Combatting Shallow Learning in a Tutor for Geometry Problem Solving

Vincent Aleven, Kenneth R. Koedinger, H. Colleen Sinclair, and Jaclyn Snyder*

HCI Institute
School of Computer Science
Carnegie Mellon University

*Langley High School / Pittsburgh Public Schools

E-mail: aleven@cs.cmu.edu, koedinger@cs.cmu.edu, colleens+@andrew.cmu.edu, jsnyder@pps.pgh.pa.us

Abstract. The PACT Geometry tutor has been designed, with guidance from mathematics educators, to be an integrated part of a complete, new-standards-oriented course for high-school geometry. We conducted a formative evaluation of the third "geometric properties" lesson and saw significant student learning gains. We also found that students were better able to provide numerical answers to problems than to articulate the reasons that are presumably involved in finding these answers. This suggests that students may provide answers using superficial (and possibly unreliable) visual associations rather than reason logically from definitions and conjectures. To combat this type of shallow learning, we are developing a new version of the tutor's third lesson, aimed at getting students to reason more deliberately with definitions and theorems as they work on geometry problems. In the new version, students are required to state a reason for their answers, which they can select from a Glossary of geometry definitions and theorems. We will conduct an experiment to test whether providing tutoring on reasoning will transfer to better performance on answer giving.

Introduction

A problem for many forms of instruction is that students may learn in a shallow way [Burton and Brown, 1982; Miller, *et al.*, submitted], acquiring knowledge that is sufficient to score reasonably well on some test items, but that does not transfer to novel situations. One manifestation of shallow learning is that students construct superficial domain heuristics that may allow them to solve some problems quite well, even though that "knowledge" is ultimately not correct. For instance, in the context of geometry, most students will learn how to find the measures of unknown quantities in diagrams. However, they may rely on the fact that certain quantities look equal rather than reason from geometric definitions and theorems. Such superficial perceptual strategies, enriched with some correct geometric knowledge, can be very serviceable and lead to correct solutions on many naturally-occurring problems. However, they fall short on more complex problems or when students are asked to discuss reasons for their answers.

Superficial strategies occur in many domains. In physics problems solving, consider a problem where students are asked to draw an acceleration vector for an elevator coming to a halt while going down. They often draw a downward arrow, assuming, incorrectly, that the acceleration has the same direction as the velocity. Also, when asked to categorize physics problems, novices tend to do so on the basis of surface level features, while experts use the deeper physics principles involved [Chi, *et al.*, 1981].

We can interpret the shallow learning problem within the ACT-R theory of cognition and learning [Anderson, 1993], as follows: In the ACT framework, learning a procedural skill means acquiring a set of production rules. Production rules are induced by analogy to prior experiences or examples. Superficial knowledge may result when students pay attention to the wrong features in those experiences or examples, features that may be readily available and interpreted, but that do not connect to deeper reasons. However, not much is known about what types of instruction are more likely or less likely to foster shallow learning.

Evaluations of cognitive tutors indicate that they can be significantly more effective than classroom instruction [Anderson, *et al.*, 1995; Koedinger, *et al.*, 1998]. In spite of this success, cognitive tutors (and other computer-based learning environments) may not be immune from the shallow learning problem. It is important to determine to what degree students come away with shallow knowledge, when they work with cognitive tutors. This may help to find out how these tutors can be designed to minimize shallow learning and be even more effective.

We study these issues in the context of the PACT Geometry Tutor, a cognitive tutor developed by our research group used in four schools in the Pittsburgh area. In a formative evaluation study, we found that the instructional approach of which the tutor is part, leads to significant learning gains. We also found evidence of a form of shallow learning: Students cannot always give a reason for their answers to geometry problems, even if the answer itself is correct. Such a reason would be, for example, a definition or theorem applied in calculating certain quantities in a diagram. We have redesigned the tutor in an attempt to remedy this kind of shallow learning. Currently, we are pilot-testing the new tutor.

In this paper, we give a brief overview of the PACT Geometry tutor. We present results from our formative evaluation study that motivated the redesign of the tutor. We describe how we have modified the tutor, and finally discuss why the changes may lead to a more effective tutor, in geometry and potentially also in other domains.

The PACT Geometry Tutor

The PACT Geometry Tutor was developed at the PACT Center at Carnegie Mellon University, as an adjunct to classroom instruction, in tandem with the PACT geometry curriculum, which covers high school geometry from a problem-solving perspective, consistent with the standards for mathematics instruction developed by the National Council of Teachers of Mathematics [NCTM, 1989]. The PACT Geometry Tutor is different from the earlier Geometry Proof Tutor [Anderson, 1993] and the ANGLE tutor [Koedinger and Anderson, 1993], reflecting changes in the geometry curriculum de-emphasizing the teaching of proof skills. As shown in Figure 1, the PACT Geometry tutor curriculum consists of four lessons (a fifth lesson on circle properties will be added soon), each divided into a number of sections. The topics in each section are introduced during classroom instruction. Students then use the tutor to work through problems, proceeding at their own pace. Usually, students spend 40% of classroom time solving problems on the computer.

The PACT Geometry Tutor provides intelligent assistance as students work on geometry problems, aimed at making sure that students successfully complete problems and that students are assigned problems that are appropriate to their skill level. In each problem, students are presented with a geometry diagram and are asked to calculate the measures of some of the quantities in the diagram, as is illustrated in Figure 2, which shows the tutor interface. The problem statement and diagram are presented in separate windows on the top left and top right, respectively. The tutor has a store of 56 problems for lesson 3.

Lesson 1. Area
1. Area of parallelograms
2. Area of triangle
3. Area of trapezoid
4. Area of circle
5. Area

Lesson 2. Pythagorean Theorem
1. Square & square root
2. Pythagorean Theorem
3. 45-45-90 right triangle
4. 30-60-90 right triangle
5. Pythagorean & area Problem

Lesson 3. Angles
1. Angles
   - Linear pair angles
   - Vertical angles
   - Complementary angles
   - Supplementary angles
   - Angle addition
   - Angle Subtraction
   - Triple angles sum
   - Linear pair formula
   - Vertical angles formula
   - Complementary angles formula
   - Supplementary angles formula
2. Angle associated with triangles
   - Triangle sum
   - Isosceles triangle, base angle
   - Isosceles triangle, vertex angle
   - Exterior angle of triangle
   - Interior angle of triangle
   - Equilateral triangle
3. Parallel lines
   - Corresponding angles
   - Alternate exterior angles
   - Alternate interior angles
   - Supplementary interior angles

Lesson 4. Similar triangles
1. Similar triangles

Fig. 1. PACT Geometry tutor curriculum, organized by lessons, sections, and skills.

Students enter answers in the Table Tool, a spreadsheet-like device shown on the left, second window from the top. The table cells correspond to the key quantities in the problem, such as the givens, the target quantities, or intermediate steps. As students enter values into the table, they receive immediate feedback indicating whether their answer is correct or not. When students are stuck, they can ask for hints, which the tutor presents in a separate Messages window (see Figure 2, left, third window from the top). The hints become more specific as students repeat their request for help. Students can move on to the next problem only when they have entered correct values for all quantities in the table.

The Diagram Tool presents an abstract version of the problem diagram (Figure 2, bottom right) with cells for students to record the measures of angles, as one often does when solving geometry problems on paper. This makes it easier for students to relate the quantities in the problem to the entities in the diagram and to keep track of information. For problems that involve difficult arithmetic operations, students can use the Equation Solver (not shown, see [Ritter and Anderson, 1995]), a tool which helps students to solve equations step-by-step.

Finally, the PACT Geometry Tutor provides a skillometer, which displays the tutor's assessment of the student, for the skills targeted in the current section (see Figure 2, bottom left). The skillometer helps students keep track of their progress. The skills for which a student has reached mastery level are marked with a "✓". When students have reached mastery levels for all skills, they graduate from the current section of the curriculum.

The PACT Geometry Tutor is a cognitive tutor, an approach based on the ACT theory which has proven to be effective for building computer tutors for problem-solving skills [Anderson, 1993; Anderson, *et al*, 1995]. The PACT Geometry tutor is based on a production rule model of geometry problem solving, organized by lessons and sections as shown in Figure 1. The model, which contains 77 geometry rules and 198 rules dealing with equation-solving, is used for ***model-tracing*** and ***knowledge tracing***. The purpose of model-tracing is to monitor a student's solutions and to provide feedback and hints. When the student enters a value into the Table Tool, the tutor uses the model output as a standard to evaluate the student's answer. To provide hints, the tutor applies its production rule model to the current state of problem-solving and displays the hint messages associated with the applicable rule that has highest priority.

The purpose of knowledge-tracing is to compile detailed measures of an individual student's competence (i.e., a student model), based on that student's performance over a series of problems. The student model is an overlay on the production rule set. For the critical production rules, the tutor uses a Bayesian algorithm to estimate the probability that the student knows the skill [Corbett and Anderson, 1995]. The estimate for a given rule is updated each time the rule is applicable, taking into account whether the student's action is correct or not, or whether the student asks for help. The tutor uses this information to select appropriate remedial problems for each student and to decide when a student is ready to advance to the next section of the curriculum. The information in the student model is displayed on the screen in the skillometer. The system is implemented using the plug-in Tutor Agent architecture [Ritter and Koedinger, 1997] and the Tutor Development Kit [Anderson and Pelletier, 1991].

Formative evaluation of the angle properties lesson

In the spring of '97, we collected data to evaluate how effective the combination of classroom instruction and practice with the PACT Geometry tutor is, primarily to identify areas where the instruction or tutor can be improved. Also, we wanted to assess how well students are able to explain answers to geometry problems. The study focused on lesson 3 of the PACT geometry

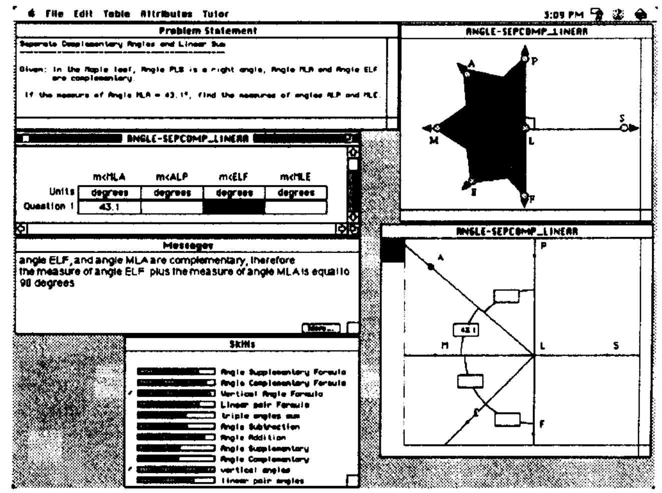

Fig. 2. The PACT Geometry Tutor provides tools and intelligent assistance

curriculum, which deals with geometric properties relating primarily to the measures of angles (see Figure 1). A total of 71 students in two schools participated. All students received classroom instruction on the topics of lesson 3 and used the tutor to work through problems related to this lesson. The students took a pre-test before working with the tutor and a post-test afterwards. The classroom instruction took place in part before the pre-test, in part in between pre-test and post-test.

Each test involved four multi-step geometry problems in which students were asked to calculate certain measures in a diagram and were asked also to state reasons for their answers, in terms of geometry theorems and definitions. Students were given a sheet listing relevant definitions and theorems and were told that they could use the sheet freely. We used two different test forms, each of which was given (randomly) to about half the students, during pre-test and post-test, to counterbalance for test form difficulty. An example question is shown in Figure 3, together with correct and incorrect reasons that students gave for correct numeric answers (e.g., correct angle measures), when they took the post-test.

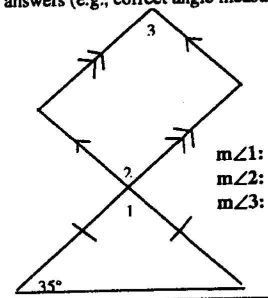

m∠1: 110° **Reason:** isosceles triangle, triangle sum
m∠2: 110° **Reason:** vertical angles
m∠3: 110° **Reason:** opposite angles in a parallellogram (or: supplementary interior angles)

	Correct	Incorrect
∠ 1	• triangle sum, isosceles triangle • third # to the sum of triangle • you subtract 180 - 35 - 35 and get 110	• only angle besides congruent angles; opposite congruent sides • linear pair • alt. Interior angles are congruent
∠ 2	• vertical pair of angles	• because it 2 is a linear pair with 1 • it is opposite ∠1 • corresponding angles are congruent
∠ 3	• Opposite Angles of a Parallelogram are equal • Interior angles on same side of transversal are supplementary • all the lines are parallel so there will be 2 pairs of equal angles	• parallel lines --> Alt. Int. angles are congruent • ANG 2 & 3 are CONG cause of parallel sides • same as m∠2

Fig. 3. Sample test question with correct answers and reasons (top) plus a sample of reasons given by students for correct answers on the post-test

The criterion for grading the reasons was whether students were able to justify their answers in terms of geometry definitions and theorems, possibly stated in their own words. For example, to calculate the measure of angle 1, one needs to apply the isosceles triangle theorem and the triangle sum rule, as shown in Figure 3, first correct reason for angle 1. Since the grading was lenient, listing only one of the two rules was deemed correct, as can be seen in the second correct reason for angle 1. Even a procedural description which did not mention any geometry rules was deemed (borderline) correct, as in the third correct reason for angle 1. We see also that some of the incorrect reasons were more incorrect than others. Some are very close to being correct (e.g., the first incorrect answer for angle 1), some are plain wrong (e.g., the first incorrect reason for angle 2 mentions the wrong theorem), some are in between.

As shown in Figure 4, students' test scores improved from pre-test to post-test. Numeric answer-finding increased from an average of 0.74 on the pre-test to 0.86 on the post-test, reason-giving improved from 0.43 on the pre-test to 0.60 on the post-test. A two-factor ANOVA with test-time (pre v. post) and action type (numeric answer vs. reason) as subjects factors, revealed significant main effects of both test-time ($F(1, 70) = 39.1$, $p < .0001$) and action type ($F(1,70) = 191.4$, $p < .0001$). This indicates that students improved significantly from pre-test to post-test and were significantly better at giving answers than at giving reasons. The two factors did not interact ($F(1, 70) = 3.3$, $p = .075$), indicating that there was as much improvement on reason-giving as they did on numeric answer finding.

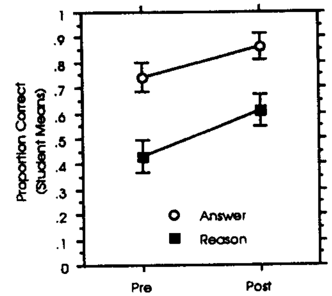

Fig. 4. Students' scores for answers and reasons (proportion correct) at pre-test and post-test

These results indicate that a combination of classroom instruction and practice with the PACT Geometry tutor, based on the PACT curriculum, is effective. The high pre-test scores may reflect the fact that the bulk of the classroom instruction took place before the pre-test. Much of the improvement in students' test scores may be due to students' working with the tutor. Students' ability to state reasons for their answers improved, in spite of the fact that the instruction did not focus on teaching them to do so. This may be due to transfer from answer-giving, the tutor's hints (which often mention the reason), or from the classroom instruction between pre-test and post-test.

Students may use shallow knowledge

The results from the formative evaluation show that students are better at finding numerical answers than at articulating the reasons that are presumably involved in finding these answers. We hypothesize that this is due to the use of shallow knowledge.

At best, students may have a fairly robust visual encoding of the knowledge involved, associating diagram configurations with inferences that can be drawn from them, but may not know the name of the definition or theorem involved. (Geometry experts organize their knowledge in this way, but they also know the name of the rules involved or at least, are able to identify corresponding rules on a reference sheet [Koedinger and Anderson, 1990].) At worst, students may draw on superficial knowledge that enables them to get the answer right in some circumstances but that may be overly general or inappropriately contextualized. This knowledge may take the form of "guessing heuristics" such as: If two angles look the same in the diagram, their measures are the same. Or even: An unknown measure is the same as that of another angle in the diagram. Or: The measure of an angle may be equal to 180• minus a measure of another angle. Such guessing heuristics can be quite helpful. For example, by looking at the diagram shown in Figure 3, one could guess correctly that the measures of angles 1, 2, and 3 are equal. Thus, once one has the measure of angle 1, one can find the measures of angles 2 and 3. But one also needs more robust knowledge to solve this problem. It seems difficult if not impossible to find the measure of angle 1 solely using guessing heuristics, without using (something like) the triangle sum rule.

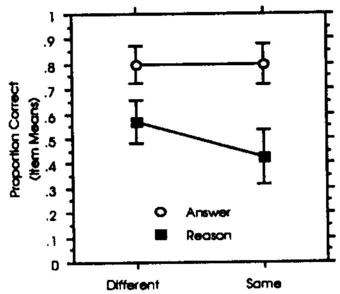

Fig. 5. Students' answer and reason scores on test items where a quantity sought is different from the quantities from which it is derived (Different) vs. scores on items where a quantity sought is equal to a quantity from which it is derived (Same)

We found evidence that students use guessing heuristics by comparing their scores on two classes of test items, namely (1) items where a quantity sought is equal to a quantity from which it is derived (in one step) and (2) test items where the quantity sought is different from the quantities in the problem from which it is derived. We found that the difference between students' answer scores and reason scores is greater on "same measure" items than it is on "different measure" items (see Figure 5). This suggests that students are, in part, relying on a heuristic that says: If angles look equal, their measures are equal.

Finally, we found evidence of the use of guessing heuristics in the logs of students' work with the tutor. For example, consider the problem shown in Figure 6. In an attempt to find the measure of ∠ELF (which is 50•), one student entered values of 90, 40, and 130 (on successive, unsuccessful attempts), evidence of the use of guessing heuristics. For example, she may have entered 40 because a given quantity in the problem is equal to 40•. She then turned to ∠ALP (leaving open the answer for ∠ELF) and found that its measure is 50•. She proceeded to enter 50 for both m∠ELF (correct) and for m∠MLE (incorrect), illustrating further use of the same heuristic. If she had entered 50 only for m∠ELF but not for m∠MLE, it would have been harder to argue that she was guessing. However, the quick entry of 50 for both quantities indicates a lack of deliberate processing characteristic of shallow learning.

This example is not unusual. It should be noted from our classroom experience (the fourth author is currently a teacher and the second author has taught high school geometry previously) that such guessing behavior is not unique to our computer tutor, but is also commonly observed in classroom dialogues.

Hypothesis: Teaching to give reasons is beneficial

We hypothesize that students acquire more reliable knowledge when they are trained consciously to interpret rules (definitions, theorems, etc. written in English) and apply them in a logical manner to find answers. By doing so, students may learn to rely less on implicit, recognition-based inferencing that results from the shallow perceptual coding of geometric knowledge. Encouraging students to reason logically from definitions and theorems may result not only in better performance on reason-giving, but may transfer to better performance on quantitative answers. First, this kind of instruction may help students to induce the correct geometry knowledge and therefore lessen the need to form shallow rules such as the ones illustrated above. The verbal encoding of the rules focuses students on the right features when inducing rules from practice examples. Consistent with this line of reasoning, studies in both a statistics and a logic domain [Holland, *et al.*, 1986] showed that students learn better from instruction using both examples and rules than either one alone. Our interpretation is that example instruction facilitates the same kind of implicit, perceptual level learning that our students are engaged in while practicing answer finding, while rule instruction facilitates the kind of explicit, verbal learning we expect to support in glossary search and reason giving.

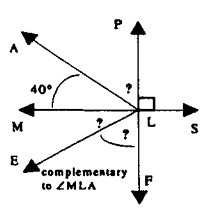

Given: ∠PLS is a right angle, ∠MLA and ∠ELF are complementary*.

If the measure of ∠MLA = 40°, find the measures of ∠ALP, ∠ELF, and ∠MLE.

Answers:
m∠ALP = 50°
m∠ELF = 50°
m∠MLE = 40°

* If two angles are complementary, the sum of their measures is 90°.

Fig. 6. Sample tutor problem

A second argument for the transfer of deliberate rule use to answer finding is based on memory research. In explaining experimental results showing that concrete words are remembered better than abstract words, Paivio (1971) posed a "dual code" theory. Concrete words are stored in memory both perceptually and verbally and these redundant codes make retrieval of them more likely than equally frequent abstract words which are stored in memory only in a verbal code. In our case, providing more practice in deliberate rule use should complement student's natural inclination toward perceptual encoding of geometric properties and encourage students toward a second verbal encoding that can enhance the probability of reliable recall. In both these arguments, we are not suggesting stamping out the use of perceptual encodings and heuristics in geometric reasoning, but rather bolstering them with complementary verbal-logical encodings.

Extensions to the tutor

We have modified and extended the PACT Geometry tutor, so as to get students to reason more deliberately with definitions and theorems. In the extended version of the tutor, students are required to provide reasons for their answers, usually the name of a definition or theorem. Students enter the answers and reasons in a new "answer sheet", shown on the left in Figure 7. To complete a problem, students must give correct answers and reasons for all items. Students can look up relevant definitions and theorems in a Glossary, shown on the right in Figure 7. The Glossary contains entries for all geometry rules that students may need in order to solve the problems. Students can cite the Glossary rules as reasons for their answers, as illustrated in Figure 7. They can click on Glossary items to see a statement in English of the definition or theorem, illustrated with an example and a diagram. Students can browse the Glossary as they see fit. Our intention is that students will consult the Glossary when they do not know how to proceed with a problem, interpret the English rules, aided by the examples, and judge which rule can be applied.

The tutor's hint messages have been changed to encourage students to use the Glossary and to help narrow the search for applicable rules in the Glossary, as is illustrated in Figure 7. When a student asks for help, the tutor initially suggests a quantity to work on and points to items in the diagram that are relevant (i.e., the tutor points to the next subgoal and to the premises to use to infer the goal quantity). If the student repeats the request for help, the tutor suggests that the student look at a certain category of rules in the Glossary (e.g., rules dealing with parallel lines). At the next level of help, the tutor highlights in the Glossary the rules that are within that category. If the student asks for even more help, the tutor states which definition or theorem to use. Finally, the tutor points out how the rule can be applied to the current problem. The tutor's production rule model has been extended to accommodate these changes. In all other respects, the tutor is the same as the previous version.

Thus, the new tutor provides opportunities to work with English descriptions of the rules and to some extent forces students to do so. We have begun pilot-testing the new tutor. Initial reactions of two geometry teachers and a student were favorable. The student remarked that he liked the tutor better than the original version, "because you have the reasons."

Discussion and conclusion

Shallow learning is a problem for many forms of instruction and current computer-based tutors may not be an exception. Students who have worked with the PACT Geometry tutor do well on problem-solving tests, but are not always able to state reasons for their answers. We hypothesize that we can reduce this kind of shallow learning by having students work more deliberately with (textual statements of) definitions and theorems as they solve geometry problems with the tutor. Towards this end, we have added a Glossary of definitions and theorems to the tutor and have modified the tutor so that students are required to state which definition or theorem was used to arrive at answers. By providing a textual representation of the geometry definitions and theorems, we aim to focus students on the right features to use in constructing the geometry production rules, as is required according to ACT theory of cognition and learning. There is also an argument that having both a visual and a textual encoding of an item may make that item easier to recall.

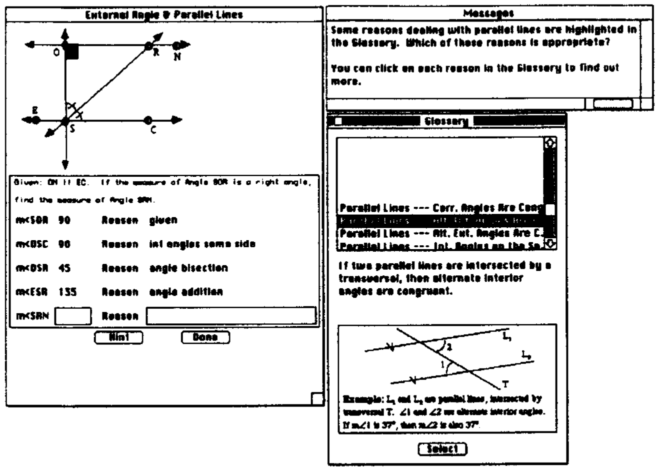

Fig. 7. Extensions to the tutor, to get students more deliberately to reason with definitions and theorems

We will carry out a controlled experiment to evaluate how effective the new tutor is, as compared to the old version of the tutor. We hypothesize that training students to reason more deliberately in this way results not only in greater ability to state reasons for quantitative answers, but also that there will be transfer to better scores on quantitative questions, especially on more difficult problems. We have designed two types of challenge problems to better discriminate the use of shallow versus deep reasoning to find numerical answers. *Misleading problems* have deceptive diagrams in which superficial perceptual strategies will yield the wrong answers. *Possibly unsolvable problems* ask students to compute unknown quantities when there is sufficient information, but to answer "not enough information" when there is not. The pre- and post-test in our evaluation will contain problems of these types, in order to test our hypothesis.

The approach to combatting shallow learning presented in this paper is not specific to geometry instruction, but could also be applied to other cognitive skills. In teaching Lisp programming, for example, it may help to have students articulate what a Lisp operator does in order to get them better to work with unfamiliar Lisp objects such as lists of lists.

If our hypothesis turns out to be true, the work will be significant in the following ways: From a practical point of view, creating instruction, computer-based or otherwise, that leads to deeper learning is an important goal. From a theoretical perspective, if it can be shown that one can effectively combat the learning of shallow knowledge by providing for both implicit perceptual and explicit textual modes of learning, that is a significant result for cognitive science generally.

Acknowledgments

Ray Pelletier and Chang-Hsin Chang contributed to the development of the PACT Geometry tutor. The PACT Geometry curriculum was developed by Jackie Snyder. The research is sponsored by the Buhl Foundation, the Grable Foundation, the Howard Heinz Endowment, the Richard King Mellon Foundation, and the Pittsburgh Foundation. We gratefully acknowledge their contributions.

References

Anderson, J. R., 1993. *Rules of the Mind.* Hillsdale, NJ: Addison-Wesley.

Anderson, J. R., A. T. Corbett, K. R. Koedinger, and R. Pelletier, 1995. Cognitive tutors: Lessons learned. *The Journal of the Learning Sciences*, 4,167-207.

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

Burton, R. R., and J. S. Brown, 1982. An Investigation of Computer Coaching for Informal Learning Activities. In D. H. Sleeman and J. S. Brown (eds.), *Intelligent Tutoring Systems*, 79-98. New York: Academic Press.

Chi, M. T. H., P. J. Feltovich, and R. Glaser, 1981. Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5, 121-152.

Corbett, A. T., and J. R. Anderson, 1995. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction*, 4: 253-278.

Holland, J. H., K. J. Holyoak, R. E. Nisbett, and P. R. Thagard, 1986. *Induction: Processes of Inference, Learning, and Discovery.* Cambridge, MA: The MIT Press.

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

Koedinger, K. R., and J. R. Anderson, 1993. Reifying implicit planning in geometry. In S. Lajoie and S. Derry (eds.), *Computers as Cognitive Tools*, 15-45. Hillsdale, NJ: Erlbaum.

Miller, C. S., J. F. Lehman, and K. R. Koedinger, 1997. Goals and learning in microworlds. Submitted to Cognitive Science.

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

Paivio, A., 1971. *Imagery and verbal processes.* New York: Holt, Rinehart, and Winston.

Ritter, S., and K. R. Koedinger, 1997. An architecture for plug-in tutoring agents. *Journal of Artificial Intelligence in Education*, 7 (3/4), 315-347. Charlottesville, VA: AACE.

Ritter, S., and J. R. Anderson, 1995. Calculation and strategy in the equation solving tutor. In Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society, 413-418. Hillsdale, NJ: Erlbaum.