

Qualitative Models in Interactive Learning Environments: An Introduction

Bert Bredeweg¹ and Radboud Winkels²

¹Department of Social Science Informatics (S.W.I.)

University of Amsterdam, Roetersstraat 15

1018 WB Amsterdam (The Netherlands)

Telephone: +31 20 525.6788

Telefax: +31 20 525.6896

E-mail: bert@swi.psy.uva.nl

<http://www.swi.psy.uva.nl/usr/bert/home.html>

²Department of Computer Science & Law

University of Amsterdam, PO Box 1030

1000 BA Amsterdam (The Netherlands)

Telephone: +31 20 525.3485

Fax: +31 20 525.3495

E-mail: winkels@lri.jur.uva.nl

<http://www.lri.jur.uva.nl/~winkels>

Publication details: Introduction to a special issue of *Interactive Learning Environments* on "the use of Qualitative Reasoning Techniques in Interactive Learning Environments", pages 1-18, volume 5, number 1-2, 1998. (Guest edited by Radboud Winkels and Bert Bredeweg)

Introduction

It is a generally held position that the process of learning will improve when learners are given computer-based tutoring programs that allow for interactive access tuned to the specific needs of each individual learner. Computer artifacts for learning should therefore be both interactive and articulated. Interactive learning environments can be seen as engines for education that facilitate learning by having learners interact with a simulation of the subject matter.

Designing, diagnosing and controlling the behavior of 'physical' systems is an important feature of daily human activities, both in professional and non-professional situations. Interacting with physical systems requires comprehension of their behavior, in particular how manipulations of some aspect of the system will effect its behavior. In order to teach these behavioral characteristics, quantitative simulations are often used in computer-based learning environments. However, some behavioral features are hard to communicate by a computer program that is based on such a quantitative simulation. Among others, generating causal explanations of the systems' behavior, reasoning from structure (i.e., deriving the behavior from a given structural description), and qualitiveness in general (i.e., a vocabulary for reasoning about behavior in qualitative terms) cannot be dealt with adequately. A large part of the research on qualitative reasoning originated from efforts trying to cope with the limitations that followed from using quantitative simulators for teaching purposes.

This introduction presents an outline of how qualitative models can be used for interactive learning environments. The first section will discuss in more detail the kind of learning that the contributions in this special issue are concerned with. The second section will elaborate on the typical characteristics of qualitative models. During the past decade a large number of promising results have been achieved by the qualitative reasoning community, whereas at the same time the limitations of the current techniques are well understood. The third section presents a range of research topics that are interesting to pursue with respect to using qualitative models in interactive learning environments. This section is followed by a short overview of the contributions in this special issue, which is split into two separate volumes. The introduction ends with some concluding remarks.

Learning about System Behavior

Learning environments are particularly interesting as artificial worlds with which learners can interact in order to learn about the behavior of certain systems. In this section we will describe in more detail the ideas underlying this approach. First, the notion of 'interacting with the (physical) environment' will be explained. Second, some ideas on what it means to learn about system behavior will be discussed, in particular, the ability to predict and postdict behavior. The third subsection discusses the benefits of using individualized instruction in combination with simulation models. The last subsection explains why qualitative models are particularly useful for this purpose.

Interacting with the Physical Environment

Humans have to interact with their physical environment. They have to drive their cars, operate VCR's, light matches, use telephones, turn on lights, use elevators, open doors, and so on. In order to deal with these different systems humans have to learn about the behavior of each of them. They have to learn how the breaks of a car work, how the buttons must be used in an elevator and how their use differs from buttons on a telephone or a VCR, that matches can be lit by striking them along specific parts of the box, that doors usually have a knob that has to be turned (counter) clockwise or pushed down, and so forth. During their lives people spend a lot of time learning about the behavior of the huge amount of systems they have to interact with. Sometimes this learning takes many years of hard work, such as learning how to fly an airplane or how to operate a power plant.

Researchers in the area of knowledge-based systems have been investigating the different ways in which humans can interact with physical systems for many years now. It turns out that three main categories can be identified, namely: (1) controlling and operating, (2) designing and constructing, and (3) diagnosing and repairing (see for example Breuker & van de Velde, 1994)). In the case of controlling, the goal of the human agent is to interact with the system in such a way that it performs a specific kind of, apparently desired, behavior. Characteristic for this situation is also the fact that the physical structure of the systems does not change, at least not significantly. It is only manipulated in order to have it perform the intended behavior. In the case of designing and constructing, the structure of the system will change. In fact, that is the key notion in this situation. When designing a system, the goal is to build a physical structure such that the desired behavior will be performed by it. Many of the systems that surround

humans to date are indeed constructed by humans. It is common in this respect to make a distinction between systems designed and constructed by humans (artifacts) and systems not made by humans (natural systems). Finally, in the case of diagnosing and repairing a system, the system does not show the desired behavior anymore. The goal of this interaction with a system is to find the cause for the malfunctioning and to repair it. The latter may require manipulating the structure of the system, such as replacing a broken part.

The difference presented here between natural systems and artifacts is of course just a first step in a classification hierarchy. Usually both categories can be further divided according to domain specific features. There are for example biological systems, chemical systems, social systems, economic systems, and so forth. Often the names given to activities concerned with interacting with these systems differ across these domains. Also, there are many possible classifications of systems. For this introduction it is sufficient to know that there are many different systems out there in the world with which humans have to interact in one way or another.

Learning about Systems and their Behavior

The large number of systems poses a big problem to humans, for they are born without knowledge of these systems, and yet, at some point in their lives they have to interact with these systems. One approach would be to have humans interact individually with these systems and let them discover the crucial insights by themselves. This would of course not be a very efficient approach. First of all, it would be very expensive and maybe even traumatic. Think of all the children that know nothing about cars and would get run over by them while trying to discover 'what kind of behavior this system might produce'. Secondly, people can reason about the behavior of a system as it evolves in time. This is of course where education comes in: Teachers spend a lot of time teaching learners about the behavior of systems in the physical world and, more specifically, how to interact with them. The basic idea is that learners have to acquire appropriate models of physical systems and their behavior. These models provide the basis for successful interaction with these systems (cf. Clancey, 1986). From a pragmatic point of view, 'an appropriate model' implies at least two important notions, namely prediction and postdiction (or explanation) of the behavior of some system (cf. Klee, 1984; Forbus, 1984). In order to perform these reasoning capabilities a person has to be able to identify some physical structure as a stand-alone unit (a system) with its own individual behavior. A person should then be able to identify the behaviors of this system that are important to him or her, and be able to either predict how these will change in the near

future, or explain how they came about following some previous behavior. It is important to realize that both prediction and postdiction require causal models of the system's behavior that enable someone to relate some set of behaviors at time t_1 to some set of behaviors at time t_2 . Also important is the fact that the whole notion of an 'appropriate model' is a relative one. Different goals require different models. For someone driving a car, it is sufficient to observe a red traffic light and be able to predict that it will turn green at some point in the near future. If, however, the red light turns off without the green one being lit, the driver should be able to explain that this is probably due to some power failure or to the light bulb being broken. Usually there will be more behavioral cues in the environment to disambiguate the possible interpretations. If for example other traffic lights are still red or green then the chance of a power failure is less likely. The electrician who has to repair the traffic lights has different goals compared to the car driver and therefore uses a different set of (more detailed) models. In an educational context, this means that depending on the goals to be achieved, specific models have to be learned, or taught.

Using Computers and Simulation Models

Today computers allow learners to learn about system behavior in a way that advances the traditional classroom oriented approach in a number of ways. First, there is the notion of individualized instruction. Learning will improve when learners are given computer-based tutoring programs that allow for interactive access tuned to the specific needs of each individual learner (cf. Wenger, 1987). Next to using individualized tutoring, the use of simulations of real-world systems has a number of advantages in itself (cf. Jong, 1991). If, for example, the real system cannot be accessed by the learner then computer models can provide interactive learning environments which are full of genuine stimuli that closely resemble the important characteristics of the original system, particularly if multimedia features are included (cf. Schank & Cleary, 1995). In addition, manipulations may be carried out with the model of the system that are undesirable (too expensive or too dangerous) under normal conditions. Safety critical operations, as for example required for operating power plants, can be carried out as many times as needed without making many additional costs. Also, in the case of a nuclear driven power plant, a series of meltdowns does not provide any danger to the outside world. Simulation models can also be used to build environments that could never exist in reality. Often these impossible worlds, for example no gravity on earth, can be very illustrative and therefore helpful for learners to acquire crucial insights (e.g. ARK, Smith *et al.*, 1987). Computer simulations also allow one to manipulate time, and by doing so speed up or slow down the behavior of some system. This allows learners to

access more global notions of how the behavior of systems evolves in time (e.g. global climate changes or pollution processes) or to study complex and rapidly changing phenomena step by step and in close detail (e.g. chemical reactions or electrical phenomena).

What kind of Models?

Research shows that simulations are only effective when the actions of the learners are monitored by a teacher (human or computer) and guidance is provided (cf. Elsom-Cook, 1990; Hulst, 1996). As we are concerned with computer simulations the question is how to relate tutoring activities to ingredients of the simulation model. In order to connect the two, the simulation model has to provide handles by which it can be accessed for tutoring purposes. This is an old issue already faced in early programs such as Steamer (Hollan *et al.*, 1987) and Sophie (Brown *et al.*, 1982), and very much the basic problem that gave rise to fundamental research areas such as qualitative reasoning (Bobrow, 1984). What does it take to build articulated models (see e.g. Forbus & Falkenhainer, 1992)? Although on the one hand it is easy to understand the essence of building articulated models, this is often not understood by engineers and other highly trained experts in physics and related areas. It is only in automated tutoring situations, when the computer program has to generate and provide feedback to the learner by itself, that one realizes what is missing in a quantitative simulation model for that purpose. It turns out that there is a whole vocabulary and a corresponding reasoning strategy that experts use, which is not available by itself from the quantitative model (Kleer, 1990). Something has to be added in order for the computer to have access to that kind of knowledge. Before going into detail about what has to be added and how, let us first point out the problem by means of a simple example.

In Fig. 1 a set of containers is shown. Each container is closed by a piston and contains an amount of gas.



Figure 1: A container-piston situation

Using Boyle's law, $V \cdot P = n \cdot r \cdot T$, we can easily compute the values for Pressure given different values for Volume and a fixed value for $n \cdot r \cdot T$ as shown in Table 1.

	V	P	$n \cdot r \cdot T$
st 1	1	10	10
st 2	2	5	10
st 3	5	2	10

Table 1: Quantitative values for $V \cdot P = n \cdot r \cdot T$

In fact, all kinds of calculations can be made as long as only one of the values is unknown. Looking at Table 1, we human beings can easily see that increasing values for V are followed by decreasing values for P . However, in the formula $V \cdot P = n \cdot r \cdot T$ there is nothing that captures this notion explicitly. Something has to be added to this equation in order to derive that P and V have some kind of monotonic relationship. At least some procedure is required that compares the values of P and V and comes up with this notion (which is by the way a rather difficult machine learning problem). But for tutoring purposes we do not want to depend on quantitative values, per se. Not only are they unavailable in many domains and specific situations, such an approach also ignores the fact that there exists this rich vocabulary that people use to communicate about the behavior of (physical) systems. In the case of the container-piston example, we would like to have access to a set of primitives represented in a computer program in such a way that it allows the following kind of utterances: "if V increases then P decreases". Moreover, we would like this language to be general and reusable for many different domains, including non-physics domains such as for example economics (Berndsen, 1992) and ecology (Salles *et al.*, 1996). This is where qualitative reasoning comes in. Qualitative reasoning provides a vocabulary (an ontology if one likes) by which computer programs can reason about the behavior of systems in such a way that these computers can communicate about the behavior of these systems with humans. We refer to this notion as 'knowledge communication about system behavior'.

Characteristics of Qualitative Models

During the past decade many important ideas have been presented by the qualitative reasoning community (see e.g. Weld & Kleer, 1990). It is far beyond the scope of this section to come up with a complete overview of all that research. Instead this section will point out some of the main characteristics that are of interest for using qualitative models in interactive learning environments (for more details see also Bredeweg, 1992).

Reasoning from Structure

An important starting point for qualitative reasoning is the notion of 'reasoning from structure'. This means that the behavior of some system is derived by analyzing its structural appearance. An essential step in the construction of a qualitative model is therefore to determine how entities from the physical reality are represented in the model. Two types of abstraction have been given much attention: (1) modeling the physical world as components connected by conduits (Kleer, 1984), and (2) modeling the physical world as a set of physical objects that interact via processes (Forbus, 1984). Each of these abstractions provides specific guidelines according to which systems must be modeled. These guidelines can in addition be used for developing general purpose libraries. As soon as a model has been constructed for a certain part of the physical world, this model can be stored in a library and used again in new situations.

Quantities and changing Behavior

In a qualitative model, behavior is typically represented by quantities, which can be assigned certain qualitative values and have a derivative that specifies a direction of change. The latter may effect the former to change and by doing so represent changing behavior. The values a quantity can have are represented as a quantity space (a set of values, usually consisting of alternating points and intervals). The most general quantity space defines three values for a quantity: negative [-], zero [0] and positive: [+] (see Fig. 2).




Figure 2: A general quantity space

This set of values is applicable to many quantities in many different situations. Take for example an amount of water: A_W . Qualitatively speaking $A_W=[+]$ could mean that there is an amount of water, and $A_W=[0]$ that there is no water ($A_W=[-]$ could then mean that there is a shortage of water). The interpretation of a qualitative value usually depends on the kind of quantity that has been assigned the value. Consider for example a pressure difference between the input and output of a valve. $P_{in-out}=[+]$ could mean that the pressure at the input of the valve is higher, $P_{in-out}=[0]$ could mean that there is no

pressure difference ($P_{in}=P_{out}$) and $P_{in-out}=[-]$ could mean that the pressure at the output is higher.

The general quantity space $\{-,0,+\}$ may not always capture the typical characteristics of a domain in sufficient detail. In those situations a domain specific quantity space is required. This can for example be true for the quantity temperature when it refers to the temperature of some substance. Typically a quantity space as shown in Fig. 3 will then be required in order to model the different aggregation phases (qualitative states) of the substance.




Figure 3: A typical quantity space for physics problems

Choosing the appropriate quantity space for a quantity is often a difficult problem. Obviously, if we refer to the human body temperature we do not want the model to represent solid, liquid and gas. A more useful quantity space for the quantity temperature in a medical domain would be: below normal (low), normal, above normal (high), as for example shown in Fig. 4.




Figure 4: A typical quantity space for medical domains

As mentioned before, derivatives are used to represent the dynamic aspects of system behavior, but again only using qualitative terms. For example: $A_W=[+]$ & $\partial A_W=[-]$, means that in the current situation the amount of water is decreasing. Termination rules are then used to determine whether the behavior in the current situation may change and lead to a new state of behavior. In this situation, the amount of water may become $[0]$ in the next state. This specific inference is based on the limit rule (cf. Kleer, 1984) and reads as follows: IF a quantity has a value and it is decreasing ($\partial=[-]$) THEN this quantity will reach the next lower value from its quantity space. More complex representations, such as using higher order derivatives, have also been discussed (see for example (Kleer & Williams, 1991).

Dependencies and Causality

An important characteristic of qualitative models is the notion of dependencies between quantities and the causality that can be modeled by these. Well understood dependencies are influences, proportionalities (Forbus, 1984) and regular qualitative (in)equalities (cf. Simmons, 1986; Williams, 1988). A simple example as shown in Fig. 5 may help to explain some of the basic ideas. It describes how an energy flow will restore the equilibrium between two objects that differ in temperature.




Figure 5: A simple heat flow process

Forbus refers to this specific partial model (or model fragment) as a heat flow process (notice that a liquid flow process would look the same except for using different quantities: Pressure, Amount, and Flow of liquid). This process is stored in a library of model fragments and applies to any situation (structural description or scenario) in which there appear two objects that differ in temperature (see Fig. 6a). If the description applies, the inequality between the temperatures causes a flow of energy between the objects that increases the amount of energy in the colder object and decreases the amount of energy in the warmer object. These dependencies are modeled by influences (I- & I+). The changes that are caused by these influences are further propagated via the proportional dependencies (P- & P+) that exist between the temperatures and the heats within each object. These changes in temperature will then effect the inequality between the temperatures and lead to a new state of behavior in which the temperatures have become equal.

A heat flow process may apply to many different physical situations as for example shown in Fig. 6 (6a: two objects differing in temperature that are moved towards each

other, 6b: a container-piston assembly containing a gas that is being heated, and 6c: a kettle containing water that is being heated).




Figure 6: Three heat exchanging situations

Depending on that specific situation the effects of the heat flow may be different. It may, for example, be the case that more than one process effects some quantity. In that case the resulting change in this quantity may be ambiguous, or not, depending on the kind of influences introduced by the processes and the magnitude of these influences. See for example Forbus (1990) for more details on this matter. Another aspect is concerned with the type of change that may occur once quantities are increasing and decreasing. In the example of the two objects (Fig. 6a), an inequality between the two quantities changes to an equality. But in some situations, as for example the kettle heating situation in Fig. 6c, this may not happen. It is very likely that the boiler stays much hotter than the water that is being heated by it. Therefore, the water temperature will not become equal to the temperature of the heat source, but instead it will change its value from being in between freezing point and boiling point to being at the boiling point (see also Fig. 3). In the new behavioral state the water will start to boil.

Other ways of deriving causal interpretations have also been proposed. Causal ordering (Iwasaki & Simon, 1986) is probably the best known in this respect. Instead of having specific causal interpretations attributed to certain types of dependencies, this method uses the order in which (in)equalities are used by the equation solver as the basis for the causal interpretation. In Top & Akkermans (1991) this is referred to as mathematical causality as opposed to physical causality discussed above.

Constructing a Running Model

Knowledge about the behavior of partial physical structures, such as the above described collections of objects and their heat exchanges, can be stored in a library of model fragments. Also the rules for determining state changes can be stored. Together

with an initial structural description (scenario) they can be presented as input to a qualitative simulator (prediction engine) and by doing so have this simulator construct a 'running model' of the system. This is shown in Fig. 7. Provided with a structural description the qualitative simulator will try to find model fragments that apply to that situation. In the figure this is referred to as specification (or classification).

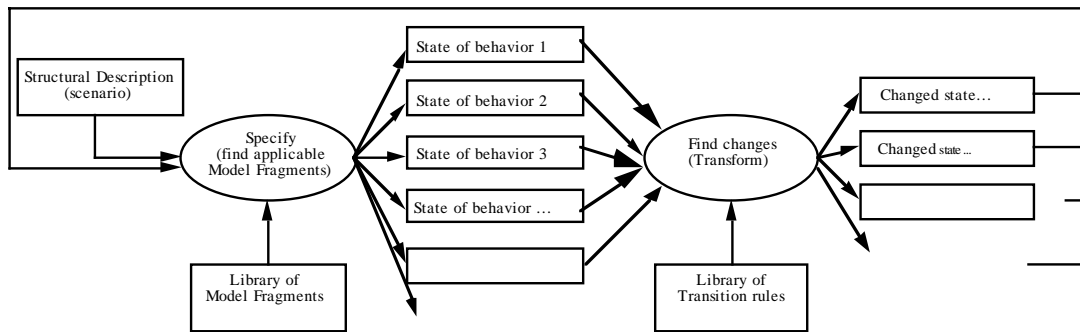


Figure 7: Inferring states of behavior

Possibly this inference process may lead to more than one state of behavior for the system. After this the qualitative engine will look for changes in the current state(s) of behavior. Given these changes, some of the applicable model fragments may not apply anymore and therefore the specification inference has to be done again. This inference cycle may continue until no new changes are found by the qualitative simulator. A possible output of a qualitative simulator for the kettle heating scenario is shown in Fig. 8.

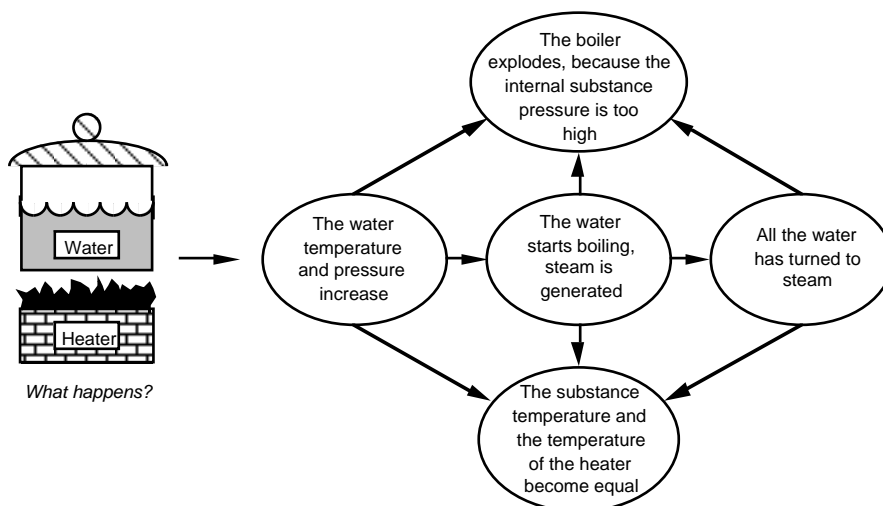


Figure 8: Behavioral states for the kettle heating scenario

For a tutoring situation a 'full' prediction of all possible behaviors of a system is usually not needed. Instead, specific trajectories of possible behaviors are often more useful. In the literature on qualitative simulation this issue has not been given much attention. It is

one of the problems that has to be tackled in order to use qualitative models for tutoring purposes.

Using QR models in ILE

Qualitative techniques provide powerful means to construct cognitive engines for education. This section discusses some of the main research areas that have to be addressed in order to successfully use qualitative techniques for that purpose.

Selecting and sequencing subject matter

One set of problems that has to be tackled concerns the things the learning environment will present to the learner(s) to interact with. This may include issues such as global curriculum planning, model selection for each of the topics that the learner has to acquire knowledge about, and detailed specifications of the exercises or assignments that the learner has to carry out. Relevant from a qualitative reasoning point of view is that all the decisions on these issues, and the automated procedures following from that, will be grounded in terms of a qualitative ontology as much as possible. In order to clarify what is meant by this, imagine a causal model that may be generated by a qualitative engine for the container-piston assembly shown in Fig. 6b. Usually such a model is too big to be comprehensible for a learner in one go. What is needed in such situations is a procedure that guides the learner through the model in smaller and understandable steps. For this purpose the ontology of qualitative models can be used. It can be regarded as a kind of meta language (cf. Harmelen, 1991) in which the steps can be defined. Consider the following procedure:

1. Focus on the quantities of objects and assemblies
2. Follow the causal paths
3. Focus on processes (and follow causal paths)
4. Consider transition to successive states of behavior (and repeat steps)

The rationale behind this procedure would read as follows: learners must first identify the objects in the simulated system and their important quantities. Next, or while considering a specific object, they have to learn about the causal dependencies for each object and aggregate. Depending on the specific causal structure for the container-piston assembly, learners will be confronted by the environment on issues concerning the quantities (and objects) relevant to: (1) the gas in the container closed by a piston, (2)

the air surrounding this assembly and (3) the heat source. The third major step concerns the introduction of the processes. In the container-piston assembly there are at least four processes. One heat flow process from the heat source to the gas in the container, one from this gas to the surrounding air, one from the heat source to the surrounding air, and one 'move' process effecting the position of the piston. The typical steps for discussing each of the processes with the learner are:

1. Find the appropriate inequality between the involved quantities
2. Discuss the flow (in the example: energy) and the initial changes (via influences) that follow from that
3. Use the causal paths to explain how the initial changes are further propagated (via proportionalities) and possibly restore the equilibrium

The last step in dealing with processes relates to the fourth step of the previous procedure. The changes represented by the derivatives of the quantities may lead to new states of behavior either because inequalities between quantities may change or because quantities take on different values from their quantity spaces.

This procedure of selecting and sequencing subject matter is only a first step in showing how the qualitative ontology can be used as a basis for driving the interactive learning environment. Many additional problems have to be solved, such as determining how to automatically generate questions and assignments to present to the learner. In dealing with these issues we can also draw from other, sometimes more traditional, approaches to instructional design. The project SMISLE (Jong *et al.*, 1994) for example, presents a set of general assignments that can be used for different situations. However, many of these approaches are not automated, i.e. similar to traditional CAI, the interaction with the learner has to be specified over and over again for each new situation. An interesting question is how we can ground these approaches, and thus automate them, by using a qualitative ontology.

Some interesting publications in this direction have been published. White & Frederiksen (1990) present their causal model progression as an approach to specifying the notion of problem complexity in terms of a qualitative model. Falkenhainer & Forbus (1991) try to solve another part of the same puzzle with their compositional modeling approach. Here the question is how to find the minimum set of model fragments needed for simulating the behavior of a system and to answer adequately a specific question. Sime (1995) tackles the problem from yet a different angle. She discusses how a specific approach to learning can be operationalised in terms of a qualitative ontology (see also this issue).

It is important to realize that not all didactic aspects can be grounded in a qualitative ontology. For example, when proceeding through successive states of a simulation, how often should a specific causal structure that reappears in a number of states be discussed with a learner? After a certain number of times the learner may have understood the phenomena and may want to move on. If the learning environment keeps bringing up this specific causal structure the learner will get bored or irritated and lose interest for the environment. However, people tend to forget things and, in order to prevent this, important issues have to be rehearsed sufficiently. It is not likely that the qualitative ontology will provide much leverage for determining how to deal with these issues. Research from other areas such as cognitive and learning sciences have to be applied for this purpose. However, from a pragmatic point of view we can circumvent the problem by giving an amount of navigation control to the user. For example, after having correctly processed a causal structure once, the learner may skip it in successive states of behavior. In fact, finding an optimal point in using both the learner's own control and the assignments set by the environment provides an interesting research question.

Cognitive Diagnosis

When the learning environment is not one of free exploration, but meant to guide the learner towards some kind of understanding and mastering of the domain, the system will have to be able to detect missing knowledge or misconceptions on the part of the learner. Relying on the learner himself to ask for additional information is usually not enough (they are often unaware of their knowledge gaps), but even in the case of a request for help, the system will in most cases have to do some reasoning in order to decide what exactly it is the learner needs to know. In order to detect or infer learners' needs for additional information, the system first of all needs to know the 'correct' conceptions and skills; usually this is represented in the domain model. Next, it has to use this 'correct' or *norm* model to either interpret the learner's request for help, or to compare the learner's answer or action to it, to see whether that deviates from the norm. When the learner's behavior deviates from the norm, it is an indication for missing knowledge or a misconception on the part of the learner (especially when accompanied by a request for help). In that case the system could try to pinpoint the exact piece of knowledge that is missing or incorrect. This process is usually referred to as (cognitive) diagnosis. The outcome of this diagnostic process should be either a lack of knowledge or a misconception of the learner (cf. Winkels, 1992; Bredeweg & Winkels, 1994). For diagnostic methods we can turn to the field of model based diagnosis, as some people have suggested (cf. Self, 1992; Bredeweg & Breuker, 1993).

In model based diagnosis, one can distinguish two processes: generating hypotheses and testing them (cf. Davis & Hamscher, 1988). Most teaching systems that do diagnosis use decision tree-like structures, where tests are directly associated with hypotheses. In the case of an explicit and cognitively plausible domain model, for instance a qualitative model, we can use this norm model for generating hypotheses in a more systematic way. There are basically two ways to do that, depending on (our understanding of) the domain: by ‘decomposition’ and by ‘specialization’. The first method assumes a hierarchical structure of parts or components at the domain level. In principle the diagnostic process is very simple: partition the system according to its decomposition and eliminate the correctly functioning parts by testing. The second approach, generating hypotheses by ‘specialization’, requires less structure in the domain, and works by descending taxonomies (e.g. as in heuristic classification; Clancey, 1985). Hypotheses can be tested against the data of the actual behavior of the learner, previously acquired data about the learner's knowledge and competence at the task (as reflected by a learner model), or by presenting a new problem to the learner that will discriminate between competing hypotheses. In principle, there are two solutions for picking the next problem: either selecting the best (most discriminating) problem from a stored set, or constructing a critical test on the basis of the current hypothesis set. The outcome of this diagnostic process is a ‘faulty component’. What can be wrong with each of these elements? In analogy with troubleshooting of artifacts it can be said that either the particular component does not work at all (the learner does not possess or cannot retrieve the particular knowledge), or it functions incorrectly (the learner has a misconception). The first case is relatively easy to check (we are not concerned with computational tractability for the moment). The second case is more difficult. If the correct version of the knowledge is replaced by something else, what is it replaced with? A pragmatic solution would be to provide the system with ‘fault models’ that reflect (common) misconceptions (as e.g. proposed by Koning *et al.*, 1996).

Take the container-piston situation in Fig. 6b as an example. A learner is asked to predict what will happen. Suppose he or she predicts the piston will move outwards (to the left), i.e. the volume of the gas will increase. “And next?” we might ask. “It will keep on moving outwards until we turn off the heat source, or it drops from the container”, the learner replies. This is certainly a likely possibility, but only one of the two. Another possibility the qualitative reasoning engine comes up with, is that the piston will stop moving at a certain moment, because the flow of energy from the source to the gas will be equal to the flow of energy from the gas to the outside world. The pressure of the gas will therefore not increase anymore, and the piston will come to a hold. One possible explanation (hypothesis) of why the learner does not see this

possibility, is that the learner misses the heat flow process between the gas and the world in his or her model altogether. Another hypothesis is that the learner just forgot about the possibility that the two flows of energy might be equal. When the learner model does not help us to distinguish between the two possibilities, we may decide to ask the learner a question, or to present a new problem that will resolve the ambiguity, or perhaps even reveal or trigger the missing part of his or her model to the learner. We might ask the learner to tell or show us all heat flow processes in the situation. If the learner only indicates the one between the source and the gas, we may try to direct the learner's attention to the increasing temperature of the gas and the difference with the outside temperature. We may also opt for the presentation of a new 'problem'. It would have to focus on the relation between the two heat flow processes, for instance one in which the heat source is very small.

Whatever the solution, the important thing here is that the qualitative 'norm' model enables us to detect a possible lack of knowledge or misconception on the part of the learner, and can be used to suggest likely candidates for them as well. Many problems still remain in this diagnostic process, e.g. finding the right level of abstraction for hypotheses, ensuring the cognitive tractability and plausibility of the norm models (cf. Koning & Bredeweg, this issue), managing the search space for the diagnostic engine, generating or finding fault models, constructing critical tests, etc.

Generating Explanations

Qualitative domain models facilitate the generation of explanations in several ways. First and foremost, as was mentioned before, they provide a vocabulary and conceptual framework to talk and think about the domain. They deal with 'quantities', 'components', 'processes', 'influences', etc., instead of with mathematical functions. Most qualitative models also provide direct access to the structure of the system it represents. This can be used to describe the physical structure in terms of components and sub-components, and their function, to the learner. As a mechanism for generating explanations for these static model elements, one could use schemata of rhetorical predicates, as described by McKeown (1985) for database objects. She identified several answer schemata for specific question types of users about these objects. An example is the 'identification' schema for providing definitions, that includes the description of 'type' and 'constituency' relations, 'concept attributes', and will give an example. For instance, for the kettle example in Fig. 6c, an identification schema could be used to describe the concept of a 'contained liquid' when the learner asks about it,

such as: “A contained liquid is a liquid that is contained by a container. An example is contained-liquid1 that is contained by container1.”.

Besides the static model elements, the reasoning on the basis of the models (pre- and postdiction) can be used to explain ‘how’ the system works and ‘why’ it behaves as it does (causality). Provided that the qualitative model has (some) cognitive plausibility (i.e. can serve as *mental model*, cf. Gentner & Stevens, 1983), these causal explanations can be used in teaching (cf. Bredeweg & Schut, 1991). For instance, the model of Fig. 5 can be used to explain what will happen to the temperatures of the two objects: a flow of energy will occur which will cause an increase of energy in the colder object and a decrease of energy in the warmer object (influences), which in turn will lead to a rising temperature of the colder object and a dropping temperature of the warmer object (proportionalities), until both temperatures are equal. For such a relatively simple model, the process of explaining a causal chain is straightforward, but for more complex models the reasoning may contain ambiguities, and the chains will be far too long to communicate to learners. Causal chains at higher levels of abstraction (a larger grain size) are then needed. This could for instance be achieved by a chunking process, where paths that do not branch (i.e. there are no alternatives for what can happen) are collapsed into fewer, or even one step. In the example given above, one could decide to take the influence and the proportionality in one go, and state that the flow of energy will ‘cause’ a drop of temperature in the warmer object, and a rising temperature in the colder object. One step further would be to just state that the flow of energy will ‘cause’ the temperatures to become equal. Later on, one can always expand (parts of) the causal chain to a more detailed level. What is needed for these types of explanations are more dynamic mechanisms that are driven by communicative or instructional goals and *intentions* (c.f. Winkels, 1992; Vadillo *et al.*, this issue), rather than the static, *content* driven approach of McKeown (1985).

An interesting use could be made of the library of model fragments for explanations, namely creating *analogies* (cf. Gentner, 1983). When a learner has been confronted with the heat flow process of Fig. 5, we could use structure mapping to find the analogous process of liquid flow and use it either to explain the heat flow (e.g. when the learner already understands the liquid flow process), or to introduce the liquid flow process once the learner understands the heat flow process. Suppose for instance that a learner holds the belief that a heat flow process, as depicted in Fig. 5, terminates when Heat1 and Heat2 are equal, instead of when the temperatures are equal. This may show up in an example where one object is much larger than the other object. One could use an analogy with the liquid flow process to show the consequences of that belief. Will there be a flow of water in the situation of a large container with water connected through a

tube to a much smaller container of water, where the fluid heights are equal in both containers (i.e. the pressures are equal)?

Of course, the story is not as simple as the example seems to indicate. Not all analogies are based on structure mapping, and not all analogies one can find will be useful in an educational setting. Further research will have to unravel other mechanisms for finding analogies and deciding on their appropriateness (cf. Hofstadter, 1995).

Authoring and Model Construction Support

Authoring for tutoring purposes, and model construction in general, is an area of research that is largely ignored by the qualitative reasoning community (cf. Schut & Bredeweg, 1996). As a result there is hardly any 'easy-to-use' software available to support authors during the process of model construction. Two areas of research can be pointed out. One concerns the development of interfaces that support the task of constructing qualitative models by an author (and the learning environment that should accompany this). This line of research is closely related to studies in human computer interaction (cf. Preece *et al.*, 1994). The second area concerns the construction of tools that automate certain subtasks of the overall authoring task. Consider for example tools that help the author to diagnose and repair a buggy model. Also machine learning techniques can be used for this purpose (cf. Bratko *et al.*, 1992; Mozetic *et al.*, 1990).

Improving Qualitative Simulators

Many researchers in the qualitative reasoning community are concerned with improving the reasoning capabilities of the qualitative simulators. This is typically what the majority of the papers presented at the annual qualitative reasoning workshop deal with (cf. Iwasaki & Farquhar, 1996; Bredeweg, 1995). An interesting topic that gets much attention is the integration of qualitative and quantitative simulators (cf. Forbus & Whalley, this issue). Lately, within the community an awareness has emerged concerning the importance of task-level reasoning and other more goal directed and applied research questions. The community is more aware of the fact that the anticipated use of the qualitative techniques to a large extent determines the specific improvements that are required. Some extensions specific for tutoring situations have been described (see for example: Bredeweg *et al.*, 1995; Falkenhainer & Forbus, 1991; Weld, 1988). There is however still a large area of research to be covered.

How to Teach?

In this introduction to using qualitative techniques as the basis for intelligent simulation environments we have not made commitments to a specific style of teaching. Using a qualitative model does in itself not require a specific choice on this matter. On the one hand it is possible to construct an interactive environment which allows learners to freely explore the subject matter. On the other hand, more restricted approaches, such as guided discovery, or fully tutoring are also possible. Also possible are co-operative and distributed forms of learning and alternative styles of teaching, such as having learners design certain behavioral artifacts or diagnose errors. It would in fact be very interesting to have more research focusing on questions related to how the qualitative ontology can be used to support these different styles of learning and teaching.

Contributions in the two issues

This issue contains eight contributions, divided over two separate volumes. All contributions deal with some of the research topics described in the previous section. We will shortly introduce them.

Issue number 1 (this issue)

Forbus & Whalley describe CyclePad, a learning environment for analyzing and designing thermodynamic cycles. CyclePad is focused on quantitative analysis of thermodynamic cycles, but qualitative reasoning is used to rule out physically impossible designs. It is a good example of a working and efficient system that combines quantitative and qualitative modeling for teaching by simulation.

Frederiksen & White argue that causal, typically qualitative models are of particular importance for obtaining transferable expertise. These models are general enough to be applicable across certain domains, and yet specific enough to be powerful and allow for successful and efficient reasoning. Frederiksen & White demonstrate the use of such causal models in ILEs for electronic troubleshooting, and end with some instructional implications for the design of computer based learning environments.

Hartley describes how learners acquire knowledge in certain domains of physics using computer-based tools. In particular, how learners can be guided in changing their mis-

or pre-conceptions about some phenomena, by having them construct qualitative explanatory models. An experimental study shows that learners did change their ideas when the interaction was focused on causal reasoning.

Koning & Bredeweg go on to examine whether qualitative models provide the right vocabulary and reasoning for a dialogue between teachers and learners. In other words, are qualitative models good candidates for normative, prescriptive models to be used for teaching? It turns out that their model forms a good basis to start from, but needs to be extended to allow for task level reasoning which is required for tutoring purposes.

Issue number 2

Michau et al present an interesting use of QR techniques, which is slightly different from the typical 'reasoning from structure' view held usually. It deals with deriving qualitative features from graphs showing quantitative behavior simulation of a system. It is essential for control engineering to understand such graphs in terms of critical qualitative features. These curve classifications are then used in teaching qualitative estimations of process performance. Their AUTO-DIDACT learning environment contains a module (ANAIS) that is able to do qualitative curve interpretation and can show the result to the learner, or guide the learner through the interpretation process.

Ploetzner & Spada hold the position that both qualitative and quantitative models are necessary for successful problem solving, at least in physics. They suggest the use of qualitative problem representations to facilitate the construction of quantitative problem representations in two ways. First, they can enable the derivation of additional quantitative information, and secondly, they can provide constraints to be met by quantitative representations. Their program SEPIA demonstrates this by solving quantitative problems in the domain of classical mechanics successfully and more efficiently by first forming a qualitative representation of them. The cognitive simulation model SEPIA also suggests how and where qualitative misconceptions effect problem solving.

Sime describes her work on MS-PRODS, a learner-centered learning environment to promote better understanding of a process rig. Emphasis is on the use of different domain models, both quantitative and qualitative to achieve that understanding. She introduces seven dimensions to classify the different domain models, and a mechanism to progress through the models, based on these dimensions. The system could use several strategies for model progression. Sime's interest is mainly in the Cognitive Flexibility Theory (CFT) (Spiro & Jehng, 1990) to guide the sequencing. Contrary to

other work, e.g. QUEST (White & Frederiksen, 1990), CFT suggests the alternating use of different models to foster better understanding of the domain, instead of moving from simple to more complex models.

Vadillo *et al.* concentrate on the potential use of qualitative models for generating explanations to help users or learners to learn a domain. They present an extension of the INTZA system, a tutoring environment for industrial training situations. In order to provide good behavioral explanations for simulations in INTZA, they extend domain models with a qualitative, causal viewpoint. The causal model is obtained by applying causal ordering (Iwasaki & Simon, 1986) to the set of differential equations that describes the system. When the user asks for help, or when an error of the user in performing a task with the simulation has been detected, an explanation involving a behavioral description of the system is often required. The causal model is used to generate the content of that behavioral description, and domain independent explanation strategies are used to present that information to the user in a way similar to that used in Eurohelp (Winkels, 1992).

Concluding Remarks

We have argued that it is useful for people to learn to understand and handle systems in the physical world through interacting with computer based simulations. Traditionally, these simulations are based on mathematical models, that are very efficient, but have several shortcomings when it comes to explaining them to relative novices. This formed one of the inspirations for the research in qualitative reasoning. Qualitative domain models have a number of features that make them very interesting for use in Interactive Learning Environments (ILEs). They provide us with a conceptual framework to talk and reason about the physical world. The models are explicit and articulated representations of systems and their behavior, and can therefore be used for causal explanations. We have addressed several research topics related to interactive learning environments, and tried to show that qualitative models have great potential, notably in the areas of selecting and sequencing subject matter, cognitive diagnosis, and explaining domain models and their behavior to learners. We have outlined some ideas of how these tasks could be aided when the interactive simulation is based on a qualitative model, and suggested research questions to be addressed in the future. Furthermore, we have mentioned other interesting areas of research related to the use of qualitative models in ILE. The contributions in this special issue will explore some of the issues introduced in more depth.

Acknowledgment

We would like to thank Joost Breuker, Ken Forbus, and Barbara White for helping us in reviewing the submissions to this special issue. We would also like to thank Kees de Koning and Paulo Salles for proofreading this introduction; as always, all remaining mistakes are ours.

References

- Berndsen, R. J. (1992). *Qualitative Reasoning and Knowledge Representation in Economic Models*. PhD thesis, University of Brabant, Tilburg.
- Bobrow, D.G. (ed.) (1984). *Qualitative Reasoning about Physical Systems*. Elsevier Science, Amsterdam, The Netherlands.
- Bratko, I., Muggleton S. & Varsek, A. (1992). Learning qualitative models of dynamic systems. In S. Muggleton (ed.). *Inductive logic programming*, pages 437-452, Academic Press, London.
- Bredeweg, B. (ed.) (1995). *Working papers of the Ninth International Workshop on Qualitative Reasoning*. University of Amsterdam, Amsterdam, The Netherlands.
- Bredeweg, B. (1992). *Expertise in qualitative prediction of behaviour*. PhD thesis, University of Amsterdam, Amsterdam, The Netherlands.
- Bredeweg, B. & Breuker, J.A. (1993). 'Device Models' for Model-Based Diagnosis of Student Behaviour. In *Proceedings of AI-ED-93*, pages 441-448.
- Bredeweg, B., Koning, K. de & Schut, C. (1995). Modelling the Influence of Non-Changing Quantities. In J. Wainer and A. Carvalho (eds.). *Advances in Artificial Intelligence (SBIA'95)*, pages 131-140, Springer-Verlag, Berlin.
- Bredeweg, B. & Schut, C. (1991). Cognitive plausibility of a conceptual framework for modeling problem solving expertise. *Proceedings of the 13th Conference of Cognitive Science Society*, Lawrence Erlbaum, Hillsdale, New Jersey, pages 473-479.
- Bredeweg, B. & Winkels, R.G.F. (1994). Student Modeling through Qualitative Reasoning. In J.E. Greer & G. McCalla (Eds.). *Student Modelling: The Key to Individualized Knowledge-Based Instruction*. Springer Verlag, Berlin, pages 63-97.
- Breuker, J.A. & Velde, W. van de (eds.) (1994). *The CommonKADS Library for Expertise Modelling*. IOS Press, Amsterdam, The Netherlands.
- Brown, J. S., Burton, R. R. & Kler, J. de (1982). Pedagogical, natural language and knowledge engineering techniques in SOPHIE I, II and III. In D. Sleeman & J. S. Brown (Eds.), *Intelligent Tutoring Systems*. Academic Press, New York, pages 227-282

- Clancey, W.J. (1986). Qualitative Student Models. *Annual Review of Computer Science*, 1:381-450.
- Clancey, W.J. (1985). Heuristic Classification. *Artificial Intelligence*, 27:289-350.
- Davis, R. & Hamscher, H. (1988). Model-Based Reasoning: Troubleshooting. In H.E. Shrobe (ed.). *Exploring Artificial Intelligence*. Morgan Kaufmann, San Mateo, CA, pages 297-346.
- Elsom-Cook, M. (1990). Analysis of a tutoring dialogue. In M. Elsom-Cook (ed.). *Guided discovery learning: A framework for ICAI research*, pages 113-131. Chapman, London.
- Falkenhainer, B.C. & Forbus, K.D. (1991). Compositional Modeling: Finding the Right Model for the Job. *Artificial Intelligence*, 51, 95-143.
- Forbus, K.D. (1990). The Qualitative Process Engine. In D.S. Weld and J.H. (Eds.). *Readings in Qualitative Reasoning about Physical Systems*,. Morgan Kaufmann, San Mateo, California, pages 220-235.
- Forbus, K.D. (1984). Qualitative process theory. *Artificial Intelligence*, 24:85-168.
- Forbus, K.D. & Falkenhainer, B. (1992). Self-explanatory simulations: Scaling up to large models. In *Proceedings of AAAI-92*, pages 685-690.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.
- Gentner, D., & Stevens, A. (1983). *Mental models*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Harmelen, F. van (1991). Meta-level Inference Systems. *Research Notes in AI*. Pitman, Morgan Kaufmann, London, San Mateo California.
- Hofstadter, D. (1995). *Fluid Concepts and Creative Analogies*. Basic Books.
- Hollan, J.D., Hutchins, E.L., & Weitzman, L. (1987). STEAMER: An interactive inspectable, simulation-based training systems. In G.Kearsley (ed.). *Artificial intelligence and instruction: applications and methods*. Addison-Wesley, Reading (Mass), pages 113-134.
- Hulst, A. van der (1996). *Cognitive Tools. Two exercise in non-directive support for exploratory learning*. PhD thesis, University of Amsterdam, Amsterdam, The Netherlands.
- Iwasaki, Y. & Farquhar, A. (Eds.) (1996). *Qualitative Reasoning: The Tenth International Workshop*. AAAI Technical Report WS-96-01, Stanford Sierra Camp, California, USA.

- Iwasaki, Y., & Simon, H A. (1986). Causality in device behavior. *Artificial Intelligence*, 29, 3-32.
- Jong, T. de (ed.) (1991). Computer simulations in an instructional context. *Education and Computing (Special issue)*, 6.
- Jong, T. de, Joolingen, W. van, Hoog, R. de, Lapied, L., Scott, D.& Valent, R. (1994). SMISLE: System for Multimedia Integrated Simulation Learning Environments. In: T. de Jong & L. Sarti (Eds). *Design and Production of Multimedia and Simulation-based Learning Material*. Kluwer Academic Publishers.
- Kleer, J. de (1990). Qualitative physics: A personal view. In D.S. Weld and J. de Kleer (Eds.). *Readings in Qualitative Reasoning about Physical Systems*, pages 1-8. Morgan Kaufmann, San Mateo, California.
- Kleer, J. de & Brown, J.S. (1984) A Qualitative Physics Based on Confluences. *Artificial Intelligence*, 24, 7-83.
- Kleer, J. de & Williams, B.C. (1991). Qualitative reasoning about physical systems 2. *Artificial Intelligence*, 51.
- Koning, K. de, Bredeweg, B. & Breuker, J. (1996). Interpreting Student Answers: More than Diagnosis Alone. In *Proceedings of Euro AI-ED*, pages 233-239.
- McKeown, K.R. (1985). Discourse strategies for generating natural-language text. *Artificial Intelligence*, 27:1-41.
- Mozetic, I., Bratko, I. & Urbancic, T. (1990). Varying Levels of Abstraction in Qualitative Modelling. In J.E. Hayes, D. Michie and E. Tyugu (Eds.). *Machine Intelligence 12:259-280*, Clarendon Press, Oxford.
- Preece, J., Rogers, Y., Sharp, H., Benyon, D., Holland, S. & Carey, T. (1994). *Human-Computer Interaction*. Addison-Wesley, Reading Massachusetts, USA.
- Salles, P.S.B.A, Pain, H. & Muetzelfeldt, R.I. Qualitative Ecological Models for Tutoring Systems: a Comparative Study. In P. Brna, A. Paiva and J. Self (eds). *Proceedings of EuroAIED95*, Fundacao Calouste Gulbenkian, Lisbon, pages 226 - 232.
- Schank, R.C. & Cleary, C. (1995). *Engines for education*. Lawrence Erlbaum, Hillsdale, New Jersey.

- Schut, C. & Bredeweg, B. (1996). An overview of Approaches to Qualitative Model Construction. *The Knowledge Engineering Review*, 11(1): 1-25, 1996.
- Self, J. (1992). Cognitive diagnosis for tutoring systems. In *Proceedings of ECAI-92*, pages 699-703.
- Sime, J.A. (1995). Model Progressions and Cognitive Flexibility Theory. In J. Greer (ed) *Proceedings of AIED95*. AACE: Charlottesville, VA, USA, pages 493-500.
- Simmons, R. (1986). Commonsense arithmetic reasoning. *Proceedings of the AAAI*, pages 118-124.
- Smith, R.B., O'Shea, T. & Scanlon, E. (1987). Building and Using Alternative Realities for Physics Education. In *Abstracts of the 3rd AIED Conference*, page 52.
- Spiro, R.J. & Jehng, J-C. (1990) Cognitive Flexibility and Hypertext: Theory and Technology for the Nonlinear and Multi-dimensional Traversal of Complex Subject Matter. Chapter 7 in D. Nix & R.J. Spiro (Eds.) *Cognition, Education and Multi-Media: Exploring Ideas in High Technology*. Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- Top, J.L. & Akkermans, J.M. (1991). Computational and physical causality. In *Proceedings of the IJCAI*: Morgan Kaufmann, USA, pages 1171-1176.
- Weld, D. (1988). Comparative Analysis. *Artificial Intelligence*, 36:333-373.
- Weld, D., & Klerer, J. de (Eds.) (1990). *Readings in qualitative reasoning about physical systems*. Palo Alto, CA: Morgan Kaufmann Publishers.
- Wenger, E.. *Artificial intelligence and tutoring systems. Computational and cognitive approaches to the communication of knowledge*. Morgan Kaufmann, Los Altos, California.
- White, B. & Frederiksen, J. (1990). Causal model progressions as a foundation for intelligent learning environments. *Artificial Intelligence*, 42, 99-157.
- Williams, B.C. (1988). MINIMA, A Symbolic Approach to Qualitative Algebraic Reasoning. In *Proceedings of the AAAI-88*, pages 264-270.
- Winkels, R.G.F. (1992). *Explorations in Intelligent Tutoring and Help*. IOS Press, Amsterdam, The Netherlands.