



# Learning from Explanations: Extending One's Own Knowledge During Collaborative Problem Solving by Attempting to Understand Explanations Received from Others

Rolf Ploetzner, Eric Fehse

## ► To cite this version:

Rolf Ploetzner, Eric Fehse. Learning from Explanations: Extending One's Own Knowledge During Collaborative Problem Solving by Attempting to Understand Explanations Received from Others. International Journal of Artificial Intelligence in Education (IJAIED), 1998, 9, pp.193-218. hal-00257112

HAL Id: hal-00257112

<https://telearn.archives-ouvertes.fr/hal-00257112>

Submitted on 18 Feb 2008

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

## **Learning from Explanations: Extending One's Own Knowledge during Collaborative Problem Solving by Attempting to Understand Explanations Received from Others**

**Rolf Ploetzner, Eric Fehse**, *Department of Psychology University of Freiburg 79085 Freiburg, Germany*

**Abstract.** On the basis of an experimental study, we propose a cognitive simulation model of collaborative problem solving and learning. In the experimental study, we investigated how qualitative and quantitative problem representations in classical mechanics are acquired and successively interrelated during collaborative problem solving. Two students, who were taught different aspects of classical mechanics, collaborated on problems which were beyond the competence of each of them individually. Students successfully learned to interrelate qualitative and quantitative problem representations. Furthermore, students who initially were taught qualitative aspects of classical mechanics gained significantly more from the information provided by their quantitatively instructed partners than the other way round. The model simulates collaborative problem solving and learning under the conditions set up in the experimental study. On the basis of the model it was possible to reconstruct the main results of the experimental study.

### **INTRODUCTION**

Two important questions in research on collaborative problem solving and learning are the questions of (a) which arrangements make collaboration both instructive and efficient and (b) which interaction processes underlie collaboration (for a recent overview see Dillenbourg, Baker, Blaye & O'Malley, 1996). Two factors which correspond to these questions and which seem to positively influence collaborative problem solving and learning are (a) task structures that promote mutual interdependence among the collaborating students and (b) communication structures that promote the construction of high-level questions and explanations.

With respect to the first factor, it has been demonstrated that students' understanding advances the most during collaboration if the students initially possess different pre-knowledge about or conceptual perspectives on the application domain under scrutiny (e.g., Howe, Tolmie & Rodgers, 1990, 1992; for an overview see Knight & Bohlmeier, 1990). Therefore, students are sometimes taught different but complementary aspects of the application domain before collaborative problem solving takes place (e.g., Aronson, Balney, Stephan, Sikes & Snapp, 1978; Slavin, 1995). Subsequently, problems are posed to the students which are designed in such a way that the students mutually depend upon each other's competencies in order to be successful. Because this approach demands that each student contributes to the solution of a posed problem, it gives the students various opportunities to learn from each other during the collaboration.

With respect to the second factor, Webb (1989) revealed by means of a meta-analysis that the success of collaborative problem solving and learning depends largely on the level of elaboration of the information exchanged between the collaborating students. It may range from the mere exchange of achieved results, over the exchange of single specific or general pieces of information, to the exchange of detailed and coherent specific or general explanations, for example.

While learning on the basis of the mere exchange of achieved results seems not to be very far reaching, the exchange of detailed and coherent explanations might considerably improve learning under the conditions that (a) the explanations are relevant to the explainees questions, (b) the explanations match the level of help needed, (c) the explanations are provided without major delay, (d) the explanations match the explainees level of understanding and (e) the explainee has the opportunity to apply the information provided by the explanations to the posed problem.

In accord with these findings, Graesser and Person (1994) observed that successfully learning students predominantly raise questions during collaborative problem solving which ask for detailed explanations. Less successfully learning students, in contrast, ask mostly questions which result in rather uninformative answers.

In this paper, we propose a cognitive simulation model of collaborative problem solving and learning that takes both factors described above into account. The model has been developed on the basis of an experimental study (cf. Ploetzner, Fehse, Spada, Vodermaier & Wolber, 1996). In the experimental study, we investigated how qualitative and quantitative problem representations in classical mechanics are acquired and successively interrelated during collaborative problem solving. The study was made up of two main phases. In the first phase, students were taught either qualitative or quantitative aspects of classical mechanics by means of two different instructional units.

By teaching different aspects of classical mechanics to the students, we intentionally gave rise to a systematic variation in the students' pre-knowledge about the application domain. In the second phase, dyads were formed with students who had received different instructional units and therefore possessed systematically different pre-knowledge. The dyads collaboratively worked on problems which were beyond the competence of each student individually.

Before and after the instruction as well as after the collaborative problem solving, students had to work individually on multi-component tests. An analysis of variance of the multi-component tests revealed that the students successfully learned to interrelate qualitative and quantitative problem representations during collaborative problem solving. Furthermore, students who initially were taught qualitative aspects of classical mechanics gained significantly more from the information provided by their quantitatively instructed partners than the other way round.

The model simulates selected aspects of collaborative problem solving and learning under the conditions set up in the experimental study. It rests on three basic assumptions. The first assumption is that collaboration is especially beneficial to the collaborating partners if they initially possess different pre-knowledge about the application domain. The second assumption is that the ability to direct felicitous questions to one's partner frequently depends on adequate self-diagnoses of shortcomings of one's own competence. The third assumption is that explanations received from one's partner provide a valuable source of information for extending one's own knowledge.

The model comprises a cognitive as well as a meta-cognitive level. The cognitive level was realised by taking advantage of a cognitive simulation model of individual problem solving in classical mechanics (cf. Ploetzner, 1995). In correspondance to the first assumption described above, the simulation model of individual problem solving was duplicated. While the knowledge available to one copy was restricted to knowledge about qualitative aspects of classical mechanics, the knowledge available to the other copy was restricted to knowledge about quantitative aspects.

The meta-cognitive level was implemented by making use of meta-programming techniques. It comprises two domain-independent mechanisms. The first mechanism corresponds to the second assumption described above. It simulates the construction of questions directed to one's partner on the basis of deductive self-diagnoses. The second mechanism corresponds to the third assumption described above. It simulates learning by attempting to understand explanations received from one's partner on the basis of one's own pre-knowledge. By means of the model it was possible to reconstruct the main results of the experimental study.

The paper is organized as follows. In the following section, the distinction between knowledge about qualitative and quantitative aspects of classical mechanics is set forth. A summary of the cognitive simulation model of individual problem solving is provided next. Afterwards, the design and the results of the experimental study are described. Thereafter, the cognitive simulation model of collaborative problem solving and learning as well as the results of a simulation study are delineated. A discussion and conclusion complete the paper.

## **KNOWLEDGE ABOUT QUALITATIVE AND QUANTITATIVE ASPECTS OF CLASSICAL MECHANICS**

The considered application domain is made up of standard textbook physics problems which involve one-dimensional motion with constant acceleration (e.g., Halliday & Resnick, 1988). The focus is on problems which ask for a precise quantitative problem solution. For example:

A block of mass  $m = 15$  kg starts from rest (initial velocity  $v = 0$  m/s) down a frictionless (coefficient of friction  $f = 0$ ) plane inclined at an angle  $\alpha = 30^\circ$  with the horizontal. What is the block's velocity  $v$  after the time  $t = 2$  s?

The knowledge under scrutiny is knowledge about qualitative and quantitative aspects of various concepts in dynamics (e.g., gravitational force, normal force and friction force) and kinematics (e.g., time, position, displacement, velocity and acceleration). Knowledge about qualitative aspects encodes information such as the conditions under which concepts can legitimately be applied, the attributes possessed by concepts and the values concept attributes might have. It might specify, for example, that there is a kinetic friction force  $F_f$  on an object, whenever there is a normal force  $F_n$  on the object and the object is moving on a surface which is not frictionless.

Knowledge about quantitative aspects encodes information about interaction and motion laws. It relies on mathematical formalisms to describe definitions of and functional relationships between concepts by means of algebraic and vector-algebraic equations. It might specify, for instance, that the magnitude of the kinetic friction force  $F_f$  on an object equals the magnitude of the normal force  $F_n$  on the same object times the coefficient of friction  $f$ :  $F_f = F_n * f$ . As de Kleer (1975) points out, knowledge about qualitative and quantitative aspects cannot be separated at a clear-cut boundary. Rather, this distinction refers to the ends of a continuum with a considerable body of knowledge between them (e.g., Ploetzner, Spada, Stumpf & Opwis, 1990; van Joolingen, 1994; White & Frederiksen, 1990).

We conceptualize qualitative and quantitative problem representations as complementary representations based on knowledge about qualitative and quantitative aspects, respectively. Qualitative problem representations include information about essential problem features to be taken into account and important distinctions to be drawn. While quantitative problem representations help to resolve ambiguities frequently inherently involved in qualitative problem representations, qualitative problem representations seem to be a necessary prerequisite for the appropriate construction and use of quantitative problem representations.

## **A COGNITIVE SIMULATION MODEL OF INDIVIDUAL PROBLEM SOLVING**

The development of the cognitive simulation model of individual problem solving started from the assumption that successful and efficient problem solving in formal sciences such as physics demands the construction and coordination of qualitative and quantitative problem representations. The simulation model and its implementation is described in detail in Ploetzner (1995). It has been realised with two main goals in mind. The first goal was to investigate the knowledge structures which underlie qualitative and quantitative problem representations in classical mechanics. The second goal was to examine the reasoning mechanisms which allow one to construct and coordinate both kinds of problem representations.

The model has been implemented in Prolog as a knowledge-based system. It comprises four major components: (1) an interpreter, (2) a knowledge base for knowledge about qualitative and quantitative aspects of classical mechanics, (3) a knowledge base for qualitative and quantitative vector-algebraic knowledge and (4) a knowledge base for geometric and algebraic knowledge.

The interpreter employs domain-independent as well as domain-specific control knowledge to enable the construction of qualitative and quantitative problem representations during problem solving. The domain-specific control knowledge includes, for instance, a procedure made up of six steps to construct so-called free-body diagrams: (1) identify the object whose motion has to be analysed, (2) determine all the forces on the object, (3) draw an arrow for each force on the object, (4) choose a suitable coordinate system as a reference frame, (5) resolve the forces for their components along each coordinate axis and (6) apply Newton's second law to the resultant force along each axis. Though this procedure aims at the construction of more and more complete qualitative problem representations, it leaves nevertheless unspecified which knowledge actually needs to be applied to achieve each step.

In the model, knowledge about qualitative and quantitative aspects of classical mechanics concepts has been implemented by means of an equation-based, relational representation language similar to the representation language employed by VanLehn, Jones and Chi (1992). Each expression which encodes knowledge about qualitative or quantitative aspects of classical mechanics concepts is made up of a left- (i.e., the condition) and a right-hand side (i.e., the conclusion). An expression's left-hand side is a conjunction of a possibly empty set of atomic sentences  $\{S_1 \dots S_k\}$ . An expression's right-hand side consist of exactly one atomic sentence  $S$  which is always an equation. The left-hand side and the right-hand side of an expression are connected by the implication operator:  $S_1 \wedge \dots \wedge S_k \Rightarrow S$ .

In the case of knowledge about qualitative aspects, arithmetic operators must not occur in the expression's left- or right-hand side. Knowledge about qualitative aspects of classical mechanics is formalized by means of two kinds of expressions. The first kind of expression states the conditions under which a concept can legitimately be applied. The second kind of expression constrains the values of the various attributes a concept has. Knowledge about quantitative aspects of classical mechanics is formalized by means of expressions which describe algebraic equations. Time-dependency of concepts is represented by means of an extensional temporal operator "value(S, P)" which denotes the value of a parameter "P" in situation "S". The situation might refer to a point in time or an interval of time (cf. Davis, 1990).

**Table 1.** How knowledge about qualitative and quantitative aspects of the kinetic friction force has been formalized

---

<b>Knowledge about qualitative aspects:</b>	
F <sub>f1</sub> :	Value(S, instance(Object1, body)) = true $\wedge$ value(S, instance(force(Object1, Object2, fn), normal_force)) = true $\wedge$ Value(S, instance(Object2, plane)) = true $\wedge$ Value(S, moves_on(Object1, Object2)) = true $\wedge$ $\neg$ (value(S, frictionless(Object2)) = true) $\Rightarrow$ value(S, instance(force(Object1, Object2, ff), friction_force)) = true
F <sub>f2</sub> :	Value(S, instance(force(Object1, Object2, ff), friction_force)) = true $\wedge$ Value(S, instance(Object2, plane)) = true $\Rightarrow$ Value(S, inclination(force(Object1, Object2, ff))) = Value(S, inclination(velocity(Object1)))
F <sub>f3</sub> :	Value(S, instance(force(Object1, Object2, ff), friction_force)) = true $\wedge$ Value(S, instance(Object2, plane)) = true $\Rightarrow$ Value(S, sense(force(Object1, Object2, ff)) = Opposite(value(S, sense(velocity(Object1))))
<b>Knowledge about quantitative aspects:</b>	
Eq:	Value(S, magnitude(force(Object1, Object2, ff))) = value(S, magnitude(force(Object1, Object2, fn))) * Value(S, friction(Object1, Object2))

---

Table 1 shows how knowledge about qualitative and quantitative aspects of the kinetic friction force is formalized in the model. Whenever there is an object which is an instance of a body and there is a normal force on the object due to a plane and the object is moving on the plane which is not frictionless, then there is a kinetic friction force on the object due to the plane (cf. Expression  $F_{f1}$  in Table 1). If there is a kinetic friction force on an object due to a plane, then the inclination of the kinetic friction force equals the inclination of the object's velocity (cf. Expression  $F_{f2}$  in Table 1) and the sense of the kinetic friction force is opposite to the sense of the object's velocity (cf. Expression  $F_{f3}$  in Table 1). The algebraic equation states that the magnitude of the kinetic friction force on an object equals the magnitude of the normal force on the same object times the coefficient of friction (cf. Expression Eq in Table 1).

The model is applied to problem descriptions encoded in the same representation language which has been employed to formalize qualitative and quantitative aspects of classical mechanics. How the problem described above is encoded is shown in Table 2.

**Table 2.** How problem descriptions are encoded

---

<b>Problem description:</b>
A block of mass $m = 15$ kg starts from rest (initial velocity $v = 0$ m/s) down a frictionless (coefficient of friction $f = 0$ ) plane inclined at an angle $\alpha = 30^\circ$ with the horizontal. What is the block's velocity $v$ after the time $t = 2$ s?
<b>Encoded problem description:</b>
given(value(1,instance(block_1, block)) = true)
given(value(1,instance(plane_1, plane)) = true)
given(value(1,moves_on(block_1, plane_1)) = true)
given(value(1,mass(block_1)) = 15::kg)
given(value(1,magnitude(velocity(block_1))) = 0::m/s)
given(value(1,sense(velocity(block_1))) = down)
given(value(1,inclination(plane_1)) = 30)
given(value(1,frictionless(plane_1)) = true)
given(value(1 $\rightarrow$ 2, duration(displacement(block_1))) = 2::s)
sought(value(2, magnitude(velocity(block_1))))

---

Table 3 summarises - in a more familiar notation - how the problem is solved by the model. Initially, qualitative reasoning ascertains that there are two different forces on the block. Subsequently, a free-body diagram is constructed. The model constructs such a free-body diagram only in terms of symbolic descriptions. Graphical elements are not involved.

The free-body diagram enables the resultant force on the block to be specified in qualitative terms. Afterwards, this specification is expressed in algebraic terms by applying vector-algebraic knowledge and knowledge about Newton's second law. The resulting algebraic expressions extend the quantitative information available to the model. Additional qualitative reasoning further restricts the knowledge about quantitative aspects to be applied. Finally, the problem's solution is derived by successively applying the relevant dynamics and kinematics laws.

The entire knowledge formalized in the model makes up a general domain model. If the model is applied to a specific problem description, a qualitative and a quantitative problem representation are constructed by copying and instantiating parts of the general domain model. Both problem representations are successively constructed, coordinated and modified until they yield the problem's solution. In the model, the main emphasis is on how the information included in a qualitative problem representation can be utilized to enable and to guide the subsequent construction of an appropriate quantitative problem representation.

During problem solving, qualitative and quantitative problem representations can be coordinated in two different ways. Firstly, by drawing on vector-algebraic knowledge, the information included in a qualitative problem representation is partially expressed in algebraic terms to construct additionally required quantitative information not available to the model beforehand. Secondly, the information included in a qualitative problem representation is exploited to constrain the use of already available quantitative information. A similar mechanism for coordinating qualitative, conceptual information and quantitative, numerical

information has been employed earlier by Ohlsson and Rees (1991) in the domain of arithmetic problem solving, for example.

**Table 3.** How qualitative and quantitative problem representations are coordinated

Qualitative Reasoning	Quantitative Reasoning
the object whose motion has to be analysed is the block the forces on the block are $F_g$ and $F_n$	
direction of $F_{gx}$ equals direction of $v$ Magnitude of resultant force $\Sigma F_x$ equals magnitude of $F_{gx}$ direction of $F_{gy}$ is opposite to direction of $F_n$ magnitude of resultant force $\Sigma F_y$ equals difference between magnitude of $F_n$ and magnitude of $F_{gy}$ motion with constant acceleration and without initial velocity magnitude of $a_y$ equals zero	$\Sigma F_x = F_{gx} = m * a$ $\Sigma F_y = F_n - F_{gy} = m * a_y$ $v = a * \Delta t$ $a_y = 0 \text{ m/s}^2$ $a_x = a = F_{gx} / m$ $= (m * g * \sin \alpha) / m$ $= g * \sin \alpha$ $v = g * \sin \alpha * \Delta t$

## EXPERIMENT

### Design

The design of the experimental study is shown in Table 4. Two experimental groups and one control group were formed. Twentyfour students from two different schools were equally distributed among the three groups. Each group was made up of eight students. All students had to work on a pretest, an intermediate test and a posttest.

**Table 4.** The design of the experimental study

Group	Pretest	Intro- duction	Instruction	Intermediate test	Collaborative problem solving	Posttest
<b>Qualitative group</b>			qualitative			
<b>Quantitative group</b>			quantitative			
<b>Control group</b>						

Between the pretest and the intermediate test, the students in the experimental groups worked on instructional units. While the students in the first experimental group worked on an instructional unit on qualitative aspects of classical mechanics, the students in the second experimental group worked on an instructional unit on quantitative aspects of classical mechanics.

Between the intermediate test and the posttest, dyads were formed with students who had received different instructional units. They worked collaboratively on five classical mechanics problems which were beyond the competence of each of them individually and which demanded the coordinated use of knowledge about qualitative and quantitative aspects of classical mechanics.

By teaching different aspects of classical mechanics to the students, we intentionally gave rise to a systematic variation in the students' pre-knowledge about the application domain. In various earlier studies, it has been demonstrated that students' understanding advances the most during collaboration if they initially possess different pre-knowledge about the application domain under scrutiny (e.g. Howe, Tolmie & Rodgers, 1990, 1992; Knight & Bohlmeyer, 1990). Because this approach demands that each student contributes to the solution of a posed problem, it gives the students various opportunities to learn from each other during the collaboration.

## Materials

### *Instructional Units*

By taking advantage of the knowledge encoded in the simulation model of individual problem solving, two instructional units on classical mechanics were constructed (cf. Ploetzner, Fehse, Spada, Vodermaier & Wolber, 1996). One unit described qualitative aspects of classical mechanics and one unit described quantitative aspects of classical mechanics. Both units comprised sections on (a) coordinate systems and vectors, (b) resolution and addition of vectors, (c) velocity and acceleration and (d) forces. The qualitative and quantitative aspects of the addressed concepts were described by means of concept maps (e.g., Novak, 1990). One or more concept maps were followed by several examples and exercises. The solutions to the exercises were also presented.

### *Problems*

On the basis of the simulation model of individual problem solving described, five different classical mechanics problems for collaborative problem solving were set up. They are shown in Table 5. In order to necessitate collaboration during problem solving, the problems were designed in such a way that their solutions require the coordinated application of knowledge about both qualitative and quantitative aspects of classical mechanics.

**Table 5.** Problems for collaborative problem solving

---

<b>Problem 1:</b> A coin of mass $m = 0.03$ kg is tossed straight up into the air with the velocity $v = 7$ m/s. After which distance $r$ is the coin's velocity reduced to $v = 3$ m/s?
<b>Problem 2:</b> A block of mass $m = 15$ kg starts from rest (initial velocity $v = 0$ m/s) down a frictionless (coefficient of friction $f = 0$ ) plane inclined at an angle $\alpha = 30^\circ$ with the horizontal. What is the block's velocity $v$ after the time $t = 2$ s?
<b>Problem 3:</b> What is the minimum stopping distance for a car of mass $m = 820$ kg travelling along a flat horizontal road with the velocity $v = 12$ m/s, if the coefficient of friction $f$ between tires and road equals 0.8?
<b>Problem 4:</b> A block of mass $m = 10$ kg is projected up an inclined plane with the velocity $v = 5$ m/s. The plane is inclined at an angle $\alpha = 15^\circ$ . Which distance $r$ up the plane does the block go, if the coefficient of friction $f$ between block and plane equals 0.3?
<b>Problem 5:</b> A block of mass $m = 72$ kg moves down an inclined plane with constant velocity $v$ . What is the coefficient of friction $f$ between block and plane, if the plane is inclined at an angle $\alpha = 20^\circ$ ?

---



### *Multi-Component Tests*

Before and after the instruction as well as after the collaborative problem solving, students had to work on parallel tests. The tests were made up of two components which assess knowledge about qualitative and quantitative aspects of classical mechanics. Six items of the first component assessed knowledge about qualitative aspects and six items assessed knowledge about quantitative aspects. On the basis of the simulation model of individual problem solving, the items were designed in such a way that the solution to each item required the use of information presented in the instructional unit on qualitative or quantitative aspects, respectively.

The six items of the second component were designed in such a way that the solution to each item required the interrelation of information presented in both instructional units. These items required the ability to coordinate the application of knowledge about qualitative and quantitative aspects of classical mechanics. Thus, while half of the items of the first test component should be solvable after the instruction took place, items of the second test component should only be solvable to a larger extent after the collaborative problem solving took place.

The pretest, the intermediate test and the posttest were made up of parallel items. In order to avoid the assessment of abilities which are specific to the use of concept maps, concept maps were not included in the tests. Parallel items were designed in such a way that the same knowledge needed to be applied to solve them. However, surface features such as the objects involved and the numerical values were varied across parallel items. Within each test, items were arranged in random order.

### **Students**

Twentyfour female tenth graders from two different high schools volunteered for the study. The students were between 16 and 17 years old. They were paid for their participation. All of the students had attended classes on basic aspects of mechanics. In these classes, the concepts of time, position, displacement, velocity, acceleration and force had been introduced. The interrelations between these concepts had been described by means of position-time-, velocity-time- and acceleration-time-diagrams, for example. However, none of the students had attended classes on more advanced, Newtonian aspects of mechanics as they were in the foreground of this study.

### **Procedure**

The students in the two experimental groups were investigated in pairs. Each pair was composed of students from different schools and was investigated for four days running. On the first day, the students worked on the pretest, an introduction into the structure of concept maps and the first sections of the instructional units. While one student received instruction on qualitative aspects of classical mechanics, the other student received instruction on quantitative aspects. While they worked on the pretest and the instructional units, students had to work on their own and were not allowed to exchange information. However, they were allowed to work collaboratively on the introduction into the structure of concept maps.

On the second day, the students worked on the remaining sections of the instructional units and on the intermediate test. Again, students had to work on their own and were not allowed to exchange information. On the third day, the students had to solve three problems (Problem 1, Problem 2 and Problem 3 as shown in Table 5) which demand the coordinated use of knowledge about qualitative and quantitative aspects.

Problem solving took place in two phases. In the first phase, the students had to work individually on a problem. In the second phase, starting from their individual problem solving attempts, the students were allowed to work collaboratively on the same problem. During collaborative problem solving, students were allowed to exchange information at their will.

The individual problem solving phase served to encourage the students to seriously approach a problem, to identify their own knowledge gaps and to come up with reasonable questions to be raised during the subsequent collaborative problem solving phase. After the students finished their collaborative problem solving attempts, they were given feedback about the correctness of their solution. In the case that the solution was incorrect, they were told where an error was made. The solution to the problem, however, was not told to the students. At the end of the collaborative problem solving, students were allowed to reread selected sections in the unit which they had received during instruction.


On the fourth day, the students collaborated on the remaining two problems (Problem 4 and Problem 5 as shown in Table 5) under the same conditions as described above. Finally, each student worked individually on the posttest.

The students in the control group were investigated in groups. They worked on three different days on the pretest, the intermediate test and the posttest. Between two days of testing there was one day without testing. During the study, all students were allowed to use a ruler and a calculator. The students were not given feedback about their answers to the different tests.

## Results

The average relative solution frequencies of all groups in the first test component of the pretest, intermediate test and posttest are shown in Figure 1. A two-way analysis of variance with repeated measurements Group  $\times$  Test was computed. Across groups, the students improved significantly from the pretest to the posttest ( $F(2, 42) = 70.57, p < .001$ ). The students in the experimental groups gained significantly more from the pretest to the posttest than the students in the control group ( $F(4, 42) = 13.05, p < .001$ ). At this level of analysis, the students in the experimental groups did not differ significantly in their gains ( $F(2, 28) = .68, p = .51$ ).

In their tendency, however, the students who initially were taught qualitative aspects of classical mechanics gained more from the information provided by their quantitatively instructed partners during collaborative problem solving than the other way round. While the correctness of the qualitatively instructed students increased on average by 36% from the intermediate test to the posttest, the correctness of the quantitatively instructed students increased only by 25%.




**Figure 1.** The average relative solution frequencies of all groups in the first test component of the pretest, intermediate test and posttest

To make this difference between the students in the experimental groups even more visible, the items which assess knowledge about qualitative aspects of classical mechanics and the items which assess knowledge about quantitative aspects of the first test component were analysed separately. Again, a two-way analysis of variance with repeated measurements was computed. The corresponding average relative solution frequencies are shown in Figure 2.


The students did not differ significantly with respect to how well they were able to consolidate and to extend the knowledge about qualitative or quantitative aspects they had acquired during the instruction. However, as predicted, the students differed significantly with respect to how well they were able to take advantage of the information provided by their partners during collaborative problem solving ( $F(1, 14) = 5.3, p < .05$ ). The students who initially were taught qualitative aspects of classical mechanics gained significantly more from the information provided by their quantitatively instructed partners than the other way round.

The average relative solution frequencies of all groups in the second test component of the pretest, intermediate test and posttest are shown in Figure 3. As expected, the items of the second test component were only solvable to a larger extent after the collaborative problem solving took place.



**Figure 2.** The average relative solution frequencies of the experimental groups in the first test component of the intermediate test and posttest with respect to the test items which assess either knowledge about qualitative or knowledge about quantitative aspects

A further two-way analysis of variance with repeated measurements Group x Test was computed. Because in the pretest there was no variance at all observable, we did not compute such an analysis with respect to the pretest and the intermediate test. Students in both experimental groups were able to solve a few items of the second test component in the intermediate test. This might be due to pre-knowledge in connection with knowledge acquired during the instruction, for example. The students in the experimental groups improved significantly from the intermediate test to the posttest ( $F(1, 14) = 9.6, p < .01$ ).



**Figure 3.** The average relative solution frequencies of all groups in the second test component of the pretest, intermediate test and posttest

### Summary of the Empirical Results

The experimental study described above yielded two main results. The first result is that students with initially different but complementary representations of the domain under scrutiny learned to appropriately interrelate their representations during collaborative problem solving when the students were confronted with problems which required such an interrelation. Because the initially different domain representations demanded that each student contributes to the solution of a problem, students were given various opportunities to exchange information and to learn from each other's explanations.

The second result is that during collaborative problem solving, qualitatively instructed students gained more from their quantitatively instructed partners than the other way round. In accord with research on differences between novices and experts (for overviews see Chi, Glaser & Rees, 1982; VanLehn, 1996), this result indicates that qualitative problem representations form not only a good starting point for the subsequent construction of quantitative problem representations during problem solving but also a beneficial starting point for learning quantitative problem representations.

### A COGNITIVE SIMULATION MODEL OF COLLABORATIVE PROBLEM SOLVING AND LEARNING

To better understand which interaction and learning processes might account for the results of the experimental study, we developed a further cognitive simulation model. The model simulates selected aspects of collaborative problem solving and learning under the conditions set up in the experimental study. With respect to the interaction processes, the model rests on the assumption, that the success of collaborative problem solving and learning largely depends on the questions and explanations exchanged between the collaborating students (cf. Webb, 1989).

It is especially assumed that the ability to direct felicitous questions to one's partner frequently depends on adequate self-diagnoses of shortcomings of one's own competence. With

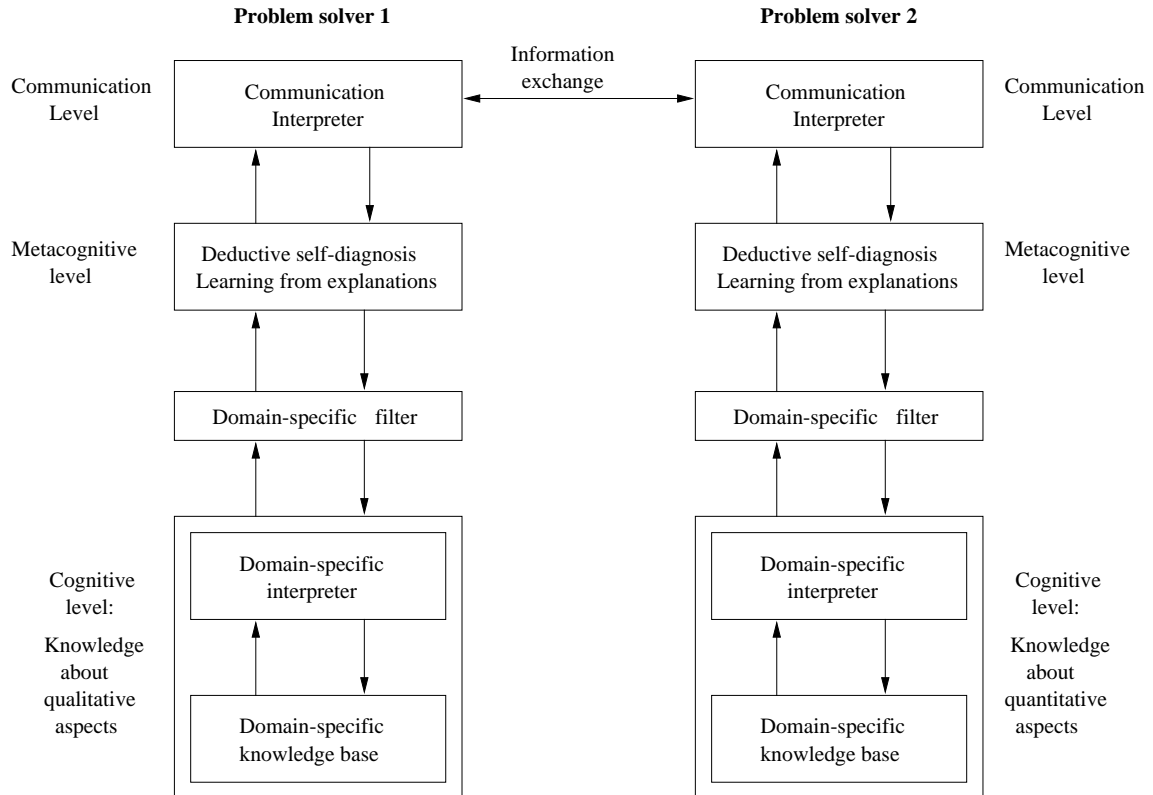
respect to the learning processes, the model rests on the assumption that explanations received from one's partner provide a valuable source of information for extending one's own knowledge. It is especially assumed that attempting to understand explanations on the basis of one's own pre-knowledge encourages learning.

**Architecture**

The model comprises two simulated problem solvers (cf. Figure 4). Each problem solver is made up of a cognitive level, a meta-cognitive level and a communication level. The cognitive level was realised by taking advantage of the cognitive simulation model of individual problem solving. Essentially, the simulation model of individual problem solving was duplicated. This lead to two problem solvers. They differ only with respect to the physics knowledge they embody on the cognitive level.

While the physics knowledge available to one problem solver was restricted to knowledge about qualitative aspects of classical mechanics, the physics knowledge available to the other problem solver was restricted to knowledge about quantitative aspects. Both problem solvers were equipped with the same knowledge about geometry, algebra and vector-algebra. In addition, both problem solvers make use of the same mechanisms to interpret the knowledge available to them.

The meta-cognitive level operates on the cognitive level. It comprises two domain-independent mechanisms. The first mechanism has been implemented by making use of meta-programming techniques. It simulates the construction of explanations as well as the construction of questions by means of deductive self-diagnoses (cf. Ploetzner, Fehse, Hermann & Kneser, 1997). The second mechanism simulates learning by attempting to understand explanations. Between the cognitive level and the meta-cognitive level a domain-specific filter was set up to specify which parts of the cognitive level are to be interpreted on the meta-cognitive level and which are not.



**Figure 4.** The components of the cognitive simulation model

The exchange of information between the two problem solvers takes place on the communication level. If one of the two problem solvers is not able to solve a posed problem, a

question is constructed on the meta-cognitive level. Subsequently, the question is forwarded to the other problem solver. The other problem solver attempts to answer the question. If the other problem solver is able to construct an answer to the question, not only the answer but also an explanation is sent back to the problem solver which raised the question.

### Implementation

The cognitive, meta-cognitive and communication level of the two simulated problem solvers were implemented in Prolog, a logic-oriented programming language (cf. Shoham, 1994; Sterling & Shapiro, 1994). The implementation of the cognitive level is described in detail in Ploetzner (1995). In this section, the implementation of the meta-cognitive and communication level is delineated.

### The Construction of Explanations and the Construction of Questions by Means of Deductive Self-Diagnoses

If a student possesses only incomplete knowledge with respect to a posed problem, he or she commonly reaches impasses during an attempt to solve the problem (cf. VanLehn, 1988a). We conceptualise those cognitive processes as self-diagnoses which lead to (a) the identification where one reaches an impasse during the attempt to solve a posed problem and (b) the determination of which information would allow one to overcome the impasse and possibly to resume problem solving. Obviously, such processes provide an informative basis for posing questions to one's partner during collaborative problem solving.

In Prolog, problem solving corresponds to the attempt to construct a formal proof by means of the knowledge available on the cognitive level. Such a proof delineates how a problem solving goal can be derived from the knowledge available. A proof tree forms an externalisation of a proof. It can serve as an explanation of how a problem solving goal can be derived from the knowledge available.

The first purpose of complementing the cognitive level with a meta-cognitive level was to provide a domain-independent means that enables the construction of explanations on the basis of proof trees. The second purpose was to provide a domain-independent means that enables the construction of questions on the basis of deductive self-diagnoses (cf. Ploetzner, Fehse, Hermann & Kneser, 1997).

**Table 6.** A standard meta-interpreter for "pure" Prolog programs

---

```

interpreter((Goal, Goals)) :-
    !,
    interpreter(Goal),
    interpreter(Goals).

interpreter(Goal) :-
    system_predicate(Goal),
    !,
    call(Goal).

interpreter(Goal) :-
    clause(Goal, Goals),
    interpreter(Goals).

```

---

The meta-cognitive level was implemented by taking advantage of meta-programming techniques which have been developed in the framework of student modelling (Beller & Hoppe, 1993; Hoppe, 1994). Its implementation started from a standard meta-interpreter for so-called "pure" Prolog programs (cf. Table 6). Pure Prolog comprises only those aspects of the Prolog

system which are in accord with formal logic. Within such an interpreter, it has to be ensured that built-in system-predicates are not processed by the meta-interpreter but by the Prolog-interpreter (cf. the second rule in Table 6).

By labelling predicates on the cognitive level as system-predicates, it is possible to specify which parts of the cognitive level are to be processed on the meta-cognitive level and which are not. For instance, on the cognitive level, the two problem solvers comprise various predicates for solving algebraic equations. Because the emphasis of the model is on physics problem solving but not on algebraic problem solving, these predicates were labelled as system-predicates.

By extending a standard meta-interpreter with two more arguments, it is possible to construct and externalise generalised proof trees (cf. Table 7). While the first argument is always instantiated with problem specific goals, the second argument is always instantiated with generalised goals. Because generalised goals are copied (for an implementation of the copy mechanism see Kedar-Cabelli & McCarty, 1987) before instantiation with problem specific information takes place (cf. the third rule in Table 7), generalised proof trees are constructed and externalised in the third argument.

A generalised proof tree is available for further processing after a problem solving goal has been derived from the knowledge available on the cognitive level. However, if the knowledge available on the cognitive level is incomplete with respect to a posed problem, no proof tree can be constructed and problem solving fails. To allow for the construction of (partial) proof trees even when only incomplete knowledge is available on the cognitive level, the meta-interpreter was further extended to simulate deductive self-diagnoses. The functionality of the extended meta-interpreter can be summarised as follows:

Given

- a problem solving goal and
- a knowledge base,

do the following:

- if the knowledge base is complete with respect to the problem solving goal, then construct a complete proof tree;
- otherwise construct a partial proof tree in which maximally specific subgoals which cannot be derived from the knowledge base are labelled as unprovable.

**Table 7.** A meta-interpreter for constructing generalised proof trees

---

```

interpreter((Goal, Goals), (G_Goal, G_Goals), (G_Proof, G_Proofs)) :-
    !,
    interpreter(Goal, G_Goal, G_Proof),
    interpreter(Goals, G_Goals, G_Proofs).

interpreter(Goal, G_Goal, system_predicate(G_Goal)) :-
    system_predicate(Goal),
    !,
    call(Goal).

interpreter(Goal, G_Goal, (G_Goal :- G_Proof)) :-
    clause(G_Goal, G_Goals),
    copy((G_Goal, G_Goals), (Goal, Goals)),
    interpreter(Goals, G_Goals, G_Proof).

```

---

To achieve this functionality, the meta-interpreter was extended by additional rules and a fourth argument turning the meta-interpreter into a "fail-safe meta-interpreter" (cf. Table 8). The most important additional rule succeeds whenever the proof of a goal fails (cf. the last rule in Table 8). If no complete proof tree can be constructed, then, in general, the notion of a goal that cannot be proved is ambiguous: if the proof of a certain goal fails, then the proof of every goal in the (partial) proof tree that is "above" the failing goal fails too. However, there exist always "lowest", maximally specific goals within the (partial) ordering of failing goals. These maximally specific goals form the "leaves" of the partial proof tree which cannot be proved. They allow for the construction of maximally specific questions.

**Table 8.** A fail-safe meta-interpreter for constructing complete as well as maximally specific partial generalised proof trees

---

```

interpreter((Goal, Goals), (G_Goal, G_Goals), (G_Proof, G_Proofs), provable) :-
    interpreter(Goal, G_Goal, G_Proof, provable),
    interpreter(Goals, G_Goals, G_Proofs, provable).

interpreter((Goal, Goals), (G_Goal, G_Goals), (G_Proof, G_Proofs), unprovable) :-
    interpreter(Goal, G_Goal, G_Proof, provable),
    interpreter(Goals, G_Goals, G_Proofs, unprovable).

interpreter((Goal, Goals), (G_Goal, G_Goals), (G_Proof, terminated(G_Goals)),
    unprovable) :-
    interpreter(Goal, G_Goal, G_Proof, unprovable).

interpreter(Goal, G_Goal, system_predicate(G_Goal), provable) :-
    system_predicate(Goal),
    !,
    call(Goal).

interpreter(Goal, G_Goal, (provable(G_Goal) :- G_Proof), provable) :-
    not(conjunction(Goal)),
    clause(G_Goal, G_Goals),
    copy((G_Goal, G_Goals), (Goal, Goals)),
    interpreter(Goals, G_Goals, G_Proof, provable).

interpreter(Goal, G_Goal, (unprovable(G_Goal) :- G_Proof), unprovable) :-
    not(conjunction(Goal)),
    not(system_predicate(Goal)),
    clause(G_Goal, G_Goals),
    copy((G_Goal, G_Goals), (Goal, Goals)),
    interpreter(Goals, G_Goals, G_Proof, unprovable).

interpreter(Goal, _, unprovable(Goal), unprovable) :-
    not(conjunction(Goal)),
    not(system_predicate(Goal)),
    not(clause(Goal, _)).

```

---

The fourth argument ensures the completeness of the meta-interpreter with respect to the construction of complete proof trees: if it is possible to construct complete proof trees, then all of them are constructed before any partial proof tree is constructed. However, to avoid nonterminating computations, the meta-interpreter is not complete with respect to the



construction of partial proof trees (cf. the third rule in Table 8). That is, there might exist partial proof trees which cannot be constructed by the meta-interpreter.

### **Learning from Explanations**

We conceptualise those cognitive processes as learning from explanations by means of which one attempts to understand explanations received from others on the basis of one's own pre-knowledge. Learning from explanations aims at (a) identifying new information in explanations and (b) extending and complementing one's own pre-knowledge by taking advantage of the new information. Learning from explanations can in part be considered as being inverse to deductive self-diagnosis:

Given

- a complete or partial proof tree and
- a knowledge base,

do the following:

- identify information in the proof tree which cannot be derived from the knowledge base and complement the knowledge base with the identified information.

That is, in order to simulate learning from explanations, a possibly partial proof tree is successively decomposed and components which cannot be derived from the knowledge available on the cognitive level are added to the knowledge base. To achieve this functionality, the meta-cognitive level was extended by additional rules which determine components of a proof tree which are not already included in the knowledge base and complement the knowledge base with the new information.

### **Communication**

Questions and generalised proof trees are exchanged between the two problem solvers on the communication level. If a problem solver cannot solve a posed problem, the meta-interpreter passes a partial proof tree to the communication interpreter. The communication interpreter filters out those leaves of the partial proof tree which were unprovable and forwards them as questions via TCP/IP (Transmission Control Protocol/Internet Protocol) to the other problem solver. The problem solver which receives questions considers them as new problems and attempts to solve them. If solutions to the problems can be found, the problem solver sends the solutions as well as generalised proof trees back to the problem solver which raised the questions. If no solutions to the problems can be found, the problem solver nevertheless sends partial generalised proof trees back to the problem solver which raised the questions.

In its current implementation, the communication interpreter allows only for non-overlapping problem solving and asymmetric communication. Non-overlapping problem solving refers to the constraint that only one problem solver can be active at a time. Asymmetric communication refers to the constraint that a problem solver which attempts to answer a question cannot ask the other problem solver for further information.

The second constraint especially necessitates that the problem solvers assume defined roles during collaborative problem solving. While one problem solver attempts to solve a posed problem, the other problem solver attempts to answer questions raised by the first problem solver. However, the roles assumed by the problem solvers may be switched when a new problem has to be solved.

### **Performance**

In Table 9, the performance of the two problem solvers is exemplified with respect to Problem 2 in Table 5. The different problem solving steps shown in Table 9 are traced on the cognitive

level. The different questions, answers and explanations shown in Table 9 are traced on the communication level.

**Table 9.** Performance of the qualitative (white areas) and quantitative (grey-shaded areas) solver

<b>Trying to solve problem ...</b>	1
<b>Question:</b> $\text{value}(2, \text{magnitude}(\text{velocity}(\text{block}_1))) = ?$	2
<b>Trying to solve problem ...</b>	3
trying equation $v=v_0+a*t$ for $\text{value}(2, \text{magnitude}(\text{velocity}(\text{block}_1)))$	4
trying equation $fr=m*a$ for $\text{value}(1, \text{magnitude}(\text{acceleration}(\text{block}_1)))$	5
<b>Answer:</b> unprovable	6
<b>Explanation:</b>	7
$(\text{unprovable}(\text{solve\_equation}(v=v_0+a*t,$	8
$\text{value}(S2, \text{magnitude}(\text{velocity}(\text{Object}))) =$	9
$\text{value}(S1, \text{magnitude}(\text{velocity}(\text{Object}))) + \text{value}(S2, \text{magnitude}(\text{acceleration}(\text{Object}))) *$	10
$\text{value}(S1 \rightarrow S2, \text{duration}(\text{displacement}(\text{Object})))) :-$	11
$(\text{unprovable}(\text{solve\_equation}(fr=m*a,$	12
$\text{value}(S, \text{magnitude}(\text{force}(\text{Object}, \text{resultant\_force}, fr))) =$	13
$\text{value}(S, \text{mass}(\text{Object})) * \text{value}(S, \text{magnitude}(\text{acceleration}(\text{Object})))) :-$	14
$\text{unprovable}(\text{value}(1, \text{magnitude}(\text{force}(\text{block}_1, \text{resultant\_force}, fr))))$	15
<b>New knowledge acquired:</b>	16
equation( $v=v_0+a*t$ , $\text{value}(S2, \text{magnitude}(\text{velocity}(\text{Object}))) =$	17
$\text{value}(S1, \text{magnitude}(\text{velocity}(\text{Object}))) + \text{value}(S2, \text{magnitude}(\text{acceleration}(\text{Object}))) *$	18
$\text{value}(S1 \rightarrow S2, \text{duration}(\text{displacement}(\text{Object}))))$	19
<b>New knowledge acquired:</b>	20
equation( $fr=m*a$ , $\text{value}(S, \text{magnitude}(\text{force}(\text{Object}, \text{resultant\_force}, fr))) =$	21
$\text{value}(S, \text{mass}(\text{Object})) * \text{value}(S, \text{magnitude}(\text{acceleration}(\text{Object}))))$	22
<b>Trying to solve problem ...</b>	23
trying equation $v=v_0+a*t$ for $\text{value}(2, \text{magnitude}(\text{velocity}(\text{block}_1)))$	24
trying equation $fr=m*a$ for $\text{value}(1, \text{magnitude}(\text{acceleration}(\text{block}_1)))$	25
derived $\text{value}(1, \text{instance}(\text{mass}(\text{block}_1), \text{mass})) = \text{true}$	26
derived $\text{value}(1, \text{instance}(\text{force}(\text{block}_1, \text{earth}, fg), \text{gravitational\_force})) = \text{true}$	27
derived $\text{value}(1, \text{instance}(\text{force}(\text{block}_1, \text{plane}_1, fn), \text{normal\_force})) = \text{true}$	28
...	29
derived $\text{value}(1, \text{arrows}(\text{block}_1)) = [\text{arrow}(\text{force}(\text{block}_1, \text{earth}, fg), 90, \text{down}),$	30
$\text{arrow}(\text{force}(\text{block}_1, \text{plane}_1, fn), 120, \text{up}), \text{arrow}(\text{velocity}(\text{block}_1), 30, \text{down})]$	31
derived $\text{value}(1, \text{instance}(\text{coordinate\_system}(\text{block}_1, [\text{axis}(x, 30), \text{axis}(y, 120)]),$	32
$\text{coordinate\_system})) = \text{true}$	33
derived $\text{value}(1, \text{instance}(\text{free\_body\_diagram}(\text{coordinate\_system}(\text{block}_1,$	34
$[\text{axis}(x,30), \text{axis}(y, 120)]), [\text{arrow}(\text{force}(\text{block}_1, \text{earth}, fg), 90, \text{down}),$	35
$\text{arrow}(\text{force}(\text{block}_1, \text{plane}_1, fn), 120, \text{up}), \text{arrow}(\text{velocity}(\text{block}_1), 30, \text{down})], \text{fbd})) =$	36
$\text{true}$	37
derived $\text{value}(1, \text{magnitude}(\text{force}(\text{block}_1, \text{resultant\_force}, fr))) =$	38
$\text{value}(1, \text{magnitude}(\text{projection}(\text{force}(\text{block}_1, \text{earth}, fg), \text{axis}(x, 30))))$	39
derived $\text{value}(1, \text{magnitude}(\text{projection}(\text{acceleration}(\text{block}_1), \text{axis}(x, 30)))) *$	40
$\text{value}(1, \text{mass}(\text{block}_1)) =$	41
$\text{value}(1, \text{magnitude}(\text{projection}(\text{force}(\text{block}_1, \text{earth}, fg), \text{axis}(x, 30))))$	42
$\text{value}(1, \text{mass}(\text{block}_1)) =$	43
<b>Question:</b> $\text{value}(1, \text{magnitude}(\text{force}(\text{block}_1, \text{earth}, fg))) = ?$	44
<b>Trying to solve problem ...</b>	44
trying equation $fg=m*g$ for $\text{value}(1, \text{magnitude}(\text{force}(\text{block}_1, \text{earth}, fg)))$	45
derived $\text{value}(1, \text{magnitude}(\text{force}(\text{block}_1, \text{earth}, fg))) = 15::\text{kg} * c(g)$	46
<b>Answer:</b> $\text{value}(1, \text{magnitude}(\text{force}(\text{block}_1, \text{earth}, fg))) = 15::\text{kg} * c(g)$	47
<b>Explanation:</b>	48
$(\text{provable}(\text{solve\_equation}(fg=m*g,$	49
$\text{value}(1, \text{magnitude}(\text{force}(\text{Object}, \text{earth}, fg))) = \text{value}(1, \text{mass}(\text{Object})) * c(g)) :-$	50
$(\text{provable}(\text{known\_value}(\text{value}(1, \text{mass}(\text{Object})))) :-$	51
$\text{system\_predicate}(\text{given}(\text{value}(1, \text{mass}(\text{Object})) = \text{Value})))$	52
<b>New knowledge acquired:</b>	53
equation( $fg=m*g$ , $\text{value}(1, \text{magnitude}(\text{force}(\text{Object}, \text{earth}, fg))) =$	54
$\text{value}(1, \text{mass}(\text{Object})) * c(g)$	55
<b>Trying to solve problem ...</b>	56
...	57
derived $\text{value}(1, \text{magnitude}(\text{projection}(\text{force}(\text{block}_1, \text{earth}, fg), \text{axis}(x, 30)))) =$	58
$15::\text{kg} * c(g) * \sin(30)$	59
derived $\text{value}(1, \text{magnitude}(\text{acceleration}(\text{block}_1))) = \sin(30) * c(g)$	60
derived $\text{value}(2, \text{magnitude}(\text{velocity}(\text{block}_1))) = \sin(30) * c(g) * 2::s$	61
<b>Solution:</b> $\text{value}(2, \text{magnitude}(\text{velocity}(\text{block}_1))) = \sin(30) * c(g) * 2::s$	62

The explanations shown in Table 9 comprise only information which results from applying domain-specific knowledge on the cognitive level. Information which results from applying interpretation mechanisms on the cognitive level is neglected.

In the following, the problem solver which is initially exclusively equipped with knowledge about qualitative aspects of classical mechanics is named the "qualitative solver". The problem solver which is initially exclusively equipped with knowledge about quantitative aspects is named the "quantitative solver". The problem solving steps, questions and explanations of the quantitative solver are grey-shaded in Table 9.

Initially, both problem solvers are initialized with the given and sought quantities of the posed problem (cf. Table 2). In the arrangement underlying the performance shown in Table 9, the qualitative solver attempts to solve the problem. The quantitative solver remains waiting for questions from the qualitative solver. On the cognitive level, problem solving essentially relies on backward-chaining. Because the qualitative solver has no knowledge available that applies to the unknown quantity, it asks the quantitative solver for the block's velocity (cf. Line 2 in Table 9).

In a first step, the quantitative solver considers a kinematic equation to determine the block's velocity (cf. Line 4). Because acceleration is a further unknown quantity in this equation, Newton's second law is considered next to determine the acceleration (cf. Line 5). However, because the quantitative solver comprises no knowledge about qualitative aspects, it cannot determine the resultant force on the block. Because no further kinematic equations are available, the quantitative solver cannot determine the block's velocity on its own.

In this case, the "lowest" problem solving goal that fails on the cognitive level of the quantitative solver is the goal that aims at determining the resultant force on the block. Consequently, the quantitative solver replies to the qualitative solver that it cannot answer the question (cf. Line 6). In addition, the quantitative solver forwards a partial proof tree to the qualitative solver (cf. Lines 8 to 15).

After the qualitative solver received the proof tree, it successively decomposes the proof tree and tries to identify components which cannot be derived from its own knowledge base. As a result, the qualitative solver extends its own knowledge base with the two equations the quantitative solver failed to apply (cf. Lines 17 to 22).

Afterwards, the qualitative solver again attempts to solve the problem. It initially makes use of the newly acquired equations (cf. Lines 24 and 25) and then proceeds to construct a free-body diagram (cf. Lines 26 to 37) to determine the resultant force on the block. The qualitative solver establishes that the resultant force on the block is made up of the projection of the gravitational force on the block onto the x-axis (cf. Lines 38 and 39). Consequently, the determination of the block's acceleration demands the determination of the gravitational force on the block (cf. Lines 40 to 42). Therefore, the qualitative solver again asks the quantitative solver (cf. Line 43).

The quantitative solver determines the gravitational force on the block by taking advantage of a force law (cf. Lines 45 and 46). Subsequently, it forwards the answer as well as the complete proof tree to the qualitative solver (cf. Lines 47 to 52). The qualitative solver decomposes the proof tree and extends its own knowledge base with the force law the quantitative solver successfully applied (cf. Lines 54 and 55). On the basis of the extended knowledge base, the qualitative solver finally determines the problem's solution (cf. Lines 57 to 62).

## **SIMULATION STUDY**

The performances of the two problem solvers were compared in a simulation study. The study was made up of two phases: collaborative problem solving and individual testing. During collaborative problem solving, the problem solvers were applied "collaboratively" to five problem descriptions which encoded the problems posed to the students in the experimental study (cf. Table 5). Two arrangements were considered. In the first arrangement, the qualitative solver attempted to solve the problem and the quantitative solver remained waiting for questions

from the qualitative solver. In this arrangement, the qualitative solver learned from explanations constructed by the quantitative solver. In the second arrangement, the quantitative solver attempted to solve the problem and the qualitative solver remained waiting for questions from the quantitative solver. In this arrangement, the quantitative solver learned from explanations constructed by the qualitative solver.

If no restrictions are applied to the learning capabilities of the two problem solvers, they are able to decompose explanations of arbitrary size and to acquire an arbitrary number of new pieces of knowledge included in these explanations. Because the assumption of unrestricted learning capabilities appears to be highly implausible, the learning capabilities of the two problem solvers were restricted in such a way that the problem solvers were capable of acquiring a maximum of seven new pieces of knowledge included in an explanation. Such new pieces of knowledge can be any expressions which encode knowledge about qualitative or quantitative aspects of classical mechanics (cf. Table 1 and Table 9).

We do not claim that this restriction adequately reflects restrictions of human learning capabilities. In a more adequate model, theories about the human working memory and its role in problem solving and learning (e.g., Sweller, 1994; Sweller & Chandler, 1994) would have to be taken into account, for example.

With respect to the first arrangement, the number of explanations the qualitative solver received from the quantitative solver during collaborative problem solving is shown in the upper part of Table 10. Across all problems, the qualitative solver received 12 explanations from the quantitative solver. These explanations included 29 pieces of information. Because not all pieces of information were new to the qualitative solver, six new pieces of knowledge were acquired.

**Table 10.** The number of explanations exchanged and the number of pieces of new knowledge acquired during collaborative problem solving

Problem solver		Problem					
		1	2	3	4	5	All
Qualita- tive	Explanations received	1	1	0	7	3	12
	Pieces of information included	3	3	0	17	6	29
	Pieces of new knowledge acquired	2	1	0	3	0	6
	Pieces of new knowledge not acquired	0	0	0	0	0	0
Quanti- tative	Explanations received	2	4	4	3	4	17
	Pieces of information included	20	38	35	0	36	129
	Pieces of new knowledge acquired	8	9	7	0	0	24
	Pieces of new knowledge not acquired	4	3	0	0	0	7

With respect to the second arrangement, the number of explanations the quantitative solver received from the qualitative solver during collaborative problem solving is shown in the lower

part of Table 10. Across all problems, the quantitative solver received 17 explanations from the qualitative solver. These explanations included 129 pieces of information. While 24 new pieces of knowledge were acquired, seven new pieces of knowledge were not acquired due to the restriction of the learning capabilities.

During individual testing, the problem solvers were applied separately to 18 problem descriptions which encoded the items of the post-test posed to the students in the experimental study. Before collaborative problem solving, the qualitative solver solved the six items (33%) of the first test component which assess knowledge about qualitative aspects of classical mechanics. Correspondingly, the quantitative solver solved the six items (33%) of the first test component which assess knowledge about quantitative aspects. The performance of the two problem solvers after collaborative problem solving is shown in Table 11. While the qualitative solver solved 16 items (89%), the quantitative solver solved 12 items (67%).

**Table 11.** The performance of the qualitative and quantitative solver on the multi-component test after collaborative problem solving

Problem solver	Successfully solved test items			
	Qualitative	Quantitative	Interrelated	All items
Qualitative	6 (100%)	5 (83%)	5 (83%)	16 (89%)
Quantitative	5 (83%)	6 (100%)	1 (17%)	12 (67%)

The results of the simulation study roughly correspond to the main findings of the experimental study. However, although the learning capabilities of the two problem solvers have been restricted, their learning performance is still above the average learning performance of the students in the experimental study. In the following, we discuss various idealisations realised in the model which might be responsible for the model’s superior learning performance.

## DISCUSSION

In this paper, we described a cognitive model that simulates selected aspects of collaborative problem solving and learning under the conditions set up in an experimental study. The model rests on three basic assumptions. The first assumption is that collaboration is especially beneficial to the collaborating partners if they initially possess different pre-knowledge about the application domain. The second assumption is that the ability to direct felicitous questions to one’s partner frequently depends on adequate self-diagnoses of shortcomings of one’s own competence. The third assumption is that explanations received from one’s partner provide a valuable source of information for extending one’s own knowledge.

In accord with the conditions identified by Webb (1989), the exchange and use of explanations as simulated by the cognitive model should lead to successful collaborative learning, because (a) the explanations are always relevant to the explainees questions, (b) the explanations commonly match the level of help needed, (c) the explanations are provided without major delay, (d) the explanations match the explainees level of understanding and (d) the explainee has the opportunity to apply the information provided by the explanations to the posed problem.

We made use of the cognitive model in a simulation study. By means of the simulation study it was possible to roughly reconstruct the main findings of the experimental study, namely, that qualitatively instructed students gained more from their quantitatively instructed

partners than the other way round. The results of the simulation study give rise to a potential account for this finding.

With respect to the problems which were posed to the students in the experimental study as well as to the qualitative and quantitative solvers in the simulation study, the qualitative solver constructed not only more but also larger explanations than the quantitative solver. Especially, the construction of free-body diagrams by the qualitative solver required much knowledge about qualitative aspects of classical mechanics to be taken into account.

On the one hand, quantitatively instructed students might have constructed rather limited explanations. Their qualitatively instructed partners were possibly able to understand these explanations without severe difficulties and to take advantage of them to complement their own knowledge. On the other hand, because explanations of qualitative reasoning had often to be extensive, qualitatively instructed students might have constructed either (a) incomplete explanations which informed their quantitatively instructed partners only partially or (b) complete explanations which overburdened their quantitatively instructed partners.

This hypothesis is supported by a discourse analysis of the verbal exchange of information of three selected pairs of students (cf. Kneser, 1997; Ploetzner & Kneser, 1998). Among other results, the discourse analysis revealed that (a) while the qualitatively instructed students gradually made fewer requests for information about quantitative aspects, the number of requests for information about qualitative aspects gradually raised from the first to the last problem and (b) explanations of qualitative aspects were commonly more extensive than explanations of quantitative aspects.

Together, the results of the empirical study as well as the results of the simulation study suggest that qualitative problem representations form not only a good starting point for the subsequent construction of quantitative problem representations during problem solving (e.g., Chi, Feltovich & Glaser, 1981; Larkin, 1983; Ploetzner, 1994) but also a beneficial starting point for learning quantitative problem representations.

The results also suggest that the complementarity of knowledge about qualitative and quantitative aspects of classical mechanics is not a symmetric complementarity. While knowledge of qualitative physics seems to encode a rather expanded and related set of qualitative aspects, knowledge about quantitative physics seems to encode a rather limited and only weakly related set of quantitative aspects. Thus, from an instructional point of view, knowledge about quantitative aspects complements knowledge about qualitative aspects rather than the other way round.

Though by means of the cognitive model it was possible to roughly reconstruct the findings of the experimental study, the model is nevertheless based on various simplifications or idealisations. These idealisations address

- the exchange of information between the problem solvers,
- the specificity of questions raised during the collaboration,
- the generality of the explanations exchanged between the problem solvers,
- the coherence of the explanations exchanged between the problem solvers,
- the correctness of the explanations exchanged between the problem solvers,
- the knowledge representation language available to the problem solvers and
- the number of partners involved in the collaboration.

Each idealisation and how it might be overcome is briefly discussed in the following.

### *Exchange of information*

In its current implementation, the model allows only for non-overlapping problem solving and asymmetric exchange of information. Non-overlapping problem solving refers to the constraint that only one problem solver can be active at a time. Asymmetric exchange of information

refers to the constraint that a problem solver which attempts to answer a question cannot ask the other problem solver for further information.

In the empirical study, in contrast, commonly both students were active at a time. Furthermore, students which were attempting to answer questions very often asked their partners for further information. In order to allow for a symmetric exchange of information, it is intended to equip the model in the future with a more flexible communication interpreter based on a KQML protocol (Knowledge Query and Manipulation Language; Finin, Labrou & Mayfield, 1997).

#### *Specificity of questions*

During collaborative problem solving, the model always begins with the construction of maximally specific questions by means of deductive self-diagnoses. These questions refer to problem solving goals which are the "lowest" unprovable problem solving goals in a partial proof tree. Only when no answers to these questions are provided, are new questions constructed which are "above" the "lowest" unprovable problem solving goals.

In the empirical study, however, students occasionally began with the construction of questions which directly referred to a sought quantity. Only when no sufficient answers to these questions were provided, were more specific questions constructed. By making use of an additional parameter on the meta-cognitive level, both strategies as well as alternations between both strategies during problem solving could be simulated by the model.

#### *Generality of explanations*

In its current implementation, the model always constructs generalised explanations. As a consequence, the problem solver which attempts to understand explanations on the basis of its own pre-knowledge acquires generalised pieces of knowledge. Subsequently, these generalised pieces of knowledge can be applied to new problems. In the empirical study, however, the students frequently constructed explanations which were specific to the problems under consideration. Thus, it was up to the students who received the explanations to successively generalise the newly acquired pieces of knowledge.

In the model, the construction of problem-specific explanations could be simulated by means of a minor modification on the meta-cognitive level. The successive generalisation of specific pieces of knowledge acquired by the problem solver which receives problem-specific explanations could be simulated by taking advantage of machine learning techniques. The newly acquired pieces of knowledge could be considered as training examples. Subsequently, on the basis of an inductive learning mechanism, they would have to be generalised in such a way that the general pieces of knowledge subsume the training examples (e.g., Michalski, 1983).

#### *Coherence of explanations*

Although the model does not only construct complete but also partial explanations, it nevertheless always constructs coherent explanations. In the empirical study, in contrast, the students frequently constructed merely fragmentary explanations. To understand and to learn from fragmentary explanations is possibly much more difficult than to understand and to learn from coherent explanations.

In some cases, a student who receives fragmentary explanations might nevertheless be able to understand them. In some cases, fragmentary explanations might lead to further questions in order to receive clarifying information. In other cases, fragmentary explanations might be memorised in the hope that later one will be able to take advantage of them. In some cases, however, fragmentary explanations might completely be ignored, because a student who receives the explanations might not be able to make any sense of them.

*Correctness of explanations*

The simulated problem solvers only comprise correct knowledge about qualitative and quantitative aspects of classical mechanics. As a consequence, only correct explanations are constructed and only correct pieces of knowledge are acquired by the problem solver which receives the explanations.

Humans, however, very often possess incorrect knowledge especially about qualitative aspects of physics (e.g., Clement, 1982; McCloskey, 1983; for a bibliography see Pfundt & Duit, 1994). If the simulated problem solvers were to construct explanations on the basis of incorrect knowledge, these explanations would also be incorrect. As a consequence, in its current implementation, the simulated problem solver which receives the explanation would acquire incorrect knowledge.

Even if the simulated problem solver which receives the incorrect explanation were already to possess the corresponding correct knowledge, it would nevertheless acquire the incorrect knowledge. In order to overcome this deficiency, the model would have to be equipped with mechanisms capable of identifying inconsistencies between an explanation and its own pre-knowledge. The identified inconsistencies could subsequently be made use of to oppose an explanation, for example. Such an opposition could prompt collaboration processes to clarify which aspects are correct and which are incorrect.

*Knowledge representation language*

Currently, the simulated problem solvers take advantage of the same knowledge representation language and the same interpretation mechanisms. This idealisation largely simplifies the communication between the simulated problem solvers. The information exchanged - at least in principle - can always be interpreted by both simulated problem solvers. Especially, knowledge acquired by learning from explanations can straightforwardly be applied to new problems. The representation languages and interpretation mechanisms used by humans, however, very often differ from each other.

As a consequence, a human who receives an explanation from another human might (a) be able to only partially interpret it and (b) interpret it in a completely different way than the human who constructed it. Therefore, the human who receives the explanation might ask for clarifying information, for example. Or the human who constructed the explanation might realise that he or she has been misunderstood.

The realisation of mutual misunderstandings could subsequently prompt negotiation processes to clarify the meaning as well as to reach a common understanding of the information in question (e.g., Baker, 1994). The detailed reconstruction and simulation of the involved negotiation processes, however, remain serious challenges to research on collaborative problem solving and learning for the years to come.

*Number of partners involved*

In its current implementation, the model is restricted to two collaborating problem solvers. Collaboration, however, frequently involves more than two partners. If more than two simulated problem solvers were involved, then each individual problem solver would not only require a (partial) model of the application domain under scrutiny but also adequate models of the knowledge possessed by its partners, for example.

On the basis of these models, questions can be posed to competent partners and information can be provided to partners which are in need of them. The construction of models of the knowledge possessed by one's partners makes up an inductive learning task of its own. Its accomplishment could in part be simulated by taking advantage of modelling techniques which have been developed in the framework of student modelling (e.g., VanLehn, 1988b; Self, 1994).



## CONCLUSIONS

We described a cognitive simulation model of collaborative problem solving and learning. Our intention was not to present a full-fledged model of how humans mutually exchange information during collaboration. Instead, our intention was to present two domain-independent mechanisms which allow for the simulation of (1) constructing questions and (2) learning by attempting to understand explanations. Both mechanisms were implemented by taking advantage of meta-programming techniques and are easy to reimplement.

On the basis of the model it was possible to reconstruct the main result of an experimental study of collaborative problem solving and learning. The qualitative solver gained more from the quantitative solver than the other way round. In accord with the findings of a discourse analysis, the model suggests that explanations of qualitative aspects frequently overburden the receiver of the explanations with new information.

In order to facilitate the acquisition of knowledge about qualitative aspects by the quantitative solver, the role of a tutor could be assigned to the qualitative solver from time to time. During tutoring, the foremost goal of the qualitative solver should not be to advance the solution of a problem but to explicitly teach a small and selected set of qualitative aspects of classical mechanics to the quantitative solver.

On the one hand, the successful reconstruction of the main results of the experimental study by the model suggests that the presented mechanisms might indeed be relevant to collaborative problem solving and learning. On the other hand, we are fully aware of the fact that the model simulates only a few aspects which might be relevant to collaborative problem solving and learning. The development of more complete, cognitively adequate and educationally fruitful simulation models of collaborative problem solving and learning remains a challenge.

However, we strongly believe that one important aim of research on artificial intelligence in education is to develop computational methods which can incrementally and flexibly be used and extended. We hope that the presented mechanisms contribute to this aim.

## Acknowledgments

The research reported in this paper has been supported by the German National Research Foundation (DFG) under contracts SP 251/9-1 and PL 224/2-2. The design and the results of the experimental study have been published earlier in the *Zeitschrift für Entwicklungs-psychologie und Pädagogische Psychologie* [Journal of Developmental Psychology and Pedagogical Psychology] (Ploetzner, Fehse, Spada, Vodermaier & Wolber, 1996). The mechanism for simulating the construction of questions on the basis of deductive self-diagnoses has been published earlier in the Proceedings of the Eighth World Conference on Artificial Intelligence in Education (Ploetzner, Fehse, Hermann & Kneser, 1997).

## References

- Aronson, E., Balney, N., Stephan, C., Sikes, J., & Snapp, M. (1978). *The jigsaw classroom*. Beverly Hills, CA: Sage.
- Baker, M. (1994). A model for negotiation in teaching-learning dialogues. *Journal of Artificial Intelligence in Education*, 5, 199-254.
- Beller, S., & Hoppe, H. U. (1993). Deductive error reconstruction and classification in a logic programming framework. In P. Brna, S. Ohlsson & H. Pain (Eds.), *Proceedings of the World Conference on Artificial Intelligence in Education* (pp. 433-440). Charlottesville, VA: Association for the Advancement of Computing in Education.
- Chi, M. T. H., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5, 121-152.

- Chi, M. T. H., Glaser, R., & Rees, E. (1982). Expertise in problem solving. In R. J. Sternberg (Ed.), *Advances in the psychology of human intelligence* (pp. 7-75). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Clement, J. (1982). Students' preconceptions in introductory mechanics. *American Journal of Physics*, 50, 66-71.
- Davis, E. (1990). *Representations of commonsense knowledge*. San Mateo, CA: Morgan Kaufmann.
- Dillenbourg, P., Baker, M., Blaye, A., & O'Malley, C. (1996). The evolution of research on collaborative learning. In P. Reimann & H. Spada (Eds.), *Learning in humans and machines: Towards an interdisciplinary learning science* (pp. 189-211). Oxford: Pergamon.
- de Kleer, J. (1975). *Qualitative and quantitative knowledge in classical mechanics* (Technical Report No. AI-TR-352). Boston, MA: Massachusetts Institute of Technology.
- Finin, T., Labrou, Y., & Mayfield, J. (1997). KQML as an agent communication language. In J. Bradshaw (Ed.), *Software agents*. Cambridge, MA: MIT Press.
- Graesser, A. C., & Person, N. K. (1994). Question asking during tutoring. *American Educational Research Journal*, 31, 104-137.
- Halliday, D., & Resnick, R. (1988). *Fundamentals of physics* (3rd ed.). New York, NY: Wiley.
- Hoppe, H. U. (1994). Deductive error diagnosis and inductive error generalization for intelligent tutoring systems. *Journal of Artificial Intelligence in Education*, 5, 27-49.
- Howe, C., Tolmie, A., & Rodgers, C. (1990). Physics in the primary school: Peer interaction and the understanding of floating and sinking. *European Journal of Psychology of Education*, 4, 459-475.
- Howe, C., Tolmie, A., & Rodgers, C. (1992). The acquisition of conceptual knowledge in science by primary school children: Group interaction and the understanding of motion down an incline. *British Journal of Developmental Psychology*, 10, 113-130.
- Kedar-Cabelli, S. T., & McCarty, L. T. (1987). Explanation-based generalization as resolution theorem proving. In P. Langley (Ed.), *Proceedings of the Fourth International Workshop on Machine Learning* (pp. 383-389). Los Altos, CA: Morgan Kaufmann.
- Kneser, C. (1997). *Kooperatives Problemlösen: Der Zusammenhang von Dialogstruktur und Lernerfolg* [Cooperative problem solving: The relation between dialogue structure and learning success] (Unpublished Diplom Thesis). Freiburg: Department of Psychology at the University of Freiburg.
- Knight, G. P., & Bohlmeier, E. M. (1990). Cooperative learning and achievement: Methods for assessing causal mechanisms. In S. Sharan (Ed.), *Cooperative learning: Theory and research* (pp. 1-22). New York: Praeger.
- Larkin, J. H. (1983). The role of problem representation in physics. In D. Gentner, & A. L. Stevens (Eds.), *Mental models* (pp. 75-98). Hillsdale, NJ: Lawrence Erlbaum Associates.
- McCloskey, M. (1983). Naive theories of motion. In D. Gentner, & A. L. Stevens (Eds.), *Mental models* (pp. 299-324). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Michalski, R. S. (1983). A theory and methodology of inductive learning. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning - An artificial intelligence approach* (pp. 83-129). Palo Alto, CA: Morgan Kaufmann.
- Novak, J. D. (1990). Concept mapping: A useful tool for science education. *Journal of Research in Science Teaching*, 27, 937-949.
- Ohlsson, S., & Rees, E. (1991). The function of conceptual understanding in the learning of arithmetic procedures. *Cognition and Instruction*, 8, 103-179.
- Pfundt, H., & Duit, R. (1994). *Bibliography: Students' alternative frameworks and science education* (4th ed.). Kiel: Institute for Science Education.
- Ploetzner, R. (1994). *The integrative use of qualitative and quantitative knowledge in physics problem solving*. Frankfurt/Main: Lang.
- Ploetzner, R. (1995). The construction and coordination of complementary problem representations in physics. *Journal of Artificial Intelligence in Education*, 6, 203-238.
- Ploetzner, R., Fehse, E., Hermann, F., & Kneser, C. (1997). Modeling the knowledge based exchange of information during collaborative problem solving on the basis of deductive

- self-diagnoses. In B. du Boulay, & R. Mizoguchi (Eds.), *Proceedings of the Eighth World Conference on Artificial Intelligence in Education* (pp. 223-230). Amsterdam: IOS Press.
- Ploetzner, R., Fehse, E., Spada, H., Vodermaier, A., & Wolber, D. (1996). Physiklernen mit modellgestützt konstruierten Begriffsnetzen und durch Problemlösen zu zweit [Learning physics by means of model-based concept maps and collaborative problem solving]. *Zeitschrift für Entwicklungspsychologie und Pädagogische Psychologie*, 28, 270-293.
- Ploetzner, R., & Kneser, C. (1998). Steps towards the acquisition of expertise: Shifting the focus from quantitative to qualitative problem representations during collaborative problem solving. In M. A. Gernsbacher & S. J. Derry (Eds.), *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society* (pp. 842-847). Mahwah, NJ: Lawrence Erlbaum Associates.
- Ploetzner, R., Spada, H., Stumpf, M., & Opwis, K. (1990). Learning qualitative and quantitative reasoning in a microworld for elastic impacts. *European Journal of Psychology of Education*, 4, 501-516.
- Self, J. (1994). Formal approaches to student modeling. In J. E. Greer & G. I. McCalla (Eds.), *Student modelling: The key to individualized knowledge-based instruction* (pp. 295-352). Berlin: Springer.
- Shoham, Y. (1994). *Artificial intelligence techniques in Prolog*. San Francisco, CA: Morgan Kaufmann.
- Slavin, R. E. (1995). *Cooperative learning: Theory, research, and practice* (2nd ed.). Boston, MA: Allyn and Bacon.
- Sterling, L., & Shapiro, E. (1994). *The art of Prolog: Advanced programming techniques* (2nd ed.). Cambridge, MA: MIT Press.
- Sweller, J. (1994). Cognitive load theory, learning difficulty, and instructional design. *Learning and Instruction*, 4, 295-312.
- Sweller, J., & Chandler, P. (1994). Why some material is difficult to learn. *Cognition and Instruction*, 12, 185-233.
- van Joolingen, W. R. (1994). QMaPS: Qualitative reasoning for simulation learning environments. *Journal of Artificial Intelligence and Education*, 5, 177-197.
- VanLehn, K. (1988a). Towards a theory of impasse-driven learning. In H. Mandl & A. Lesgold (Eds.), *Learning issues for intelligent tutoring systems* (pp. 19-41). New York, NY: Springer.
- VanLehn, K. (1988b). Student modeling. In M. C. Polson & J. J. Richardson (Eds.), *Foundations of intelligent tutoring systems* (pp. 55-78). Hillsdale, NJ: Lawrence Erlbaum Associates.
- VanLehn, K. (1996). Cognitive skill acquisition. *Annual Review of Psychology*, 47, 513-539.
- VanLehn, K., Jones, R. M., & Chi, M. T. H. (1992). A model of the self-explanation effect. *The Journal of the Learning Sciences*, 2, 1-59.
- Webb, N. M. (1989). Peer interaction and learning in small groups. *International Journal of Education Research*, 13, 21-39.
- White, B. Y., & Frederiksen, J. R. (1990). Causal model progressions as a foundation for intelligent learning environments. *Artificial Intelligence*, 42, 99-157.