

Intelligent Cognitive Tutors as Modeling Tool and Instructional Model

Position Paper for the NCTM Standards 2000 Technology Conference
June 5-6, 1998

Dr. Kenneth R. Koedinger
Pittsburgh Advanced Cognitive Tutor (PACT) Center
Human-Computer Interaction Institute
School of Computer Science
Carnegie Mellon University

What gets measured, gets done.
If you don't measure results, you can't tell
success from failure.
If you can't recognize failure, you can't
correct it.
If you can't see success, you can't reward it.
If you can't see success, you can't learn from
it.

From "Reinventing Government"
By David Osborne and Ted Gaebler

Abstract

Effective technologies for learning and doing mathematics should be based on sound cognitive theory, be empirically tested against alternatives, and be primarily addressed at mathematics as a modeling language. I will discuss the status and promise of research-based "intelligent" systems that support students in mathematical modeling and tool use. Such systems are beginning to reshape the mathematics classroom, the way teachers teach, and what and how students learn.

Introduction

The goals of the NCTM Standards 2000 Technology Conference were stated as follows:

1. What are appropriate/effective technologies for learning/doing mathematics?
2. What is the mathematics that is relevant in a technological world?
3. What does research say about uses of technology in mathematics education?
4. What is the technological infrastructure that one can assume most schools and teachers will have access to in years 2000-2005? Equity issues?

This paper focuses primarily on the first and third of these questions. In particular, I believe that effective technologies for learning and doing mathematics should be based on sound cognitive theory, be empirically tested against alternatives, and be primarily addressed at mathematics as a modeling language. I illustrate these points in the context of an educational technology we call "Cognitive Tutors" (Anderson, Corbett, Koedinger, Pelletier, 1995). Cognitive Tutors are based in computer science research on artificial intelligence techniques and cognitive psychology research on the nature of human learning and performance. Cognitive Tutors have been created to help students learn in a variety of mathematics and computer programming domains and have been subject to laboratory and classroom evaluations that demonstrate the potential for dramatic learning gains from appropriate use of this technology. My comments will focus primarily on a cognitive tutor for algebra called the Pump Algebra Tutor or PAT (Koedinger, Anderson, Hadley, & Mark, 1997; Koedinger & Sueker, 1996). PAT is in regular use now in over one hundred classrooms at the middle school, high school, and college levels. Many of the 40 schools using PAT are in big cities, involve average teachers, and include a large number of disadvantaged minority or learning disabled students. Thus, my comments will also be relevant to the question of equity issues.

Before turning to a summary of Cognitive Tutors in general and PAT in particular, I want to make a few general comments about the role of cognitive theory and empirical testing in the development of technology-enhanced learning innovations and the role of mathematical modeling as the appropriate core focus of innovations for mathematics.

Why empirical tests against alternatives?

Why is it important that we perform empirical tests of our innovations in comparison with alternatives? If the intuitions and beliefs that guide our design of learning environments were fully informed and perfect, there would be no need for such experiments. Unfortunately, our intuitions and beliefs about learning and instruction are limited and are not always accurate. One problem is that our intuitions are based largely on our conscious learning experiences, but a great fraction, perhaps the majority, of what we learn is at a level below our awareness. The grammar rules of our first natural language, English in my case, are an excellent example. We learn these rules, in the sense that they determine our behavior in language comprehension and production, well before we are consciously aware of them. To use that old twist of phrase, as early language learners we go from "not knowing we don't know" to "not knowing we know" without going through the intermediate states of conscious learning: "knowing we don't know" and "knowing we know". As we get older, of course, conscious learning processes play a greater role. However, it is a mistake to think conscious learning takes over. In fact, there is ample evidence from cognitive psychology research that our brains continue to engage in implicit learning processes.

Our intuitions about learning are biased by limited information -- overly influenced by our memories of our conscious learning experiences. We are subject to what I call "Expert Blindspot" -- as experts in a domain we are often poor judges of what is difficult and challenging for learners. Perhaps few would disagree about the importance of evaluating our educational innovations to better understand how they do or do not improve on current practice. Nevertheless, I think it is worth emphasizing the danger of being biased by expert blindspot and lured by our personal intuitions into assuming that

our educational innovations and reforms will necessarily be for the better.

Why design systems based on cognitive theory?

Not every experiment can be run comparing alternative features of instruction and their interactions. Thus, we need a way to guide the generation of new instructional designs. Such a guide should help us prune design ideas not likely to enhance learning and inspire new ideas that will. Cognitive theory also provides a way to accumulate reasons for past successes and failures to inform future practices.

Why address math as a modeling language?

Although technology has mastered calculation of various kinds, arithmetic, graphic, symbolic, logical, humans are the only masters of translating problems into mathematics, building theories, producing communicative forms. Learning how to create mathematical models of problem situations is difficult but it is the key to mathematical success in our modern world.

An Example: Cognitive Tutors

I will summarize a technology, Cognitive Tutors, which I believe is particularly suited (though certainly not exclusively) to facilitate the three goals I've articulated: need for cognitive theory, empirical testing, and a focus on mathematics as a modeling language. Cognitive Tutors are based on the ACT theory of learning and performance (Anderson, 1983; 1990; 1993). ACT is a broad "unified theory of cognition" (Newell, 1990), however, I will highlight just a few key features that are particularly relevant to learning mathematics. The theory distinguishes between tacit performance knowledge, so-called "procedural knowledge" and static verbalizable knowledge, so-called "declarative knowledge". According to ACT, performance knowledge can only be learned *by doing*, not by listening or watching. In other words, it is induced from constructive experiences -- it cannot be directly placed in our heads. Such performance knowledge is represented in the notation of if-then production rules that associate internal goals and/or external perceptual cues with new internal goals and/or external actions. Here three examples of English version of production rules:

| | |
|------|---|
| IF | the goal is to prove two triangles congruent and the triangles share a side |
| THEN | check for other corresponding sides or angles that may be congruent |
| IF | the goal is to solve an equation in X |
| THEN | graph the left and right sides of the equation and find the intersection point(s) |
| IF | the goal is to find the value of quantity Q and Q divided by Num1 is Num2 |
| THEN | find Q by multiplying Num1 and Num2 |

It is important to note that the rules of mathematical thinking (which production rules are intended to represent) are **not** the same as the rules of mathematics (e.g., theorems, procedures, algorithms) as they appear, for instance, in textbooks. Production rules represent people's tacit knowledge of when to choose particular mathematical rules as well

as other tacit performance knowledge like plans and informal intuitions. The particular if-then notation of production rules is not so important as the features of human knowledge they represent and the implications of these features for instruction. Production rules are *modular* and this means that we can diagnose specific student weaknesses and focus instructional activities on improving these. Production rules are *context specific* and this means that mathematics instruction cannot be effective if it disconnects mathematics from its contexts of use. Students need true problem solving experiences to learn the if-part of productions, the conditions for appropriate use of mathematical rules, as well as some occasional small exercises (which are still over-emphasized in many curricula) to introduce or reinforce the then-parts of productions, the mathematical rules themselves. Production rules are of *limited generality*. In other words, cognitive research (e.g., Singley & Anderson, 1989) has shown that the performance knowledge, though general (i.e., it applies in multiple contexts), tends to be fairly narrow in its applicability and tied to particular contexts of use. Thus, we must gauge our expectations about how far student learning will transfer and construct curricula that both encourages general encodings of mathematical ideas and also provides multiple examples and activities applying these ideas in a variety of well-chosen contexts.

In applying the ACT theory to instruction, we have focused on the idea that human one-to-one assistance or tutoring is extremely effective in facilitating learning. Bloom (1984) showed that an individual human tutor can improve student learning by two standard deviations over classroom instruction. In other words, the average tutored student is performs better than 98% of students receiving classroom instruction. This results provides a sort of "gold standard" for comparing the effectiveness of educational technologies. The results of meta-analysis of hundreds of studies of traditional computer-aided instruction (CAI) suggests that CAI can leads, on average, to a significant 0.3 to 0.5 standard deviation improvement over non-computer-aided control classrooms (e.g., Kulik & Kulik, 1991). There are too few studies of multimedia and simulations at this point to provide a generic figure, though some of these studies indicate little effect, for instance, of animations (e.g., Pane, Corbett, & John, 1996) or of game-like simulations (e.g., Miller, Lehman, & Koedinger, in press). In studies of our Cognitive Tutor technology we have shown Cognitive Tutors to yield about a one standard deviation effect (Anderson, Corbett, Koedinger, & Pelletier, 1995; Koedinger, Anderson, Hadley, & Mark, 1997).

To build a Cognitive Tutor, we create a cognitive model of student problem solving in that domain by writing production rules that capture students multiple strategies and their common misconceptions. These productions are written in a modular fashion so that they can apply to a goal or context independent of how that goal was arrived at. For simplicity of illustration, I provide an example from the domain of equation solving:

| | |
|----------|--|
| Correct: | IF the goal is to solve $a(bx+c) = d$ THEN rewrite this as $bx + c = d/a$ |
| Correct: | IF the goal is to solve $a(bx+c) = d$ THEN rewrite this as $abx + ac = d$ |
| Buggy: | IF the goal is to solve $a(bx+c) = d$ THEN rewrite this as $abx + c = d$ |

The first two productions illustrate alternative strategies for the same problem solving goal. This allows the cognitive tutor to follow students down problem solving paths of their own choosing. The third "buggy" production represents a common misconception.

Such rules allow the cognitive tutor to recognize such misconceptions and provide appropriate assistance. The cognitive tutor makes use of the cognitive model to follow students through their individual approaches to a problem. It does so using a technique called "model-tracing". Model tracing allows the cognitive tutor to provide students individualized assistance that is just-in-time and sensitive to the students' particular approach to a problem.

The cognitive model is also used to trace students knowledge growth across problem solving activities. The knowledge tracing technique is dynamically updating estimates of how well the student knows each production rule. These estimates are used to select problem solving activities and to adjust pacing to adapt to individual student needs.

Cognitive Tutors have been subject to comparative evaluations in the lab and in classroom for more than 12 years (Anderson, Corbett, Koedinger, & Pelletier, 1995). A cognitive tutor for writing programs in the LISP computer language was compared to a control condition in which students solved the same programming problems without the aid of the cognitive tutor. Students in the experimental group completed the problems in *1/3 the time* with better post-test performance than students in the control group. The LISP tutor allowed students to engage in productive problem solving search, but reduced unproductive floundering. Two different Cognitive Tutors for geometry proof design were used in classroom studies compared to control classes using a traditional geometry curriculum without the cognitive tutor. In both studies, students in the experimental classes scored *1 sd better* than students in control classes. I will mention two important lessons from these studies. First, echoing results from experiments with LOGO (Lehrer, Randle, Sancillo, 1989; Klahr & Carver, 1988), we demonstrated that careful curriculum integration and teacher preparation was critical to our effectiveness results (Koedinger and Anderson, 1993). A second lesson came from a third party evaluator who studied changes in student motivation and classroom social processes as a consequence of the use of the Geometry Proof Tutor. Schofield, Evans-Rhodes, and Huber (1990) found the classroom evolved to be student centered with the teacher taking a greater facilitator role supporting students as-needed on the particular learning challenges each was experiencing. This point was repeated in a Math Teacher article by one of the participating teachers in which Wertheimer (1990) emphasized that because the cognitive tutor was effectively engaging students he was more free to provide individualized assistance to students who needed it.

The Pump Algebra Tutor (PAT)

Two experiences led us to take a different "client-centered" approach when we began to develop the Pump Algebra Tutor. First, at the time of the geometry studies, the NCTM Standards were coming out and suggesting a deemphasis on proof in high school geometry. Second, we had experienced the importance and difficulty of integrating the technology with the classroom and paper-based curriculum. Thus, in the Pump Algebra Tutor project we designed the tutor and the curriculum hand-in-hand. A high school math teacher, Bill Hadley, and a curriculum supervisor, Diane Briars, had been working on an algebra curriculum to make algebra accessible to more students, to help students make connections between algebra and the world outside of school, and to prepare students for the "world of work" as well as further academic study. We teamed up with Hadley and began evolving the curriculum and designing the tutor by sharing ideas from both

research and practice.

Functional Models of Authentic Problem Situations

Table 1 shows part of day one activities of a two day performance assessment used in the Pump Algebra curriculum. The goal of day one is not for students to produce a particular numerical answer, but rather to produce an analysis and a representation of that analysis in mathematical forms including tables, graphs, equations, and words. In other words, the students goal is to create mathematical models of the alternative situations presented in the problem. These models are then used on day two to answer specific questions about these choices and to write a full page recommendation using the models as support for the recommendations made. As Paul Goldenberg has said models provide a "general answer" to problems that can be used to answer multiple specific questions of the kind usually asked in math class. A major goal of Pump Algebra is to aid students in developing successively more sophisticated models of quantitative relationships using multiple representations each with different costs and benefits (cf., Tabachneck, Koedinger, & Nathan, 1995).

Table 1. A Performance-Based Assessment from the Pump Algebra Curriculum

My Life as a Photographer

Final - Day One

Your friend has decided he is very interested in a career as a photographer. You look up Photographer in the Pennsylvania Career Guide and find out that there are three different paths to becoming a photographer. You can enter the field upon graduating from high school, go to a technical school or attend college. You research these options and find out the following:

Option 1: High School Diploma. He can become a photographer with only a high school diploma. The average salary for these photographers is \$1115 per month.

Option 2: Technical School. Completing high school and attending a technical school for an eighteen month (1 and 1/2 years) course in photography is a second option. This technical program costs about \$18,000 and the average salary for those completing the course is \$1925 per month.

Option 3: College. Completing high school and attending a four year college program in photography is the third option. The average cost for a four year program is \$50,000 and a graduate can expect to earn about \$2745 per month.

You want to make a complete comparison of these three options for your friend. Since you are an excellent algebra student you want to use algebra to make this analysis, and then use this analysis to write a letter to your friend explaining clearly the advantages and disadvantages of each option. You also want to make a recommendation to him as to what he should do!

Learning to Model with Algebraic Symbols: The Inductive Support Strategy

To reach the goal of creating improved instructional supports to help students learn to be successful on assessments like the Photographer Career problem in Table 1, we began to research issues of mathematical modeling and the underlying competencies required. In particular, we focused on studying what students know and do not know about symbolic modeling. Skills for symbolic modeling are important because they are not currently automated, they are the entry point to using today's powerful calculation tools (e.g., graphic and symbolic calculators, spreadsheets, programming), and they are particularly

skills for students to acquire. We began to experiment with different approaches that might better aid students in learning to model with algebraic symbols.

Table 2 shows a problem from an algebra textbook. Hadley had been using a similar problem format, but put a particular emphasis on problem contexts that would be more authentic to students and that contained real data. Forester intended this problem format to illustrate the nature of an algebraic variable as truly varying in contrast to traditional algebra word problems (e.g., leave out questions 1-3 in Table 2) in which there is an unknown constant, but no variable. As I began to observe and analyze student thinking toward creating a cognitive model, I formed the simple instructional hypothesis that having students answer the concrete "result-unknown" questions 2 and 3 might facilitate their learning to translate such problems into the language of algebra (question 1).

Table 2. Textbook problem with different question types.

Drane & Route Plumbing Co.
charges \$42 per hour plus \$35 for the
service call.

1. Create a variable for the number of hours the company works. Then, write an expression for the number of dollars you must pay them.
2. How much you would pay for a 3 hour service call?
3. What will the bill be for 4.5 hours?
4. Find the number of hours worked when you know the bill came out to \$140.

Symbolization

Question

Result-
Unknown

Questions

Start-
Unknown

Question

This hypothesis followed from my prior cognitive science research on the importance of inductive experiences in the evolution of geometry knowledge (Koedinger & Anderson, 1990) and from observations of students who can successfully solve concrete result-unknown questions, like 2 and 3, but cannot produce the corresponding algebraic sentence. Such students already have effective performance knowledge or production rules for comprehending the English problem statement, for extracting the relevant quantities and quantitative relations, and for a producing a numerical answer. However, production rules for "speaking algebra", that is, for taking a problem understanding and expressing it in algebraic symbols, are either weak or missing (cf. Heffernan & Koedinger, 1997; 1998). One such production rule is illustrated below:

If the goal is express a quantity Q1 in algebraic symbols
and Q1 is result of combining Q2 and Q3 with operator Op
and the expression for Q2 is Expr2
and the expression for Q3 is Expr3
Then set a goal to write: Expr2 Op Expr3
set a goal to check for correct order of operations

This production rule characterizes tacit performance knowledge for composing algebraic

"embedded clauses" for a quantity, like $42h + 35$ for the total bill (Q1), from knowledge of simple clauses, like $42h$ for the hourly charge (Q2) and 35 for the service charge (Q3).

The rationale for having students do result-unknown questions before symbolization is this: By stepping through the arithmetic operations with concrete instances of quantities, students have a source from which to learn, by analogical induction (Anderson, 1993), production rules for producing algebraic sentences using the same operations and quantities. I call this use of concrete instances to help students induce algebraic sentences the *inductive support strategy*.

We used an early version of PAT to implement an experimental learning study in which we compared a control condition in which students solved problems in the Forester textbook format to an experimental group in which students solved same problems but with the questions rearranged so that the concrete result-unknowns (#2 and #3) would first, before the symbolization question (#1). Figure 1 shows the results of that study. Students in the Inductive Support experimental group learned significantly more from pre-test to post-test than students in the Textbook control group (Koedinger & Anderson, in press).

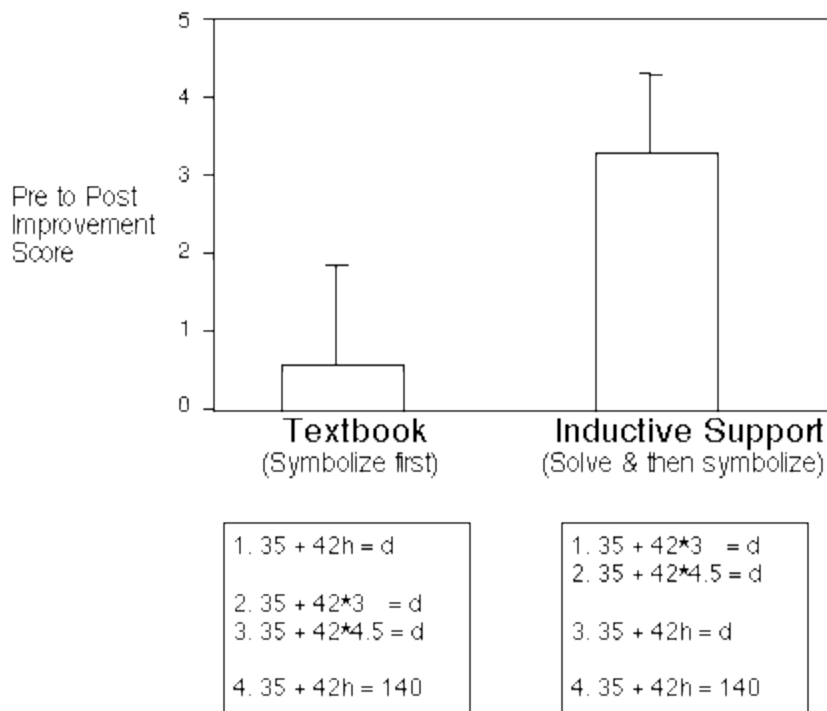


Figure 1. Improving algebraic modeling by bridging off students' existing knowledge.

The general point here, consistent with Paul Goldenberg's paper, is that students need to not only learn mathematical concepts, but also develop a mastery or fluency with mathematical modeling languages like algebra, programming, statistics notation, dynamic geometry tools, etc. Further, this experiment illustrates one effective way to assist students in developing mathematical language fluency (e.g., in algebra): Create instructional activities that help students bridge off existing fluency in known languages (e.g., arithmetic and English).

Description of PAT: A Cognitive Tutor for Practical Algebra

As part of the development of the Pump Algebra Tutor (PAT), Pittsburgh teachers wrote problem situations, like the Photography problem above, intended to be personally or culturally relevant to students. Some problem situations are of potential general interest (e.g., the decline of the condor population), while others are somewhat more specific to Pittsburgh 9th graders (e.g., making money shoveling snow or declining population in Pittsburgh after the demise of the steel mills). These problems were added to PAT using an intelligent problem authoring system in which teachers type the problem description and enter an example solution (Ritter, Anderson, Cytrynowitz, & Medvedeva, 1998). The authoring system has an incomplete and imperfect model for reading English text, but can make reasonable guesses as to how quantities in the entered solution map to phrases in the text. The author can edit and correct these guesses.

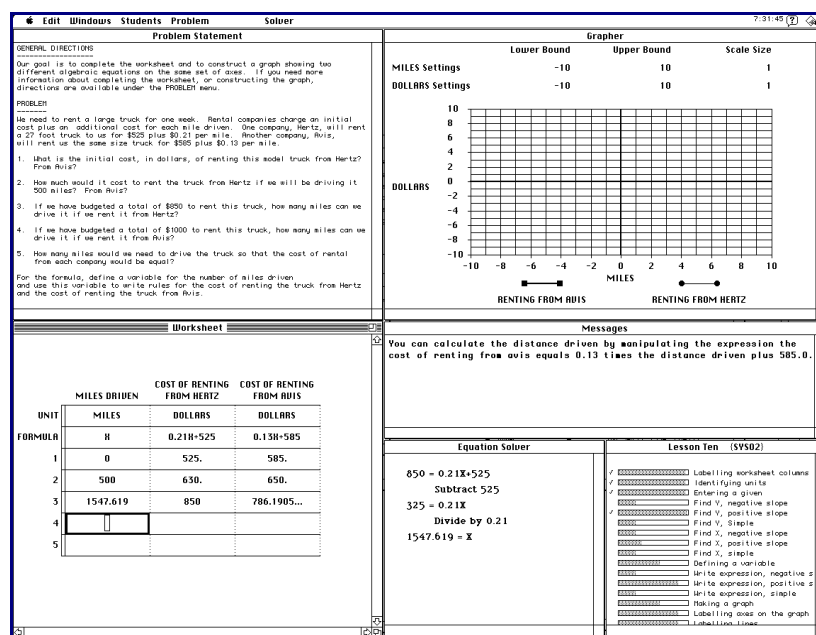


Figure 2. In the Pump Algebra Tutor (PAT), students create tabular (lower-left), graphical (upper-right), and symbolic (lower-center) models of problem situations (upper-left) with as-needed assistance (middle-right) and dynamic assessment (lower-right) from the Cognitive Tutor. ([click on image for high resolution image](#))

Students begin work on PAT problem situations by reading a description of the situation and a number of questions about it. They investigate the situation by representing it in tables, graphs, and symbols and using these representations to answer the questions. Helping students to understand and use multiple representations is a major focus of the tutor.

In Figure 2, the PAT screen shows a student's partial solution for a problem. This problem appears in later stages of the curriculum after students have acquired some expertise with constructing and using graphs and tables for single linear equations. The top-left corner of the tutor screen provides a description of the problem situation. The problem involves two rental companies, Hertz and Avis, that charge different rates for renting large trucks. Students investigate the problem situation using multiple

representations and computer-based tools, including a spreadsheet, grapher, and symbolic calculator -- in Figure 2 these are the Worksheet, Grapher and Equation Solver windows, respectively. Students construct the Worksheet (lower-left of Figure 2) by identifying the relevant quantities in the situation, labeling the columns, entering the appropriate units, entering algebraic expressions, and by answering the numbered questions in the problem description. Students construct the graph of the problem situation (upper-right) by labeling axes, setting appropriate bounds and scale, graphing the lines, and identifying the point of intersection. The Equation Solver (lower-center) can be used at any time to help fill in the spreadsheet and identify points of intersection. The student can use these representations to reason about real-world concerns, such as deciding when it becomes better to rent from one company rather than another.

Most students spend 20-30 minutes solving a problem of this type on the computer. During that time, the tutor monitors their activities, and provides feedback on what they are doing. The provision of timely feedback is one way in which the tutor individualizes instruction. Like good human tutors, PAT's tutorial interactions with students are brief and focused on individual students' particular learning needs as they come up in the context of problem solving. When a student is having trouble, PAT does not immediately provide correct answers or even detailed advice. Instead, PAT tries to maximize student opportunities to discover or reinforce appropriate concepts or skills on their own. PAT provides two kinds of assistance: just-in-time feedback on problem solving steps taken and on-demand hints on next steps needed.

When providing feedback, PAT does not make a big deal out of errors, usually only implicitly "flagging" them using ~~outline text~~. As Schofield et al. (1990), Wertheimer (1990), and many of our current teachers have observed, students do not feel the same social stigma when making errors on the computer as they do when making errors in class. Commonly occurring slips or misconceptions are recognized by "buggy" production rules and specific advice can be provided that may explain what is wrong with the current step or hint toward an appropriate correct step. Examples of buggy productions in PAT include leaving out the initial value (intercept) in a computation or formula, leaving out parentheses or otherwise violating order of operations, confusing the dependent and independent variable in a table, graph or formula, and many others.

This provision of timely feedback is a critical feature of Cognitive Tutors that leads to substantial cognitive and motivational benefits. In a parametric study with the LISP tutor, Corbett and Anderson (1991) provided a demonstration of how the timely feedback leads to dramatic reductions in the learning time needed to reach the same level of post-training performance. In addition to cognitive benefits, there are also motivational benefits of timely feedback. Much like the motivational attraction of video games, students know right away that they are making progress and having success at a challenging task.

In addition to just-in-time feedback, a second way PAT individualizes instruction is by providing context-sensitive hints. Using the model tracing mechanism described above, PAT is always following each student's particular approach to a problem. At any point in constructing a solution, a student can request a hint and PAT will provide one that is sensitive to what the student has done up to that point. The tutor chooses a hint message by using the production system to identify the set of possible next strategic decisions and ultimate external actions. It chooses among these based on the student's current focus of

activity, what tool and interface objects the student has selected, the overall status of the student's solution, and internal knowledge of relative utility of alternative strategies. Successive levels of assistance are provided so as to maximize the students' opportunities to construct or generate knowledge on their. This approach is consistent with the notions of cognitive apprenticeship (Collins, Brown & Newman, 1989; Vygotsky, 1978).

The "Message" window in Figure 2 shows the result of a student help request. The current focus of attention is based on the selection of the worksheet cell for question 4, under the column entry for 'miles driven' -- this cell is highlighted in Figure 2. Given the information in the problem about the costs of renting from Avis or Hertz, the student is asked: "If we have budgeted a total of \$1000 to rent this truck, how many miles can we drive it if we rent it from Hertz?"

An initial hint directs the student to consider information in the question that is relevant to finding a value for the distance: "You know that the cost of renting from Avis depends on the distance driven, and you are given a value for the cost of renting from Avis." By asking for help a second time, the tutor hints towards redescribing the problem statement in a solution-enabling representation: "You can calculate the distance driven by manipulating the expression the cost of renting from Avis equals 0.13 times the distance driven plus 585.0." This hint is an intermediate abstraction between the more concrete complete problem statement and a more abstract solution-enabling representation like an algebraic equation ($1000 = 0.13d + 585$) or an internal encoding of the problem constraints that facilitates a mental "unwinding" strategy (Koedinger & MacLaren, 1997). If needed, a further message provides the algebraic equation that the student needs to set up and solve. Such detailed hints are like the examples provided in textbooks illustrating a new idea. The difference is that, in PAT, these examples come in the process of problem solving when the student is better able to understand and make use of the example.

The Equation Solver window (lower-center) shows how the student solved a similar question (question 3). The student enters their own equation and solves it by indicating standard algebraic manipulations.

By keeping students engaged in successful problem solving, PAT's feedback and hint messages reduce student frustration and provide for a valuable sense of accomplishment. In addition to these functions of model tracing, PAT provides learning support through *knowledge tracing*. Results of knowledge tracing are shown to student and teacher in the Skillometer window. By monitoring a student's acquisition of problem solving skills through knowledge tracing, the tutor can identify individual areas of difficulty (Corbett, Anderson, & O'Brien, 1995) and present problems targeting specific skills which the student has not yet mastered. For example, a student who was skilled in writing equations with positive slopes and intercepts, but had difficulty with negative slope equations would be assigned problems involving negative slopes. Knowledge tracing can also be used for "self-pacing", that is, the promotion of students through sections of the curriculum based on their mastery of the skills in that section.

The activities in PAT are organized hierarchically so that related problem situations that draw on a core set of skills are organized in to "sections" and then sections that use the same set of notations, tools, and broader concepts are organized into "lessons". In the current PAT curriculum (used in schools in the 1997-98 school year) there were 22 lessons each of which contained about 4 sections on average and each section contained

about 5 required problems and about 5 additional problems. Initially, students explore common situations involving positive quantities, mostly whole numbers and some simple fractions and decimals, and represented these situations mathematically in tables and graphs. As the year progresses more complex situations are analyzed involving negative quantities and four quadrant graphing is introduced. Similarly, as situations increase in complexity, formal equation solving and graphing techniques are introduced to enable students to find solutions. Systems of linear situations and quadratics are developed through the introduction of situations in which they naturally occur, for example, modeling and comparing the price structures of two rival companies that make custom T-shirts. Modeling vertical motion and area situations provide contexts for introducing and using quadratic functions.

Classroom Context of PAT Use

Most schools using PAT are also using the Pump curriculum and text materials. The typical procedure is to spend 2 days a week in the computer lab using PAT and 3 days a week in the regular classroom. In the classroom, learning is active, student-centered, and focused primarily on learning by doing. Teachers spend less time in whole-group lecture and more time in individual and cooperative problem solving and learning. The teacher plays a facilitator role, but may do some whole-group lecture and discussion to highlight student discoveries or respond to recognized needs of the class as a whole.

In the classroom, students work together in groups or teams to solve problems similar to those presented by the tutor. Teams construct their solutions by making tables, expressions, equations, and graphs which they then use to answer questions and make interpretations and predictions. Teachers play a key role in helping students to make connections between the computer tools and paper and pencil techniques and to see how the general concepts and skills for representation construction and interpretation in each case are the same. Literacy is stressed by requiring students to answer all questions in complete sentences, to write reports and to give presentations of their findings to their peers.

The Pump curriculum uses alternate forms of assessment including performance tasks, long term projects, student portfolios, and journal writing. From the first day all answers must be written in complete sentences to be accepted. At the end of each quarter students are given a performance task as a final examination. At the end of each semester these tasks are graded by the teachers at a mini-scoring conference where teachers come together, construct a scoring rubric, and score all the student papers in an afternoon. Because each teacher scores papers from every other teachers' class as well as their own they come to have a better understanding of the objectives of the curriculum, what students know and do not know, and in what ways other teachers' students may differ.

Replicated Field Study Results

Our prior research with intelligent tutors (Koedinger & Anderson, 1993) and the work of others (Lehrer, Randle, Sancilio, 1989; Klahr & Carver, 1988) suggests that to be effective in improving student learning, educational technology must be closely integrated with curriculum goals and other learning resources like texts and teacher practices. Such integration has been a major goal from the start of the development of PAT. The consequence has been a dramatic impact on student learning and achievement. We have

demonstrated this impact in experimental field studies in city schools in Pittsburgh and Milwaukee, replicated over 3 years, and targeting both 1) higher order conceptual achievement as measured by performance assessments of problem solving and representation use and 2) basic skills achievement as measured by standardized test items, for instances, from the math SAT. In comparison with traditional algebra classes at the same and similar schools, we have found that students using PAT and the Pump curriculum perform 15-25% better than control classes on standardized test items and 50-100% better on problem solving & representation use.

Following the observations of Schofield, Evans-Rhodes, & Huber (1990) and Wertheimer (1990), we have also observed the impact of the use of PAT on changes in classroom social and motivational processes. Visitors to our classrooms often comment on how engaged students are. PAT may enhance student motivation for different reasons: 1) authentic problem situations make mathematics more interesting, sensible, or relevant, 2) students on the average would rather be doing than listening and the incremental achievement and feedback within problems has a video-game-like appeal, 3) the safety net provided by the tutor reduces potential for frustration and provides assistance on errors without social stigma, 4) the longer term achievement of mastering the mathematics is empowering.

In the computer lab, teachers are glad to essentially have a teacher's aid for every student and thus are freed to be facilitators and provide more one-on-one instruction with individual students. This experience is eye-opening for many teachers who may see new aspects of student thinking and feel the advantages of greater student-centered learn by doing.

Cognitive Tutors as Teacher Support and Change Agents

How and why does Cognitive Tutor use facilitate the spread of effective teaching principles and practices and the institution of curriculum reform? Is there something special about Cognitive Tutors that makes such spread more likely than alternative educational technologies like books, traditional CAI, or malleable tools? Unlike other educational technologies, Cognitive Tutors have a running model of student thinking and of adaptive student-centered instruction. Thus, the system provides an active "living example" of research-based principles and practices. Furthermore, by essentially serving as an extra teacher's aid for each student in the classroom, Cognitive Tutors free teachers to observe individual student thinking more often and more closely and to reflect on their instructional practices in the context of more one-to-one interaction with students. Student responses in such close interactions provide teachers with immediate and detailed feedback on the effectiveness of their practices that they can then use to adjust those practices.

A DoE-FIPSE project in which we adapted PAT for use in college-level developmental math courses has provided the opportunity to observe PAT-inspired changes in curricula and teacher practices over multiple semesters. These changes were accompanied by significant quantitative improvements in student learning. PAT was initially used at the University of Pittsburgh in the Fall semester of 1995 as an add-on to a traditional "Intermediate Algebra" course. As shown in the left of Figure 3, this supplementary use led to significant improvement in students' problem solving abilities over a traditional

course without PAT (45% vs. 30% as measured by a performance-based Cell Phone assessment much like the Photographer assessment in Table 1). In the Spring of 1996, instructors used the PAT authoring tool to create new problems better adapted to the particular needs and interests of their students. More interestingly, changes were not limited to the software. Instructors began to use PAT problems in their regular classes and began to experiment with more student-centered learn by doing outside of the computer lab. The consequence of these new practices was an increase in end-of-course achievement beyond that found in the experimental classes in the Fall (65% vs. 45%). Over the next two semesters, these practices evolved and PAT use became better integrated with further modest increases in student learning over past semesters (68%, 71%). In preparation for the Spring of 1997, the University of Pittsburgh made a decision to fully reform their intermediate algebra course to more generally target "quantitative literacy". Lecturing was deemphasized in favor of more "workshop" time in which students worked on projects in collaborative groups. The typical college level approach of a lot of instructor-centered lecture time and a little student-centered recitation time was replaced with a lot of student-centered workshop and a little whole-group reflection and targeted lectures. The consequence of these changes, in the first semester of their implementation, we yet another increase in end-of-course student performance (79%).




Figure 3. Increasing PAT-inspired reforms leads to increasing scores across five semesters.

Conclusions

I believe that effective technologies for learning and doing mathematics should be based on sound *cognitive theory*, be *empirically tested against alternatives*, and be primarily addressed at *mathematics as a modeling language*. I have argued for and illustrated these points in the context of Cognitive Tutors and in particular, the Pump Algebra Tutor (PAT). PAT is based on the ACT theory of cognition and a production rule model of student problem solving and mathematical modeling. PAT has been subject to empirical tests both in the laboratory and in the field. In a formative laboratory study with PAT we showed that we could improve students learning of symbolic modeling skills over a textbook approach by having students first solve concrete arithmetic versions of problems before abstracting these operators in an algebraic model. In field studies of the use of PAT and the Pump curriculum we demonstrated the combination leads to dramatic

increases in student learning on both standardized test items (15-25% better than control classes) and new standards-oriented assessments of problem solving and representation use (50-100% better than control classes). The approach of PAT and the Pump curriculum focus is on developing students competence in creating mathematical models of problem situations rather than on answers to isolated questions. Mathematical models provide the power to garner a deeper understanding of a problem situation such that multiple, unanticipated questions can be addressed and answered.

Cognitive Tutors, like PAT, have the potential not only to dramatically increase student achievement, but also serve a professional development function for teachers. Because of the cognitive model, Cognitive Tutors can provide a living example of effective instructional practices. Teachers working in the computer lab have more time to observe student performance on thought revealing problems and have observed learn-by-doing instruction in action. In this way, Cognitive Tutors can carry research-based practices into the classroom and serve as change agents for professional development.

References

- Anderson, J. R. (1993). *Rules of the Mind*. Hillsdale, NJ: Erlbaum.
- Anderson, J. R., Corbett, A. T., Koedinger, K. R., & Pelletier, R. (1995). Cognitive tutors: Lessons learned. *The Journal of the Learning Sciences*, 4 (2), 167-207.
- Collins, A., Brown, J. S., & Newman, S.E. (1989). Cognitive apprenticeship: Teaching the crafts of reading, writing, and mathematics. In L. B. Resnick (Ed.), *Knowing, learning, and instruction: Essays in honor of Robert Glaser*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Corbett, A. T. & Anderson, J. R. (1991). Feedback control and learning to program with the CMU LISP tutor. Paper presented at the annual meeting of the American Educational Research Association, Chicago, IL.
- Corbett, A. T., Anderson, J. R. & O'Brien, A. T. (1995). Student modeling in the ACT programming tutor. In P. Nichols, S. Chipman, & R. Brennan (Eds.), *Cognitively diagnostic assessment* (pp. 19-41). Hillsdale, NJ: Erlbaum.
- Heffernan, N. & Koedinger, K. R. (1997). The composition effect in symbolizing: The role of symbol production vs. text comprehension. In *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum. [Marr Prize Winner]
- Heffernan, N. & Koedinger, K. R. (1998). A developmental model for algebra symbolization: The results of a difficulty factors assessment. Accepted to the Eighteenth Annual Conference of the Cognitive Science Society.
- Klahr, D., & Carver, S. M. (1988). Cognitive objectives in a LOGO debugging curriculum: Instruction, learning, and transfer. *Cognitive Psychology*, 20, 362-404.
- Koedinger, K. R., & Anderson, J. R. (in press). Illustrating principled design: The early evolution of a cognitive tutor for algebra symbolization. To appear in *Interactive Learning Environments*.

- Koedinger, K. R., Anderson, J. R., Hadley, W. H., & Mark, M. A. (1997). Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education*, 8, 30-43.
- Koedinger, K. R., & MacLaren, B. A. (1997). Implicit strategies and errors in an improved model of early algebra problem solving. In *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Kulik, C. C. & Kulik, J. A. (1991). Effectiveness of computer-based instruction: An updated analysis. *Computers in Human Behavior*, Vol. 7, pp. 75-95.
- Lehrer, R., Randle, L., & Sancilio, L. (1989). Learning preproof geometry with LOGO. *Cognition and Instruction*, 6, 159-184.
- Miller, C. S., Lehman, J. F., & Koedinger, K. R. (in press). Goals and learning in microworlds. Accepted to *Cognitive Science*.
- Pane, J.F., Corbett, A.T., and John, B.E. (1996). Assessing Dynamics in Computer-Based Instruction. In *Proceedings of ACM CHI'96 Conference on Human Factors in Computing Systems*. Vancouver, pp. 197-204.
- Ritter, S., Anderson, J. R., Cytrynowitz, M., & Medvedeva, O. (1998). Authoring content in the PAT algebra tutor. *Journal of Interactive Media in Education*, 2.
- Singley, M. K., & Anderson, J. R. (1989). *Transfer of Cognitive Skill*. Hillsdale, NJ: Erlbaum.
- Schofield, J. W., Evans-Rhodes, D., & Huber, B. R. (1990). Artificial intelligence in the classroom: The impact of a computer-based tutor on teachers and students. *Social Science Computer Review*, 8:1, 24-41.
- Tabachneck, H. J. M., Koedinger, K. R., & Nathan, M. J. (1995). A cognitive analysis of the task demands of early algebra. In *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Vygotsky, L. S. (1978). *Mind in society*. Cambridge, MA: Harvard University Press
- Wertheimer, R. (1990). The geometry proof tutor: An "intelligent" computer-based tutor in the classroom. *Mathematics Teacher*, April, 308-317.