

Intelligent Tutoring Goes To School in the Big City

Kenneth R. Koedinger
Human-Computer Interaction Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
USA

ken.koedinger@cs.cmu.edu
<http://sands.psy.cmu.edu/ACT/people/koedinger.html>

John R. Anderson
William H. Hadley
Mary A. Mark
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213
USA

Abstract. This paper reports on a large-scale experiment introducing and evaluating intelligent tutoring in an urban High School setting. Critical to the success of this project has been a client-centered design approach that has matched our client's expertise in curricular objectives and classroom teaching with our expertise in artificial intelligence and cognitive psychology. The Pittsburgh Urban Mathematics Project (PUMP) has produced an algebra curriculum that is centrally focused on mathematical analysis of real world situations and the use of computational tools. We have built an intelligent tutor, called PAT, that supports this curriculum and has been made a regular part of 9th grade Algebra in 3 Pittsburgh schools. In the 1993-94 school year, we evaluated the effect of the PUMP curriculum and PAT tutor use. On average the 470 students in experimental classes outperformed students in comparison classes by 15% on standardized tests and 100% on tests targeting the PUMP objectives. This study provides further evidence that laboratory tutoring systems can be scaled up and made to work, both technically and pedagogically, in real and unforgiving settings like urban high schools.

1. Introduction

We have created an intelligent tutoring system for algebra problem solving that we call PAT. PAT stands for PUMP Algebra Tutor or Practical Algebra Tutor. PAT is practical in two ways. First, PAT is practical in its pedagogical focus. Students engage in investigations of real world problem situations and use modern algebraic tools (spreadsheets, graphers, and symbolic calculators) to express covariance relationships, to solve problems and to communicate results. Second, PAT is practical in going beyond a laboratory prototype to a fully functioning system. It is currently being used by more than 500 high school students in 3 Pittsburgh city schools.

This paper reports on a large-scale experiment in the classroom implementation of PAT. We start with a description of the system design -- a marriage of content guidance, provided by experts in mathematics pedagogy, and scientific support, provided by the ACT theory and cognitive tutoring technology (Anderson, Corbett, Koedinger, & Pelletier, 1995; Anderson & Pelletier, 1991). Next, we present the results of the first formative evaluation of this system. Students in tutor-using classes outperformed students in comparison classes by 15% on standardized tests and 100% on tests that emphasized real world problem solving and multiple mathematical representations.

2. A Curriculum and Cognitive Tutor for Practical Algebra

2.1. Client-Centered Design and the PUMP Curriculum

The PAT tutor has been developed through a collaboration between the Pittsburgh Urban Mathematics Project (PUMP) in the Pittsburgh School System and the cognitive tutoring group at Carnegie Mellon University. Critical to the success of this project has been a client-centered design approach that has matched our client's expertise in curricular objectives and classroom teaching with our expertise in artificial intelligence and cognitive psychology.

As part of the Pittsburgh Urban Mathematics Project (PUMP), local mathematics teachers have produced a more accessible algebra curriculum that focuses on mathematical analysis of real world situations and use of computational tools. The PUMP curriculum materials employ "real world" situations designed to make mathematics more meaningful and accessible to students. All students come to high school with experience of the mathematics used in "everyday" life, but many are unable to connect this to "school" mathematics (Resnick, 1987). The PUMP curriculum materials try to bridge this gap by using situations from everyday life to generate the mathematics and as a means for the students to anchor their knowledge (cf. CTG, 1990). The unifying concept of the PUMP Algebra curriculum is the use of functional models, represented variously in tables, graphs, and symbols, to analyze and explore situations. The PUMP curriculum is consistent with the new curriculum recommendations of the National Council of Teachers of Mathematics (NCTM, 1989). NCTM recommends increased attention to the use of real-world problems, use of computer utilities, mathematical communication, and making connections. They also recommend decreased attention to traditional word problems by type (e.g., coin, work, mixture), the simplification of radicals, factoring polynomials, and other paper-and-pencil techniques.

In the PUMP classroom students work on mini-projects investigating problem situations like comparing the current quantity and growth rate of old growth forest in the US to the harvest rate. Students investigate such situations by (1) addressing questions, like "Assuming these figures do not change, when will all the old growth forest be gone?", (2) creating a table to investigate the relationships between quantities, (3) scaling, graphing and identifying points of intersection, (4) using algebraic notation to concisely represent the underlying structure of the situation, and (5) using algebraic notation to compute solutions.

PAT was built to support this kind of mathematical investigation and problem solving. Most importantly, PAT was designed to help students develop algebraic skills which they can use in the context of real-life problem situations. The PAT learning environment includes a set of computational tools to aid investigation (a spreadsheet, grapher, and symbolic calculator) and an organized curriculum of problem situations. In developing PAT, we worked closely with both curriculum designers from the school system, and teachers in actual classrooms with actual students, all of whom have given us valuable information. Direct observation and protocols from tutoring sessions provide rich sources of evidence that we have drawn on to increase our understanding of students and to improve the design of PAT.

2.2. Principled Design of Cognitive Tutors

The design of PAT was also guided by theoretical principles. As a *cognitive tutor* (Anderson, et. al, 1995), PAT has the defining feature of containing a psychological model of the cognitive processes behind successful and near-successful student performance. Based on the ACT theory, this cognitive model is written as a system of if-then production rules that are capable of generating the multitude of solution steps and mis-steps typical of students. The cognitive model is the basis for two student modeling techniques: model tracing and knowledge tracing. *Model tracing* is used to monitor student's progress through a problem solution (see Anderson, Boyle, Corbett, & Lewis, 1990). This tracing is done in the background by matching student actions to those the model might generate. The tutor is mostly silent. However, when help is needed, the tutor knows where the student is and can provide hints that are individualized to the student's particular approach to the problem. *Knowledge tracing* is used to monitor students' learning from problem to problem (see Corbett & Anderson, 1992). A Bayesian estimation procedure identifies students' strengths and weaknesses relative to the production rules in the cognitive model. This assessment information is

used to individualize problem selection and optimally pace students through the curriculum.

PAT's cognitive model and general design is the consequence of basic research on mathematical cognition. Our previous research has shown that students have informal *inductive* routes to mathematical knowledge that often precede formal instruction in the deductive use of symbols (Koedinger & Anderson, 1990; 1991). Thus, contrary to popular belief, students can perform better on algebra word problems under certain circumstances than on the equivalent algebraic equations (Koedinger & Tabachneck, 1995). We applied such results in early experiments with PAT where we showed that students learned more from a theory-inspired "inductive-support" version of the tutor than from a "textbook" version based on a popular Algebra text (Koedinger & Anderson, 1996).

2.3. Description of PAT: A Cognitive Tutor for Practical Algebra

In day to day life, people deal with a wide variety of situations that cause them to draw on basic algebra and reasoning skills. Checking the amount of a paycheck, estimating the cost of a rental car for a trip, and choosing between long-distance telephone service offers from AT&T and MCI are just three examples of real-world situations in which algebraic skills are useful. As part of the development of PAT, Pittsburgh teachers wrote problem situations like these 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 more specific to Pittsburgh 9th graders (e.g., making money shoveling snow). These problems were added to PAT using a problem authoring environment in which teachers type the problem description, enter an example solution, and edit the guesses the system makes about how quantities in the solution to relate to phrases in the text.

Students work through PAT problem situations by reading a textual 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 of information is a major focus of the tutor.

The screenshot displays the PAT interface with several windows:

- Problem Statement:** Describes a rental truck problem involving Hertz and Avis. It asks for initial costs, costs at 500 miles, and the number of miles for a \$850 budget.
- Grapher:** A coordinate plane with 'MILES' on the x-axis and 'DOLLARS' on the y-axis. It shows settings for lower and upper bounds and scale size.
- Worksheet:** A table with columns for Miles Driven, Cost of Renting from Hertz, and Cost of Renting from Avis. It includes a formula row and five numbered rows for data entry.
- Equation Solver:** Shows the algebraic steps to solve $850 = 0.21x + 525$, resulting in $x = 1547.619$.
- Messages:** A message box stating: "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."
- Lesson Ten (SYS02):** A checklist of skills such as "Labeling worksheet columns", "Identifying units", and "Entering a given".

Figure 1. A screen from PAT, a cognitive tutor for practical algebra.

In Figure 1, 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 1 these are the Worksheet, Grapher and Equation Solver windows, respectively. Students construct the Worksheet (lower-left of Figure 1) 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. For the most part, the tutor is silently tracing students actions in the background. When a student makes an error, it is "flagged". For text items, flagging is achieved by putting the student's entry in **outline text**. Errors in plotting points in the grapher tool are flagged by coloring the point gray rather than black and indicating the coordinates of the incorrectly placed point so that the student can see how they differ from the intended coordinates in a Worksheet row. Often flagging is done without comment, which appears to reduce students' negative feelings associated with making errors in math class. But if the student's error is a commonly occurring slip or misconception that has been codified in a buggy production rule, a message is provided that indicates what is wrong with the answer or suggests a better alternative. Examples of buggy productions in PAT include putting a correct value for a cell in the Worksheet in an adjacent row or column, confusing the dependent and independent variable in formula writing, incorrectly entering arithmetic signs in equation solving, and confusing the x and y coordinates in graphing.

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 immediacy of feedback leads to dramatic reductions in the learning time needed to reach the same level of post-training performance -- learning time was 3 times longer in the most delayed feedback condition than in the most immediate. 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. Further, because the system does not make a big deal out of errors, students do not feel the social stigma associated with making an error in class or on homework. Errors are a private event that are usually quickly resolved and the student is then back to making progress.

In addition to error feedback, a second way PAT individualizes instruction is by giving help on request. At any step in constructing a solution, a student can ask for help. The tutor chooses help messages for presentation by using the production system to identify a desirable next activity. Choice of a desirable action is based on the student's current focus of activity, the overall status of the student's solution, and internal knowledge of interdependencies between problem-solving activities as represented in the production rules. Multiple levels of help are provided

so that more detailed information can be obtained by making repeated help requests.

The "Message" window in Figure 1 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 1. 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 student receives a more detailed description suggesting that the distance can be calculated by relating information given in the question to a particular algebraic relationship described in the problem. "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." Further messages are also available, describing in more detail the type of equation that the student needs to set up and solve. 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, Carver, and Brancolini, 1994) 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 lessons of the curriculum based on their mastery of the skills in that lesson. In the 1993-94 study, knowledge-tracing capabilities of the tutor were not fully used. Knowledge tracing controlled the selection of problems within lessons, but not the self-paced advancement of students between lessons.

Self-pacing was not used for two reasons. First, participating teachers were not certain how to coordinate students' differing rates of progress through PAT lessons with the material being addressed in the regular classroom. Teachers were already tackling a number of new challenges in using PAT and in simply using computers in classroom. Second, the researchers needed more student data to decompose domain competence into production rules that best match the grain size of algebra learning events.

Instead of self-pacing progress between lessons, students spent a fixed amount of time on each lesson, about 3-4 class periods. At the end of each such installment, all students were advanced to the next lesson whether or not knowledge tracing had judged them to have mastered the skills in that lesson. The current 1994-95 study fully utilizes knowledge-tracing capabilities of the tutor, within and between lessons.

The PAT curriculum for the 1993-94 school year contained six lessons of problems. Initially, students explored common situations involving positive quantities and graphing in the first quadrant of the Cartesian coordinate system (positive values only on the x and y axis's). As the year progressed more complex situations were analyzed that required negative quantities and graphing in the other quadrants. Similarly, as the situations increased in complexity, formal equation solving and graphing techniques were 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, two rival companies that make custom T-shirts with different price structures provided an opportunity to explore a system of two linear equations. Modeling vertical motion provided a context for introducing and using quadratic functions. Problems involving quadratic functions are part of the PUMP curriculum, but were not yet implemented in PAT.

2.4. Special Features of the PUMP Classroom

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. The transfer of the computer tools to paper and pencil techniques and the interpretation and understanding of these tools are the emphasis of the classroom. 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 project also 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 all the teachers in the project 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 this new curriculum.

3. A Large-Scale Classroom Experiment

The on-going evaluation of PAT and the PUMP curriculum is a kind of "design experiment" (Brown, 1992) on the effect of both instructional innovations in the unforgiving setting of real schools. Evolving versions of PAT have been tested in laboratory experiments following the cognitive tutoring design methodology (Koedinger & Anderson, 1996; Anderson, et al., 1995). However, the urban classroom situation is unlike the refined and controlled environment of the lab and laboratory standards cannot realistically be applied. As such, we have begun by addressing the practical question of whether the whole package, PUMP curriculum and PAT, is effective by comparing it against a traditional curriculum without PAT. By laboratory standards, this experiment is confounds two variables, a change in curriculum and the use of PAT. However, our strategy is first to establish the success of the whole package and then, if indeed it is successful, to examine the effect of the curriculum and intelligent tutoring components independently.

3.1. Method

Data reported is from the 1993-94 school year. The student population came from 3 Pittsburgh Public High Schools, Langley, Brashear and Carrick, with similar demographics and student aptitudes. These schools are about 50% African-American, 50% single-parent families, and only 15% go on to college. Students in the experimental classes received two treatments: they were taught the new PUMP curriculum and they worked with PAT for approximately 25 out of 180 of their normal class periods. Students worked on 6 lessons with the complete PAT environment and 1 lesson with the equation solving tutor module alone. The "PUMP+PAT" group consisted of 20 algebra classes that involved 470 students and 10

teachers. The 12 classes from Langley high school contained students who in prior years would have been placed into a non-academic general math class, rather than algebra. Because of satisfaction with the pilot use of the curriculum and tutor in the prior year, Langley decided to assign all 9th graders to algebra instead.

The comparison classes received a traditional curriculum and did not use PAT. There were two types of comparison classes. The matched "Comparison" group consisted of 5 algebra classes that involved 120 students and 3 teachers. These students were from roughly the same background as the experimental classes. If anything, Comparison students were somewhat better prepared as a group given the inclusion of students in the Langley experimental classes who would otherwise have been placed in a lower level math. The "Scholars-Comparison" group consisted of 2 "scholars" algebra classes involving 35 students and 1 teacher. Scholars courses are an academic track for students who are selected based on prior school success. None of the experimental classes were scholars algebra classes.

We looked at students' math grades in the previous school year to verify that there were no differences in students prior mathematical background that would put the PUMP+PAT group at an advantage. Using 1 for a D and a 4 for an A, the grade averages were lowest at 2.1 for PUMP+PAT, next at 2.4 for the Comparison group, and highest at 2.6 for the Scholars-Comparison group.

3.2. Assessment Design

Designing a fair assessment plan for an experiment involving curriculum reform is difficult. Standardized tests are often rejected for this purpose because they do not address the objectives of the new curriculum. However, we reasoned that if PUMP+PAT students did better on reform objectives yet were worse on the basic skills tapped by standardized assessments, we would have just shifted the focus of instruction. Such assessments would in effect reflect creation of a new course: not necessarily a bad goal, but not evidence of an improvement in the instructional process. We wanted to show experimental classes doing much better on the reform objectives of authentic problem solving and representational tool use, and at least as well or better on the basic skills tapped by standardized tests. Thus, we gave both types of tests.

We used two types of standardized tests: an Iowa Algebra Aptitude test and a subset of the Math SAT appropriate for 9th graders. We also designed two tests to assess reform objectives reflecting NCTM's recommendations and the PUMP curriculum. The Problem Situation Test

was created to assess students' abilities to investigate problem situations, presented verbally, that have algebraic content. The Representations Test was created to assess students' abilities to translate between representations of algebraic content including verbal descriptions, graphs, and symbolic equations.

The assessments were given over two days at the end of the spring semester during a normal 44 minute class period. All students took the Iowa on the first day of testing. On the second day, approximately half of the students took each of the other tests.

3.3. Results and Discussion

Table 1 shows the percentage correct, standard deviation (in parentheses), and N (second line) for the groups on the four tests. Note that because of the high absenteeism that is typical of city schools particularly near the end of the school year, there were a substantial number of students who missed one or both days of the assessment. For each test, a between-subjects ANOVA was run with three levels, Comparison, PUMP+PAT, and Scholars-Comparison. The addition of the Scholars-Control group provides a upper edge comparison for the intervention. The results of these tests are shown in the 5th column of Table 1.

The 6th and final column shows effect sizes in terms of standard deviation units (sigma) of the PUMP+PAT group above the Comparison group. Effect sizes provide a metric for assessing the impact of instructional interventions. The Bloom (1984) result that individual human tutors can bring students 2 sigma above normal classroom instruction sets a standard of comparison for the impact of intelligent tutors. Previous studies have shown cognitive tutors to yield as much as a 1 sigma effect over control conditions (Anderson, Corbett, Koedinger, and Pelletier, 1995; Koedinger & Anderson, 1993).

Table 1. Results of Final Testing of PUMP+PAT, Comparison, and Scholars-Comparison Classes

| | Comparison | PUMP+PAT | Scholars-Comparison | F value and significance | sigma |
|------------------------|-----------------|------------------|---------------------|-------------------------------|-------|
| Iowa Algebra Aptitude | .46 (.17) 80 | .52 (.19) 287 | .68 (.17) 34 | F(2,398) = 17.0 p < .0001 | 0.3 |
| Math SAT Subset | .27 (.14) 44 | .32 (.16) 149 | .42 (.15) 15 | F(2, 205) = 5.1 p < .01 | 0.3 |
| Problem Situation Test | .22 (.22) 42 | .39 (.33) 127 | .38 (.26) 20 | F(2, 186) = 5.3 p < .01 | 0.7 |
| Representations Test | .15 (.18) 44 | .37 (.32) 124 | .12 (.16) 18 | F(2, 183) = 13.4 p < .0001 | 1.2 |

On the Iowa Algebra Aptitude test, PUMP+PAT scores are significantly higher than the Comparison ($p < .05$), a 0.3 sigma effect. They are significantly lower than the Scholars-Comparison ($p < .01$). On the SAT subset, the PUMP+PAT scores are higher than the Comparison scores (0.3 sigma), but there is a lot of variance in this smaller sample and PUMP+PAT is not significantly higher than Comparison ($p > .05$) nor significantly lower than Scholars-Comparison ($p > .05$). The largest effects come on the new NCTM-oriented tests. On the Problem Situation test, PUMP+PAT scores are significantly and substantially higher than Comparison ($p < .01$), a 0.7 sigma effect. They match up with the Scholars-Comparison ($p > .05$). On the Representations test, PUMP+PAT scores are significantly and substantially higher than Comparison ($p < .01$), a 1.2 sigma effect, and than the Scholars-Comparison ($p < .01$).

To summarize, the PUMP+PAT classes scored about 1 sigma better on the NCTM-oriented tests that were the target of the curriculum. Their scores were about 100% better or double those of the Comparison classes. These learning gains appear to occur at no expense to basic skills objectives of standardized tests. In fact, PUMP+PAT classes scored about 15% better on these tests.

4. Conclusion

Cooperation between the Pittsburgh Urban Mathematics Project and the cognitive tutoring group at Carnegie Mellon has led to the development of the PAT tutor, and its integration into classrooms in three Pittsburgh Public Schools. As expected from past experience (Schofield, Evans-Rhodes, & Huber, 1990; Wertheimer, 1990), the tutor has been enthusiastically received by students and teachers. Teachers comment that working in the computer lab with PAT engages students who present difficulties in the normal classroom. In addition, teachers like the way that the tutor accommodates a large proportion of student questions and frees teachers to give more individualized help to students with particular needs. As one concrete example of teacher support, a teacher's enthusiastic testimonial of the program was critical in convincing the Pittsburgh school board to purchase computer labs to expand the program to two more high schools for the 1994-95 school year.

Evaluation of PAT and the PUMP curriculum is continuing. In the 1994-95 school year, the PAT curriculum expanded to include 10 lessons and 214 problem situations. Students are in the computer lab two days a week, working with PAT at a self-paced rate. Student time on the tutor

will more than double (roughly from 25 to 70 days) compared to the 93-94 school year.

Acknowledgments

This research was supported by contract MDR-92-53161 from the National Science Foundation.

Thanks to the teachers of the experimental and comparison classes. Thanks also to other key members of our Pittsburgh Advanced Cognitive Tutor center: Al Corbett, Ray Pelletier, Steve Ritter, Jackie Snyder.

References

- Anderson, J. R., Boyle, C. F., Corbett, A. T., & Lewis, M.W. (1990). Cognitive modeling and intelligent tutoring. *Artificial Intelligence*, 42, 7-49.
- 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.
- Anderson, J. R., & Pelletier, R. (1991). A development system for model-tracing tutors. In *Proceedings of the International Conference of the Learning Sciences* (pp. 1-8). Evanston, IL.
- 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, 3-16.
- Brown, A. L. (1992). Design experiments: Theoretical and methodological challenges in creating complex interventions. *The Journal of the Learning Sciences*, 2, 137-178.
- 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. (1992). Student modeling and mastery learning in a computer-based programming tutor. In *Proceedings of the Second International Conference on Intelligent Tutoring Systems*. Montreal, Canada.
- Corbett, A.T., Anderson, J.R., Carver, V.H. and Brancolini, S.A. (1994). Individual differences and predictive validity in student modeling. In A. Ram & K. Eiselt (eds.) *The Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum.

- Cognition and Technology Group at Vanderbilt (1990). Anchored instruction and its relationship to situated cognition. *Educational Researcher*, 19 (6), 2-10.
- Koedinger, K. R., & Anderson, J. R. (1990). Abstract planning and perceptual chunks: Elements of expertise in geometry. *Cognitive Science*, 14, 511-550.
- Koedinger, K. R., & Anderson, J. R. (1991). Interaction of deductive and inductive reasoning strategies in geometry novices. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*, Hillsdale, NJ: Erlbaum.
- Koedinger, K. R. & Anderson, J. R. (1993). Effective use of intelligent software in high school math classrooms. In *Proceedings of the World Conference on Artificial Intelligence in Education*, (pp. 241-248). Charlottesville, VA: Association for the Advancement of Computing in Education.
- Koedinger, K. R., & Anderson, J. R. (1996). Illustrating principled design: The early evolution of a cognitive tutor for algebra symbolization. Submitted to *Interactive Learning Environments*.
- Koedinger, K.R., & Tabachneck, H.J.M. (1995). Verbal reasoning as a critical component in early algebra. Paper presented at the annual meeting of the American Educational Research Association, San Francisco, CA.
- National Council of Teachers of Mathematics (1989). *Curriculum and Evaluation Standards for School Mathematics*. Reston, VA: The Council.
- Resnick, L. B. (1987). Learning in school and out. *Educational Researcher*, 16(9), 13-20.
- 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.
- Wertheimer, R. (1990). The geometry proof tutor: An “intelligent” computer-based tutor in the classroom. *Mathematics Teacher*, April, 308-317