

PEOPLE POWER :

A Human-Computer Collaborative Learning System

Pierre DILLENBOURG (*) and John A. SELF (**)

* TECFA, Faculté de Psychologie et des Sciences de l'Education, University of Geneva (Switzerland)

E-mail: pdillen@divsun.unige.ch

** Computing Department, University of Lancaster (UK)

E-mail: jas@computing.lancaster.ac.uk

Abstract. This paper reports our research work in the new field of human-computer collaborative learning (HCCL). The general architecture of an HCCL is defined. An HCCL system, called People Power, has been implemented in CLOS. It contains a micro-world in which the learner can create an electoral system and simulate elections. The learner's task is to infer relations between the features of the electoral system and the distribution of seats. The human learner collaborates with a computational learner. The collaboration between learners is modelled as 'socially distributed cognition' (SDC). We view a pair of learners as a single cognitive agent whose components are distributed over two brains. This model maps inter-people and intra-people communication processes and thereby proposes an explanation of how the former generates the latter: the pattern of arguments that emerge from dialogue is reused by the artificial learner when it reasons alone. Reasoning is implemented as a dialogue with oneself. We report some results of the first experiments we have conducted.

1. People Power

This work escapes from the idea that courseware must be knowledgeable in the domain to be taught. Major advances in educational computing resulted from this idea, especially the shift of focus from learning outputs to learning processes. However, since the late 'eighties, the 'tutor-as-expert' paradigm is criticized with respect to pedagogic and philosophical values. The present work explores an avenue that radically differs from the established paradigm. The premise is that the system is initially no more knowledgeable than the learner. Instead, it will attempt to learn in collaboration with the user. Apart from a few limited systems (Chan and Baskin, 1988; Self and Hartley, 1989; De Haan and Oppenhuizen, 1990), HCCL is a virgin research area. Our first task has been to refine the concept of HCCL. When Self (1986) introduced the idea of HCCL, it was suggested as a method for learner modelling, the process of inferring the learner's knowledge. This function appeared to be inadequate: with respect to learner modelling, HCCL raises more issues than it solves. Nevertheless, the HCCL idea remained worthwhile for its originality with respect to the dominant 'tutor-as-expert' paradigm.

In HCCL, a human learner and a computerized learner collaborate to learn from experience. The computerized learner will be called the *co-learner*. Both learners share the problem solving experience they acquire with a micro-world. An HCCL includes four components: (i) a micro-world; (ii) the human learner; (iii) a computerized 'co-learner'; (iv) the interface through which learners interact with the micro-world; (v) the interface between the two learners the interface. These components are represented on figure 1. To some extent, any interactive software could be viewed as a HCCL system. For instance, an

expert system request for complementary data may be considered as a collaboration act. However, a collaborative system supposes a symmetrical interaction among learners: the same range of interventions must be available for both learners.




Figure 1.: Components of an HCCL system

We have implemented a HCCL system in political science, called PEOPLE POWER. The system includes the four components presented in figure 1. The co-learner is named 'Jerry Mander'. The learners cycle of activities in the micro-world is fairly simple. Firstly, the learners build an electoral system by specifying its parties, its constituencies, its laws, and so forth. Then, they simulate the elections and analyse the results. The goal is that learners discover the features that make an electoral system more or less proportional, i.e. whether the distribution of seats among parties correspond or not to the distribution of people preferences. PEOPLE POWER is written in object-oriented Common Lisp (CLOS) and runs on a Macintosh.

2. The 'Socially Distributed Cognition' model.

Our challenge was to develop a computational model that would guide our implementation of the co-learner. This model must be *socially valid* (more than psychological validity), i.e. it must perform the social behaviours expected in collaboration, such as agreement, interrogation, etc. If we push this idea to its extreme, we could have a system like ELISA (Weizenbaum, 1966): even if co-learner does not understand anything, it could still involve the learner in important activities, e.g. by randomly asking 'why?'. However, this would weaken the HCCL idea, since it removes the *causal relationship between social activities and learning*. Our model aims to support *authentic* collaboration. The co-learner has no access to any hidden expertise and it has not hidden didactic intention: it asks questions to get the answers, not to check if the learner knows the answers.

Our model focuses on a particular feature of cognition: the *internalisation of mutual regulation* skills. Regulations mechanisms that appear during collaboration (at the group level) are later mastered by individuals (Blaye, 1988). Wertsh (1985) has established the communicative dimension of this internalisation process (Vygotsky, 1978), by showing that a change of language prepares, within the social stage, the transition to the individual stage. The central role of communication also emerges on the Piagetian side. Blaye (1988) suggests that the intensity of the socio-cognitive conflict (Doise and Mugny, 1984) is less important than the fact it generates dialogue. It is interesting to notice that the relation between communication and regulation is also a core issue of *distributed artificial*

intelligence, i.e. the "study of how a loosely coupled network of problem solvers can work together to solve problems that are beyond their individual capabilities" (Durfee, Lesser & Corkill, 1989, p. 85).

Our 'socially distributed cognition' (SDC) model is based on three postulates. It is loosely related to Minsky's (1987) view that intelligence emerges from the combination of mental agents, each responsible for a "smaller process". We henceforth use 'agent' to refer to a cognitive process and 'device' for the place where agents are 'implemented' (a human brain or an artificial system).

Postulate 1. *An individual is a society of agents that communicate. A pair is also a society, variably partitioned into devices.*

Postulate 2. *The device border determines two levels of communication: agent-agent communication and device-device communication. Inter-agent and inter-device communications are isomorphic.*

Postulate 3. *Inter-device communication is observable by each device. Therefore, inter-device communication patterns generate intra-device communication patterns*

The individual and the group share the same computational representation, a society of agents. The composition of this society does not depend on the number of devices (individual versus group), but on the problem being solved. Several researchers have observed a spontaneous and unstable distribution of task among peer members (Miyake, 1986; O'Malley, 1987; Blaye et al., 1991). The border between devices defines two types of communication: communication between agents within a device and communication between agents from different devices (i.e. between devices). Postulate 2 claims that these two levels of communication are isomorphic. The duality of the isomorphism (postulate 2) and the difference (postulate 3) between the two communication levels (social and individual) determine the mechanism by which intra-device and inter-device communication processes influence each other: *inter-device communication structures are internalised as intra-device communication structures*. This third axiom corresponds to our research objective: mutual regulation processes, encompassed in inter-device communication, create self-regulation processes.

2. A simple example

We illustrate our approach with an imaginary dialogue between two mountaineers (Fig.1). Lucil suggests carrying on to the top and Jerry suggests going down. This discussion can be represented by the 'dialogue pattern' shown in figure 3. Jerry's first utterance (#2) refutes Lucil's argument (#1). Lucil's second utterance (#3) refutes Jerry's refutation, and so forth.

Lucil (#1)>	<i>Come on, the view is great from the top.</i>
Jerry (#2)>	<i>The clouds are coming.</i>
Lucil (#3)>	<i>The top is higher than the clouds roof,</i>
Lucil (#4)>	<i>... and snow is good, the risk of avalanche is low,</i>
Jerry (#5)>	<i>Yes, but it will be late when we'll return.</i>

Figure 2: A fictional dialogue among mountaineers.

Let us imagine that Lucil later climbs another mountain, with *no snow*, and hesitates between continuing or not. She can roughly have the same discussion with herself than with Jerry. She can replay the first part of the dialogue (#1-3), but not the second part (#4-5) which is not appropriate to the new context (no snow).

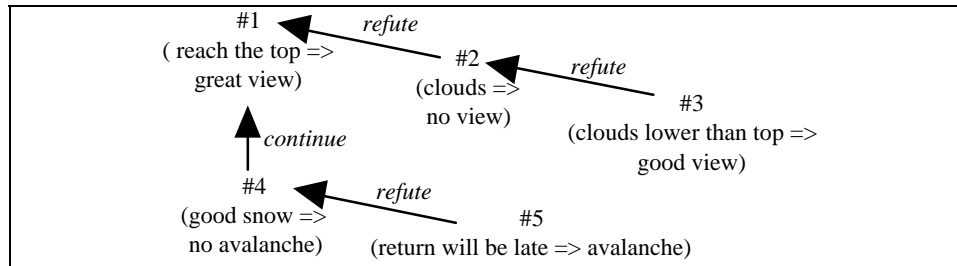


Figure 3: A simple dialogue pattern

Our approach consists of storing and replaying dialogue patterns. A pattern is a network of relationships between arguments. It is stored with some knowledge about the context in which it has been expressed. The probability that Lucil replays individually such a dialogue depends on two factors. The first one is her confidence in Jerry, hereafter, the *social sensitivity*: there are more chances that Lucil pays attention to this pattern if Jerry is a guide. The second factor is the *environmental feed-back*. If Lucil convinces Jerry to continue, and if it occurs that there is no view at all, there are more chances than Lucil pays attention to Jerry's counter-arguments. In the SDC model, these two factors modify the patterns by updating the strength of links between arguments.

3. From dialogue to monologue.

In People Power, learners play a game: for each country presented, learners have to reorganise the map (move one ward from a constituency to another) in order to gain seats for their party. Dialogue is about which ward to move and where. Then learners run the elections and check whether they have gained seats or not. We measure learning as a decrease of the number of attempts necessary to gain seats in a country. Initially, the co-learner has some naive knowledge about elections. For instance, it has a rule saying "If a party gets more votes, then it will get more seats". This rule is naive, but not basically wrong, it is only true in some circumstances. Jerry learns when it may use it.

Jerry Mander uses a single process for both searching for a solution and dialoguing with the real learner. This procedure uses two arguments, the proposer and the criticiser. If two *different learners* are respectively assigned to each argument, the procedure performs a real *dialogue*. If the *same learner* is used for both arguments, the procedure does *monologue*, i.e. reasoning. The procedure is a theorem prover: it proves that some change in the country map leads to a gain of seats for the Demagogic party. The procedure explores a tree of rules (or arguments), in depth-first search. In the monologue mode, a learner explores the tree backwards until it proves that some map change leads to gaining seats. In the dialogue mode, the proposer shows its inference path step-by-step to the criticiser. The dialogue structure is analogous to the mountaineers example. When a learner proposes an argument, the criticiser attempts to prove that this argument does not guarantee a gain of seats. If the criticiser does not find any counter-argument, the proposer

continues its argumentation. If the criticiser finds a counter-argument, it initiates a new sub-dialogue in which the proposer and criticiser roles are inverted. Figure 4 shows an example of dialogue between two artificial learners. The dialogue structure is simple and rigid. In terms of dialogues games, each learner has only two dialogues moves: *accept* (which leaves the partner continuing its explanation) or *refute* (by bringing a counter-evidence).

```

Marc > I suggest to move ward1 from Nord to Rhone-Alpes
Jerry > Why ?
Marc > If When We Remove "ward1" From Nord
Marc > The Demagogiques Get More Preferences Than Ringards In Nord
Marc > Then Demagogiques Will Take A Seat From Ringards In Nord
Jerry > OK, continue.
Marc > If Demagogiques Takes A Seat From Ringards In Nord
Marc > Then Demagogiques Will Have More Seats In Nord
Marc > And Ringards Will Lose One Seat
Jerry > OK, continue.
Marc > If Demagogiques Get More Seats In Nord
Marc > Then Demagogiques Will Have More Seats In France
Jerry > I disagree with that...
  Marc > Why ?
  Jerry > If Demagogiques Has Less Preferences In "ward1" Than In Rhone-Alpes
  Jerry > And If One Add "ward1" To Rhone-Alpes
  Jerry > Then Demagogiques Will Loose Preferences In Rhone-Alpes
  Marc > OK, continue.
  Jerry > If Demagogiques Get Fewer Preferences In Rhone-Alpes
  Jerry > Then Demagogiques Will Get Fewer Votes In Rhone-Alpes
  Marc > OK, continue.
  Jerry > If Demagogiques Party Gets Fewer Votes In Rhone-Alpes
  Jerry > Then It Will Get Fewer Seats In Rhone-Alpes
  Marc > I disagree with that...
    Jerry > Why ?
    Marc > If Demagogiques Has No Seats In Rhone-Alpes
    Marc > Then It Cannot Lose Seats
    Jerry > OK, continue.
    Marc > Let's resume where we were.
  Jerry > Let's resume where we were.
  Marc > Let's resume where we were.

```

Figure 4 : Example of dialogue between two artificial learners. The indentation indicates levels of refutation. The task was to move a ward from a constituency to another in such way that the new regrouping of votes leads the 'Demagogics' party to gain seats.

Reasoning is implemented as a dialogue with oneself, i.e the same learner plays the roles of proposer and criticiser. When the learner reaches some node (an argument) in exploring the solution tree, it tries to refute it. If it fails to refute, it continues. If it refutes its own argument, it backtracks and explores another branch of the tree. The process is recursive as in dialogue: the learner attempts also to refute its own refutation, and so forth. The main difference between dialogue and monologue is that, in dialogue, a learner refutes the other by proving that some step is wrong, while in monologue, Jerry Mander refutes itself only by using refutations that have been elaborated jointly (see next section)

4. Learning mechanisms

The learner will learn relations between arguments. A dialogue pattern is a network of links between arguments (or rules). The type of links is determined by the dialogue: a 'continue-link' relates two rules that have been consecutively verbalised in an explanation; a 'refute-link' relates two rules such that one has been verbalised to refute the other. The representation of patterns is distributed: we don't have a specific object which represents a pattern, each rule stores its links with other rules. For each link, we stored data about the context and its strength, a numeric parameter whose role will be described later on.

In monologue, the learner uses continue-links as a heuristic for exploring the tree. If rule-X has a continue-link rule-Y, this means that, after having applied rule-X, Jerry Mander considers rule-Y before any other one. If rule-X has several continue-links, they are sorted by increasing order of strength. Acquiring 'continue-links' corresponds to some incremental and context-sensitive form of knowledge compilation: the pair (rule-X rule-Y) is now some kind of 'chunk' that speeds up reasoning.

Example: Rule-6 ('If a party gets more preferences, then it will get more votes') is often followed by Rule-8 ('If a party gets more votes, then it will get more seats'). The continue link Rule-6/Rule-8 corresponds to the rule 'If a party has more preferences, it will get more seats'.

In monologue, the refute-links bring attention to rules that should be considered before continuing inference. If rule-X has a refute-link to rule-Y, Jerry will check out rule-Y before continuing. If it occurs that rule-Y is verified, Jerry will backtrack, otherwise it will continue its search. The refutation (rule-Y) may of course itself be refuted. If it is the case, Jerry may continue its inference from rule-X. Adding refute-link constitutes a special form of rule specialisation, i.e is identical to adding a rule condition. Let us imagine two rules, rule-X: $p1 \Rightarrow q$, and rule-Y: $p2 \Rightarrow (\text{not } q)$. The refute-link rule-X/rule-Y (i.e. rule-Y refutes rule-X) corresponds indeed to a specialized version of rule-X: $p1 \text{ and } (\text{not } p2) \Rightarrow q$.

Example: Rule-9 ('If a party gets fewer votes (p1), then it will get fewer seats'(q)) is refuted by Rule-13 ('If a party has no seats (p2), then it cannot loose seats' (not q)). The association Rule-9 / Rule-13 corresponds to a specialised version of Rule-9: 'If a party gets fewer votes (p1) and has seats (not p2), then it will loose seats (q)'.

Two mechanisms modify the strength of links according to the dialogue and to the results of elections. The *social sensitivity* represents the extent to which the real learner influences the co-learner (e.g. her confidence). When a new link is created or when an existing link is verbalized, the strength of this link is increased by the value of this factor. The *environmental feed-back* modifies the link strength according to the results of the simulated elections. If Jerry proposal leads to a gain of seats, the links verbalised by Jerry are strengthened and those refuting Jerry are weakened. Conversely, if some seats are lost, the continue-links are weakened and the refute-links are strengthen. This corresponds to a simplified 'learning by experimentation' strategy.

5. Experimentation

Five subjects played the 'Euro-Demago Game' during one hour. They have appreciated the micro-world as a whole and the possibility to interact with the co-learner. Some of them

expressed the feeling to be really collaborating with a partner, though this partner was not much human-like. Figure 5 shows an excerpt of human-computer discussion.

The subjects clarified the problem we anticipated: the difficulty of communication. Actually, they did not complain about the window in which they introduce explanations, which appeared to be easy to be used and fast to be learned. The bottle-neck was the need to know Jerry's rule before to be able to express oneself with these rules. This difficulty seems however to decrease after half an hour of work.

```
...
Jerry > I suggest to move ward2 from Nord to Sud
Lucil > Why ?
Jerry > If Demages Has Less Preferences In "ward2" Than In Nord
Jerry > And If One Remove "ward2" From Nord
Jerry > Then Demages Will Gain Preferences In Nord
Lucil > Continue
Jerry > If Demages Gets More Preferences In Nord
Jerry > Then Demages Will Get More Votes In Nord
Lucil > Refute
      Lucil > If When We Remove "ward1" From Nord
      Lucil > The Demages Don 'T Get More Preferences Than Mortales In Nord
      Lucil > Then Demages Won 'T Take A Seat To Mortales In Nord
      Jerry > That's not the case here !
Jerry > If Demages Gets More Preferences In Nord
Jerry > Then Demages Will Get More Votes In Nord
Lucil > Continue
Jerry > If The Demages Party Gets More Votes In Nord
Jerry > Then It Will Get More Seats In Nord
Lucil > Continue
Jerry > If Demages Gets More Seats In Nord
Jerry > Then Demages Will Have More Seats In Portugal
Lucil > Continue
Jerry > That's all
Jerry > Do you accept my proposal ?
Lucil > Yes
```

Figure 5: A excerpt from a dialogue between a subject - we renamed Lucil -and Jerry Mander.

Some users did also complain about the presence in Jerry's reasoning of rules such 'If a party has less votes, then it has not more votes'. Such a rule is necessary for Jerry's logical reasoning, but it is totally trivial for the learner and hence would not be verbalised by a learner outside the context of this experiment. Two learners pointed out that their participation in the dialogue coerced them to enter in Jerry's logics.

Because the cost of interaction was very high, subjects reduced the number of interactions. Since Jerry's learning output depends directly of the amount of interactions, Jerry has been learning slowly. All dialogue patterns elaborated by Jerry during these one-hour sessions included less than 10 inter-agents links, while dialogue patterns built in experiments with two artificial learners become efficient around 25 links. Hence, Jerry continued to provide the learners with suggestions that were not very good, and hence decreased the subjects

interest in its suggestions. The co-learner must possess some minimal skills to maintain the learner's interest in collaboration.

For improving communication between learners, we should fundamentally redesign the interface. A solution might be to provide them with an interactive table for describing electoral data. Dialogue would then be carried out by pointing out data or data transformations that illustrate, prove or disprove some statement. This graphical support to interaction goes in the direction of Roshelle's (to appear) observations.

These essays have been very insightful for us. However, because of the cost of interaction, they did not allow us to explore our model in a systematic way. Therefore, we conducted experiments with two artificial-learners (Dillenbourg, 1991). The full report of these experiments is beyond the scope of this paper. Globally, performance improved provided that learners develop a rich network of arguments, what requires intensive communication. These results converge to the computational modelling work of Hutchins (1991): cognitive effects seem to be related to the 'communication curve', i.e. the frequency of communication between subjects and how it is distributed over time. The interest of such finding is to orient empirical research with human subjects.

6. Other features of People Power

The design of People Power takes into account some issues that are not directly integrated in the SDC model.

Authenticity. Jerry Mander has no didactic intention. Mutual regulation emerges from Jerry's wish to interact, to understand and to share decisions. Jerry has no hidden knowledge. Its knowledge is directly inspectable by the real learner. Jerry's rules do not form proper expertise. They correspond to very simple principles that most people master. The acquired expertise is the organisation of these rules into patterns.

Symmetry. The interaction between learners is symmetrical: both learners have the same possibilities to change the electoral system, to make a suggestion or to interrupt the partner. There is one exception though: if Jerry rejects her suggestion, the real learner may decide to stick to her position, but, if she disproves Jerry's proposal, Jerry abandons it.

Sociocognitive conflict. As in the sociocognitive theory conflict, Jerry Mander does not spontaneously perceive internal conflict. Conflict only appears through interaction with another subject. We applied with Blaye's (1988) idea that the conflict itself is less important than the verbalisations it generates.

Sharing the working memory load. Working memory load is shared among collaborative human learners. In People Power, when the real learner is conducting the explanation, the co-learner role is restricted to the control of this explanation, through the search for refutations. The proof built by learner B for refuting A's argument zooms in on a single inference step in A's proof. Thereby B can devote its cognitive resources to controlling this particular argument while A's must maintain a broad view of the proof.

Versatility. Nothing would be more boring than a co-learner which would ask the real learner to justify all her statements. Human learners know that such behaviour would irritate their partner and refrain sometime to express their disagreement. Jerry's behaviour

can be modified by a 'systematicity' factor: if Jerry disagrees with the real learner, and if the 'systematicity factor' is 80%, there are 20 chances in 100 that Jerry will not express his disagreement.

Pre-Individualisation (Asynchronous multiple learners). Not all pairs function in an ideal way (Salomon and Globerson, 1989). Among pair efficiency conditions, researchers pay particular attention to the homogeneity of the pair. An HCCL system may provide the real learner with a co-learner appropriate to her level. The human tutor may select Jerry Mander's background knowledge (its initial set of rules). The current release only includes one knowledge set, but it would be trivial to add more sets. The user can also change Jerry's working memory limit, the ratio between the social and environmental factors and the systematicity. She may define her 'patience', i.e. a number of minutes beyond which we Jerry's inference engine is interrupted.

Extensibility (Synchronous multiple learners). People Power includes only two learners, an instance of the 'artificial learner' class and an instance of 'human learner'. The representation of learners as objects supports an extension of People Power to a multi-learner system. However, People Power only includes human/artificial learner and artificial/artificial learner dialogue procedures. A multi-learner system would require human/human learner procedures, either internal or external to the system.

Simplifications and restrictions. Developing all aspects of a collaborative learner, in a detailed way, would require solving almost all issues that have been raised in AI. We had to reduce both the range of issues and the complexity of the selected issues. For instance, learners have the same explicit goal, there is no goal negotiation, there is no dialogue repair, etc.

7. Discussion.

As often, implementing a system led us to develop a computational model. This model relates three ideas. The first idea views a group of subjects as a single cognitive system. This implies that reasoning and collaborating are two instances of dialogue: reasoning is a dialogue with oneself and collaborating a dialogue with the peer. That is the second idea. It is not new, it has been proposed by main scientists such as Vygostksy, Piaget or Mead. Because of this isomorphism, the social dialogue generates structures for the reflective dialogue. That is the third idea. The SDC model is thus based classical ideas. Our contribution has been to organise these ideas into a structure that has the consistency necessary for implementation.

The main limitation of our work lies in the simplicity and rigidity of conversation among learners. In our implementation, dialogue is restricted to agreement and disagreement utterances. Real dialogue involves the much more complex process of elaborating and maintaining a shared understanding. This process involves a set of operators, such as 'specification', 'elaboration', 'reformulation', 'restrictions' and 'reasons' (Baker, 1992). Such a rich set of operators is required to reach '*social grounding*', i.e. the participants belief that his partner has understood his utterance (up to some acceptability criterion) (Clark and Brennan, 1991). It would be interesting to adapt and test the SDC model with a larger set of dialogue operators, supporting dialogue repair. With respect to the specificity of HCCL, we should modify the interface in order to support '*physical grounding*', i.e the possibility

of peers to use the environment (the micro-world) for comparing their mutual understanding (Roschelle, to appear).

References

- BAKER, M. (1992) The collaborative construction of explanations. *Paper presented to "Deuxièmes Journées Explication du PRC-GDR-IA du CNRS"*, Sophia-Antipolis, June 17-19 1992.
- BLAYE, A. (1988) *Confrontation socio-cognitive et resolution de problemes*. Doctoral dissertation, Centre de Recherche en Psychologie Cognitive, Université de Provence, 13261 Aix-en-Provence, France.
- 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.
- CHAN, T.-W. & BASKIN, A.B. (1988) "Studying with the prince", The computer as a learning companion. *Proceedings of the International Conference on Intelligent Tutoring Systems* (pp.194-200), June 1-3. Montreal, Canada.
- CLARK, H.H. & BRENNAN S.E. Grounding in Communication. In L. Resnick, J. Levine and S. Teasley. *Perspectives on Socially Shared Cognition* (pp. 127-149). Hyattsville, MD: American Psychological Association.
- DE HAAN, A. & OPPENHUIZEN, C.J. (1990) RITS, cooperative problem solving with computers, a pragmatic approach. *Proceedings of the Second 'Congres Europeen Mutli-Media, IA et Formation'*, September 24-26. Lille, France.
- DILLENBOURG, P. (1991) Human-Computer Collaborative Learning. *Doctoral dissertation*. Department of Computing. University of Lancaster, Lancaster LA14YR, UK.
- 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, pp. 83-127). Reading, Massachusetts: Addison-Wesley.
- HUTCHINS, E. (1991) The Social Organization of Distributed Cognition. In L. Resnick, J. Levine and S. Teasley. *Perspectives on Socially Shared Cognition* (pp. 383 - 307). Hyattsville, MD: American Psychological Association.
- MINSKY, M (1987) *The society of mind*. London: William Heinemann Ltd.
- MIYAKE, N. (1986) Constructive Interaction and the Iterative Process of Understanding. *Cognitive Science*, 10, 151-177.
- O'MALLEY, C. (1987) Understanding explanation. *Paper presented at the third CeRCLC Workshop*, Ullswater, UK.
- ROSCHELLE J. (to appear) Learning by Collaborating: Convergent Conceptual Change. To appear in the *Journal of the Learning Sciences*.
- SELF, J.A. (1986) The application of machine learning to student modelling. *Instructional Science*, 14, 327-338.
- SELF, J.A. & HARTLEY J.R. (1989) Guided Discovery Learning Systems in Intelligent Computer-Aided Instruction. *Final Report of the GDLS project*. Department of Computing, University of Lancaster, Lancaster LA1 4YR, UK.
- SUCHMAN, L.A. (1987) *Plans and Situated Actions. The problem of human-machine communication*. Cambridge: Cambridge University Press.
- VYGOTSKY, L.S. (1978), *Mind in Society. The Development of Higher Psychological Processes*. Edited by M. Cole, V. John-Steiner, S. Scribner & E. Soubberman Harvard University Press. Cambridge, Massachusetts: Harvard University Press.
- WEIZENBAUM, J. (1966) ELIZA: a computer program for the study of natural language communication between man and machine. *Communications of the Association for Computing Machinery*, 29 (1), pp. 36-45.
- WERTSCH, J.V. (1985) Adult-Child Interaction as a Source of Self-Regulation in Children. In S.R. YUSSEN (Ed) *The growth of reflection in Children* (pp. 69-97). Madison, Wisconsin: Academic Press.