in D. Guin, J.-F. Nicaud et D. Py (Eds) *Environnements Interactifs d'Apprentissage* (pp. 221-232). Paris: Eyrolles.
- Dillenbourg, P., Baker, M., Blaye, A. & O'Malley, C. (1995b) The evolution of research on collaborative learning. In E. Spada & P. Reiman (Eds) *Learning in Humans and Machine: Towards an interdisciplinary learning science*. Oxford: Elsevier.
- Dillenbourg, P. , Traum, D. & Schneider, D. (Submitted) *Grounding in multi-modal task-oriented collaboration*. Submitted to the European Conference on Artificial Intelligence in Education. Lisbon, Sept. 96.
- Douglas, S.A. (1991) Tutoring as Interaction: Detecting and Repairing Tutoring Failures. In P. Goodyear (Ed) *Teaching Knowledge and Intelligent Tutoring* (pp. 123-148). Norwood, NJ: Ablex.
- Edmondson, W. (1981). *Spoken Discourse : A model for analysis*. London : Longman.
- Fischer, G., Lemke, A.C., Mastaglio, T. & Morch, A.I- (1991): Critics: an emerging approach to knowledge-based human-computer interaction. *International Journal of Man-Machine Studies*, 35, 695-721.
- Fox, B.A. (1991) Cognitive and Interactional Aspects of Correction in Tutoring. In P. Goodyear (Ed) *Teaching Knowledge and Intelligent Tutoring* (pp. 149-172). Norwood, NJ: Ablex.
- Fox, B.A. (1993). *The Human Tutorial Dialogue Project : Issues in the Design of Instructional Systems*. Hillsdale N.J. : Lawrence Erlbaum Associates.
- Frohlich, D.M. (1993) The history and future of direct manipulation, *Behaviour & nformation Technology*, 12 (6), 315-29.
- Galliers, J.R. (1989). *A Theoretical Framework for Computer Models of Cooperative Dialogue, Acknowledging Multi-Agent Conflict*. PhD thesis (unpublished), Human Cognition Research Laboratory, The Open University (GB).
- Hutchins, E. (1995) How a Cocomit Remembers its speeds. *Cognitive Science*. 19, 265-288.
- Katz, S. & Lesgold, A. (1993). The role of the tutor in computer-based collaborative learning situations. In S. Lajoie & S. Derry (eds.) *Computers as Cognitive Tools*. Hillsdale NJ : Lawrence Erlbaum Associates.
- Lerdahl, F. & Jackendoff, R. (1983). *A Generative Theory of Tonal Music*. Cambridge Mass.: MIT Press.
- Lund, K, Baker, M.J. & Baron, M. (to appear). Modelling dialogue and beliefs as a basis for generating guidance in a CSCL environment. *Proceedings of ITS-96*, Montréal.
- Mevarech, Z.R. & Light, P.H. (1992). Peer-based interaction at the computer : looking backward, looking forward. *Learning and Instruction* , 2, 275-280.
- Miyake, N. (1986) Constructive Interaction and the Iterative Process of Understanding. *Cognitive Science*, 10, 151-177.
- Moeschler, J. (1985). *Argumentation et Conversation : Eléments pour une analyse pragmatique du discours*. Paris : Hatier-Créatif.
- Roschelle, J. & Teasley, S.D. (1995). The Construction of Shared Knowledge in Collaborative Problem Solving. In O'Malley, C. (ed.) *Computer Supported Collaborative Learning*. Berlin : Springer-Verlag.
- Roulet, E. (1992). On the Structure of Conversation as Negotiation. In (*On*) *Searle on Conversation*, pp. 91-100, (eds.) Parret, H. & Verschueren, J. Amsterdam : John Benjamins.
- Sycara, K. (1988). Resolving Goal Conflicts via Negotiation. *Proceedings of the AAAI-88 Conference*, St. Paul Minnesota (USA), 245-250.
- Terveen, L.G. (1993) Intelligent systems as cooperative systems. *Intenational Journal of Intelligent Systems*

Woods, D.D. & Roth E.M. (1988) Cognitive systems engineering. In M. Helander (Ed.) *Handbook of Human-Computer Interaction* (pp. 3-43). Elsevier Science Publishers: Amsterdam.