This is a draft of Major, N., Ainsworth, S.E., & Wood, D.J. (1997) REDEEM: Exploiting Symbiosis Between Psychology and Authoring Environments, International Journal of Artificial Intelligence in Education 8(3/4), 317-340.

REDEEM: Exploiting Symbiosis Between Psychology and Authoring Environments

Nigel Major, Shaaron Ainsworth and David Wood

ESRC Centre for Research in Development, Instruction & Training, Department of Psychology,University Park, University of Nottingham, Nottingham, NG7 2RD, United Kingdom.
Email: sea@psychology.nottingham.ac.uk

Abstract

Although Intelligent Tutoring Systems (ITSs) have been shown to be effective in a number of domains, they are still rarely found outside the laboratory. Authoring environments that speed up the process of ITS development provide a partial solution to this problem. However, we propose that to be truly effective, authoring environments must be based on a psychologically plausible account of teaching. In this paper, we describe how the authoring environments can exploit the symbiotic relationship between psychology and authoring environments, which in turn leads to potential advances not just in instructional practice through the deployment of better tutoring systems, but also in instructional theory. This is illustrated by a description of the REDEEM software suite. As authoring environments are crucially dependent on teachers' expertise, we outline how REDEEM can be used to explore that expertise and even improve it.

Introduction

Turbulent Times for ITS

Intelligent Tutoring Systems have long promised significant improvements in learning outcomes. Their ability to model student behaviour offers the chance to give individualised instruction to students when a teacher is not available. Research has shown that individual training significantly improves student performance over classroom learning (Bloom, 1984). Individualised computer tutors are beginning to produce similar effects (*e.g.* Shute, 1990).

During the 1980s a number of important ITSs were developed (*e.g.* PROUST; Soloway & Johnson, 1984; SHERLOCK; Lesgold *et al.*, 1992) which showed impressive performance in particular domains, but as the decade came to a close the tide was already beginning to turn against this apparent miraculous use of artificial intelligence (AI) in education.

Although the reasons were never clearly defined, research teams began to turn their attention from ITS to interactive learning environments. These offer students creative and realistic software worlds, rich in multimedia experiences, in which to experiment and learn. One potential reason for this has been the difficulty in developing an ITS even in a relatively limited domain. Developing an ITS is estimated to take more than 200 hours to produce an hour of instructional material (Woolf & Cunningham, 1987). In addition, time taken to reconfigure the teaching style of a traditional ITS could approach the time needed for the original development. Consequently, to cover any reasonable school curriculum with ITS-based teaching would take many years of development effort. Even if this were possible, there has been concern amongst teachers that ITSs would not reflect their own pedagogic concern, embodying rather the beliefs of the system designers (Major, 1995).

One of the recent successes for intelligent tutoring systems has been the LEAP system (Bloom *et al.*, 1995). LEAP is an intelligent coached apprenticeship environment which offers multimedia training in corporate customer contact. Having been successfully evaluated, LEAP was reported to be on the point of full-scale deployment in a large telecommunications company. The authors however summarised their work by pointing out that LEAP needs authoring tools, usable by subject matter experts, which would allow general courses to be built and could change the style of training.

A partial panacea to the difficulties of ITS did indeed begin to emerge with the advent of authoring environments. These allowed appropriate domain material to be developed together with the ability to have teachers create their own teaching strategies. The early authoring environments included the Instructional Design Environment (Russell *et al.*, 1988), GTE (Van Marcke, 1992), KAFITS (Murray & Woolf, 1992) and COCA (Major, 1995). An evaluation of the authoring tools in COCA (Major, 1994) showed that despite offering considerable power to teachers, there remained a gap between the kinds of interfaces teachers would be prepared to use and the AI-based representations that authoring tools required them to manipulate. Consequently, REDEEM reduces the teacher's opportunities to modify low level instructional behaviour in favour of improving the ease of authoring to reduce the overhead for the teacher.

Even supposing these practical difficulties can be overcome, flexible tools can allow too many questions to be asked. So the right tools must be able to ask the right questions. Accordingly, these tools should be based on a scientifically and psychologically motivated model. In this paper we present the REDEEM tools, which offer a principled solution to the problems we have been considering.

How Psychologically Motivated Authoring Tools Might Help Through Reuse

Authoring tools can clearly help in terms of reducing ITS development time. They usually contain a shell which allows reuse of the essential decision-making machinery which delivers adaptable instruction. Further, they also offer reusable instructional components of varying degrees of complexity.

At the domain level, reusability is no less important, but much more difficult. The older ITSs would have had expert systems operating in the domain, with problem solutions found by interrogating that expert system. But expert system development is complex and costly, and so does not offer easy reusability. In contrast, traditional computer-based training (CBT) alone may be essentially incapable of meeting a teacher's learning objectives. Without instructional knowledge, CBT doesn't easily adapt to teachers' specific learning goals. Ideally, the same domain material should be able to be used for a wide range of objectives and users. Sometimes, teachers may wish to keep a large amount of control over the type and order of the material presented (when learners are new to a domain) but allow much more freedom when students are revising. To solve these problems, an ideal compromise between ITSs with domain level intelligence and CBT is to move the expertise to the instructional level. This offers a more intelligent use of an ITS's domain level and allows courseware to be configured to suit a wider range of instructional objectives.

In order to exploit the advantages of this approach to the creation of ITSs, the task of authoring must be simple enough that a trained teacher without experience in artificial intelligence would be able to use the system effectively. Although relatively little research has been undertaken to assess the difficulties of these tasks for different groups of authors (*e.g.* expert and trainee teachers), what evidence we have suggests that even experienced teachers who are highly computer literate find these demands hard (Major, 1994). To achieve usability at the instructional level, the authoring system must offer instructional control which provides the teacher with the key decisions in a way that can easily be manipulated. It must provide the critical dimensions of a strategy without offering so many instructional choices that the task becomes unreasonably complex. Psychological research which has examined the factors that influence effective instruction provides a foundation for identifying these essential dimensions.

Default instructional behaviour must be psychologically motivated so that it supports the overlaid instructional components as effectively as possible. This is necessary if the tools do not allow maximum flexibility. Additionally, by providing teachers with the opportunity to manipulate critical instructional variables, they are given the opportunity to prescribe the teaching strategy they prefer. These dimensions must therefore capture a large degree of variation between different approaches. A good model of the teaching process will permit substantially different ITSs through a small set of configurations.

While psychological theories of instruction can form the basis of the development of authoring environments for ITSs, there is another connection that should be exploited between psychology and ITS. No current theory is general enough to provide a basis for all the design decisions that have to be made to construct working systems. Authoring systems offer a means of speeding up scientific investigation. By providing ITSs whose teaching strategies can be easily and systematically altered whilst keeping the domain material constant, we are in an ideal position to test theories of instruction. Thus, the relationship between psychology and reusable ITSs should be a symbiotic one. Current research can inform the design of ITSs, which can then be used to test the theories embedded within it, which in turn can inform developing theories of instruction and learning.

How Can Authoring Environments Benefit From Psychology

Two important aspects that must be considered when designing ITS's are decisions about the teaching sequence (*i.e.* what to teach) and the nature of instruction (*i.e*. how to teach).

What to teach

Gagné (1965) pioneered the application of psychological task analysis to the decomposition of complex learning tasks. Since then the value of information processing approaches to the design of instructional systems has become well established. In this paper, we are less concerned with subject matter of the domain as this will be provided by the courseware. The major issue for the authoring tools and the resulting ITS is the decision about how to sequence instruction within a given learning domain. Research into curriculum planning holds that no single, fixed sequence of learning tasks or 'routes' through a domain is optimal for all learners (Resnick, 1976). Investigations into both learning 'styles' and aptitude by treatment interactions (ATIs) indicate that, to approach something like optimal teaching environments, it is necessary to plan for different learning routes through a domain. Thus, ITSs produced with REDEEM can provide both teachers and learners with various ways of approaching the domain.

A number of different aspects of REDEEM's model of teaching provide for this flexibility. One dimension that is crucial to achieving this is the variation in the level of control provided to the learner. At one extreme, learners must move through the material in the way prescribed by the teacher's specification and the shell's default behaviour. At the other extreme, the learner can be given complete control over the order in which they tackle material. A third approach provided by REDEEM is to provide learners with a choice of optimal pages and allow a selection to be made amongst these.

Furthermore, these suggested routes are highly configurable by teachers. One important aspect of these sequences is the specification of pre-requisite knowledge. Thus, teachers can ensure that (under conditions of low learner control) students can't access pages which contain material that would not be understood without some previous learning objectives having been met. Higher level objectives can also be configured. Some aspects of domains may best be tackled by providing examples before moving to general principles. On other occasions, it may be best to provide an over-arching framework before illustrating it. These two approaches represent different pedagogical views, as reported by Sutherland *et al.* (1996). Thus, REDEEM allows teachers to describe material as either general or specific and then allows this information to be incorporated into teaching strategies.

Domain material is likely to be provided with varying degrees of difficulty. Information about the complexity of a unit of material will be crucial to ensure sensible sequences of instruction. Thus, domain material must be described along several dimensions: is this material likely to be difficult, familiar, introductory, general or specific. This information is used to ensure plausible approaches to material in the absence of any more specific instructions. For example, when learners are faced with new material, then the default is to start with easy introductory material. The shell can also exploit these descriptions in relation to information about learner performance provided by the student model. Thus, students judged as having little problem with the material could be provided with more difficult units.

How to teach

REDEEM provides ITSs whose default instructional behaviour can be described as contingent. The aim of contingent support for learning is (a) that all learners can complete the learning activities successfully and (b) that they are continually provided with problems that are neither too easy nor too hard. Thus, in addition to making decisions about sequences of instruction, ITSs have to support the learners activities whilst they achieve some goal. One critical aspect of this involves rules about how any help is to be given. Wood *et al.* (1978) have formulated two seemingly simple principles which constrain the effectiveness of such decisions When a learner gets into difficulty, the tutor should offer more specific help immediately. When a learner succeeds after a given level of help, then the tutor should provide less specific instruction at the next intervention. Adherence to these principles predicts learning rates in a number of domains across a wide age range.

Within REDEEM, the provision of instruction (when asked for by the author of the system) will follow the principles of contingency by default. Changes to the defaults will only be encouraged if the author wishes to test alternative hypotheses about the rules for effective instruction.

In addition to default contingent behaviour, a teacher may specify various dimensions that configure the different teaching styles of REDEEM. This provides REDEEM with two main advantages (1) teachers can configure teaching strategies on an ITS along important dimensions to meet specific objectives and (2) the ITS will be able to use a variety of teaching styles during a single lesson to meet its objectives.

The different dimensions of instructional strategy provided in REDEEM allow different ITSs to be produced. This behaviour is crucial for allowing teachers to achieve specific objectives. Differences in instructional variables can be accommodated when ITSs are capable of being configured to support different lessons and students.

The evidence for different types of students learning better under differing conditions is grounded in the research on aptitude by treatment interactions conducted in the mid seventies (*e.g.* Cronbach & Snow, 1977; Snow & Yalow, 1982). A number of consistent effects were found (*e.g.* high aptitude students are found to work better in unstructured environments, with rapidly paced material; lower ability students with small step sizes, redundant text, *etc*.). More recently, research centrally concerned with computer-based learning, suggests that the nature of the learning environment will interact with the type of knowledge supported (*e.g.* declarative or procedural), the learner's cognitive style and the domain (Kyllonen & Shute, 1989). Consequently, REDEEM allows teachers to specify variables concerned with both instructional treatments and student aptitude. Dimensions of the instructional strategy include the amount of testing, when testing should occur, the extent of learner control, and whether answers should be deduced or given. The teacher can also develop a range of categories in which to place students. Once individual students have been placed in the categories, different strategies may be associated with the different categories.

One common dimension used to classify ITSs is the amount of learner control. This ranges from systems high in constraint such as Anderson's LISP tutor (Anderson & Reiser, 1985) to free discovery and programming environments such as Logo. Yet learner control is neither good nor bad. The success of high learner control depends on factors such as prior knowledge, motivation, instructional goals and cognitive skills (Steinberg, 1989). A fairly robust finding from the research on ATIs is that learners with lower aptitude should be given less control over how they learn. Another dimension associated with differential outcomes depending upon learner style and knowledge type is whether answers to questions should be deduced or given. Shute (1992) found interactions between question answering and aptitude. Higher ability subjects learned more declarative knowledge in rule induction (answer deduced) environments and lower ability subjects in rule application environments (answers given). Another of REDEEM's dimensions allows statements of the perceived level of difficulty. The ideal level of difficulty will be one that falls within the learner's zone of proximal development, (Vygotsky, 1978). This is the region of activity in which learners can perform successfully given the aid of supporting context. Allowing the ITS to choose material at different levels of difficulty allows us to cater for a wider range of understanding.

In addition to allowing ITSs to be configured to produce specific instructional strategies for particular learning goals, our approach allows one ITS to employ multiple teaching styles. This allows for much greater flexibility of teaching than has traditionally been the case in ITSs. Adapting the form (macro-adaptation) as well as the content (micro-adaptation) of teaching has been identified as crucial for successful tutoring systems (Ohlsson, 1986). Certainly, human teachers often subtly alter their strategies in response to student behaviour (Major, 1993; Wood *et al.*, 1976).

One way of providing an ITS with multiple strategies is to simply consider what makes teaching styles distinguishable. This approach need not entail consideration of educational or psychological validity (*e.g.* as in Dominie; Elsom-Cook & Spensley, 1990). Another approach recommended by Ohlsson (1986) is to identify how expert teachers adapt their teaching styles and then employ this knowledge in an expert system component of an ITS. Our approach draws even more heavily on teachers' expertise. Teachers are allowed to directly configure different teaching strategies and to describe, using a simple interview tool, when they would change some aspect of their strategy in response to a learner's performance.

How Psychology Can Benefit From Authoring Environments

Although psychological research can inform much of the design of ITSs, generalised principles are few and far between. We believe that by basing an authoring environment on the identified critical dimensions that underlie these existing principles, we provide the basis for the other aspect of this symbiotic relationship; how ITSs can in turn inform psychological theory.

The first essential element of this approach is reusability. With reusability of courseware, we can explore how various teaching strategies are differentially effective with the same domain material or with different categories of learners. Further, because REDEEM has been designed to exploit critical instructional variables, it becomes an ideal testing ground for controlled experiments that examine teaching strategies. These instructional approaches can either be based on existing literature or could be created by individual teachers. This allows us to explore the convergence between theoretical objectives and teachers' practical approaches. We can relate these different approaches to learning outcomes.

Teaching strategies

We reviewed above some of the more robust findings from the research on ATIs. However, research on many instructional variables had revealed inconsistent effects (*e.g.* diagrams, inserted question; Snow &

Yallow, 1982). Using authoring environments provides opportunities to revisit many of these issues, but in a more systematic and controlled fashion. Evaluation using authoring environments has the advantage that different teaching styles can be applied consistently and to the same course material, thus eliminating other sources of variance.

Additionally, studies done in the past with ITSs investigating teaching strategies have been restricted to particular kinds of students, particular domains and particular types of problem. With REDEEM we will be able to consider what strategies are appropriate for different kinds of learners (*i.e.* different ages, abilities) in different kinds of domains and with different kinds of problems.

We wish to go beyond the sort of evaluations that are currently possible with ITSs. Individual systems have implemented particular strategies before (see Lesgold *et al.*, 1992) and such systems have been evaluated. REDEEM provides control over specific elements of strategies. Consequently, evaluation studies can be much more focused than broad comparisons between strategies. The same course can be taught with the same strategy but with only a single variable altered, such as the amount of testing. Thus, we will be in a position to examine the effectiveness of teaching strategies in a fine-grained manner.

The design space of possible ITSs is large. This is particularly true when teaching strategies have been considered in the detailed manner we are prescribing. Consequently, evaluation of ITSs to inform psychological theory will be costly and time-consuming. A partial solution is provided by authoring environments which contain psychologically motivated variables that can be manipulated in a controlled fashion. REDEEM does not offer a solution to the time taken to run experiments. However, it does substantially reduce the time needed for ITS re-implementation.

We have already highlighted the need for macro and micro adaptation for effective learning. Although there is some research that examines when human teachers change their instructional approach (Wood *et al.*, 1976), little is known about when it is appropriate to change teaching strategy in a computer-based environment. A few systems have employed multiple strategies, albeit in an unprincipled way. But we have lacked the tools that would have allowed precise experiments to be carried out. REDEEM offers the opportunity to explore the effects of both coarse and fine-grained changes of strategy on learning. The results of this exploration will inform a meta-strategy knowledge base, which can, in turn, guide the design of future generations of ITSs.

The REDEEM Software Suite

The REDEEM tools consist of three main pieces of software: the courseware catalogues; the authoring tools and the instructional shell. This software has been developed in Asymetrix Multimedia ToolBook and will run on Windows 3.1 or higher.

The primary interface for the teacher is through the authoring tools, in which courses are described, teaching strategies constructed and students classified. The shell then uses this knowledge that the teacher has prescribed with the authoring tools, together with its own default teaching knowledge, to interpret the courseware in such a way as to deliver intelligent instruction.

Courseware Catalogues

Domain material in REDEEM is based on the idea of a courseware catalogue. That is, it is able to interpret pages of computer-based training developed in a standard authoring package, ToolBook. REDEEM cannot interpret general ToolBook courses because they are likely to include much hard-coded control knowledge. Rather, it is able to use ToolBook courses which offer discrete pages of material showing the different aspects of the domain at varying levels of difficulty. These pages form the content of the courseware catalogue. The material on each page is pre-prepared which allows greater reusability but, of course, can limit the flexibility of the resulting instruction to some extent. If the course has pre-existing control knowledge coded into a page, this information will have to be manually removed before authoring. Much of this control information can be re-represented using the authoring tools.

The genetics course

The first significant domain we have developed for use with REDEEM is a course in genetics for 15 year old high school students, built with the assistance of a high school biology teacher. The course is based around the National Curriculum in biology, which is taught in all UK schools. The material developed would normally take around 12 hours of teaching time to cover in normal classroom based teaching, and so we have a fairly significant amount of material.

The course itself does not look like a typical CBT course, because it is developed with the courseware catalogue idea in mind. Thus, it contains a certain amount of repetition as various topics are treated at different levels of difficulty, thereby giving the teacher and hence REDEEM the choice of different approaches to teaching the same information.

The material includes a full range of multimedia presentations, with still photographs, text, graphics, sound, animation and simulation. This has been done in consultation with the teacher in order to offer a rich and motivating environment. Figure 1 shows part of the course itself, together with one of the authoring tools in which the teacher is describing the course.

Authoring Tools

REDEEM's authoring tool suite consists of five main tools. These allow the teacher to describe the different courses at both page and section level, to construct teaching strategies, to categorise students and to describe how strategies should be refined during a lesson. In addition, there is a simple tool which allows the teacher to assign different strategies to various student categories.

Describing course material

The first part of the authoring process involves the description of the course material. For each page, the teacher describes the teaching material upon a number of dimensional ratings. These can be seen in Figure 1. Rapid elicitation of course features is ensured by graphical manipulation. These dimensions were generated from interviews with teachers, psychological research, and previous research with COCA. Teachers make judgements along these dimensions for use by the shell during teaching. The positions of the sliders do not take into account any changes that may occur during the lesson itself.

Figure 1: Describing a page of course material in REDEEM

Limited interactivity may be added to otherwise fairly static course material by the addition of appropriate questions. Teachers can define relevant questions, specify answers (including second best where appropriate) and give five hints for each question that follow the principles of contingent help. Figure 2 shows the tool for eliciting multi-choice questions. The shell will monitor the student's performance on these questions. Pages may also contain interactive problem-solving tasks.

Whilst REDEEM does not at present monitor student performance on these types of tasks, it can direct the student to these pages as appropriate.

Figure 2: Defining a multi-choice question

Teachers also select between a maximum of six relations between pages, such as, which pages are prerequisite to the current page, or whether one page is analogous to another. We intend to investigate which of these relations prove to be most useful to teachers.

The teacher combine pages into sections. Thus all pages that address a similar topic can be combined into one section even if they are distant within the courseware catalogue. These sections are themselves characterised independently of the pages. The dimensions available are the same as the first four sliders used to describe the page (*i.e.* introductory, difficulty, familarity and specificity of material). It is also possible for teachers to specify explicit relations between sections. An example from the Genetics course is shown in Figure 3.

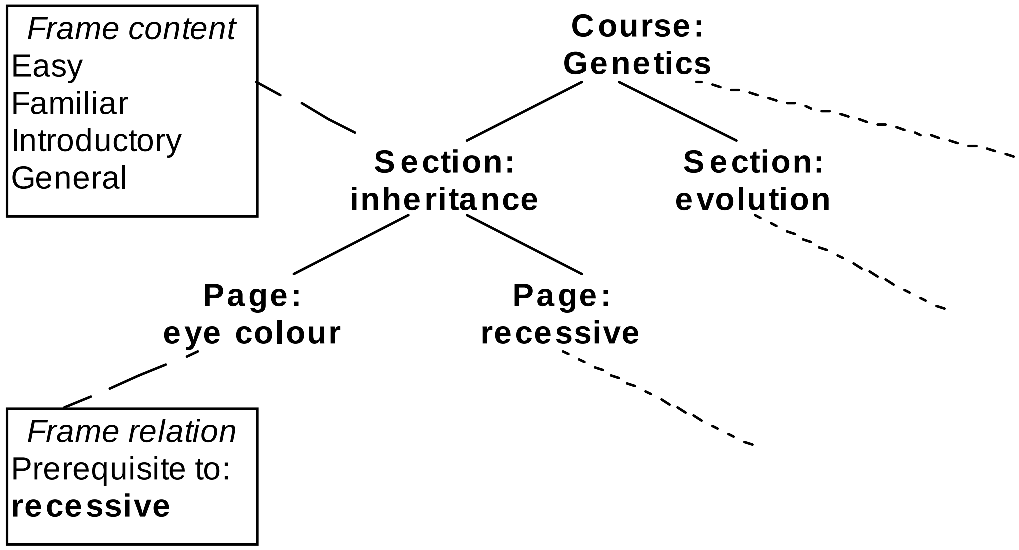

Figure 3: A part of the semantic network describing the Genetics course

These tools provides information that the system uses as a semantic network describing the structure of the teaching material. This network expresses information about the properties of page and section contents and relations between these pages and sections. This enables the shell to make default decisions about adapting content and to implement teachers' preferred sequences and routes through material.

Constructing teaching strategies

The second important part of the authoring task is the definition of a number of teaching strategies as the basic repertoire of the ITS shell. Figure 4 shows the teacher defining a strategy that has been called *practice*. A graphical interface allows teachers to dimensionally scale and rank the critical aspects of the strategy and to choose a number of potential teaching and testing styles. Thus the process of constructing teaching strategies is a very simple enterprise in terms of the interface. The simpler authoring necessitates a certain loss of control for the teacher that is transferred to the teaching shell, but there is still considerable control available.

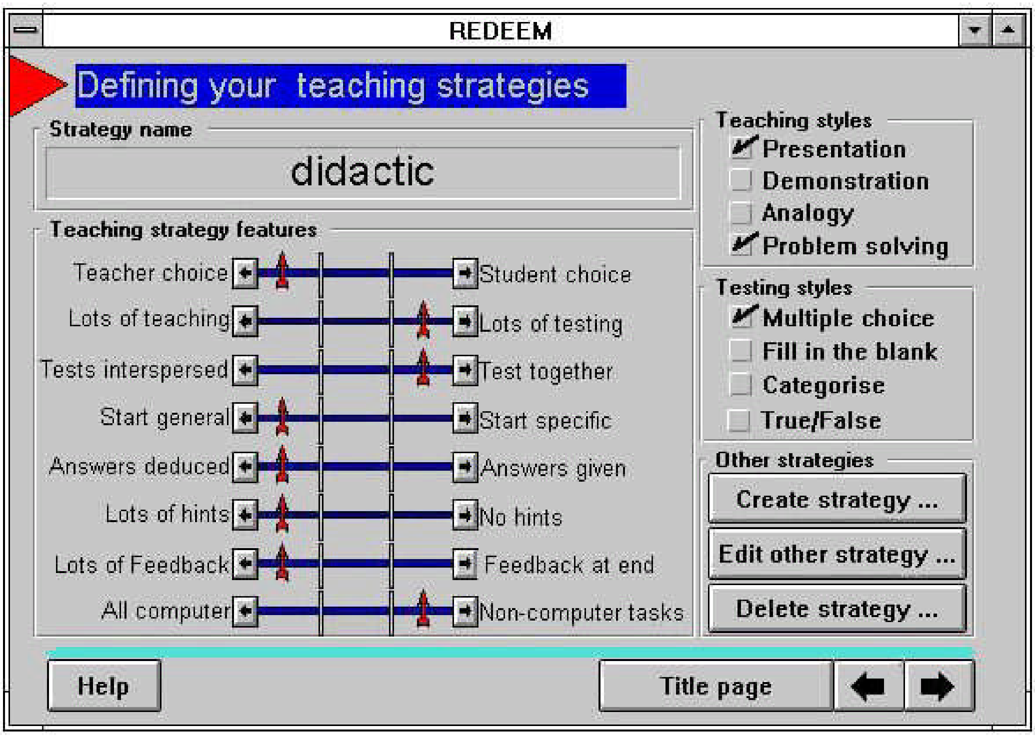

Figure 4: Screen showing strategy authoring

Different instructional principles can be embodied in various strategies by manipulating the sliders. Each slider in Figure 4 has three discrete positions that result in different instruction. Consequently teachers are free to create as many strategies as can be composed from the various instructional attributes shown above. Table 1 provides a brief summary of the effects of the different slider positions. Additionally, the teaching styles inform the shell that it can exploit built-in facilities of the course material. For example, the analogy style seeks to prefer a new page that has an analogous relationship to the current one if such a page is available, and the problem-solving style will seek to present an interactive problem to solve. The testing styles inform the shell that the teacher prefers questions to be asked in certain ways, such as a multiple choice or a fill in the blank.

	Slider Left	Slider Centre	Slider Right
Teacher choice	No student choice	Choice of shell's preferred pages	Free choice of all pages
Lots of teaching	Offer no tests	One test limit per page	All tests available
Tests interspersed	Test after each page	Test after each section	Test after course
Start general	Prefer general pages first		Prefer specific pages first
Answers deduced	Right answer provided only when no further options available	Right answer provided upon second error	Right answer provided upon error
Lots of hints	Help on request and error	Help on error	No help
Lots of interruption	Summarise after section and question	Summarise after section	No summary
All computer	No non-computer tasks		All non-computer tasks

Table 1. The effects of the slider positions on teaching behaviour

Describing students

Students can be categorised into a set of teacher-defined categories. These categories are used by the shell to decide which strategy is most appropriate at any given moment. Figure 5 provides an example of one teacher's decisions, where the categories chosen are based upon performance ratings.

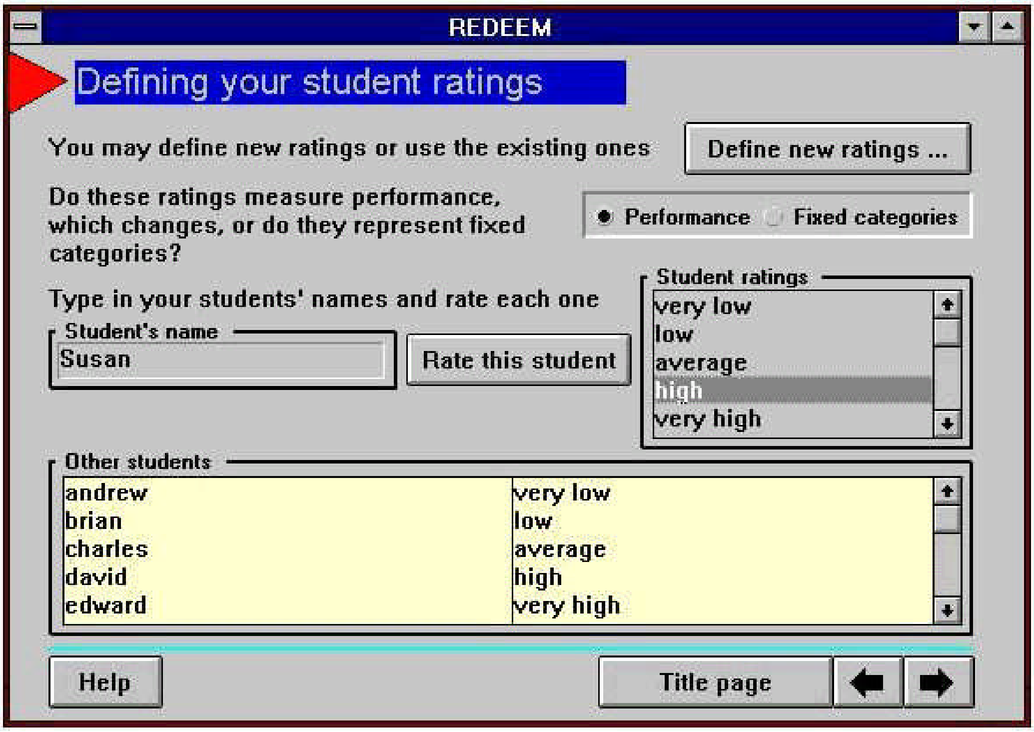

Figure 5: Describing students

If the teacher wishes, the validity of categories can be evaluated against performance. If this is the case, then the shell will automatically change the category as the overall standard of the student (as defined in the shell's student model) changes. This could result in a new teaching strategy. Thus, student categories also have influence over strategy changes. Alternatively, these categories could represent dimensions other than performance, such as learning styles. Teachers have suggested literacy, numeracy and cognitive style as potential student ratings. REDEEM uses these student categories, as well as information gained during a student's interaction with the shell, to modify its instructional behaviour.

Knowledge about strategy refinement

The student categories that the teacher defines can be assigned to particular teaching strategies, thus offering the REDEEM shell coarse-grained knowledge about when to change strategy. In addition to this facility, there is another authoring tool which leads the teacher through a number of multiple-choice questions, eliciting information about the circumstances in which particular aspects of the current strategy might change. The teacher may decide, for example, that if the student's performance improves to a certain level, then the strategy should offer more learner control. Figure 6 shows one of the interview screens.

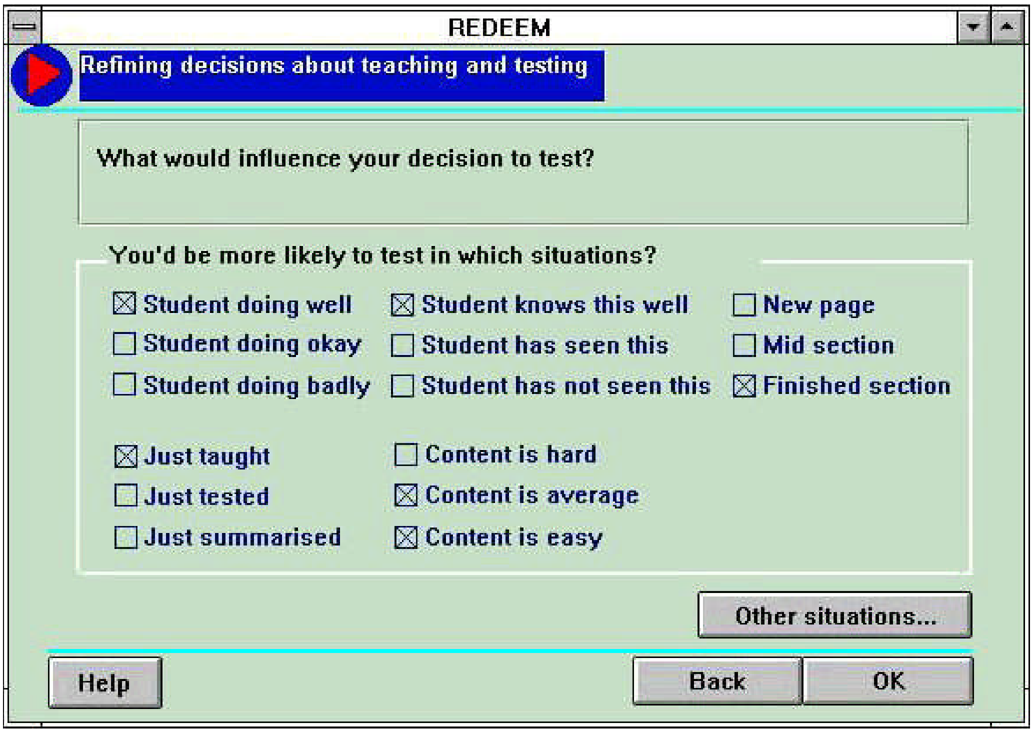

Figure 6: REDEEM tools eliciting knowledge about strategy refinement

The original model of teaching employed in COCA was analysed so that all the essential decisions were organised in hierarchical fashion. The interview asks questions of the teacher which enables the shell to select the appropriate model at any given time, somewhat analogous to traversing a decision tree. More details of the approach taken in this tool are given in Major & O'Hara (1995). An interviewing tool has been implemented although, as yet, there is limited functionality in the shell to exploit the results of the interview.

ITS Shell

The ITS shell delivers the courseware according to the instructions generated by teachers using the authoring tools in combination with its predetermined defaults. In order to achieve adaptive instruction, it must handle a number of other roles.

Delivering adaptive instruction

The main role for the ITS shell is to deliver the course material to each student in the manner that the teacher has specified using the authoring tools. The shell is given the teacher's strategy description contained in the dimensional ratings, and the semantic network created from the teacher's descriptions of the course material which provides a powerful engine for achieving adaptive behaviour in the final ITS. At a fine-grained level, a tutorial action available to the shell is the choice of appropriate domain material, for example should the next page be more general, more difficult or less familiar. By using the semantic network relating pages, default tutorial knowledge will ensure prerequisites are mastered before new material is introduced or that examples or analogies are used appropriately. If there is a non-computer based task associated with a chosen page, then the strategy information controls whether that particular task is offered to the student, and monitors whether this has been achieved before allowing the student to progress. At a more coarse-grained level, the shell also selects the next tutorial activity, that is, do we teach next by presenting new material, or do we test, and if so should we test after the page or the completion of the next section? The shell will also consider whether it is appropriate to hand control over to the student, allowing either free choice of all material or a limited selection of the shell's preferred pages.

Default tutorial behaviour

Although substantial amounts of the shell's teaching and meta-strategic knowledge will have been defined by the teacher in the authoring environment, there is still a requirement for default tutorial behaviour by the shell. The major function of the default knowledge is to ensure sensible routes through the domain. This sequencing behaviour follows teachers' commonly held principles of teaching obtained by interview and known experimental results on instructional design. Rules used to govern these sequences are to prefer easy material before difficult material, introductory before final or familiar before unfamiliar. This is done both at the section and page level. Other rules specify that pages in the same section should be taught together and that pre-requisite pages are taught in the correct order. The experimental results from contingent instruction (Wood *et al.*, 1978) are used to determine the help students receive. This states that when students are failing then offer more help, if they are succeeding, then offer less help.

Student modelling

REDEEM employs a basic overlay model which records the system's understanding of the students' knowledge of an area. The values of the model change over the course of a session both as the student sees new or existing material and as the student answers questions.

The basic course material unit being modelled is the page. This provides a degree of flexibility on the modelling process not usually found in a conventional ITS. Depending on the nature of the material on a page, this basic unit could correspond to an individual declarative fact, or to a step in a procedure. As the lesson proceeds the shell will update the model entries for each page.

Strategy refinement

The aim of meta-strategic knowledge is to alter, where appropriate, the rules for making these decisions about teaching. In order to achieve macro-level adaptation, the model of the teaching strategy used at each cycle of interaction with the student will not necessarily be the same as the models used in the previous or subsequent cycles. The shell uses knowledge elicited in the authoring tools to offer more sophisticated changes of strategy during run-time instruction. These fine-grained changes allow any of the critical dimensions of the current strategy to be altered.

Student history

An important role for the shell is to maintain a student history in addition to the student model. The shell keeps a trace of all modules taken, including pages visited, questions that were asked and their answers, number of hints offered, scores and time on tasks. This history is sufficiently detailed that it is possible to recreate a student's session artificially. This information is used for a number of purposes - it limits repetition of pages during teaching, it serves as the basis of the report given to the teacher and provides useful information for research purposes.

Now that we have described the REDEEM authoring tools, we will consider how they can be used to examine teachers' knowledge of learning and instruction - the expertise that authoring environments depend upon.

Exploring Teachers' Expertise

REDEEM allows us to replace the responsibility for designing instruction back in the hands of the experts - teachers and instructors - rather than in the hands of knowledge engineers and programmers. This approach is contrasted with traditional ITSs which have a fixed instructional strategy. However, in order for authoring environments to be successful, teachers' pedagogical knowledge must be capable of generating effective ITSs. Our future research is centrally concerned with addressing this question. REDEEM provides the opportunity to explore two distinct but related research agendas. The first is to explore the theoretical question as to whether involving teachers in ITS development produces the promised practical advantages. The second is to examine how authoring tools could be used to capture the nature of teachers' expertise.

Firstly, we propose and wish to test the following practical advantages for including teachers in ITS design: (a) improved teaching, (b) increased teacher acceptance, (c) knowledge transmission and (d) professional practice. We will now consider each of these advantages in turn.

We propose that as a consequence of allowing teachers to author ITSs, systems which employ a more effective instructional approach will result. Authoring systems allow us to exploit the knowledge bases of those people with expertise and responsibility for teaching and training - teachers themselves.

With regard to teacher acceptance, teachers report that they would trust computer-based tutoring more if they could modify general and student specific teaching strategies in addition to modifying course material (Major, 1993).

Knowledge transmission can occur by comparing the lessons produced by educational experts in different schools using similar materials. Virtual communities of teachers could develop and share resources. Teacher trainers would be in a position to produce model lessons for their trainees to experience. Within schools, knowledge transmission would occur as teachers with particular expertise produce environments to be shared with other teachers.

In addition to experiencing lessons designed by other teachers, authoring environments offer teachers two additional ways to improve in their professional practice. Firstly, REDEEM offers unique opportunities for teachers to reflect upon pedagogy by requiring users to make pedagogical judgements in the context of student needs and by providing pre-existing domain material. Feiman-Nemser (1990) has identified a major component of expertise as the teacher's ability to focus on students' needs rather than on themselves or the domain. Additionally, a characteristic of newly qualified teachers' instruction is an overemphasis on selection and presentation of content at the expense of developing and assessing pedagogy. REDEEM provides the subject material and hence should support teachers' reflection upon general pedagogical reasoning. Secondly, REDEEM allows teachers to experience their own teaching from the learners' perspective. This experience is something it would be nearly impossible to create without tools such as REDEEM.

A second focus of our research with teachers will be to explore aspects of teacher expertise. Although, a substantial body of research has attempted to analyse the components of good teaching, this has rarely been related to learning outcomes (Desforges, 1995). REDEEM can be used to examine the lessons produced by experienced teachers in comparison with those produced by novice teachers. These environments could be related to any differences in learning outcomes.

Based upon the existing research on knowledge acquisition with novices and experts (*e.g.* LaFrance, 1989), the cognitive expertise literature (*e.g.* Chi, Glaser & Farr, 1988) and novice/expert teacher research (*e.g.* Carter *et al.,* 1988), we propose a number of predictions about the expert/novice teacher differences in relation to authoring. For example, experts will have more proficient plans than novices and so will be able to author more quickly and effectively and experts will differentiate their instruction more finely.

We propose to examine how differences such as these will be reflected in the way the authoring tools are used. In addition to examining these differences between teachers' lessons, we are also looking for points of convergence. For example, what are the common categories that teachers use to classify students; how many categories are normally described; do teachers agree that individual children categorised as similar should receive similar lessons?

Future Developments

The research on teacher expertise should additionally identify factors that contribute to the successful use of authoring environments. These will help us to design both the facilities of new authoring tools and the ways such tools should be supported. Three of these factors are discussed below.

Cut Down Version For Schools

We anticipate that having used REDEEM to experiment on what are the most effective instructional features, we will consequently be able to identify the crucial dimensions that teachers wish to manipulate, and the critical relations used to link pages. Having done so, the authoring task may further be reduced by just offering these most commonly used features. Such a minimalist version is more likely to find acceptance in situations where authoring time is short. Another alternative for teachers who haven't the time for authoring is to have educational publishers produce courses and course descriptions that can be used almost immediately with the teacher's strategies.

Web-Based Teaching

The internet itself is too dynamic for REDEEM to use directly as courseware; the content of pages changes too often. It is, however, an ideal communication medium for sharing the reusable components that REDEEM uses. Thus, material taken off the net could be kept either on local web pages, which, being more stable could become the basis of a REDEEM course, or could be imported into ToolBook itself. A small change to the implementation of REDEEM would allow it to interpret HTML.

There is considerable scope for children and students to use REDEEM themselves. Once provided with a small library of strategies, children could themselves pull course material off the net. Another possibility is that groups of students could develop ITSs for younger students or indeed for one another. A subject area could be assigned to each student, who then, having learned the subject area, develops an ITS using REDEEM and exchanges it for the other ITSs developed by the others in the group.

Training Implications

If authoring is to become the task of a qualified teacher, we need to have good models showing how they use such environments. We need to evaluate both instructor satisfaction and time taken to author the environment in relation to the amount of instructional material produced. In addition, we wish to use a video-based case study method to examine the process of authoring itself. Examining how teachers behave when authoring ITSs using REDEEM will be a first step to specifying the necessary skills. This information will be used when considering subsequent re-implementation of the tools. Additionally, it will serve to inform us about how we should train teachers when they are confronted with authoring for the first time.

Summary

This paper has considered the relationship between psychological theories of learning and instruction and the development of authoring environments. We started by considering why ITSs have not been the success that was hoped for, and showed how authoring environments might resolve many of these difficulties when they are based on a model that is grounded in psychological research. We proposed that this relationship is not uni-directional but that theories of instruction could be enhanced by the results of experiments performed with such an authoring environment. We then presented REDEEM, a set of tools which encapsulate this potential symbiotic relationship. We moved on to see how teachers' expertise informs REDEEM based ITSs and, in turn, how teachers might benefit in their professional practice by using REDEEM. Finally, we looked at possible future developments.

Initial user trials have been undertaken and we are now in a position to test the utility of this approach. We propose that this harnessing of psychology and authoring environments offers a unique solution to effective computer-based instruction.

Acknowledgements

The work was supported by the ESRC at the ESRC Centre for Research in Development, Instruction and Training. The authors would like to acknowledge all the other members of the centre for providing a stimulating and supportive atmosphere in which to conduct research.

References

Anderson, J.R. & Reiser, B.J. (1985) The LISP tutor. *BYTE*, *10*(4), 159-175.

Bloom, B.S. (1984) The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher*, *13(*6), 4-16.

Bloom, C.P., Meiskey, L., Sparks, R., Dooley, S. & Linton, F. (1995) Putting Intelligent Tutoring systems into Practice: A Study in Technology Extension and Transfer. *Machine Mediated Learning, 5,* 13-41.

Carter, K., Sabers, D., Cushing, K., Pinnegar, S. & Berliner, D. (1987) Processing using information about students; A study of expert, novice and postulant teachers *Teaching and Teacher Education, 3*, 147-157.

Chi, M.T.H., Glaser, R. & Farr, M.J. (1988) *The Nature of Expertise*. Hillsdale, NJ: LEA.

Cronbach, L.J. & Snow, R.E. (1977) *Aptitudes and Instructional Methods*. New York: Irvington.

Desforges, C.W. (1995) How does experience affect theoretical knowledge for teaching? *Learning and Instruction, 5*, 385-400.

Elsom-Cook, M. & Spensley, F. (1990) A tutor for procedural skills. In M. Elsom-Cook (Eds.), *Guided Discovery Tutoring: A Framework for ICAI Research* (pp. 57-69). London: Paul Chapman Publishing Ltd.

Feiman-Nemser, S. (1990) Teacher preparation: Structural and conceptual alternatives. in W.R. Houston (Ed.) *Handbook for research on teacher education*. New York Macmillan.

Gagné, R.M. (1965) *The Conditions of Learning.* New York, Chicago, San Francisco, Atlanta, Dallas, Montreal, Toronto, London, Sydney, Tokyo: Holt, Rinehart and Winston.

Kyllonen, P.C. & Shute, V.J. (1989). A Taxonomy of Learning Skills. In P. L. Ackerman, R.J. Sternberg, & R. Glaser (Eds.), *Learning and Individual Differences: Advances in Theory and Research* (pp. 117-163). New York: Freeman.

LaFrance, M. (1989) The quality of expertise: Implications for expert-novice differences for knowledge acquisition. *Sigart* April 1989.

Lesgold, A., Lajoie, S., Bunzo, M. & Eggan, G. (1992) Sherlock: A coached practice environment for an electronics troubleshooting job. In J. Larkin & R. Chabay (Eds.), *Computer Based Learning and Intelligent Tutoring* (pp. 202-274). Hillsdale, NJ: LEA.

Major, N.P. (1993) Teachers and Teaching Strategies. *Proceedings of the Seventh International PEG Conference*, Heriot-Watt University, Edinburgh.

Major, N.P. (1994) Evaluating COCA - What do teachers think? *Proceedings of the World Conference on Educational Multimedia and Hypermedia - EDMEDIA 94*. AACE Press.

Major, N.P. (1995) Modelling Teaching Strategies. *Journal of Artificial Intelligence in Education, 6,* 117-152.

Major, N.P. & O'Hara, K. (1995) Grounding DTMs: An Interview Tool For Acquiring Meta-Strategic Teaching Knowledge. In Hallam, J. (Ed.) *Hybrid Problems, Hybrid Solutions*. Amsterdam: IOS Press.

Murray, T. & Woolf, B. (1992) Tools for Teacher Participation in ITS Design. In Frasson, C., Gauthier, G. & McCalla, G.I. (Eds.) *Proceedings of Intelligent Tutoring Systems 92*. Berlin: Springer Verlag.

Ohlsson, S. (1986) Some principles of intelligent tutoring. *Instructional Science, 14,* 293-326.

Resnick, L.B. (1976) *The Nature of Intelligence*. LEA, Hillside, N.J.

Russell, D.M., Moran, T.P. & Jordan, D.S. (1988) The instructional design environment. In Psotka, J., Massey, S.A. & Mutter, A. (Eds.) *Intelligent tutoring systems: Lessons learned*. Hillsdale, NJ: LEA.

Shute, V.J. (1990) Rose Garden Promises of Intelligent Tutoring Systems: Blossom or Thorn *Presented at the Space Operations, Applications and Research Symposium,*

Shute, V.J. (1992) Aptitude-treatment interactions and cognitive skill diagnosis. In J. W. Reigan & V.J. Shute (Eds.), *Cognitive approaches to automated instruction.* Hillsdale, NJ: LEA.

Snow, R.E. & Yalow, E. (1982) Education and Intelligence. In R.J. Sternberg (Eds.), *A Handbook of Human Intelligence* Cambridge: Cambridge University Press.

Soloway, E.M. & Johnson, W.L. (1984) Remembrance of blunders past: a retrospective on the development of PROUST. *Proceedings of the Sixth Cognitive Science Society Conference*. Boulder, Colorado.

Steinberg, E. (1989) Cognition and learner control. A literature review. *Journal of Computer Based instruction, 16*(4), 117-121.

Sutherland, R., Rojano, T., Mochon, S., Jinich, E. & Molyneux, S. (1996) Mathematical Modelling in the Sciences through the Eyes of Marina and Adam. In *Proceedings of PME-20*, University of Valencia, Valencia.

Van Marcke, K. (1992) Instructional Expertise. In Frasson, C., Gauthier, G. & McCalla, G.I. (Eds.) *Proceedings of Intelligent Tutoring Systems 92.* Berlin: Springer Verlag.

Vygotsky, L.S. (1978) *Mind in Society: The Development of Higher Psychological Processes.* Harvard University Press, Cambridge: MA.

Wood, D.J., Bruner, J.S. & Ross, G. (1976) The role of tutoring in problem solving. *Journal of Child Psychology and Psychiatry, 17, 2, 89-100*.

Wood, D.J., Wood, H.A. & Middleton, D.J. (1978) An experimental evaluation of four face-to-face teaching strategies. *International Journal of Behavioral Development, 1*, 1.

Woolf, B.P. & Cunningham, P.A. (1987) Multiple knowledge sources in intelligent teaching systems. *IEEE Expert*, Summer 1987.