

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
1211 Geneva 4 (Switzerland)
E-mail: pdillen@divsun.unige.ch

** Computing Department
University of Lancaster
Lancaster LA1 4YR(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): a theoretical framework is proposed, design principles are elaborated and a working system is described. The general architecture of an HCCL is defined. The design of the computational learner concentrates the mechanisms by which learners internalise mutual regulation processes. 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 can generate the latter.

An HCCL system, called People Power, has been implemented in LISP on the Apple Macintosh. 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 implementation of the co-learner is based on the SDC model. The pattern of arguments that emerge from dialogue is reused by the co-learner when it reasons alone. Reasoning is implemented as a dialogue with oneself. We tested People Power in two settings: with a human and an artificial learner or with two artificial learners. The common findings of these two experiments is that the SDC model is valid provided that the discussion is intensive, i.e. that many arguments are brought into dialogue.

Introduction

This work escapes from the idea that courseware must be knowledgeable in the domain to be taught. Twenty years ago, this very idea constituted the promise of research in 'Artificial Intelligence and Education'. 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 problem solving experience. 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 microworld; (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. The system includes the four components presented in the previous section: a microworld plus its interface and a computerized learner plus its interface. The whole system is named PEOPLE POWER. The co-learner is named 'Jerry Mander'. The learners cycle of activities in the microworld 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.

The 'Socially Distributed Cognition' model.

Our challenge was to develop a computational model that would guide our implementation of the co-learner. The role of the model within the HCCL systems generates constraints of *social validity* rather than psychological validity. A model is socially valid if it can perform the social behaviours expected in collaboration, such as agreement, interrogation, etc. The idea to implement this model is too ambitious, our model focuses on a particular feature of cognition: the *internalisation of mutual regulation* skills. Blaye (1988) observed that regulations mechanisms that appear during collaboration (at the group level) are later mastered by individuals. Wertsh (1985) has established the communicative dimension of the 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 emerge 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* (Gasser, 1991), 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). The designer of a DAI system knows the trade-off between the need for central control (regulation) and the agents' ability to convince each other through argumentation.

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

The model's foundation is rather trivial: to analyse the relation between individual and social cognition, we consider a group of subjects as a single subject. A individual who is solving some problem X and a pair working jointly at X can be represented as two instances of some 'cognitive systems' class. This 'mitosis' axiom considers cognition as a socially distributed process. Psychologists have observed a spontaneous and implicit distribution of task among peer members. The set of running cognitive processes at some stage of problem solving seems to be more determined by the problem being solved than by the number of problem solvers.

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

There are of course differences between learning alone and learning with a peer. The second axiom brings some nuances to the first one. It points out that the concrete form of communication processes distinguishes individual and social cognition. We discriminate two kinds of communication: communication among processes within the brain of a single learner (or device) and communication between devices. It postulates that the communication processes are isomorphic, i.e. that they use the same patterns of arguments. Therefore, we merged our inference engine and our dialogue algorithms into a single 'dialogue engine'. This procedure searches in a tree of arguments (rules). Search operators are expressed by the simplest set of dialogue moves: continue (go down) and refute (backtrack). The dialogue engine arguments specify two agents. If these are two different learners, the engine runs a real dialogue. If the same learner is provided as both arguments, we get a dialogue between the learner and himself, i.e. the learner's reflective thinking.

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

The third axiom refines the second one. Intra-device and inter-device communication patterns differ by the fact the latter are observable by each agent. They can therefore be perceived as any object of the environment, for instance as input to inductive processes. This axiom claims that inter-device communication patterns generate intra-device patterns.

In our implementation, the pattern is a trace of the dialogue engine. It is recorded when it is verbalised and stored in a distributed way, as a network of links between agents. These links memorize the context in which they have been created. The relative strength of links guides the selection of agents by the inference engine. This strength changes when the interaction proceeds. The changes depend on the ratio between two parameters, the social and the environmental factors. The social factor indicates the attention that one learner pays to her partner ideas, the extent to which one learner influences the other. The environmental factor refers to the extent to which a learner changes her mind on the basis of the indirect feed-back provided by the environment. In our system, the environmental feed-back is the result of the simulation of the elections within the electoral system built through collaboration. .

```
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 To Ringards In Nord
Jerry > OK, continue.
Marc > If Demagogiques Takes A Seat To Ringards In Nord
Marc > Then Demagogiques Will Have More Seats In Nord
Marc > And Ringards Will Loose One Seat
Jerry > OK, continue.
Marc > If Demagogiques Gets 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 Gets Fewer Preferences In Rhone-Alpes
    Jerry > Then Demagogiques Will Get Fewer Votes In Rhone-Alpes
    Marc > OK, continue.
    Jerry > If The 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 The Demagogiques Has No Seats In Rhone-Alpes
        Marc > Then It Cannot Loose 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 2 : Example of dialogue between two artificial learners, Jerry and Mark. The indentation indicate levels of refutation. The task was to move a ward from a constituency to another such way that that the new regrouping of votes leads the 'Demagogics' party to gain seats.

Design Features

The SDC model that has been described and implemented does not exhaust the list of issues related to HCCL. The SDC model concentrates on our main research objective, the internalisation of mutual regulation skills. Nevertheless, the design of Jerry Mander takes into account some issues that are not directly integrated in the SDC model.

Authenticity. Jerry Mander has no didactic intention. Mutual regulation does not result from a systematic control activity. It emerges from Jerry's wish to interact, to understand and to share decisions. Jerry has no hidden knowledge. Its various knowledge bodies are directly inspectable by the real learner: its inference rules, the hierarchy of electoral concepts, and the representation of the current electoral system. Jerry's rules can be read by the real learner in the 'enter a rule' window. Those rules do not form proper expertise. They correspond to very simple principles that most people master. The expertise result from the organisation of rules into patterns.

Symmetry. The interaction between learners is symmetrical: both learners have the same possibilities to change the system, to make a suggestion or to interrupt the partner. There is one exception though. If A can interrupt B exactly the same way B can interrupt A, the dialogue might recurse infinitely. Therefore, we give the final word to the human learner. This means that when Jerry rejects her suggestion, the real learner may decide to stick to her position. If the human learner 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. The dialogue between the real learner and Jerry Mander leads to verbalise links between rules and these links are only stored when they are verbalised .

Multiplicity of learning strategies. Our approach covers various learning strategies: induction, through the effect of pattern frequency on strength; deduction since the whole process is based on 'proof' similar to explanation-based learning (but whose generalisation is controlled by dialogue) and analogy, since dialogue patterns guide new inferences on the basis of previous ones. Learning in itself is neither inductive, nor deductive, nor analogical. A learning session takes a particular colour, inductive, deductive or analogical, according to the context and to the learning activities.

Psychological plausibility. Psychological plausibility remains desirable, provided that it is not achieved at the expense of social validity. We intentionally use 'psychological plausibility' instead of 'psychological validity'. Plausibility means that we applied general principles of psychology to the design of Jerry Mander. For instance, human working memory has a capacity limited to around seven elements (± 2). Similarly, Jerry's working memory (the set of explicit facts) is limited by the values of the engine argument 'max-depth'. The working memory parameter defines a depth threshold beyond which the inference engine backtracks. The nuance between validity and plausibility lies in the fact that the value for the working memory boundary has been tuned empirically rather than extracted from protocol analysis. The chosen boundary has no absolute value.

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 peer knowledge. An HCCL system may indeed provide the real learner with a co-learner appropriate to her level. Jerry Mander's background knowledge (its initial set of rules and actions) is stored in an explicit way. The human tutor may select the relevant background knowledge set in the 'tune-collaboration' window. The current release only includes one knowledge set, since we did not focus on that issue, but it would be trivial to add more sets. We refer to this aspect of our system as 'pre-individualisation'. The collaboration mechanisms defines some kind of individualisation: since Jerry's behaviour is highly dependent on the human learner's behaviour, the system reactions are indeed tailored to the particular user. By selecting the co-learner, the human learner pre-determines the range of later individualisation.

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. People Power could for instance work with four learners, two human learners working in remote places, each of them being assisted by an artificial learner. 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 .

Sharing the working memory load. Working memory load is shared among collaborative human learners. Although it does not focus on that phenomena, the SDC model may partially reproduce it. 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 attention to controlling this particular argument while A's must maintain a broad view of the whole proof.

Versatility. Nothing would be more boring than a co-learner which would ask the real learner to justify all its statements, to explain everything, etc. Human learners know that such behaviour would irritate their partner and hence refrain sometime to express their arguments. Jerry's behaviour can be modified by a 'systematicity' factor. It is applied when Jerry verifies an action or a rule proposed by the real learner. For instance, if the 'systematicity factors' is 80% , there are 20 chances in 100 that Jerry does not express his disagreement.

Tuning collaboration. The design and implementation of Jerry Mander compelled us to choose values for various parameters that influence Jerry's reasoning. This set of factor values constitutes Jerry's profile and can be modified by the user. They include: Jerry's prior knowledge (naive rule base), Jerry's working memory limit, the social-environmental ration and the systematicity. Another parameter is very important for Jerry, although it describes the human learner: its patience. The human learner is invited to define her 'patience' as a number of minutes during which she will wait for the co-learner's answer. This value is used to stop the inference engine when it runs a too long time.

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. We did not implement mechanisms by which ambiguities in communication recursively generate disambiguation dialogues. It is clear, though, that these ambiguities constitute a major source of dialogue in human-human collaboration.

Experimentation

The goal of the experimentation was to get a feed-back on the quality of our model in general and our system in particular. We asked to five subjects to play the 'Euro-Demago Game' during one hour. The goal is that the learner's party gains one seat in each EEC country, by moving wards from a constituency to another. Learners have appreciated the microworld 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 totally human-like. Figure 3 shows an excerpt of human-computer discussion.

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

The subjects clarified the problem we anticipated: the difficulty of explanation and refutation processes. Actually, they did not complain about window in which they introduce refutations, 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.

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. This can be perceived either as a limitation of HCCL possibilities or as a new way to perceive the utility of HCCL.

Because the cost of interaction with Jerry 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. This statement raises the question: what are the minimal skills that the co-learner must possess 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 a interactive table for describing electoral data. Dialogue would then be carried out by pointing out data or data transformations that illustrate, proof or disproof some statement. This graphical support to interaction goes in the direction of Roshelle's (1990) symbolic mediation idea.

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 completed the experimentation by a computer / computer experimentation, i.e. to observe collaboration between two artificial-learners (Dillenbourg, 1991). The full report of these experiments is beyond the scope of this paper. However, these experiments confirmed the human-machine experiments: performance improved provided that learners develop a rich network of arguments, what requires intensive communication among learners.

7. Discussion.

The design of the co-learner has been influenced by two sets of constraints. The first and primordial set corresponds to our research objectives, implementing a computational account of how mutual regulation skills may be internalised. The second set of constraints is due to the integration of the co-learner model within a collaborative learning environment. We paid much attention to the microworld interface, especially to provide good support for reflection. By contrast, the design of the learners interface is more basic. This constitutes a weakness of People Power. However, although we do not deny the

defaults of our interface, we want to point out that its simplicity improves - to some extent - the transparency of Jerry Mander. Jerry uses a few pre-defined messages such as 'I disagree with that'. The frequent use of these messages makes dialogue boring. We then replaced each message by a small library of synonymous messages within which Jerry randomly selected the one to be displayed. This trivial trick made conversation more pleasant. But we abandoned this practical advantage due to its scientific cost: the structure of dialogue was not legible any more. This could easily lead the observer to believe that Jerry's reasoning is much more complex than it actually is. This may be interesting from a pedagogical point of view, but not from a scientific perspective.

REFERENCES

BLAYE, A. (1988) *Confrontation socio-cognitive et resolution de problemes*. Doctoral dissertation, Université de Provence, Aix-en-Provence, France.

CHAN, T.-W. & BASKIN, A.B. (1988) "Studying with the prince", The computer as a learning companion. *Proceedings of the ITS-88 Conference* (pp.194-200). Montreal, Quebec.

DE HAAN, A. & OPPENHUIZEN, C.J. (1990) RITS, cooperative problem solving with computers, a pragmatic approach. *Proceedings of Applica 90, Second Congress European Multi-Media, IA et Formation* (pp. 1-9)

DILLENBOURG P. (1991) Human-Computer Collaborative Learning. Unpublished doctoral thesis. University of Lancaster. Department of Computing.

DILLENBOURG, P. & SELF, J.A. (to appear) Designing human-computer collaborative learning. In C.E. O'Malley (Ed), *Computer-Supported Collaborative Learning*. Wiley & Sons

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.

GASSER L. (1991) Social conceptions of knowledge and action: DAI foundations and open systems semantics. *Artificial Intelligence*, 47, pp. 107-138.

MAES, O. (1988) Issues in Computational Reflection in P.Maes & D. Nardi (Eds) *Meta-level Architectures & Reflection* (pp. 21-36). Elsevier. Amsterdam..

ROSHELLE, J. (1990) Designing for Conversations. *Paper presented at the AAAI Symposium on Knowledge-Based Environments for Learning & Teaching*, Stanford, CA.

SALOMON, G. & GLOBERSON, T. (1989) When teams do not function the way they ought to. *International journal of Educational research*., 13 (1), 89-100.

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.

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.