Expertise, motivation and teaching in learning companion systems

Jorge Adolfo Ramirez Uresti, Benedict du Boulay

Publication date
01-04-2004

Published in
International Journal of Artificial Intelligence in Education

Expertise, motivation and teaching in learning companion systems

Article (Unspecified)

Expertise, Motivation and Teaching in Learning Companion Systems

Jorge Adolfo **RAMIREZ URESTI**[1] and Benedict du Boulay[2]

[1]Deparment of Computer Science, ITESM-CEM,
Apdo. Postal 50, Módulo de Servicio Postal,
Atizapán, Edo. Mex., MEXICO 52926.
juresti@itesm.mx
http://research.cem.itesm.mx/juresti/

[2]School of Cognitive and Computing Sciences,
University of Sussex, Falmer,
BRIGHTON BN1 9QH, UK.
bend@cogs.susx.ac.uk

Abstract

This paper describes work carried out to explore the role of a learning companion as a teachable student of the human student. A LCS for Binary Boolean Algebra has been developed to explore the hypothesis that a learning companion with less expertise than the human student would be beneficial if the student taught it. The system implemented two companions with different expertise and two types of motivational conditions. An empirical evaluation was conducted. Although significant differential learning gains between the experimental conditions were not observed, differences in learner behaviour between these conditions were. In particular students in the motivated condition with a weak companion taught it many more times than in the other experimental conditions and in general worked harder. Finally, the experiment also suggested that learning companions might be confusing for students if they try to resemble human behaviour, i.e. if they do not perform exactly as they are told.

1 Introduction

Recent research on Intelligent Tutoring Systems (ITSs) is exploring the benefits of having human students collaborate with computerised agents. The issues being studied range from the external representation of such agents (Rickel & Johnson, 1999) to the selection of their internal characteristics (Aïmeur, Dufort, Leibu, & Frasson, 1997; Hietala & Niemirepo, 1998). Among all of these systems, Learning Companion Systems (LCSs) extend the traditional model of ITSs by adding computerised agents whose aim is to provide a peer for the human student. This kind of agent is called a *Learning Companion* (Chan, 1991).

In principle this companion could be helpful to the student in a number of ways. For example, the companion could be a role model for the student (Okamoto & Kasai, 1999); both students could collaborate or compete as equals (Chan & Baskin, 1990); the companion could be a source of advice (Hietala & Niemirepo, 1998); the companion could be a student of the human student (Chan & Chou, 1997). This last role, being the student of the human student, has recently started to be explored (Chan & Chou, 1997; Jun, 1997). The rationale for such a selection of role is that by teaching the companion, the student should be able to reflect on her own knowledge and thus learn more effectively (Berliner, 1989; Goodlad & Hirst, 1989; Michie, Paterson, & Hayes-Michie, 1989; Nichols, 1994).

This paper describes work carried out to study this role: a learning companion as a student of the human student. A LCS for Binary Boolean Algebra has been developed to explore the hypothesis that a learning companion with less expertise than the human student would be beneficial for the student in her learning. The system was empirically evaluated in a study with 32 subjects. The results gave some support to the hypothesis that the weak companion is beneficial for the student's learning as it encourages teaching. However, for teaching to occur the student must be strongly motivated to interact with the weak companion. A trend was also found for the strong companion to encourage a passive attitude to learning in the student — it might be used as a tool that provides solutions. Finally, as an unexpected and interesting result, the experiment suggested that learning companions might be confusing for students if they try to resemble human behaviour, i.e. if they do not perform exactly as they are told.

The rest of this paper describes the research done. Section 2 discusses issues related to the Learning Companion and provides the rational for the experiments performed. Section 3 describes the research objectives and hypotheses of this work. Section 4 explains the system implemented to explore the hypotheses — LECOBA. Section 5 describes the subjects and tasks used for the empirical evaluation. Section 6 discusses the results obtained from the experiments undertaken with LECOBA. Section 7 provides a discussion of the results obtained. Finally, Section 8 presents the conclusions of this work and provides suggestions for future work.

2 The LC as a Student

Work on LCSs has increased in the last few years but much remains to be done to explore the full capabilities and possibilities of these systems. Self and his colleagues proposed ITSs which offered collaboration with the student rather than instruction (Gilmore & Self, 1988; Cumming & Self, 1989). But it was Integration-Kid (Chan, 1991), the first system built as an LCS, which introduced the idea of a Learning Companion (LC). The LC in this system was capable of collaborating or competing with the student. Integration-Kid proved the feasibility of LCSs and demonstrated that they stimulated learning interactions which are not possible with a computer tutor only. This work started a discussion about the characteristics that learning companions should have: What expertise should the LC possess? How should it be implemented, using machine learning or simulation? What role(s) could it take? How will the student communicate with the LC? How will the computer tutor relate to the LC? In fact, Chan's (1996) comment in a review of these systems is that 'there has been no serious evaluation of the effectiveness of having such a learning companion in a learning environment'.

From the above issues perhaps the most important are the expertise and the role that such a companion agent should have if it is going to be of educational value to the student. Integration-Kid explored the roles of collaborator and competitor, and sub-optimal expertise for the LC. These are only two of the roles and one expertise level that a learning companion could have. There are many other roles and expertise levels that these agents could take. The limits in the number of roles and expertise that learning companions could be programmed to represent are only defined by human peer learning.

2.1 Role and Behaviour

At this moment it is pertinent to define the concepts of role and behaviour in LCSs. In most of the LCS literature the role and the behaviour of a LC seem to refer to the same concept, i.e. there appears to be no difference between the concept of *role* and the concept of *behaviour* of a companion. Both concepts tend to be described at the same time. For example, in Integration-Kid (Chan, 1991) the companion collaborates via responsibility sharing. It is difficult to decide here if this description of the LC refers to its role, behaviour or both. Even though the concepts of role and behaviour are very similar, there is a difference between them which should be made explicit when discussing LCs.

The *role* of a LC is the function which it performs in the system. For instance, a LC could have one or various of the following roles: competitor, collaborator, advisor, student, etc. On the other hand, the *behaviour* is a description of the way in which the LC performs its role. For example, a collaborative companion could collaborate by working together with the student on a problem or by discussing how to solve a problem. The behaviour of the LC should also qualify the attitude in which the LC is performing its role: actively, passively, only when problems emerge, enthusiastically, etc.

A role could have various behaviours associated with it. In the previous example of the collaborative companion there are two behaviours: 1) working together with the student on a problem and 2) discussing how to solve a problem. Therefore, differentiating between roles and behaviours helps to create a categorisation of LCs.

2.2 Expertise

One of the most important issues to explore in LCSs is the question of the expertise level that the learning companion should possess in order to be of educational value to the student interacting with it. Most of the systems developed so far have dealt in one way or another with this issue. However, only Hietala and Niemirepo (1997, 1998) have designed their system in order to study explicitly the expertise level of the companion. Their interest was to select the companion's expertise to maintain the student's motivation to collaborate with the LC. They classified companions as *weak* or *strong* based on their expertise. For them a weak LC is one that has minimal expertise whereas a strong LC has almost an expert-like expertise. Their results showed that, in general, students preferred strong companions, specially when tasks got harder.

Students in Hietala and Niemirepo's experiment faced a LCS which provided four learning companions. Only one of them was active at any one moment. Two of them were strong and the other two were weak. The companions were not labelled as 'weak' or 'strong' so students had to establish their expertise by experimenting with them. Students had always the possibility to select at any time the companion they preferred to collaborate with. These subjects sought to collaborate with the strong learning companions when tasks got harder. A probable reason is that students saw the task as something that had to be completed. Collaborating with a strong learning companion was a way to complete the task as quickly as possible. On the other hand, weak learning companions were used by students specially at the beginning of the interaction, when tasks were easy. Later on, when tasks got harder, these companions were still used but this was not the general tendency. In general, students preferred strong companions.

Subjects in Hietala and Niemirepo's system did find the LC they wanted. But, was it the best learning companion to learn with?

Students, in general, were more comfortable with a strong companion at the end of the interaction, when tasks got harder. This was a very effective way to complete the task on hand, by asking someone who knows more for the answer. This is also one of the disadvantages of having a strong learning companion: if it is almost an expert in the domain, it could easily be confused with a tutor. Also, excessive use of strong LCs may make students not pay attention to the teaching they are receiving. And even worse, if the companion would answer or do all the student tells it, the student could end up by directing the companion to do all the work. This use of the companion may promote in the students a passive attitude towards their own learning. In any of these scenarios, there might be few benefits in having an LC, quite the opposite, there might be disadvantages instead of advantages.

On the other hand, subjects in the experiment used weak learning companions at the beginning of the interaction. Most probably this was because companions were not labelled with their expertise level — students did not know which type of companion they were using. This unclarity was deliberately introduced into the system to force students to test all companions until they had found the one they preferred. The use of weak learning companions was therefore due to a search for the best companion for a given individual, the one preferred. Although, it must be said, some students did prefer weak learning companions, increased their use and continued using them until the end, but this was not the most common choice. So, in general, weak companions were used mainly in the search for the preferred companion but were not good enough to collaborate with for serious tasks. This type of companion was also found to be disappointing, lazy and irritating. This is the main disadvantage of weak companions, that students find them a nuisance and decide not to use them anymore. However, interacting with a weak companion could potentially benefit the student's learning more than a strong companion could. A weak companion may allow a student to explain and teach to it. Seen in this perspective, a weak learning companion would give many benefits to a collaboration with an LC. There is, therefore, a need to understand what type of expertise should learning companions posses to be of pedagogical use to students.

2.3 Role: Learning by Teaching

Besides the issue of the expertise level of LCs, another important characteristic, which has been explored since the beginning of research in LCS, is their role.

In particular the role of the companion as a student of the user has just begun to be explored in the LCSs by Chan and Chou (1997) and Jun (1997). The results of the experiments conducted with these systems were not encouraging. Students did not benefit, as it was expected, by a 'learning by teaching' interaction with the companion. Subjects in Chan and Chou's experiments took a passive attitude towards teaching the learning companion. The tool provided for this purpose, called Diagnosis-Hint-Tree (DHT), allowed students to watch the learning companion's work if they wanted, or to make a minimal effort while teaching — as the tool gave suggestions on what to teach if the student could not teach correctly. In consequence, when teaching the LC, the student was not making any effort to understand what the LC was doing, losing in this way many of the benefits of teaching a peer. Subjects in Jun's work may have been discouraged to teach the companion by the tutor's constant interruptions while they were teaching. If the student was incorrectly teaching, the tutor would immediately intervene to rectify the situation. Therefore, the passive attitude of students in both systems towards teaching the learning companion may have been the cause of their failure to benefit from this kind of interaction.

However, research has found evidence to support the notion that Learning by Teaching can be a facilitator for learning. Students who teach other students learn more and better (Berliner, 1989; Goodlad & Hirst, 1989; Michie et al., 1989; Nichols, 1994). A student who needs to teach other people will have to revise, clarify, organise and reflect on her own knowledge in order to be able to teach, i.e. the student will need to master the knowledge. Besides, recent work on Teachable Agents has shown that students find teaching a virtual human agent interesting (Brophy, Biswas, Katzlberger, Bransford, & Schwartz, 1999). This agent captured the students' attention and motivated them to teach it. Teaching it was an active enterprise as students had to research and study beforehand. In consequence, the students in the experiments with this system showed high learning gains. Also, work by Scott and Reif (1999) has found that students who coached (taught) a computer tutor benefited as much as those subjects who had personalised tutoring from expert teachers.

In conclusion, there is an opportunity to investigate the application of the paradigm of learning by teaching into ITSs, in general, and in LCSs, in particular. Although empirical results from LCSs used to investigate if teaching a machine would produce similar benefits as those reported in Learning by Teaching have been contradictory, the drawbacks experienced seem to have been the result of the design of systems. More research with learning by teaching systems may find out that certainly learning by teaching a computer is beneficial for students (Ramírez Uresti, 1998).

2.4 Motivation

Motivating students to use the learning companion(s) has been an issue in LCSs since the beginning of the area. Unfortunately, it has not been studied enough.

Early LCSs relied simply on the *presence* of the companion as the pure motivational factor (Chan, 1991). However, not all people are motivated by interacting with another learner. Given this fact, it cannot be assumed that students will take advantage of the LCs just because they happen to be present in the system.

Proposals to solve this problem have been varied. But, the one that is most related to this work is the one suggested by Hietala and Niemirepo (1997, 1998). As mentioned before, Hietala and Niemirepo's interest in their experiments was to select the companion's expertise to maintain the student's motivation to collaborate. In this case, the selection of the LC's expertise is itself the 'motivating' factor through which students may decide to collaborate or not with an LC. Hietala and Niemirepo concluded that this motivating factor, the degree of expertise, was successful because students' interest for the companions was kept 'alive for the collaboration'.

The expertise as a motivational factor has been shown to be good for this purpose. However, for some cases it is not possible to use it as a motivational factor. When the objective is to motivate students to *collaborate with a particular type of learning companion*, which might be either weak or strong, Hietala and Niemirepo's motivational factor (the degree of expertise) is in fact a key element of the LC's required characteristics. It cannot be used as a motivational factor.

At this point, one may well be thinking why is it necessary to motivate students to collaborate with learning companions? To answer this question consider the case of a student using a system with a weak learning companion. A weak learning companion has poor knowledge of the domain so its comments, suggestions, etc. are often incorrect. The student interacting with this LC may at the beginning interact with the LC just out of curiosity. However, once the student realizes that the LC does not possess a good knowledge of the domain, she may well decide not to use the LC anymore, or worse, to stop using the system, wasting the benefits of having a companion.

A similar case can be envisaged for an interaction with a strong companion. Once students realize that a strong companion knows a lot about the domain, they can decide to let the companion do all the work — as in Hietala and Niemirepo's research where strong LCs were found to be the ones preferred by students due to difficult tasks. In this case motivating students to use LCs is not the issue, the issue is to motivate students to work more by themselves.

There is therefore a need to motivate students to interact with learning companions which have a predefined expertise. Or more generally, to motivate students to collaborate with companions without using required elements of the companion's characteristics as the only motivational factor. The motivational factor should strive to get the most beneficial educational collaboration with the companion depending on its characteristics.

2.4.1 Motivational Degrees

Different kinds of system will need different kinds of motivation. When deciding how to motivate a student to collaborate with the LC the degree of the motivation is a factor, i.e. how much emphasis will the system put on encouraging a student to interact with the LC? The degree of the motivational factor employed can be somewhere in a continuum of motivational degrees. On one end the encouragement can be very strong, even compulsory, putting a lot of emphasis on the collaboration with the LC. On the other end, it can be very weak, even non-existent, without any emphasis on the collaboration. For example, a very strong motivational factor to interact is via bi-directional dialogues. In this kind of interaction each agent has to contribute to the task in order to be able to continue with it. The motivation to collaborate is so strong, in fact compulsory, that the system will just not continue with the next activity if the student refuses to collaborate. People Power (Dillenbourg & Self, 1992) is a good example of a system with this kind of very strong motivation based in bi-directional dialogues.

At the other end of the possible degrees of motivation, an example of a very weak motivational factor has already been discussed: the mere presence of the companion. In this scenario collaborating with the companion is an option of the student. If she wants to collaborate, she does — there is no penalty for not collaborating. An example of a system with this kind of very weak motivational factor is the one developed by Hintze (1991).

When developing LCSs, the motivational degree and the way the student will be motivated to interact must be carefully decided. Depending on the aims, domain, task or stage, the motivation will vary from system to system. Furthermore, it may vary from task to task inside a system.

3 Research Objectives

As argued in the previous sections, there is a need to understand how the expertise of learning companions affects the students' learning, how students can be motivated to collaborate with such learning companions, and whether learning by teaching is an effective role for learning companions. Given these needs, the aims of the research described here were to explore:

1. The effect on the students' learning when interacting with weak and strong companions.
2. The motivation of students to interact with these two types of LCs, and
3. The effect of a teaching window as a reflection tool for learning.

3.1 Experimental Conditions

The issues of expertise, motivation and learning by teaching role were explored with a **LE**arning **CO**mpanion system for binary **B**oolean **A**lgebra (LECOBA — described in the next section). The system provided two types of learning companions (Weak and Strong) and two types of motivational conditions (Motivated and Free — see Table 1). In consequence, LECOBA was used in four different conditions which allowed us to investigate the effects of expertise and motivation in the students' learning. These conditions are summarised in Table 2.

3.1.1 Hypotheses

For each of the conditions in Table 2 students were expected to learn Binary Boolean Algebra (BBA). However, each condition was expected to influence the students' learning

Motivational Conditions	
Motivation	Description
Motivated	Motivation is strongly encouraged by means of scores. Both the student and the LC have their own score, plus a total score of their joint performance. Students must try to get the maximum joint score.
Free	Motivation is weakly encouraged by reminding students to collaborate with the LC. Tutor reminds students that collaborating with the LC is beneficial for them. Students are expected to at least be curious about the LC.

Table 1: Motivational conditions offered in LECOBA.

to different degrees.

Condition 1 (Weak/Motivated) was expected to be the best one of the four conditions for helping students learn. In this situation a student would be strongly motivated to interact with the companion and thus have many opportunities to learn by teaching the LC (*Hypothesis 1*). This was the main hypothesis of this work.

Students in Condition 2 (Strong/Motivated) would also be strongly motivated to collaborate, and were expected to learn by watching and collaborating with the LC. However, their learning should be less than in Condition 1 because there would not be much of a necessity to teach the companion (*Hypothesis 2*).

Students in Condition 3 (Weak/Free) would be weakly encouraged to interact with the LC. They were expected to lose interest in the weak LC and thus make use of the system in an ITS-like fashion. Their learning should be less than in Condition 1 due to little or no interaction with the companion (*Hypothesis 3*).

Finally, Condition 4 (Strong/Free) was expected to be the worst of the four conditions for learning. In this situation students were also weakly motivated to collaborate. Once a student realized that the LC is a strong one and can do most tasks by itself, it was expected that she would let the LC do all the work. If this were the case, students would become passive learners hampering their own learning (*Hypothesis 4*).

4 LECOBA

The LECOBA system consists of three agents: a tutor, a learning companion and a student (Ramírez Uresti, 2001). In this work users of LECOBA are referred to as 'students' or as 'users', whereas the pair made up of the student and the learning companion is referred to as such or as 'learners'.

A screen dump of LECOBA is shown in Figure 1. The figure shows the system at the moment when the learners are beginning to work on a problem. The windows shown here are: 1) the tutor's window at the top left corner, 2) the LC's window at the top middle of the screen, 3) the student's window at the top right corner, 4) a tool for the student to solve problems (Simplification Tool) at the bottom left corner, and 5) a window for the

Experimental Conditions		
	Expertise	
Motivation	Weak	Strong
Motivated	CONDITION 1 Student gets a lot of opportunities to teach the LC. *Student learns more than in the other conditions.*	CONDITION 2 Student gets few opportunities to teach the LC. *Less learning than in Condition 1.*
Free	CONDITION 3 Student ignores opportunities to teach the LC. *Student gets bored and may stop interaction with the LC.*	CONDITION 4 Student lets the LC do all the work. *Student hampers her own learning.*

Table 2: Experimental conditions under LECOBA. Hypotheses are shown in italics.

student to give suggestions to the companion (Student - Suggestion) near the middle of the screen.

Students using LECOBA follow the same cycle of activities. This cycle is presented in Figure 2. The system is designed to be used at several different times by the same user. It keeps a record of each user so they continue from where they quit in their previous session. When the student logs in for the first time, the tutor starts by teaching a lesson – teaching is done just by displaying in the tutor's window the text that students must read for that lesson. Immediately after the lesson, the tutor gives examples of the concepts just seen. Problems to be solved by the student and the learning companion follow. Finally, after they have finished solving the problem, the tutor comments on their performance. This cycle continues until the end of the curriculum.

When the tutor presents a problem to the learners it makes it clear that they both have to work on it as a team. There is flexibility in the division of labour within the team. It is up to the student to decide who will actually solve the problem and if she will interact with the LC or not. The student is always free to choose not to interact with the companion if she wants.

4.1 Motivation

LECOBA can motivate the student to collaborate with the companion in two modes: *Motivated* and *Free*. Whether the student will be assigned to a Motivated interaction or a Free interaction depends on the experimental condition being tested. The tutor will remind the student about collaborating with the companion every time it reaches the end of the comments stage, i.e. it will use the motivational condition selected to remind students.

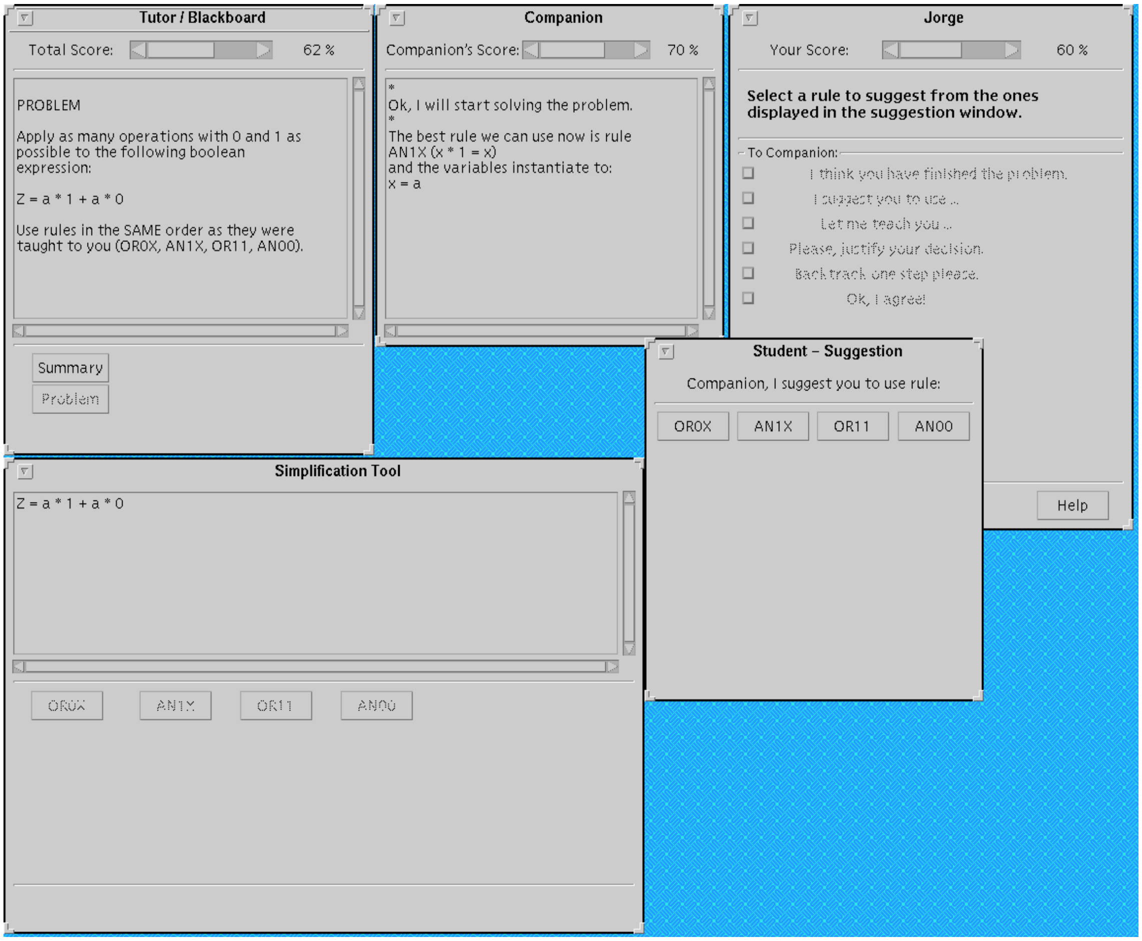

Figure 1: LECOBA. The learning companion is solving the problem and the student is about to give a suggestion to the companion.

4.1.1 Motivated Interaction

In Motivated interaction LECOBA has a built-in mechanism to motivate the student to collaborate with the learning companion: the Score Mechanism. This mechanism can be enabled or disabled before the student starts the first session with the system.

Figure 1 shows LECOBA with the Score Mechanism enabled. The scores are displayed at the top of the tutor's, companion's and student's windows as horizontal bars with their corresponding percentages to the right. Each score ranges between 0% and 100%. The scores in the companion's and student's windows are based on the performance of each respective learner. Scores are designed to challenge the student to interact more with the LC or to work more by herself. The challenge is to obtain the maximum score in the `Total Score` at the tutor's window. This score is determined not only by the student's performance, but also by the companion's performance. The companion's score can be improved mainly by teaching the companion. The student's score improves as a mixture of her own work and her involvement in the problem resolution when the LC is working. While the `Total Score` has not reached its maximum value, students will continue studying the same topic in the curriculum and will not advance to the next level.

A student is encouraged to interact with a weak learning companion, especially to let it solve problems so she could have the opportunity to teach it. If she has not been letting the companion solve problems, the companion's score will be decremented by a quarter of its total until the companion has solved at least the same number of problems

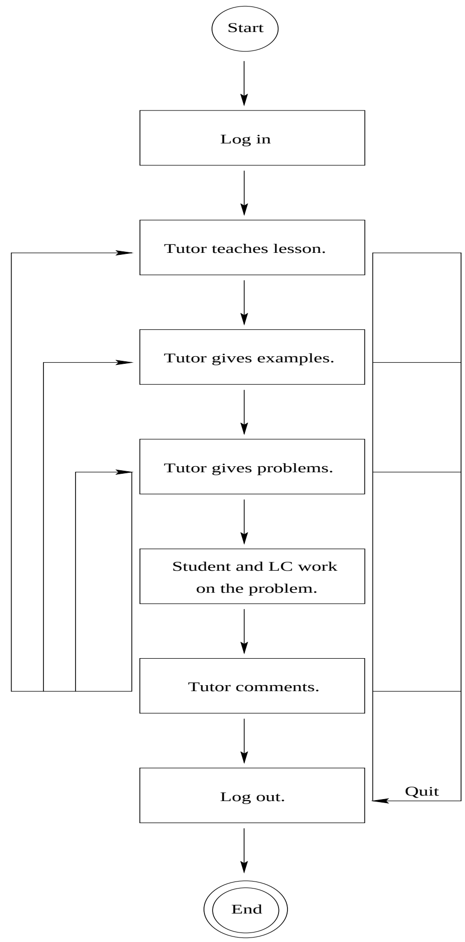

Figure 2: Cycle of Activities in LECOBA.

as the student. This action will impact on the `Total Score` preventing it from reaching its maximum value. Therefore, if the student wants to get a `Total Score` of 100%, she will have to let the companion solve more problems; in doing so, she will have the opportunity to teach the companion. A similar situation occurs with a strong companion. A student is encouraged to work more by herself when interacting with a strong companion. If she has not been solving problems by herself, her score will be decremented by a quarter of its total until she has solved at least the same number of problems as the companion. The student will have to solve more problems to reach the maximum score. Solving problems by herself allows her to practice more with the rules seen so far, preventing her from directing the companion to do all the work.

4.1.2 Free Interaction

In a Free interaction there is no Score Mechanism present. The tutor's, companion's and student's windows are displayed without the scores. The student is only mildly encouraged by the tutor to collaborate with the companion. At the end of the comments stage, the tutor will tell the student that it is beneficial for her learning to collaborate with the learning companion. *This is the only motivation for the student to collaborate with the companion in a Free interaction.*

The learners' progression in the curriculum does not depend on the student interacting or not with the companion, nor on the type of LC, and nor on their current knowledge of the rules. It depends on the number of consecutive correct problems solved by either learner in the current level. It is up to the student to decide whether to collaborate with the companion or let it solve any problems.

4.2 Teaching Window

While the learning companion is in charge of solving the problem, and the student is collaborating with it, the student has the opportunity to teach to the companion. Teaching the LC is done via a window called the *Teaching Window* (Figure 3).

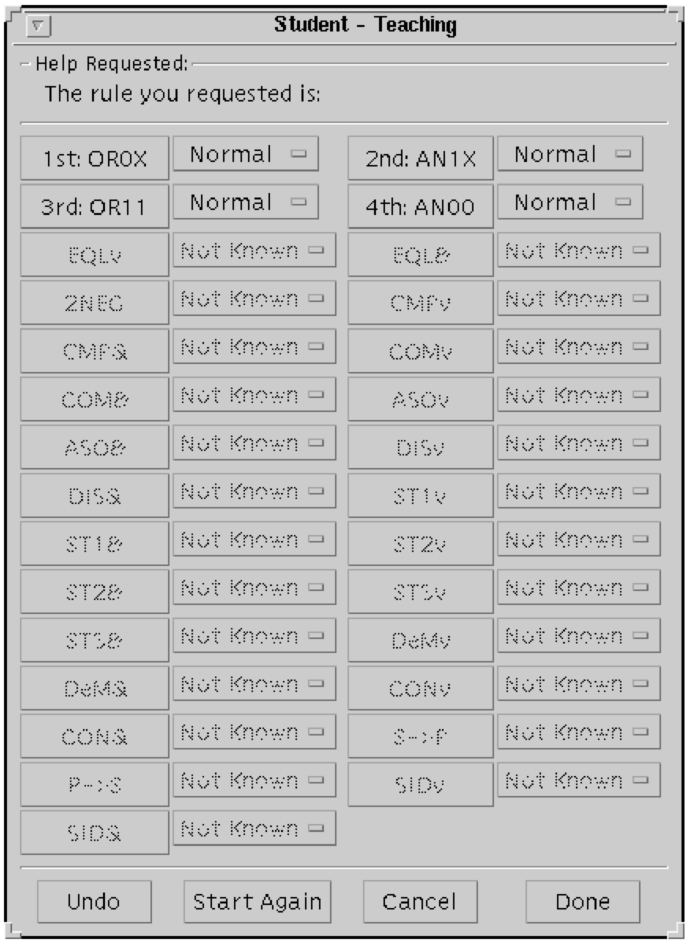

Figure 3: The Teaching Window in LECOBA.

The Teaching Window is based on the idea of inspectable student models (Bull, Brna, & Pain, 1995a; Bull, Pain, & Brna, 1995b). It presents to the student the companion's understanding of the domain at a specific moment during the interaction — i.e. its 'student model' (Companion's Expertise — CE). The objective is to let the student see exactly what the companion knows when trying to solve a problem. In Figure 3, the companion's knowledge is represented by a series of buttons and menus — each button with its corresponding menu. In this case only, the first four buttons and their menus are enabled. These represent the boolean rules which the learners have studied so far. Rules are ordered by the priority that the companion will try to apply them to a boolean expression. For instance, rule `OROX` is labelled as `1st`, this means that this rule is the first one which the companion will try to use when solving a problem. Rule `AN00` is the last one the companion would consider. The order in which the companion uses rules can be changed by clicking on their corresponding buttons and then swapping position with another rule.

The menus allow the student to tell the companion how to use a rule. She can tell it to use rules in a specific mode[1] or not to use them at all.

To teach the companion students must change its knowledge (CE) to make it more suitable for the task. Through the Teaching Window students can enable, disable, change the order and modify the way in which rules are used by the companion – i.e. students must click on the buttons and menus to select a new configuration of the companion's knowledge (CE). All of these changes are recorded while the student is performing them in the Teaching Window. They are considered as 'explanations' which the student is giving to the LC. Once the student has taught the companion, the learning mechanism decides how much the LC has understood from the student's teaching. Depending on the companion's current knowledge of the domain, a proportion of concepts is selected from the explanations given to the LC — selection is done from the first change in priority to the last one — and the rest of the new CE is obtained from the old CE. The companion automatically modifies its behaviour to immediately reflect its new understanding of the domain based on the student's teaching. The effect of this learning mechanism is that, if the companion did not 'understand' all the student taught it, when the student opens the Teaching Window on a subsequent occasion, she will notice that the learning companion did not 'learn' all it was taught by her the last time, i.e. that some of the changes in the CE made by the student have not been kept.

The Teaching Window can be seen as a reflection tool for the student. It encourages her to reflect on her own learning before deciding what to teach to the companion. This is the main objective of this window, to promote reflection before teaching to the companion. When teaching the companion the student will need to modify its knowledge in a way she considers to be better to solve the current problem. To select the companion's new knowledge the student will need to understand why the LC is using that particular combination of rules or heuristic. In order to try to understand the companion's knowledge, she will first need to think about her own knowledge of the domain, i.e. what knowledge does she use to simplify expressions and why. In summary, the student will need to revise, clarify, organise and reflect on her own knowledge before she can teach the learning companion effectively.

5 Empirical Evaluation

5.1 Design

The experiment conducted was a 2x2 factor, between-subjects design. There were two independent variables: expertise and motivation, and one main dependent variable: learning gain.

5.2 Subjects

The experiment was conducted in June 1999. Thirty two (32) undergraduate engineering students in their first and second years at the University of Sussex volunteered to take part in the experiment. Students in their first year had attended a course, where Binary Boolean Algebra (BBA) was taught, the term before experiments took place. LECOBA provided them with an opportunity to use a tutoring system to revise before the examination later in that same term. Students in their second year had taken the same course, as first year students, in their first year. LECOBA allowed them to review Boolean Algebra before they attended a course where previous knowledge of Boolean Algebra was needed.

[1]Some rules (the basic laws) can be used in different modes depending on the curriculum level being studied. This means that at higher levels the student must also decide the best mode for a rule to be applied to an specific problem.

Subjects in the experiment were volunteers and were paid for their participation. They were assigned randomly to each one of the four conditions (see Table 2).

5.3 Tasks

Each subject taking part in the experiment attended two sessions. Each of the sessions lasted approximately 1 hour.

In the first session students were given a pre-test on Binary Boolean Algebra. The time to answer the pre-test was restricted to a maximum of 15 minutes. Immediately thereafter the subject was introduced to LECOBA through a small demonstration. After the demonstration the student started the interaction with the system. This interaction was set to a maximum of 30 minutes due to time restrictions.

The second session should ideally have taken place the day after the first session to limit the effect of people having revised the topic with a different method. However, it was difficult to get all subjects back on the next day. Some subjects had their second session on the same day as the first session and some a few days after. The second session started with a 30 minutes interaction with LECOBA. Afterwards, students had to perform a post-test on Binary Boolean Algebra, which lasted 15 minutes as for the pre-test. Both the pre-test and the post-test consisted of a series of questions to asses the rules and strategies students employed to simplify Binary Boolean expressions (see Appendix). Jointly with the post-test, a questionnaire was given to students to asses their perception of LECOBA as a tool for their learning. Students were requested to fill in this questionnaire after completing the post-test.

All the students, without taking into consideration any previous knowledge of the domain, were requested to start the interaction with LECOBA from the beginning of the curriculum. The sessions were logged by the system.

5.3.1 Demonstration

Subjects had a small demonstration of LECOBA during their first session. The aim of this demonstration was to allow subjects to familiarise themselves with the system so they were able to make the most of their interaction with it.

The demonstration was presented by the experimenter and was of an informative nature. First, the experimenter presented the four windows displayed all the time in the system (tutor's, companion's, student's and Simplification Tool) and explained their functionality. Next, subjects were told the different stages of the interaction with LECOBA. The path followed during the demonstration of these stages was always to select the learning companion to solve the problem. This action permitted the experimenter to introduce subjects to the Suggestion Window and to the Teaching Window. Finally, regardless of the condition subjects were placed in for the experiment, the demonstration was given using the strong learning companion and the Motivated motivational condition.

During the demonstration the following facts about the interaction were emphasised:

First, that both, herself and the learning companion, had to collaborate to solve the problem presented to them as the learning companion was her peer.

Second, that the learning companion was also learning BBA.

Third, that she (the student) was always *free* to decide who will solve the problem — either herself or the learning companion.

Fourth, that the learning companion could reject or not her suggestions to it.

Fifth, that the learning companion could understand or not when it was taught.

Finally, that the scores may not be present when she used the system.

The subjects were not informed that there were two types of learning companion and two types of motivational conditions.

6 Results

6.1 Learning Gain

Table 3 presents the average scores of the pre-test and the post-test, and the difference between them (learning gain). In general subjects improved their performance in the post-test after the interaction with LECOBA. On a scale from 0 to 100 the average improvement was of 11.79 points.

Test Scores					
		Pre	Post	Diff.	Std. Dev.
Total		47.41	59.20	11.79	15.58
By LC	Weak	46.60	59.21	12.60	12.32
	Strong	48.21	59.19	10.98	18.68
By Motivation	Motivated	45.65	58.50	12.85	18.64
	Free	49.17	59.89	10.73	12.33
By Condition	1) Weak/Motivated	44.42	56.42	12.00	7.91
	2) Strong/Motivated	46.87	60.58	13.71	26.09
	3) Weak/Free	48.79	62.00	13.21	16.18
	4) Strong/Free	49.54	57.79	8.25	7.06

Table 3: Averaged scores. Scores range from 0 to 100. Diff. is the difference between the Pre and the Post scores. The standard deviation (Std. Dev.) presented is of the difference (Diff.).

After determining that the assumptions of the Analysis of Variance (ANOVA) had been met, a one-way independent measures ANOVA was run to determine if subjects in the four conditions were of similar expertise in the pre-test. Expertise of the LC (Weak/Strong) and motivation (Motivated/Free) were the independent variables. Pre-test scores (0 to 100) was the dependent variable. For an $F(3,31) = 0.0889, p = 0.9656 >> 0.05$ there was not enough evidence to suggest that subjects came from different populations, i.e. all the subjects had an initial similar expertise.

The pre- and post-test scores of subjects were used as a measurement to see if they had learnt from the interaction with LECOBA. A two-way mixed-design Analysis of Variance (ANOVA) was run after determining that its assumptions had been met. There were two between-subjects variables: expertise (Weak/Strong) and motivation (Motivated/Free), and one within-subjects variable with two levels: test scores (pre/post). The effect of the expertise ($F(1,28) = 0.08, p = 0.779 >> 0.05$) and of the motivational condition ($F(1,28) = 0.14, p = 0.714 >> 0.05$) were not statistically significant, neither their interaction ($F(1,28) = 0.34, p = 0.566 >> 0.05$). However, there was a highly significant effect of the tests (pre and post) on the subjects ($F(1,28) = 16.87, p = 0.000 << 0.01$). In other words, there was a very marked overall difference on the scores of the pre- and the post-tests. This is corroborated by the subjects impression on how much they learnt from the interaction with LECOBA. In a scale from 1 (A lot) to 5 (Nothing) the mean was 2.47. Subjects felt that they had learnt some BBA with the system.

6.1.1 Problems F & G

Because there was not a significant difference in the learning gain between the four conditions, the pre- and post-tests were analysed in more detail. Of the nine simplification problems subjects faced in the tests, problems F and G were the ones which, to be solved, needed almost specifically the application of the rules that students studied in the five levels implemented in LECOBA (see Part 2. Simplification Problems of the pre- and post-tests in the Appendix.) In fact, only problem G needed the use of a commutative rule (commutation in `OR`), which none of the subjects had problems with — as they demonstrated when solving other problems in the tests. Therefore, the analysis of these problems gives a more precise idea of the effect of each of the four conditions on the subjects.

Figure 4 shows the difference between the pre-test and the post-test scores of problems F and G, and their joint difference. From the figure it is clear that those subjects interact-

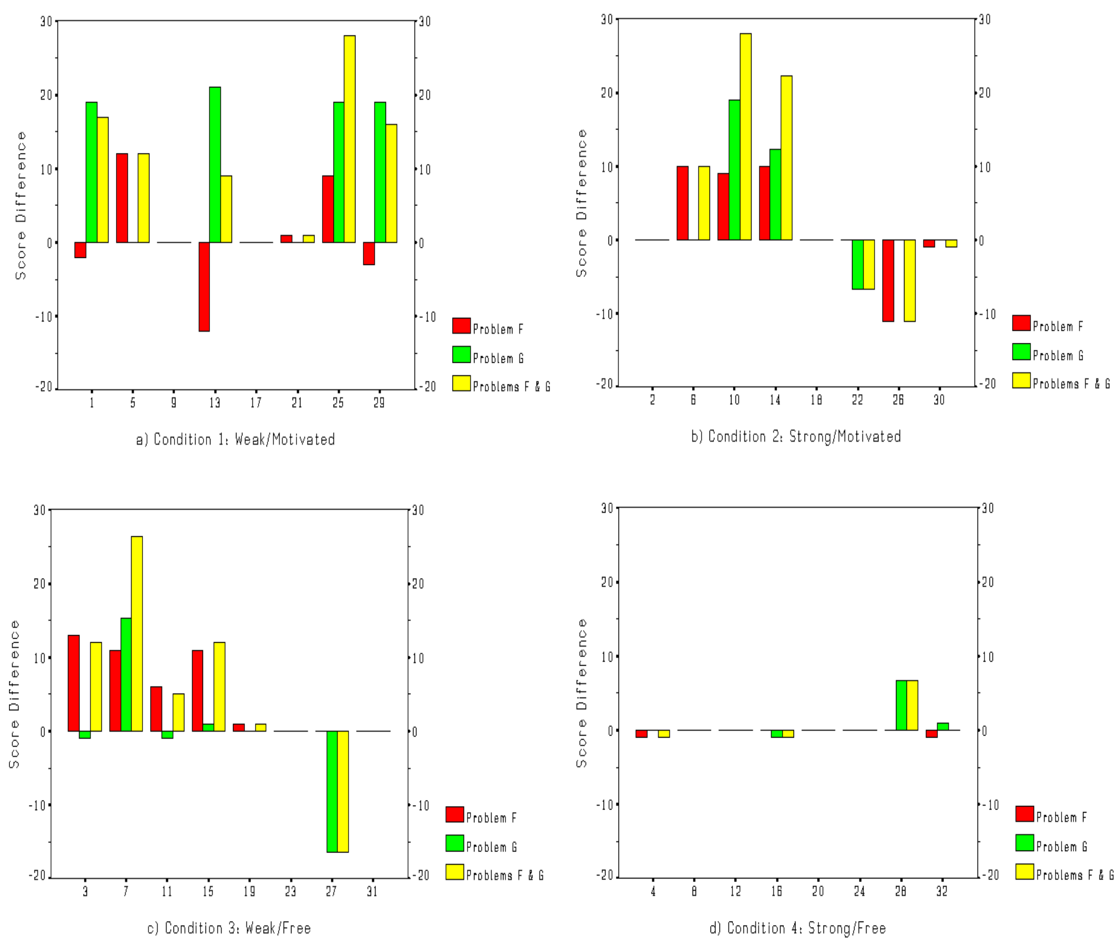

Figure 4: Difference between the pre-test and the post-test scores of problems F and G. The numbers below the bars are the subject identity numbers.

ing with a weak companion improved more in these problems than subjects with a strong companion. And from these two conditions, subjects in Condition 1 (Weak/Motivated) improved most. The improvement is more clearly seen in Figure 5, which shows the improvement percentage for each subject in problems F and G, and their joint improvement percentage.

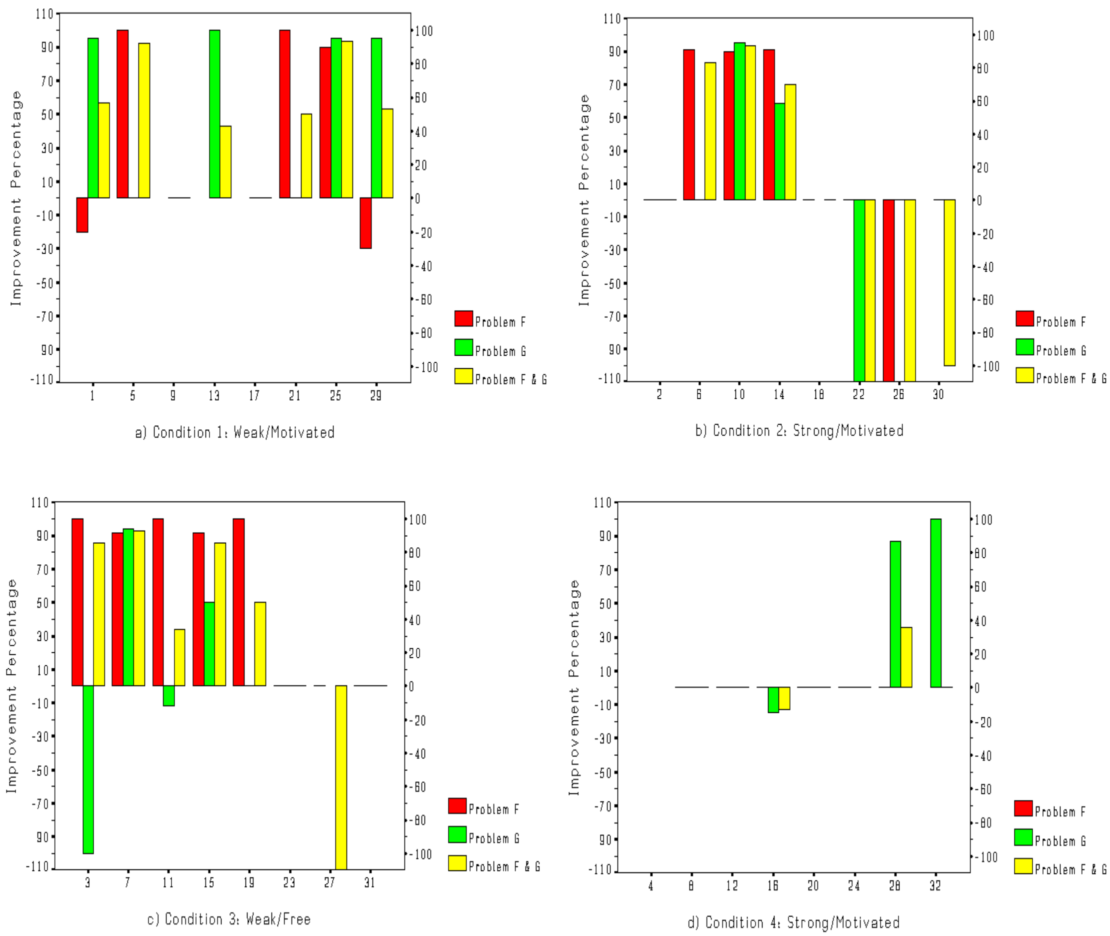

Figure 5: Improvement percentage in problems F and G. The percentage is based on how much a subject could have improved from the pre-test to the post-test regardless of the amount of points — for example, if a student had a pre-test score of 70 and the maximum mark was 100, he could have improved 30 points from the pre-test to the post-test. The numbers below the bars are the subject identity numbers. Those bars which reach the lower limit go in fact much lower.

6.1.2 Effect of Teaching and Suggesting

Teaching was hypothesised as an activity that would help students to benefit more from the interaction with LECOBA. To analyse the effect of teaching on the subjects' learning gains, a correlation between the score improvement and the number of teaching incidents was performed for the total learning gain and for the learning gain of problems F and G. Table 4 presents this analysis.

For the total learning gain there was a small positive correlation that was only significant with the presence of outliers — and without the outliers was very close to being significant. In the case of problems F and G, there was as well a small positive correlation which was not significant either with outliers or without them – but again, the correlations were near to being significant.

In a similar manner, an analysis of correlation was performed for the number of suggesting incidents (see Table 5). Giving a suggestion to the companion may have also been helpful for subjects as this action allows for some reflection before actually giving the suggestion. For the total learning gain there was a small positive correlation that was significant with and without outliers. However, for the learning gain of problems F and

G, the correlation was also small and positive but it was not significant — although it was close to being significant. These results give some support to the claim that teaching or giving a suggestion to the companion may have been beneficial for the students learning.

Correlation Analyses for Teaching					
		Kendall's		Spearman's	
Data Type	Number of Data	Correlation	Significance	Corr.	Sig.
All Questions	Outliers (32)	0.2513	0.030	0.3350	0.030
	No outliers (30)	0.1814	0.096	0.2307	0.110
F and G	Outliers (32)	0.1823	0.094	0.2337	0.099
	No outliers (31)	0.2171	0.061	0.2742	0.068

Table 4: Correlation between teaching incidents and learning gain. The data was analysed for the total number of questions in the tests and for problems F and G only. The correlation tests used were: Spearman's rank correlation and Kendall's tau-b, both for non-parametrical data. The significance for both tests is one-tailed.

Correlation Analyses for Suggesting					
		Kendall's		Spearman's	
Data Type	Number of Data	Correlation	Significance	Corr.	Sig.
All Questions	Outliers (32)	0.3207	0.008	0.4329	0.007
	No outliers (31)	0.2939	0.015	0.3979	0.013
F and G	No outliers (32)	0.1738	0.102	0.2164	0.117

Table 5: Correlation between suggesting incidents and learning gain. The data was analysed for the total number of questions in the tests and for problems F and G only. The correlation tests used were: Spearman's rank correlation and Kendall's tau-b, both for non-parametrical data. The significance for both tests is one-tailed. There were no outliers for F and G problems.

6.2 Motivational Condition

Figure 6 shows the average number of problems subjects solved. The average was 12.81. Overall the student solved (8.81) more than twice the problems that the companion solved (4.00). With a weak companion students solved (9.50) almost two times more problems than the LC (5.13). With the strong companion students also solved (8.13) more problems than the LC (2.88) — almost three times more. In both Motivated and Free interactions the student solved more problems than the LC. As expected the overall number of problems solved in the Motivated condition (16.44) was higher than in the Free condition (9.19), the student worked more in the Motivated condition (11.44) than in the Free condition (6.70), and the LC worked more in the Motivated condition (5.00) than in the Free condition (3.00).

6.2.1 Motivated

Most subjects in the Motivated motivational condition responded in the affirmative when asked if they were encouraged to collaborate with the LC because of the scores — only two responded in the negative. However, when asked if they could have advanced faster in the curriculum if the scores were not present, about 40% responded in the affirmative. This is

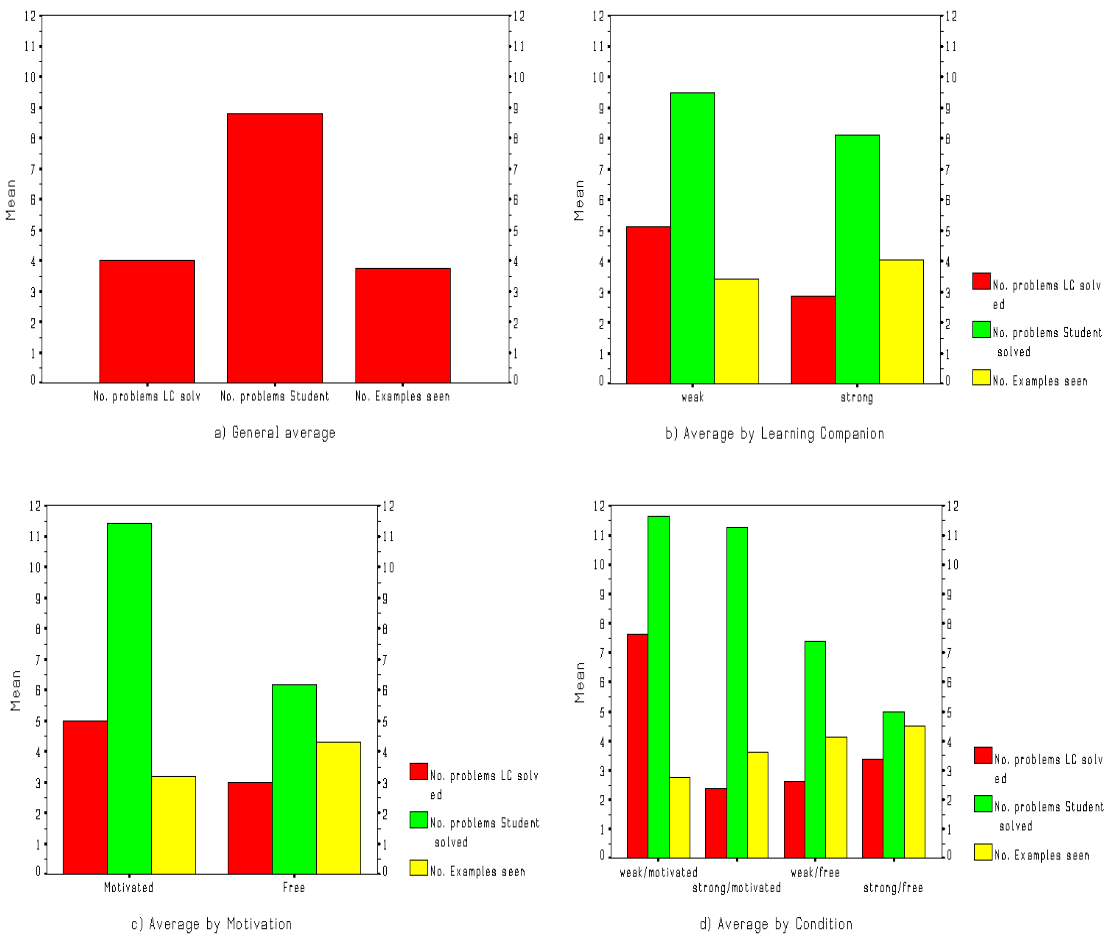

Figure 6: Average number of problems and examples seen.

a sign that some subjects may have found the scores confusing and difficult to understand. Subjects with the weak companion solved considerably more problems (11.63) than the LC (7.62). This was not expected. The expectation was that either the LC would have solved more problems than the student or that both would have solved a similar number of problems. On the other hand, as expected, subjects with the strong companion solved (11.25) almost five times more problems than the LC (2.38).

6.2.2 Free

Most subjects (11 out of 16) in the Free motivational condition responded in the affirmative when asked if they had collaborated with the companion as a result of the tutor encouraging them so to do. Of the five subjects who responded in the negative, four interacted with the weak companion and only one with the strong. Therefore, half of the subjects who interacted with the weak companion did so because they were encouraged by the tutor and the other half because they wanted to. Nevertheless, even though most subjects said that they were motivated to collaborate with the companion, subjects solved more problems than the LC. As expected, subjects with the weak companion solved almost three times more problems (7.38) than the LC (2.63). But contrary to expectations, subjects with the strong companion solved more problems (5.00) than the LC (3.38). The expectation was that the strong LC would have solved more problems than the student.

6.2.3 Physical Activity

An alternative measure of the amount of work performed by the subjects in the interaction is the number of button presses they made while using LECOBA. Subjects with the weak companion had a higher average of button presses (440.63) than subjects facing the strong companion (341.56). For the motivational condition, as expected, subjects in the Motivated condition pressed the buttons more times (449.94) than subjects in the Free condition (332.25). Per condition (Figure 7), as expected subjects in Condition 1 (Weak/Motivated) were by far the most active of the four conditions (547.13), followed by subjects in Condition 2 (Strong/Motivated, 352.75), Condition 3 (Weak/Free, 334.13) and Condition 4 (Strong/Free, 330.38). This last result emphasises the fact that students in Condition 3 did not interact much with the companion or its physical activity would have been higher due to teaching the LC.

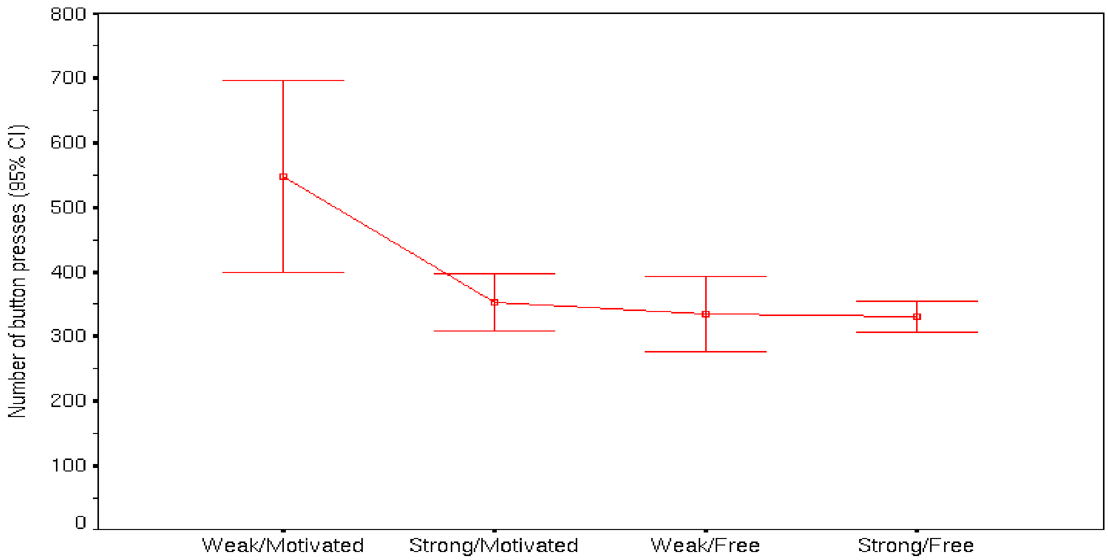

Figure 7: Average number of button presses per condition.

6.3 Teaching the LC

Subjects taught the LC an average of 4.44 times. As expected the weak companion (8.00) was taught much more than the strong companion (0.87) — more than nine times as much. Also as expected, the companion was more taught in the Motivated interaction (7.19) than in the Free interaction (1.69) — more than four times as much.

Considering the four conditions separately (Figure 8 a), Condition 1 (Weak/Motivated) had by far the highest number of teaching incidents (14.00). The smallest number of teaching incidents occurred in Condition 2 (Strong/Motivated) with an average of 0.38 incidents. Comparing the teaching incidents between the subjects who interacted with the weak companion, as expected in the Motivated interaction (14.00) there were seven times as many incidents as in the Free condition (2.00). Also, as expected the difference in the teaching incidents between the four conditions was significant (see Table 6). Subjects in Condition 1 taught significantly more than subjects in any other condition. Figure 8 b) shows clearly the effect of the type of LC and the motivational condition in the number of teaching incidents.

Finally, most subjects (24) responded that they believed they had learnt as a result of teaching the LC. From these subjects eleven (11) were with the weak companion and

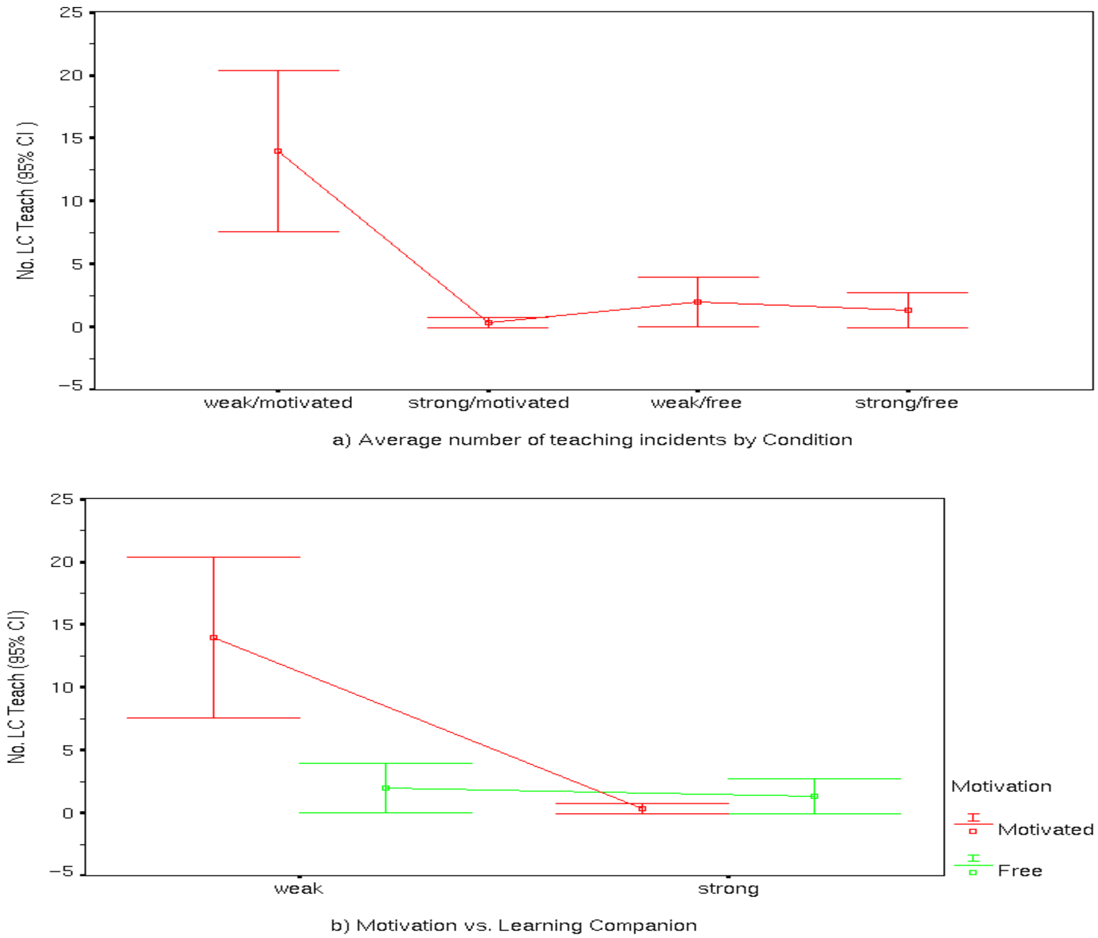

Figure 8: Average number of times subjects taught the LC per Condition.

thirteen (13) with the strong one. By motivational condition, thirteen (13) were in the Motivated condition and eleven (11) in the Free condition. By condition, six (6) were in Condition 1 (Weak/Motivated), seven (7) in Condition 2 (Strong/Motivated), five (5) in Condition 3 (Weak/Free) and six (6) in Condition 4 (Strong/Free). It is interesting to see that, although by little, more subjects who were with the strong companion thought that they had learnt by teaching the LC. In any case, teaching the companion was seen by subjects as a beneficial activity for their own learning.

7 Discussion

7.1 Learning Gain

Subjects improved their scores significantly from the pre-test to the post-test. This improvement must have been the result of the teaching administered by LECOBA, the subjects' teaching and suggesting to the learning companion, the subjects' previous expertise in the domain, and their experience of solving the tests. Teaching the learning companion and giving suggestions to it might have had some effect on the students' learning, as the correlations found for the these activities suggest. Although the effect of these activities was expected to be higher, the engagement in them might account for some of the improvement in the post-test scores. Also, that subjects improved their scores as a result of their previous expertise in the domain was expected.

Statistical Analyses for the number of Teaching incidents			
Test	Independent	Decision Value	Significance
1)	All conditions	$\chi^2_{.05}(3)$, $H = 18.4749$	$p = 0.0004 << 0.01$
2)	Conditions 1 and 2	$Z = -3.4242$	$p = 0.0006 << 0.01$
2)	Conditions 1 and 3	$Z = -3.1282$	$p = 0.0018 << 0.01$
2)	Conditions 1 and 4	$Z = -3.2725$	$p = 0.0011 << 0.01$

Table 6: Statistical analyses for the number of teaching incidents. The dependant variable for all the analyses is `Number of Teaching Incidents`. The type of tests are: 1) Kruskall-Wallis for non-parametrical data, corrected for ties, and 2) Mann-Whitney U test for non-parametrical data, corrected for ties.

It was also observed that many subjects during the pre-test took more time solving it than the post-test[2]. When answering the pre-test, even though subjects were told that they had only 15 minutes for this task, they were progressing slowly and many of them were surprised to realize that the time for the test had finished 'so soon'. Therefore, in the post-test subjects were more aware of the 15 minutes to answer the test. The main effect of this behaviour was that most subjects answered more questions in the post-test that in the pre-test. Taking this fact into consideration, the analysis of only the problems F and G of the tests makes more sense. They were two of the questions which were solved by almost every subject in both tests.

Unfortunately the effect of the different learning companions' expertise and of the motivational condition in the improvement of the subjects' scores was not distinguishable, as the differences between the conditions were not significant. However, there were some trends in the data in some of the conditions which allow for conclusions based on the expertise and on the motivational condition.

Condition 1 (Weak/Motivated) could be said to have a trend of being the most beneficial for the subjects' learning. It had a very low dispersion of the learning gain, i.e. learning under this condition was similar for all subjects in the condition. And it was the condition where subjects improved most when problems F and G were considered separately from the other simplification problems. These two problems, F and G, were the only ones which almost exclusively considered the rules and concepts that subjects would see in the five levels implemented in LECOBA. Had all the test been based only on these rules and concepts, the effect of the companion's expertise and motivational condition may have been significant. This trend in Condition 1 of subjects improving consistently and more than in the other conditions — although not significant, gives some support to this work main hypothesis: that a weak LC would help the student to learn by teaching it.

On the other hand, *Condition 4 (Strong/Free) seems to have a trend of being the worst condition for learning.* In this condition subjects learning gain was very similar but had the worst improvement when all the problems of the tests were taken into consideration and when only problems F and G were used. Therefore, subjects in Condition 4 had a consistent low improvement. Besides these results, when the tests were analysed for improvement in the understanding of BBA concepts, subjects in this condition were the only ones with a trend of low beneficial effect on their understanding of strategies. Some subjects even had a negative effect on their learning of the strategies. These subjects may have been confused by the new heuristics. They did not practice them enough as they preferred to solve problems by themselves, not to collaborate with the companion and were allowed by the tutor to progress easily in the curriculum. All of these situations

[2]The exact times to solve the tests were not recorded. The observation was taken from the notes of the experimenter.

allowed them to solve problems without having to use the heuristics taught by the tutor. In summary, subjects in Condition 4 were the ones who least benefited from the interaction with LECOBA. This trend gives some support to the Hypothesis 4: subjects in this condition would be the worst of the four conditions.

7.2 Motivation

The Motivated motivational condition was designed to encourage subjects to collaborate with the weak companion and to work by themselves when facing the strong companion. On the other hand, the Free motivational condition was designed to let subjects decide how much collaboration they would have with the weak and the strong companions.

7.2.1 Motivated

Subjects in the Motivated motivational condition were the ones who worked hardest. They spent more time working with the system than subjects in the Free condition. Many more problems were solved in the Motivated condition than in the Free condition. The physical activity of subjects was higher in the Motivated condition than in the Free condition. To fully appreciate the impact of the Motivational condition in the subjects' work, the number of messages they sent to the companion must be taken into consideration. When the LC was working on the problem, subjects sent it more than twice as many messages as in the Free condition. And when the student was in charge of solving the problem, the number of messages to the LC was also higher than in the Free condition. Therefore, as all of these results show, *subjects in the Motivated condition were the ones who had to work hardest in the interaction with LECOBA.*

With the strong companion, as expected, subjects solved more problems than it did. The difference between the number of problems solved by the student and the LC is very big. It is possible that these subjects had an inclination of solving problems by themselves and that the scores encouraged them to do so more frequently than if the scores had not been present. In any case, *the scores worked as expected with the strong companion (Condition 2, Strong/Motivated).* On the other hand, contrary to expectations, with the weak companion subjects also solved more problems than it. But, if all but two subjects in the Motivated condition said that they had been encouraged by the scores to interact with the LC, why did subjects with the weak companion solve more problems than it? From an analysis of the logs of the sessions it is clear that 3 subjects facing the weak companion markedly preferred to solve problems by themselves. For long periods of time they chose to solve problems by themselves mixing them with a few selections of the LC. Therefore, they were not using the scores all the time to decide who should solve the problems. This behaviour had two effects. First, these subjects could not advance to the next level — as they were requested to collaborate with the learning companion. Second, it increased considerably the mean of the number of times students facing the weak companion decided to solve problems by themselves. The logs show that the other subjects (5) in Condition 1 (Weak/Motivated) did mix the number of selections of themselves and the LC. They collaborated more evenly with the LC as a result of the scores. So, *in Condition 1 there is a trend, though not a clear indication, that the scores worked as expected.*

A high percentage of subjects thought that they would have advanced faster if the scores had not been present. Some subjects expressed discomfort with the scores because they did not understand clearly how they worked. An example, which describes the feelings of some of the subjects towards the scores, is the comment given by subject 17: "could not seem to progress even when answered lots of questions correctly". The result of the boolean expressions this subject had solved may have been correct but not the heuristic used, so he had to solve more problems. Consequently, the scores must be improved to provide the students with more specific indications of their progress in the curriculum. They must let them decide efficiently who should solve a problem, either the LC or themselves.

In summary, the scores worked as expected. There was a clear indication that with the strong companion subjects preferred to work by themselves. With the weak companion the trend was that subjects selected the LC to solve problems as a result of the scores. Subjects in the Motivated condition were the ones who worked most. Finally, subjects had a better impression of the LC in the Motivated interaction than in the Free interaction. This may have been due to the higher exposure to the LC in the Motivated condition.

7.2.2 Free

Subjects facing the weak companion in the Free interaction (Condition 3) solved more problems than the LC. This result was expected. The tutor's mild encouragement to interact with the weak companion may not have been really necessary. Half of these subjects mentioned that they had collaborated with the weak companion because they had wanted to. Nevertheless, the interaction with the weak LC was minimal as the number of messages to the weak companion is almost the lowest when the LC was working on the problem and the lowest when the student was in charge. Furthermore, all of the subjects finished the interaction having decided to solve most problems by themselves and not letting the companion work. In conclusion, *subjects may interact with the weak companion just out of curiosity or because the tutor suggested it, but once they realise that its knowledge is poor, they are very likely to stop interacting with it if they are allowed.* This result gives support to Hypothesis 3.

In the case of the strong companion subjects also decided to solve more problems than it did. The difference between the number of problems solved by the student and the LC was not much. This result was not expected. The expectation was that in Condition 4 (Strong/Free) subjects would select the LC to solve most problems once they had realized that it was strong. It would seem that subjects preferred to solve problems entirely by themselves. However, this was not completely true. Subjects had a lot of collaboration with the strong companion — the number of messages sent in this condition is almost the highest of the four conditions both when the LC was in charge of the problem and when the student was solving it. Most of these messages were related to the strong LC's suggestions. And the logs of the sessions reveal that there is a tendency in this condition for students to ask for advice from the LC while they were working. Most of the subjects increased the number of requests over the interaction — only two asked for very little advice. Therefore, even though students chose to solve more problems by themselves, they were not solving them only by themselves. They were really seeking the support of the strong companion to solve problems. Consequently, *there is a trend of using the strong companion as a tool that will provide solutions to the problems.*

In summary, subjects under the Free motivational condition solved more problems than the weak LC and eventually decreased their interaction with it. With the strong companion subjects also solved more problems than the LC. But they were not solving the problems entirely by themselves, they were seeking the companion's suggestions to solve them.

7.3 Suggestions to the LC

A very interesting result was the subjects' reactions to the learning companion's answer when it was given a suggestion. The learning companion accepts or rejects suggestions in accordance with the following rule: the more it knows about a rule, the less probable it will accept the suggestion (unless it is the same rule it suggested).

Mainly with the strong companion, but sometimes with the weak companion as well, subjects were confused to see that the companion was rejecting their suggestion. If the companion was strong, they were more amazed to read that the companion was telling them that they were wrong and that it would not follow the suggestion given. Subjects were not happy to have their suggestions to the learning companion rejected by it. One of them called the experimenter because he believed that he had found a bug in the system as the companion was rejecting all of his suggestions. He said to the experimenter: "How can I force the companion to do what I want?". This comment clearly demonstrates how subjects were not prepared to be rejected by a machine and wanted it to do as it was told. The experimenter had to explain to the subject that collaborating with the companion should be similar as collaborating with another person. Even after this explanation the subject was still unsure that the machine should be 'convinced' to accept his suggestions. In any case, subjects were expecting the companion to accept the suggestion that it had been given. They were not prepared to be rejected by a machine. Eventually most students got used to this kind of behaviour of the companion.

Programming a learning companion to behave as a human peer seems not to have been very well received by the users and, as a consequence, caused unwanted effects. The companion's behaviour to reject the user's suggestion when it 'thought' that its own rule selection was correct damaged the collaboration between the learning companion and the student. Subject 24 expressed this unwanted effect very clearly: "Often when I know the companion is saying something wrong, I did suggest something, the companion outright rejects it. Thus the user (as I) gets irritated and justs (sic) solves the problem himself". *In LECOBA, that the companion rejected the student's suggestion was a source of confusion and for some subjects it may have generated a damaging effect to their interaction with the companion.* More research needs to be done on the behaviours that learning companions can exhibit as credible for users.

7.4 Teaching the LC

Teaching the learning companion also proved to be a source of confusion and irritation for subjects, specially for those interacting with the weak companion. The learning companions in LECOBA were programmed not to understand some of the concepts that were being taught to them.

Subjects commented that the companion was very difficult to teach. If the subjects had very carefully taught the companion and it happened not to understand what was taught to it, subjects would have had to teach it all the concepts again. Subjects were confused to see this behaviour of the learning companion as they expected it to memorise all that was taught to it. It also created some reluctance to teach the companion. In the sessions' logs it was clear that subjects in Condition 3 (Weak/Free) decreased the clarity and carefulness of their teaching after they had observed that the companion would not always learn all that had been taught to it. They did not like it that the companion was not learning. Some subjects perceived this process of repeated teaching as an activity which was detrimental, not for their learning, but for their speed to solve problems. Subject 27 commented on this issue: "[Teaching the companion] slowed me down, the companion was difficult to teach". However, a high speed in solving problems does not necessarily mean good learning.

Teaching the companion all the concepts every time it had to be taught is the issue that seems to have irritated subjects most. Actually, the companion did not have to be taught all the concepts every time subjects wanted to teach it. The companion had only to be taught those concepts that it did not understand. A weak companion, therefore, had to be taught more concepts than a strong one. But in any case, subjects had to check carefully that the companion understood all the concepts. It is likely that this checking activity was most irritating for subjects. It constituted a difficult and time consuming task, and if the companion had to be taught a lot of concepts two or three times, because it did not understand the first time it was taught, then subjects would get irritated as they would have had to check again and re-teach. The inability of the companion to understand quickly was deliberately programmed to encourage in the students more and careful teaching. However, clear and careful teaching decreased for those subjects who were in Condition 3 (Weak/Free). A possible improvement to the teaching of the companion could be to also have the possibility to teach by example, as subject 9 suggested. He pointed out: "the teaching window is not as someone teaches". Teaching by example would eliminate the need of having to teach 'all' the concepts every time teaching is wanted retaining the benefit of having to teach some concepts several times — as this may help students to think more carefully about what they are teaching.

Having to teach the companion several times also posed, at least to one of the subjects, an interesting dilemma in that subjects may know how to solve a problem with another person but not with a learning companion. Subject 29 asked: "the LC does not understand and I do not want to tell it that it has finished [because it was not true]. What can I do?". The solution was to teach it again as this activity was the one being encouraged in the subjects. It is very likely that this particular subject may have just given up with a human peer but LECOBA did not allow the user to react in this way. An important question arises then: if the learning companion is behaving in a manner similar to a human student, does the system need to provide the user with all or most of the possible reactions that she could have with a human peer? A tentative answer is that it depends on the aims of the system. In LECOBA, in Condition 1 (Weak/Motivated), the aim was to encourage the user to teach the companion and, therefore, the student should not have been given the possibility not to teach the companion if there was a need to do it. Teaching the learning companion was essential for the system's aims.

In conclusion, a learning companion that does not understand all that it is taught is confusing for subjects because they expect the machine to understand everything that is taught to it. Having to teach the companion several times can slow down the user but it can also give her many benefits. As subject 8 put it: "The process requires a lot of thinking and personal effort so things may take longer than they usually would but they stay in your mind for *longer*!".

8 Conclusion

This paper has explored two types of expertise level for learning companions (weak and strong) and two types of motivational condition (Motivated and Free). While differential learning gains were not found between conditions, differences in learning activity were observed. In particular it was found that students with a weak companion in the motivated condition taught their companion many more times that in the other three conditions and in general worked harder.

Given the nature of the pre and post-test, it is not possible to say whether the lack of differential learning gains was a result of the imprecision of the tests, the relatively short exposure time of the subjects, the small number of subjects or other factors associated with the system which masked any effects of teaching. A future experiment would need to make the activity of teaching the companion more transparent and less irritating. It would also need to make solution strategies for these problems more explicit rather than learning them implicit simply as rule orderings.

Finally, future work on LCSs could explore the effect of the relations between the agents of an LCS in the student. LECOBA implemented two types of interactions between the agents of the system. A private interaction between the tutor and the learning companion, and public interactions between the three agents of the system (see Figure 9). Only the

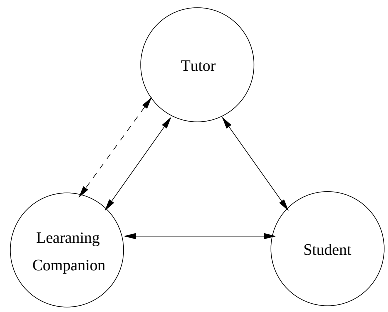

Figure 9: Agent architecture of LECOBA. The tutor controls the LC. Solid arrows represent public interactions. Dashed arrows represent private interactions.

agents in a private interaction were aware of its existence whereas all agents were able to 'see' a public interaction. These interactions represent relations which exist between the agents of the system. In LECOBA's case, the tutor and the companion had a relation between them that the student was not aware of. This notion of private and public interactions as representing relations could be generalised to a model which includes all the private and public interactions that may exist between the agents of an LCS. Figure 10 presents an abstract representation of this model for LCSs with three agents.

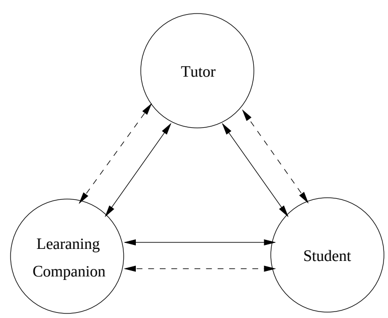

Figure 10: Relationship model for LCSs with three agents. Solid arrows represent public interactions. Dashed arrows represent private interactions.

Not all the systems resulting from implementing the possible interactions between the agents may be viable or interesting. For instance, having a system in which the only interaction is a private one between the tutor and the learning companion would certainly not be beneficial for the student, she would not be able to observe anything in a system like that. Some other interconnections would resemble systems already implemented. For example, having public interactions between the three agents would resemble the original LCS proposed by Chan (1991). Systems that may be interesting to explore would be those which include private interactions between the student and the tutor or the learning companion. Those kind of systems could explore if there is a benefit in providing private interactions between the student and the tutor and between the student and the companion. For instance, consider a system based on Pask's idea of TEACHBACK (Pask, 1976a, 1976b; Entwistle, 1978). The student could be asked to teach the learning companion the concepts that the tutor teaches her. In a system like that a private interaction between the student and the tutor may be used to help the student teach the learning companion. The interaction would need to be private as, in principle, if this interaction were public, the companion may 'learn' just by 'hearing' what the tutor is commenting to the student about how to teach the companion. This second scenario would also be interesting to explore: how would the student react to agents who learn just by watching what she is doing if she is not aware of this fact? This model of relationships between the agents in LCSs could provide interesting ideas to explore in future research.

Acknowledgements

The first author would like to thank the Mexican people who through CONACyT of México funded this research.

References

Aïmeur, E., Dufort, H., Leibu, D., & Frasson, C. (1997). Some Justifications for the Learning by Disturbing Strategy. In du Boulay, B., & Mizoguchi, R. (Eds.), *Artificial Intelligence in Education: Knowledge and Media in Learning Systems. Proceedings of AI-ED 97 World Conference on Artificial Intelligence in Education*, Vol. 39 of *Frontiers in Artificial Intelligence and Applications*, pp. 119–126. International AIED Society, IOS Press.

Berliner, D. (1989). Being the teacher helps students learn. *Instructor*, *98*(9), 12–13.

Brophy, S., Biswas, G., Katzlberger, T., Bransford, J., & Schwartz, D. (1999). Teachable Agents: Combining Insights from Learning Theory and Computer Science. In Lajoie, S. P., & Vivet, M. (Eds.), *Artificial Intelligence in Education. Open Learning Environments: New Computational Technologies to Support Learning, Exploration and Collaboration*, Vol. 50 of *Frontiers in Artificial Intelligence and Applications*, pp. 21–28. International AIED Society, IOS Press.

Bull, S., Brna, P., & Pain, H. (1995a). Extending the Scope of the Student Model. *User Modelling and User-Adapted Interaction*, *5*, 45–65.

Bull, S., Pain, H., & Brna, P. (1995b). Mr. Collins: A collaboratively constructed, inspectable student model for intelligent computer assisted language learning. *Instructional Science*, *23*, 65–87.

Chan, T.-W. (1991). Integration-Kid: A Learning Companion System. In Mylopolous, J., & Reiter, R. (Eds.), *Proceedings of the Twelfth International Conference on Artificial Intelligence (IJCAI-91)*, Vol. 2, pp. 1094–1099. Morgan Kaufmann Publishers Inc.

Chan, T.-W. (1996). Learning Companion Systems, Social Learning Systems, and the Global Social Learning Club. *Journal of Artificial Intelligence in Education*, *7*(2), 125–159.

Chan, T.-W., & Baskin, A. B. (1990). Learning Companion Systems. In Frasson, C., & Gauthier, G. (Eds.), *Intelligent Tutoring Systems: At the Crossroads of Artificial Intelligence and Education*, chap. 1, pp. 6–33. Ablex Publishing Corporation, Norwood, New Jersey.

Chan, T.-W., & Chou, C.-Y. (1997). Exploring the Design of Computer Supports for Reciprocal Tutoring. *International Journal of Artificial Intelligence in Education*, *8*, 1–29.

Cumming, G., & Self, J. (1989). Collaborative Intelligent Educational Systems. In Bierman, D., Breuker, J., & Sandberg, J. (Eds.), *Artificial Intelligence and Education.*

Proceedings of the 4th International Conference on AI and Education, Vol. 3 of *Frontiers in Artificial Intelligence and Applications*, pp. 73–80 Amsterdam, Netherlands. IOS Press.

Dillenbourg, P., & Self, J. A. (1992). PEOPLE POWER: A Human-Computer Collaborative Learning System. In Frasson, C., Gauthier, G., & McCalla, G. I. (Eds.), *Second International Conference on Intelligent Tutoring Systems (ITS'92)*, Vol. 608 of *Lecture Notes in Computer Science*, pp. 651–660. Springer-Verlag.

Entwistle, N. J. (1978). Symposium: Knowledge Structures and Styles of Learning: a Summary of Pask's Recent Research. *British Journal of Educational Psychology*, *48*, 1–11.

Gilmore, D., & Self, J. (1988). The application of machine learning to intelligent tutoring systems. In Self, J. (Ed.), *Artificial Intelligence and Human Learning, Intelligent Computer-Aided Instruction*, chap. 11, pp. 179–196. Chapman and Hall Computing.

Goodlad, S., & Hirst, B. (1989). *Peer Tutoring: A Guide to Learning by Teaching.* Kogan Page, London.

Hietala, P., & Niemirepo, T. (1997). Collaboration with software agents: what if the learning companion agent makes errors?. In du Boulay, B., & Mizoguchi, R. (Eds.), *Artificial Intelligence in Education: Knowledge and Media in Learning Systems. Proceedings of AI-ED 97 World Conference on Artificial Intelligence in Education*, Vol. 39 of *Frontiers in Artificial Intelligence and Applications*, pp. 159–166. International AIED Society, IOS Press.

Hietala, P., & Niemirepo, T. (1998). The Competence of Learning Companion Agents. *International Journal of Artificial Intelligence in Education*, *9*, 178–192.

Hintze, D. (1991). Guided Discovery Learning: The Computer as a Cooperative Learning Assistant. In Birnbaum, L. (Ed.), *Proceedings of the 1991 International Conference on the Learning Sciences*, pp. 231–239. Association for the Advancement of Computing in Education (AACE).

Jun, Y. (1997). Linear Kid: A mathematical software designed as a computer-based peer tutoring system. In *Electronic Proceedings of the Second Asian Technology Conference in Mathematics (ATCM'97)*. http://www.cs.runet.edu/ atcm/EPATCM97/ATCMP008/paper.html.

Michie, D., Paterson, A., & Hayes-Michie, J. (1989). Learning by Teaching. In *2nd Scandinavian Conference on Artificial Intelligence 89*, pp. 307–331. IOS Press.

Nichols, D. (1994). Issues in Designing Learning by Teaching Systems. In Brusilovsky, P., Dikareva, S., Greer, J., & Petrushin, V. (Eds.), *Proceedings of the East-West International Conference on Computer Technologies in Education (EW-ED'94)*, Vol. 1, pp. 176–181. http://www.comp.lancs.ac.uk/computing/users/dmn/docs/crimea.html.

Okamoto, T., & Kasai, T. (1999). Students' Strategies in Collaborative Learning Environment with Multi Companion Agents. In Baggott, L., & Nichol, J. (Eds.), *Intelligent Computer and Communications Technology: Teaching & Learning for the 21st Century. PEG99, Ninth International PEG Conference*, pp. 45–53. School of Education, University of Exeter, England.

Pask, G. (1976a). Conversational Techniques in the Study and Practice of Education. *British Journal of Educational Psychology*, *46*, 12–25.

Pask, G. (1976b). Styles and Strategies of Learning. *British Journal of Educational Psychology*, *46*, 128–148.

Ramírez Uresti, J. A. (1998). Teaching a Learning Companion. In Ayala, G. (Ed.), *Proceedings of the International Workshop "Current Trends and Applications of Artificial Intelligence in Education". The Fourth World Congress on Expert Systems*, pp. 83–89 Monterrey, México. ITESM.

Ramírez Uresti, J. A. (2001). The LECOBA Learning Companion System: Expertise, Motivation and Teaching. *International Journal of Continuing Engineering Education and Life-Long Learning*, *11*(3), 216–228. Special Issue of the Best of PEG99 Conference.

Rickel, J., & Johnson, W. L. (1999). Virtual Humans for Team Training in Virtual Reality. In Lajoie, S. P., & Vivet, M. (Eds.), *Artificial Intelligence in Education. Open Learning Environments: New Computational Technologies to Support Learning, Exploration and Collaboration*, Vol. 50 of *Frontiers in Artificial Intelligence and Applications*, pp. 578–585. International AIED Society, IOS Press.

Scott, L. A., & Reif, F. (1999). Teaching Scientific Thinking Skills: Students and Computers Coaching Each Other. In Lajoie, S. P., & Vivet, M. (Eds.), *Artificial Intelligence in Education. Open Learning Environments: New Computational Technologies to Support Learning, Exploration and Collaboration*, Vol. 50 of *Frontiers in Artificial Intelligence and Applications*, pp. 285–293. International AIED Society, IOS Press.

Appendix

This appendix presents the pre-test and the post-test used during the experimental sessions with LECOBA. The time to solve each of the tests was 15 minutes. With each of them a separate sheet containing all the Laws and Theorems of Binary Boolean Algebra was provided to the subjects.

Pre-test

Learning Companion Systems Research

Binary Boolean Algebra Experiment

First of all thank you for participating in this experiment. Please feel free to leave at any time if you do not wish to continue. However, your cooperation will be much appreciated.

This stage of the experiment should take **15 minutes** in total.

1. Aims.

The aim of this stage is to measure knowledge in Binary Boolean Algebra. In particular, the number of rules mastered and the strategies used to simplify an expression.

2. Task Description.

This stage is divided into two parts and consists of a series of questions. For most of the questions you will need to consult the separate sheet containing the Laws and Theorems of Binary Boolean Algebra. The rules on this sheet have a **name** you will need. For example, in rule:

OR0X: `X + 0 = X`

the name of the rule is OR0X.

Part 1 is about strategies. Part 2 contains four simplification problems. **Please start on Part1 and then follow to Part 2. If you do not know the answer to a question, you may continue with the next one or go back to an earlier one. If you have any questions about the task, please let me know**.

The next page contains the first part of this stage. **Please let me know when you are ready to start and I will inform you when your time is over**.

General Information

1. Have you studied Binary Boolean Algebra Before? YES / NO

   If YES:

   (a) When did you study it? (Choose one only)

      i. yesterday
      ii. last week
      iii. last month
      iv. last term
      v. last semester
      vi. last year
      vii. More than a year ago (please state the number of years).

   (b) To what extent? (Choose one only)

      i. a day
      ii. a week
      iii. a month
      iv. a term
      v. a semester
      vi. More than a semester.

Part 1. Strategies

1. Supposing you wanted to simplify the following Boolean expression:

   ```
   1 + a * 1 + 0 + a * 0
   ```

   (a) Which Boolean rule would you apply to it as the FIRST STEP in the simplification process? (**State clearly the name of the rule**).

   (b) Why?

2. Supposing you wanted to simplify the following Boolean expression:

   ```
   a'' * 0 + 0 + a + a' + 1
   ```

   (a) Which Boolean rule would you apply to it as the FIRST STEP in the simplification process? (**State clearly the name of the rule**).

   (b) Why?

3. In the following list [table] of Laws and Theorems of Binary Boolean Algebra, enumerate the rules from the one you think is the most useful to simplify an expression (1) to the one you think is the least useful (29) — **Try to enumerate at least five rules**. For example, if you think that rule COM& is the most useful to simplify an expression, put a 1 in its corresponding space.

Laws and Theorems of Binary Boolean Algebra		
Number	Name	Rule
	OR0X	X + 0 = X
	AN1X	X * 1 = X
	OR11	X + 1 = 1
	AN00	X * 0 = 0
	EQLv	X + X = X
	EQL&	X * X = X
	2NEG	(X')' = X
	CMPv	X + X' = 1
	CMP&	X * X' = 0
	COMv	X + Y = Y + X
	COM&	X * Y = Y * X
	ASOv	(X + Y) + Z = X + Y + Z
	ASO&	(X * Y) * Z = X * Y * Z
	DISv	X + Y * Z = (X + Y) * (X + Z)
	DIS&	X * (Y + Z) = X * Y + X * Z
	ST1v	X * Y + X * Y' = X
	ST1&	(X + Y) * (X + Y') = X
	ST2v	X + X * Y = X
	ST2&	X * (X + Y) = X
	ST3v	X + X' * Y = X + Y
	ST3&	X * (X' + Y) = X * Y
	DeMv	(X + Y)' = X' * Y'
	DeM&	(X * Y)' = X' + Y'
	CONv	X * Y + X' * Z + Y * Z = X * Y + X' * Z
	CON&	(X + Y) * (X' + Z) * (Y + Z) = (X + Y) * (X' + Z)
	S->P	X * Y + X' * Z = (X + Z) * (X' + Y)
	P->S	(X + Y) * (X' + Z) = X * Z + X' * Y
	SIDv	X * Y + X * Z = X * (Y + Z)
	SID&	(X + Y) * (X + Z) = X + Y * Z

Part 2. Simplification Problems

1. For the following problems simplify the Boolean expression using the rules provided. At each step **state clearly the name of each rule** being used.

   (a) `a * a`

   (b) `c + 1`

   (c) `b * 0`

   (d) `a + a'`

   (e) `(a + b')'`

   (f) `c + 0 * 1`

   (g) `a + b + a' + b'`

   (h) `(a + b) * b' + b' + b * c`

   (i) `(a * b' + (a' + b) * c') + c`

Post-test

Learning Companion Systems Research

Binary Boolean Algebra Experiment

First of all thank you for participating in this experiment. Please feel free to leave at any time if you do not wish to continue. However, your cooperation will be much appreciated.

This stage of the experiment should take **15 minutes** in total.

1. Aims.

The aim of this stage is to measure knowledge in Binary Boolean Algebra. In particular, the number of rules mastered and the strategies used to simplify an expression.

2. Task Description.

This stage is divided into two parts and consists of a series of questions. For most of the questions you will need to consult the separate sheet containing the Laws and Theorems of Binary Boolean Algebra. The rules on this sheet have a **name** you will need. For example, in rule:

OR0X: `X + 0 = X`

the name of the rule is OR0X.

Part 1 is about strategies. Part 2 contains four simplification problems. **Please start on Part1 and then follow to Part 2. If you do not know the answer to a question, you may continue with the next one or go back to an earlier one. If you have any questions about the task, please let me know**.

The next page contains the first part of this stage. **Please let me know when you are ready to start and I will inform you when your time is over**.

Part 1. Strategies

1. Supposing you wanted to simplify the following Boolean expression:

   `1 + d * 1 + 0 + d * 0`

   (a) Which Boolean rule would you apply to it as the FIRST STEP in the simplification process? (**State clearly the name of the rule**).

   (b) Why?

2. Supposing you wanted to simplify the following Boolean expression:

   `d'' * 0 + 0 + d + d' + 1`

   (a) Which Boolean rule would you apply to it as the FIRST STEP in the simplification process? (**State clearly the name of the rule**).

   (b) Why?

3. In the following list [table] of Laws and Theorems of Binary Boolean Algebra, enumerate the rules from the one you think is the most useful to simplify an expression (1) to the one you think is the least useful (29) — **Try to enumerate at least five rules**. For example, if you think that rule COM& is the most useful to simplify an expression, put a 1 in its corresponding space.

Laws and Theorems of Binary Boolean Algebra		
Number	Name	Rule
	OR0X	X + 0 = X
	AN1X	X * 1 = X
	OR11	X + 1 = 1
	AN00	X * 0 = 0
	EQLv	X + X = X
	EQL&	X * X = X
	2NEG	(X')' = X
	CMPv	X + X' = 1
	CMP&	X * X' = 0
	COMv	X + Y = Y + X
	COM&	X * Y = Y * X
	ASOv	(X + Y) + Z = X + Y + Z
	ASO&	(X * Y) * Z = X * Y * Z
	DISv	X + Y * Z = (X + Y) * (X + Z)
	DIS&	X * (Y + Z) = X * Y + X * Z
	ST1v	X * Y + X * Y' = X
	ST1&	(X + Y) * (X + Y') = X
	ST2v	X + X * Y = X
	ST2&	X * (X + Y) = X
	ST3v	X + X' * Y = X + Y
	ST3&	X * (X' + Y) = X * Y
	DeMv	(X + Y)' = X' * Y'
	DeM&	(X * Y)' = X' + Y'
	CONv	X * Y + X' * Z + Y * Z = X * Y + X' * Z
	CON&	(X + Y) * (X' + Z) * (Y + Z) = (X + Y) * (X' + Z)
	S->P	X * Y + X' * Z = (X + Z) * (X' + Y)
	P->S	(X + Y) * (X' + Z) = X * Z + X' * Y
	SIDv	X * Y + X * Z = X * (Y + Z)
	SID&	(X + Y) * (X + Z) = X + Y * Z

Part 2. Simplification Problems

1. For the following problems simplify the Boolean expression using the rules provided. At each step **state clearly the name of each rule** being used.

   (a) `b * b`

   (b) `g + 1`

   (c) `a * 0`

   (d) `f + f'`

   (e) `(d + c')'`

   (f) `g + 0 * 1`

   (g) `b + e + b' + e'`

   (h) `(f + c) * c' + c' + c * g`

   (i) `(d * e' + (d' + e) * g') + g`