Intelligent Tutoring and High School Mathematics

John R. Anderson

Department of Psychology, Carnegie Mellon University
Pittsburgh, PA 15213, USA

1 The State of American Mathematics Education

The situation with respect to low educational achievement has been raised to the status of a national crisis in America. For no subject is the problem felt more acutely than mathematics education. American students enter school scoring somewhat lower than students of most first-world countries and leave school scoring much worse. This is not just a matter of the fact that American schools have to teach more difficult to teach students. In a comparison of fifth grade classrooms, [14, 15] no classroom was found in Minneapolis (USA) with a better math score than any classroom in Taipei (Taiwan) or Sendai (Japan). The average Japanese twelfth grader scores higher than the average of the top 5 percent of American students [9]. The comparisons between America and Japan define the extremes with most other first-world countries coming somewhere between. The contrast between America and Japan is also much more extreme for mathematics achievement than it is for other subjects such as reading.

In the context of an international conference, it is interesting to discuss the possible causes of America's low international standing in mathematics education. Here I will present some views of where we stand from the American perspective leaving to discussion the perspectives from other countries. The popular explanation among American mathematics educators is to point to the American mathematics curriculum and call for reform. However, one could argue that curriculum reform is part of the problem and not the solution. American mathematics education has been in a constant state of reform ever since the new math movement. American teachers face the need to change what they teach much more often than their counterparts in other countries and complain bitterly about the fact [14]. However, one does need to recognize that technology is changing the nature of mathematics and every country is facing the need to change its curriculum to face this fact. The need for curriculum reform is a necessary consequence of changes in our society but it seems extremely implausible that it will change the achievement level of American students. Without other changes American students will be as bad at the new curriculum as they were at the old.

One analysis, The Underachieving Curriculum [9] written by mathematics educators, considers four other popular explanations which it rejects before voting for the golden road of curriculum reform. One is class size which is the popular explanation of teachers. It points out that there is little relationship between national class size and achievement. For instance, the average class size in Japan is 41 and in America 26. Studies within America have found little relationship between class size and achievement. Stigler and Perry [15] noted one reason why. As class size decreases, American teachers are more and more tempted to individualize instruction to a single student or subgroups of students, interacting with those students while other students are receiving no instruction. In contrast in large classes teachers will spend more time teaching to the whole class. This means that the actual time the student is instructed is often less in the smaller class.

Their second popular explanation is that America aspires to teach a larger portion of its students than do other countries. There is some truth to this. Their figure is that only 17 percent of English students complete school whereas 82 percent of American students. However, there are some troubling counterexamples to this explanation. For instance, Japan retains 92 percent of its students.

The third explanation they consider is perhaps the most popular among the American public: the quality of the teachers. McKnight et al. [9] point out that, contrary to popular belief, American mathematics teachers are relatively well trained compared to their counterparts in most countries including Japan when measured in terms of number of mathematics courses taken.

The fourth explanation they consider is that American students receive much less time on mathematics instruction. To debunk this explanation they report statistics comparing seventh grade Japanese and eighth grade American students reporting 144 hours of instruction for American students and 101 hours for Japanese students. However, this is a misleading comparison. There is only one year in the 12 years that Japanese students spend this little time in mathematics. They spend 175 hours in most of elementary school (where they are receiving twice the class time devoted to mathematics as American students), 140 hours in the rest of junior high, and more hours again in senior high [17]. In addition most Japanese students spend after school hours in juku classes where they get further intensive tutoring. Moreover, time is spent much more efficiently in the Japanese class. Students are receiving instruction 90 percent of the time in a Japanese classroom while only 46 percent of the time in the American classroom [15].

Speaking a psychologist who studies learning, I have to say that the first variable of human learning is time on task. This is not to say that it is not important how one spends that time and it is not to say that we cannot spend time learning rather useless things, but all other things being equal amount learned is roughly proportional to amount of time spent learning. I also have to say that the second variable of human learning is time away from task. This points a finger of suspicion at the long summer vacation in America and the fact that some important topics like algebra are dropped for a year in the curriculum (typically geometry intervenes between algebra I and algebra II) which create ideal opportunities for forgetting to do its dirty work. Certainly, teachers complain about all the reteaching they have to do after summer vacation.

The situation goes beyond time on task. There are different attitudes that impact on the amount of time and effort that is given to mathematics learning. American students and parents believe that mathematics achievement is a matter of ability while Japanese students and parents think it is a matter of effort [14]. While Americans are dissatisfied with their nations mathematics achievement, they are more often satisfied with their own child's education than are Japanese parents who are proud of their nation's educational performance but think their child could do more. The intensive effort that Japanese students put into what is a 12 year competition for colleges is an anathema to all but the most striving of American parents. It seems unlikely that we are going to see American education become a mirror of the Japanese system. However, the impossibility of transforming society should not blind us to the fact that the probable explanation of the achievement differences really is effective time on task. Changing curriculum, lowering class size, and improving teacher quality will do little as long there is less effective time spent learning.

In contrast to the general lack of effective educational interventions, it is well known to the intelligent tutoring community that there is one intervention which can produce enormous achievement gains with time on task fixed. This is private tutoring [4]. In contrast to typical efforts to reduce class size, if we could reduce class size to one the student can get much more effective instruction. If a private tutor is extended to monitoring homework, that time can also be spent with maximum efficiency. The effectiveness of private tutoring does not contradict the importance of time on task but rather reinforces it. What private tutors do is manage the microstructure of learning time to assure it is effectively spent. In contrast to other proposals that requiring changing societal attitudes in America, personalized instruction is in keeping with the American belief in the uniqueness of the individual. Unfortunately, a private human tutor is rather too expensive for the average American and is not conceivable in a public education system.

The promise of computer-based tutors is that they can make the benefits of individualized instruction available to all students at affordable costs. While their promise goes far beyond American mathematics education, they do have a special promise there. Tutors may be particularly adapted to the American society. Second, computers naturally pose a platform for teaching the high-tech mathematics which will be the mathematics of the future where there will be little emphasis on mastering mechanistic skills like the long-division algorithm or symbol manipulation in algebra, and much greater emphasis on using powerful mathematical software packages effectively and with understanding.

2 Review of Research on Model-Tracing Tutors

We have been working with a style of tutor which we think is particularly well designed to deliver individual computer-based instruction in the mathematics classroom. Our research on tutoring systems began with the completion of the ACT* theory of cognition [1]. That theory proposes that human problem solving is enabled by a set of production rules. The theory describes how these rules are learned and how they are executed in the process of solving a problem. The theory makes strong claims about how problem-solving skills like those in mathematics are learned. The theory can be turned into a set of strong prescriptions for instruction. In 1984 it seemed that the obvious vehicle for delivering these prescriptions was the intelligent tutoring paradigm which had been evolving in artificial intelligence as a way of getting computers to interact with students much as private human tutors interact with students. Merging our cognitive models with the intelligent tutoring methodology had the promise of providing a demanding test of our cognitive theory and making substantial educational contributions.

Over the eight years that we have been working on the topic we have developed a highly articulate approach to tutoring which we call *model-tracing tutoring*. The basic premise of the tutoring approach is to develop a cognitive model of how the student should solve problems and use this model to interpret the student's problem-solving behavior and to guide the student through the curriculum. This cognitive model is represented as a set of production rules. This cognitive model represents an "ideal" that we want the student to achieve. It should be capable of producing any acceptable solution path for a problem. We supplement this model with some of the bugs that students are observed to make. We use this cognitive model to interpret the student's problem solving behavior. When the student makes errors we can interpret these errors and provide appropriate feedback. When the student asks for help we can propose an appropriate path of solution. The key to the model-tracing methodology is the ability to interact with the students at this step-by-step grain size and interpret their behavior in terms of cognitive rules.

We have had some success with this methodology and have followed up that success with some research trying to identify what determines its success. It seems that there are three key factors: (1) Most important is the creation of a successful cognitive model and communication of that model to the student. The tutor, if well designed, facilitates the communication of the model but we have gotten partial success communicating these models off-line with just verbal instruction. (2) Critical to minimizing learning time is to have some means of protecting the student from the potentially devastating cost of errors. In untutored environments students can spend hours on problems which can be done in a few minutes with a few well chosen pointers. (3) Critical to a successful growth of knowledge in a course is the ability to monitor the students' acquisition of individual rules in the tutor and only promote students when they have mastered these rules. The cognitive model provides a psychologically viable analysis of the skill into individual components. The process of following the growth of rule knowledge over problems we call knowledge-tracing to contrast it with model-tracing which is following the students' use of rules within individual problems. Our knowledge-tracing capacity enables individualized learning.

By 1987 [3] we had completed three computer-based tutors—one for proof skills in geometry, one for symbol-manipulation skills in algebra, and one for beginning coding skills in the computer programming language LISP. We found that these tutors could accelerate the rate of skill acquisition by as much as a factor of three [5]. This result, which has been many times replicated, remains our major finding: Our tutors dramatically accelerate the learning of a curriculum by optimizing the learning process through individualization of instruction.

A dichotomy developed in that early phase which has stayed with our research to this date. On one side was the research with high school mathematics (the algebra and geometry tutors) which was the more significant of our projects in terms of its practical implications and on the other side was our work on LISP (and later on Prolog and Pascal) which turned out to be yielding more fundamental data about tutoring and human cognition. The work on high school math tutors was addressing a much more important topic nationally, a much larger domain in content (high school mathematics is between 500 and a 1000 hours while introductory programming is between 50 and 100 hours), and a domain which for which our tutors were more applicable because expertise in high school mathematics can be more easily formalized. Our work on introductory programming yielded more information because we had ready access to the college undergraduate population for research and instruction. The advantages were both ease of access to the subject population and general cooperativeness of that subject population. So our understanding of what we can do with intelligent tutoring is largely based on studies of introductory college programming. The most important application of this knowledge is to high school mathematics.

Any attempt to get empirical feedback on our mathematics tutors was a large endeavor in dealing with administrative barriers. With some difficulty we have brought high school students into the laboratory where we could explore the consequences of various design features but we could not access the consequences for classroom achievement. The geometry and algebra tutors were both demonstrated in the Pittsburgh Public Schools over a period of three years. In contrast to the college programming course, the school environment did not let us accelerate students so we had to concentrate our measures on achievement gains and could not get learning time gains. With the geometry tutor, there were achievement gains of about one standard deviation or one letter grade which we have reported elsewhere [3]. In addition there were large positive effects on class morale that have been documented by Schofield and Evans-Rhodes [12] and a positive report from the teacher's perspective has been written by

Wertheimer [16]. These efforts also resulted in a set of positive relationships with the Pittsburgh Public Schools that are still serving us well.

The results with the algebra tutor were a substantial contrast. When we brought the tutor into the classroom we experienced a considerable difficulty because of the conflict between the symbol-manipulation techniques we had built into the tutor and the techniques the teacher wanted to use. This has been a major object lesson in the importance of having the consumer buy into the tutor. The students in the classroom were not at a disadvantage to control students and still showed advantages for one topic (factoring of quadratics) but did not show the advantages we had expected from our laboratory work.

Subsequent to that research, we have developed a algebra word-problem tutor [13] which has produced large gains in the laboratory and a new geometry tutor [8] which is now being tested in Pittsburgh Public Schools. These tutors reflect a major ongoing shift in our approach to tutoring to one which we think will prove more congruent with the needs of modern mathematics education.

Our new tutor development philosophy is hinted at in Anderson and Pelletier [2]. The major shift in our development philosophy is to focus on educator's conception of the skill rather than our own and to focus on embedding powerful problem-solving tools in our tutor environment. We are now working with educators and teachers in the Pittsburgh Public Schools trying to identify from them their conception of what should be taught. This we try to codify as a cognitive model. This is very much like expert system development where the educators serve as the experts and we as the knowledge engineers trying to codify their expertise. This can be a trying experience in two ways. First, the educators do not naturally think of the competence they are trying to teach as formalizable. Thus, it can be a struggle to extract from them the rules. Second, we may not always like the rules we extract in the sense we may believe there are better things that could be taught. While we may influence the product we have had to take an attitude that in the end the educator is always right. The result is something the teachers are happy with and which is viable in the classroom.

Having identified the competence we want the students to acquire, we try to create a powerful interface for communicating it. So, for instance, our algebra interface has built-in facilities for setting up tables, graphing functions, solving equations, etc. It is much like a modern symbol-manipulation package. For certain applications (e.g., solving equations) we may disable certain features but our goal is to teach students to express the target competence in a modern computational environment. Our tutors can fade in and out in terms of the amount of control they exert. At one extreme, they can force the student to take a specific path while at the other extreme the student is free to do anything they want in the interface and will receive no feedback. We want to produce students who are effective problem solvers on their own in such an environment. This requires some effort at teacher education as most high-school teachers are quite unfamiliar with using such mathematical tools.

We have now created a development system [2] which organizes the creation of model-tracing tutors in this mold, automatically providing many of the facilities. It supports development of production-rule models, model-tracing, development and integration of interfaces, and organization of large curriculum. It requires a rather sophisticated knowledge engineer to use but it avoids having to build complete systems for each application. It also guarantees that the interface will have a common feel over a large span of curriculum such as high school mathematics.

3 Application to the Pittsburgh Public Schools

We have entered in to a collaborative relationship with the Pittsburgh Public Schools to help revise the high-school mathematics curriculum to one that is more modern, that is computer intensive, and one that is organized around model-tracing tutors. This is largely just a matter of good intention now. We are working with the mathematics faculty of one high school where we have created a classroom of 24 Mac II computers thanks to the generosity of Apple. We are experimenting with a geometry tutor and are working on the development of other tutors. As things progress, we hope to port our tutors to other high schools in Pittsburgh, to out-of-school "learning centers" that are being created, and to middle schools. Before describing our intentions however, it is worth describing the current situation.

3.1 The Current Situation

The situation in the city of Pittsburgh is by no means among the worst in America, but it is representative. Each year over 3,000 students enter high school which extends from ninth to twelfth grade. About 300 students are designated as either scholar students or gifted students. There is a state-mandated special education program for gifted students who are defined as students with IQs greater than 130. Scholar students are non-gifted students who show high achievement. Most of these students have already taken Algebra I in middle school and will progress with little difficulty through the math curriculum. A large fraction will take calculus in high school. The only problem for these students is that they are bored and are not being challenged to perform at their potential. A good portion of the remaining students (about 1,000) have already been lost to the academic mathematics track and will take course after course of "general math" in high school which is a review of the basic arithmetic they should have learned. Those students who are in neither the general math, scholars, or gifted program (about 2,000) take Algebra I sometime in high school. 72 percent of these students either get a D or E. There is rapid atrophy of students taking successively higher courses with 2,000 taking algebra 1, 1,200 taking geometry, 700 taking algebra 2, 350 taking precalculus, and 200 taking calculus. The figures after algebra 1 are inflated by the addition of the gifted and scholar students who take algebra 1 in middle school.

Thus, the large majority of students in the school system experience failure in mathematics before they leave high school. The problem is by no means entirely with the school system. There are large attendance problems, violence and drugs in the community, and many family problems. In effect a large fraction of the students whom Pittsburgh is trying to educate have opted out of the education process for reasons quite independent of what is happening in the schools. Still over half of the students are there to learn and are not learning what they should or could. Certainly, the problem is not the amount of money spent to educate a student. It is estimated it costs $9,000 per year for each student—almost as much money as is spent at the area's most preppy private school.

What is a typical high school mathematics classroom like? It varies from almost 40 students to less than 20 with size tending to decrease as level increases. Advanced courses to advanced students are regarded as fun to teach and students are relatively motivated and achieving but such courses are the exception. The more typical class has a very sullen or disruptive character (depending on the teacher and students) with a large fraction of the students not understanding the lesson material. The teacher is in a constant state of alternating between teaching those students who understand or remediating those who do not. Exam scores are dismal and cheating is rampant. Teachers joke that you can reconstruct the seating pattern in the class by looking at the overlap in wrong answers. Teachers are well paid and they regard it as battle pay. Most teachers entered the profession with the best of intentions but those have been lost by the practicalities of surviving in the classroom. Pittsburgh schools are not the scenes of the same amount of violence that typifies the public image of the American urban school but they are also not places of learning. In past generations students walked through the schools and took a job in the local steel mills. This career option is no longer available. The city itself has largely moved to a service industry economy. Students have to learn technical and mathematical skills to survive in the job market.

3.2 The Ideal Image

Let us contrast the typical classroom today with our image of what the tutored mathematics classroom would look like. This image is based both on our previous experiments with tutored classrooms in the public schools and our success with tutored courses at Carnegie Mellon University. Physically, we imagine classrooms of about 30 modern machines with rapid processors and large two-page monitor screens. During regular school hours students come and go from these rooms according to regular class schedules but these rooms are available after school hours as well in which students find a room with an available machine and do their homework.

If one went into a regular class one would see most of the students working away on their own perhaps at very different points in the curriculum. At particular points a student might have a difficulty with the tutor or a concept. Their first reaction would be to ask a neighbor but if unsuccessful they would call upon the teacher that is monitoring the class. Thus, the teacher would shift role from person who doles out instruction and drill and practice to someone who is the expert of last resort. The teacher would also pull groups of students from their regular work and assign them group projects during which their computers become not tutors but tools. Thus, the second role of the teacher would be to orchestrate collaborative problem solving. These groups would not necessarily involve students at the same point in the curriculum and students would be encouraged to try alternative methods of solution to a problem. The product of the class would be students who are capable of using modern computational tools to formulate and solve problems.

The function of the tutor would be to train and monitor the requisite skills for such problem solving. It would provide reports to the teacher about the progress of individual students. It would select problems to help students master these problems and promote students to new material as the skills were mastered.

Students would be required to maintain a minimal pace so as to guarantee mastery of the NCTM standards [10, 11]. Students who could not maintain this pace in class would have to do after school work or even summer work to keep up. This is much what happens in our university programming classes. However, students who could would be encouraged to accelerate through the problems. After school and summer work would not be reserved for the slower students. Mastering new material would enable new opportunities. The intention is to create a community of learners in which progress is valued and enables students to do new things. Students who mastered first year college calculus before the 12th grade would be have the opportunity to attend special college courses where they could explore mathematics and related subjects.

What are the prerequisites to this ideal image? A decade ago the first thing that would come to mind would be access to the computers. This is no longer the major issue. Suitable computers are affordable and are becoming cheaper by the year. There are two major difficulties. One is that we need to codify perhaps a thousand hours of instruction. The second is that we need to institute the organizational changes to allow such tutors to succeed in the classroom. These two prerequisites are not independent as we will develop below under the issue of ownership and tutor development.

3.3 The Issue of Ownership

We have already discussed the difficulty we had when the algebra tutor, which had some success in the laboratory, entered the classroom of a teacher who had a different image of what should be taught. That same teacher had great success with the geometry tutor but had played a large role in fashioning it and found it more congenial to what he wanted to teach. It is clear that if a congruence is not achieved between the teacher and the tutor, there will not be success. The curriculum that we provide has to be one that the teachers own.

The issue of ownership exists at three levels in the Pittsburgh Public Schools. District wide there is a mathematics curriculum group which holds strong opinions about what should be taught in high-school mathematics and how it should be taught. Their opinion is strongly influenced by the NCTM standards [10, 11]. Each school has a mathematics faculty through which these ideas are filtered. Then there is the individual teacher. Needless to say the opinions are not always the same at the three levels. Frequently but not always, the influence from the teacher level is conservative reflecting resistance to change.

We have identified one high school in Pittsburgh, where the three levels are relatively in synch. This is Langley High School which has an innovative mathematics faculty and principal. Its mathematics faculty pioneered a new geometry curriculum which is being adopted city wide. In conjunction with our program a master teacher who has been working on curriculum reform is being transfer to Langley. In addition to relative consensus on the three levels, it is the case that this consensus is in favor of the NCTM standards which probably maximizes the likelihood of acceptance outside of Pittsburgh.

We have come more and more to view the process of creating a tutoring system as a similar endeavor to creating an expert system [6, 7]. It is generally regarded that the most critical factor in the practical success of an expert system is guaranteeing that the client for the system will really use it. The classic mistake is to build a system to solve a problem the developer thinks is important and then go to the client and try to convince the client that this is the problem they wanted solved. The successful systems are ones that consult with the clients at the outset and focus on the problems they perceive as important. Of course, it is critical that the clients buy into the concept that the technology you have can solve their problem. What one does not want to do, however, is to try to define for them their problem and the details of the solution. Thus our model for development will be one in which we will work with the teachers codifying what their conception is of the mathematics curriculum. Thus, it will be a curriculum they own.

A related issue is class and school reorganization. Teachers vary in their willingness to let go of their role as stand-up lecturers. They are almost universally intimidated by the prospect of individualized learning and the prospect of a classroom of 30 students at very different places. Most teachers have not ever dealt with self-paced learning or computer-based classrooms.

The issue of individualized learning also strikes a sensitive political cord in the district administration as a whole who don't want to have to deal with complaints about different identifiable subgroups achieving at different levels. They are greatly enthused about enhancing the performance of low achievers and are willing to engage in radical changes (for them) like opening up schools after hours. However, they see a real problem if high achievers achieve even higher. In their resistance to the idea they stop just short of saying they want to hold back high achieving students. They also have to deal with the constituency reflected by state-mandated programs for gifted students. Reflecting this they have encouraged us to think about making these tutors available for gifted middle school students who are bored with the instruction there.

We know from our experiments in the school and at the university that these issues can be resolved successfully. Teachers enjoy their new role once they adapt to it. Almost no one complains when they are succeeding at learning. However, our saying so does not make anyone believe there are not problems nor does it lower resistance. Our plan here is just to gradually expose the teachers and administrators to the benefits of the tutored classroom. Again the issue of ownership arises. We want them to define how to reorganize school and classroom to take advantage of the instructional material they have in effect authored. They may not come exactly to our image of the classroom but they will be drawn to a radically reorganized classroom that will respond to the reality that the tutor creates. Undoubtedly, what they discover in the process will be better than what we could suggest at the outset.

3.4 Tutor Development

An important observation is that there is nothing that high school students are expected to learn through calculus that cannot be achieved by current expert systems. It is true that there are proofs in geometry that are beyond the capacity of current expert systems as is some of the more foundational reasoning in calculus. However, these are beyond the skills required of high-school students. So if we were taking a purely expert-system development approach to high-school mathematics we know that we are guaranteed success. The severe constraint in our case is that the expertise be modelled in a human-like way so that it can serve as a target of expertise. It is something of a conjecture whether we can achieve this for all of high school mathematics but we have had no difficulties in what we have attempted so far.

We have observed that it takes at least 100 hours to do the development that corresponds to an hour of instruction for a student (an hour for the slower students, can be much less for faster students). Since we are looking at codifying material that might occupy as much as 1,000 hours we are looking at a development time of 100,000 hours. Dividing this by 2,000 hours per man year, we are looking at a development effort of 50 man years which can easily be doubled to incorporate the cost of evaluation and revision. A 100 man years of investment is not that high to revise the mathematics curriculum of a school district processing more than 10,000 students per year. Amortized over 10 years, this comes down to about two hours of development per student yearly. We are looking at an expenses that are well less than 1 percent of $9,000 per student.

References

1. J. R. Anderson: *The Architecture of Cognition*. Cambridge, MA: Harvard University Press, 1983.

2. J. R. Anderson, R. Pelletier: A development system for model-tracing tutors. In Proceedings of the International Conference of the Learning Sciences, 1-8. Evanston, IL, 1991.

3. J. R. Anderson, C. F. Boyle, A. T. Corbett, M. W. Lewis: Cognitive modelling and intelligent tutoring. *Artificial Intelligence, 42*, 7-49, 1990.

4. B. S. Bloom: The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher, 13,* 3-16, 1984.

5. A. T. Corbett, J. R. Anderson: Feedback control and learning to program with the CMU Lisp Tutor. Paper presented at the Annual Meeting of the American Educational Research Association, 1991.

6. D. Diaper: *Knowledge Elicitation.* Chichester, England: Ellis Horwood, 1989.

7. M. Greenwell: *Knowledge engineering for expert systems.* Chichester, England: Ellis Horwood, 1988.

8. K. R. Koedinger, J. R. Anderson: Abstract planning and perceptual chunks: Elements of expertise in geometry. *Cognitive Science, 14,* 511-550, 1990.

9. C. C. McKnight, F. J. Crosswhite, J. A. Dossey, E, Kifer, J. O. Swafford, K. J. Travers, T. J. Cooney: *The Underachieving Curriculum: Assessing U.S. School Mathematics from an International Perspective.* Champaign, IL: Stipes Publishing Company, 1990.

10. National Council of Teachers of Mathematics Commission on Standards for School Mathematics: *Curriculum and Evaluation Standards for School Mathematics.* Reston, VA: The Council, 1989.

11. National Council of Teachers of Mathematics: *Professional standards for teaching mathematics.* Reston, VA: The Council, 1991

12. J. W. Schofield, D. Evans-Rhodes: Artificial intelligence in the classroom: The impact of a computer-based tutor on teachers and students. *In Proceedings of the Fourth International Conference on AI and Education,* (238-243). Amsterdam, 1989.

13. M. K. Singley, J. R. Anderson, J. S. Gevins, D. Hoffman: The algebra word problem tutor, *Artificial Intelligence and Education,* 267-275, 1989.

14. H. W. Stevenson, S. Lee, C. Chen, J. W. Stigler, C. Hsu, & S. Kitamura: Contexts of Achievement, *Monographs of the Society for Research in Child Development, 55,* 1990.

15. J. W. Stigler, M. Perry: Mathematics Learning in Japanese, Chinese, and American Classrooms. In J. W. Stigler, R. A. Shweder, and G. Herdt (Eds.) *Cultural Psychology.* New York: Cambridge University Press, 1990.

16. R. J. Wertheimer: The Geometry Proof Tutor: An intelligent computer-based tutor in the classroom. *Mathematics Teacher,* 1990.

17. M. White: *The japanese education challenge: A commitment to children.* New York: The Free Press, 1987.