Intelligent Tutoring Systems are Missing the Tutor:

Building a More Strategic Dialog-Based Tutor

Neil T. Heffernan & Kenneth R. Koedinger

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
neil@cs.cmu.edu, koedinger@cmu.edu

Abstract

Graesser et. al. believe "there is something about interactive discourse that is responsible for [student] learning gains." In this paper we present *Ms. Lindquist*, an Intelligent Tutoring System (ITS) designed to carry on a more human-like interactive dialog to help students learn how to write algebra expressions given a word problem. Ms. Lindquist is able to carry on a running conversation, complete with probing questions, positive and negative feedback, follow-up questions in embedded sub-dialogs, and requests for explanation as to why something is correct. In order to build Ms. Lindquist we have expanded the traditional model-tracing paradigm so that Ms. Lindquist not only has a model of the student, but also has a model of tutorial reasoning. Ms. Lindquist has a separate tutorial model encoding pedagogical content knowledge in the form of different tutorial strategies that was partially developed by observing an experienced human tutor. We discuss aspects of this human tutor's method that can be modeled well by Ms. Lindquist and share examples of what the software is able to do. Through testing, we plan to learn about the relative effectiveness of the different tutorial strategies Ms. Lindquist uses. Ms. Lindquist is available at www.AlgebraTutor.org.

Introduction

The mission of the Center for Interdisciplinary Research on Constructive Learning Environments (CIRCLE) is 1) to study human tutoring and 2) to build and test a new generation of tutoring systems that encourage students to construct the target knowledge instead of telling it to them (VanLehn et. al. 1998). Computer Aided Instruction systems were 1st generation tutors. They presented a page of text or graphics and depending upon the student's answer, put up a different page. Model-tracing ITSs (Intelligent Tutoring Systems) are 2nd generation tutoring systems that allow the tutor to follow the line of reasoning of the student. ITS have had notable success (Koedinger et. al. 1997) despite the fact that human tutoring can look very different (Moore, 1996). One difference is that there is a better sense of a dialog in human tutoring. We, and others, think that is important; after analyzing over 100 hours of untrained tutors in naturalistic tutoring sessions, Graesser et. al. (1999) believe "there is something about interactive discourse that is responsible for learning gains."

The members of CIRCLE are working on 3rd generation tutoring systems that are meant to engage in a dialog with students to allow students to construct their own knowledge of the domain. We have built a new ITS, called Ms. Lindquist, that not only can model-trace the student's actions, but can be more human-like in carrying on a running conversation, complete with probing questions, worked examples, positive and negative feedback, follow-up questions in embedded sub-dialogs, and requests for explanation as to why something is correct. In order to build Ms. Lindquist we have expanded the model-tracing paradigm; not only does she have a model of the student but also has a model of tutorial reasoning (e.g. Clancey, 1982) for our specific domain.

McArthur et. al. (1990) criticizes Anderson's et. al. (1985) model-tracing ITS and model-tracing in general "because each incorrect rule is paired with a particular tutorial action (typically a stored message), every student who takes a given step gets the same message, regardless of how many times the same error has been made or how many other error have been made. ... Anderson's tutor is tactical, driven by local student errors (p. 200)." They go on to argue for the need for a more strategic tutor. Ms. Lindquist meets that criticism. Ms. Lindquist's model of tutorial reasoning is both **strategic** (i.e. has general multi-step plans that can be used to breakdown problems) and **tactical** (i.e. can recognize and respond to "teachable moments.") Merrill et. al. (1995) argued that human tutors are tactical in that "student-tutor dialogues were centered much more around student-initiated events, as they attempted to actively understand new instructional material and solve problems, [rather] than around tutorial presentation of material and subsequent querying of student understanding." Merrill et. al. went on to say "microanalysis of student-tutorial interactions in problem-solving situations suggests that tutors do more than simply reteach a correct procedure component when students encounter impasses or errors. Our tutors focused on guiding the error repair process rather then communicating

14

their guesses about what student's misconception." We argue that a tutor needs to be tactical in recognizing certain situations, and strategic enough to plan multi-step sequences of tutorial interaction Therefore, it is important that Ms. Lindquist can understand the student and come up with a strategy to respond with, even in the presence of multiple errors.

The task domain we are working on is symbolization, which is the task of writing an algebraic expression given a real-world problem context. (An example is shown below at line 240.) Symbolization is fundamental because if students can't translate problems into a mathematical algebraic model, they will not be able to apply algebra to solve real world problems. Fortunately, this domain makes it easy to avoid some difficult natural language issues because we can ask students to write algebraic expressions, and those expressions are easy for the computer student model to "understand". We take advantage of this property of the domain to avoid any serious natural language processing. Instead, we focus our energies on modeling tutorial reasoning which includes capturing the *pedagogical content knowledge (Shulman, 1986)* of an experienced human tutor. A good tutor not only knows the domain knowledge and general tutoring knowledge, but also has specific strategies that can help a student overcome common difficulties. We have set out to observe and model some of these strategies for our specific domain of symbolization. The following is an example of a content-specific pedagogical strategy for symbolization. It was collected and transcribed from a one hour long one-on-one tutoring session between an experienced human tutor and an eighth grade student.

240. Student: [reads the "bike-trip" problem] Cathy took a "m" mile bike ride. She rode at a speed of "s" miles per hour. She stopped for a "b" hour break. Write an expression for how long the trip took.
241. S: uhm [writes "s/m+b" but should be "m/s+b"]
242. Tutor: How do you calculate the amount of time it takes you? If you're, if you're, if you're riding at, let's make it simple. If you are riding at 20 miles per hour, OK, and you go 100 miles, how many hours did that take you?
243. S: Um 5
244. T: 5. And how did you get that 5? How did you use the numbers 100 and …
245. S: 100 miles divided by miles per hour
246. T: So you took the miles and divided it by the [garbled, but possibly "speed"]
247. S: Miles divided by s plus b equals time [writes m/s+b]
248. T: Right.[Full transcript available at www.AlgbraTutor.org]

We call the tutorial strategy displayed here the *concrete articulation strategy,* which Koedinger & Anderson (1998)

refereed to as *inductive support.* McArthur et. al. also observed that human tutors often used what he called *curriculum scripts* and *micro-plans,* which often involved a series of questions designed to remediate particular difficulties. We call these scripts *knowledge construction dialogs* to emphasis the fact that we are trying to build a tutor that encourages students to build their own knowledge rather than being told it. Below, we will show how Ms. Lindquist participates in an analogous dialog. We will also show three other tutorial strategies that Ms. Lindquist can use.

We think that if you want to build a good ITS for a domain you need to:
1) Study what makes that domain difficult, including discovering any hidden skills, as well as determining what types of errors students make.
2) Construct a theory of how students solve these problem. (We instantiated that theory in a cognitive model.)
3) Observe experienced human tutors to find out what pedagogical content knowledge they have and then build a tutorial model that, with the help of the theory of domain skills, can capture some of that knowledge.
4) Use your model to answer questions about what makes for effective tutoring.

We omit discussion here of the first and second steps (Heffernan & Koedinger, 1998 and 1999) and instead focus on the tutorial modeling effort. The fourth step is future work.

Cognitive Student Model

Our student model is a cognitive model of the problem solving knowledge that students are acquiring. The model reflects the ACT theory of skill knowledge (Anderson, 1993) in assuming that the problem solving skills can be modeled as a set of independent production rules. Our model has over 68 production rules. The cognitive model enables the tutor to trace the student's solution path through a complex problem solving space. The cognitive model for symbolization is tied to the underlying hierarchical nature of the problem situation. We model the common errors that students make with a set of "buggy" productions. From our data of 218 individual student errors on eight different problems we found that the following list of errors accounts for over 75% of the errors that students made. We illustrate the errors in the context of a problem which has a correct answer of "5g+7(30-g)".
1) Wrong operator (e.g. "5g-7(30-g)")
2) Wrong order of arguments (e.g. "5g+7(g-30)")
3) Missing parentheses (e.g. "5g+7*30-g")
4) Confusing quantities (e.g. "7g+5(30-g)")
5) Missing a component (e.g. "5g+7g" or "g+7(30-g)" or "5g+30-g")
6) Omission: correct for a subgoal. (e.g. "7(30-g)" or "5g")

7) Combinations of errors (e.g. "5g+7*g-30" has the wrong order for "g-30" and is missing parenthesis)

These "buggy" productions are used to allow us to make sense of a student's input even if there are several incorrect steps. Traditional model-tracing tutors generate their feedback from templates associated with individual production rules which is why McArthur et. al. criticized then as being to tactical. Instead, our system passes the diagnosis to the tutorial model that reasons between tactical and strategic concerns.

The Tutorial Model

Our tutorial model is informed by observation of an experienced human tutor, but it doesn't pretend to model any one individual or make claims to being the most effective model. The tutorial model is implemented with 77 production rules that decide what to say and how to say it. Many of these productions manipulate the tutorial agenda which operates usually like a push-down stack. Most of the productions are concerned with implementing the various tutorial strategies but some are concerned with more mundane aspects including maintaining conversational coherence, giving positive feedback, and dealing with issues related to how to deal with multiple student errors or incomprehensible student input. This approach is similar to Freedman's (2000). We will now look at four different strategies that Ms. Lindquist can use on any problem. The examples also show instances of more tactical behavior trigged by certain situations.

Concrete Articulation Strategy

The concrete articulation strategy usually uses three questions illustrated below in T2, T3, and T4. Here is an example of the dialog generated by Ms. Lindquist which we believe closely parallels the experienced human tutor shown above.

T1: [Doing the "bike-trip" problem] Please write an expression for the total time the trip took. [The first question always asks the student to symbolize.]

S1: s/m+b

T2: No. Let me ask you an easier question. Please compute the time actually on the bikes if the speed was 20 miles per hour, and the distance ridden was 100 miles? [ask to compute a concrete instance]

S2: 5

T3: Correct. What math did you do to get that five? [ask student to articulate in symbols the answer]

S3: 100/20

T4: Correct. Now write your answer of 100/20 using the variables given in the problem (i.e. put in 'm' and 's')? [ask student to generalize their concrete articulation]

S4: b+m/s [was expecting "m/s"]

T5: Correct.

We will walk through what Ms. Lindquist did here. First, the student model diagnosed S1 ("s/m+b") as having the quantity "s/m" with reversed arguments (bug type #2.) The diagnosis also said that the student understood that the total trip time was the time of the break ("b") plus the time actually on the bikes. This diagnosis was passed to the tutorial model chosen from competing strategies. In this case the tutorial model decided to conduct a *concrete articulation* knowledge construction dialog. Therefore, the tutor model generated questions for the *compute, explain/articulation,* and *generalization* steps. The tutor put these three questions on the agenda.

At S2, the student answers the compute question correctly, and the tutor decides to continue with its plan. It has to be sensitive to what the student typed. If the student typed 100/20 instead of 5, then both the compute and explain questions would be removed from the agenda and the tutor would skip to the generalize step. An example of this sort of flexibility is demonstrated in T5. T4 was a question to articulate the time actually on bikes. The tutor was hoping the student would type "m/s" but instead the student answered "b+m/s." Only a pedantic tutor wouldn't accept that answer as it indicates an understanding that subsumes the understanding required for the sub-goal. This is where the tutorial agenda behaves differently than a stack because Ms. Lindquist pops both of the questions off the agenda. To show you an example of how the previous situation could have come out differently, lets look at the following dialog also generated by Ms. Lindquist.

T1: Hello. [Doing the "bike-trip" problem] Please write an expression for the total time the trip took. [symbolize]

S1: help! [The student can request help. The tutor treats such a request the same way it treats an totally incomprehensible student input.; it starts breaking down the problem into steps.]

T2: Let me ask you an easier question. Please compute the time actually on the bikes if instead the speed was 20 miles per hour, the time of the break was 2 hours, and the distance ridden was 100 miles? [compare to T2 for the above example]

S2: 7

T3: Correct. What math did you do to get that 7?

S3: 2+5 [student doesn't completely articulate]

T4: Okay, but how did you get the 5? [similar to T3]

S4: 100/20

T5: Correct Now lets try this again, what math did you do to get the 7? [same question as T3 but the phrase "lets try this again" thrown in for conversational coherence]

S5: 2+100/ 20 [correct articulation]

T6:Correct. Now write your answer of '2+100/20' using the variables given in the problem? [generalize like T4 from pervious example]

16

S6: b+m/s [of course accepts "m/s+b"]
T7: Correct.

This example is interesting in the way Ms. Lindquist responds at T4. The system has bug rules to apply to questions other than symbolization questions. In this case, the student model reported that the student's answer of "2+5" is a partially complete articulation but that the student didn't spell out where the "5" came from. The tutorial model is able to "dig-down" into the student 's reasoning with an embedded sub-dialog. This involves adding the additional question of explaining how the student got the "5". We think that such tactical behavior, conditioned on particular student errors, is consistent with Merrill's et. al. analysis of what human tutors do. Ms. Lindquist can use embedded multi-step dialogs to respond to other specific situations such as parentheses errors, errors of omission, miscomprehension of the question, and argument ordering errors.

Substitution with Decomposition Strategy

A second example of a tutorial strategy is what we call the *substitution with decomposition* strategy. This strategy is motivated by previous cognitive analysis and experimentation that will not reviewed here (Heffernan & Koedinger 1998). Below we present an example of Ms. Lindquist tutoring with this strategy. (As this article went to press we had just conducted our first pilot testing with human students. The following example is verbatim from this pilot test.)

T1: Hello. Jane is "j" years old. Peg is "p" years old. Mary is Peg's age minus Tom's age. Tom is 3 years younger than Jane. Peg is "x" years older that Mark. Write an expression for Mary's age.[correct answer is "p-(j-3)"]
S1: (p-3)-(3-j)+(p+x)
T2: Hmm. ["Hmm" indicates students response was not diagnosed] No. Let me ask you and easier question. Use j and 3 to calculate Tom's age. [ask for the one subgoal]
S2: 3-j [This is a common error type that the tutor should be tactical enough to recognize. The tutor has since been modified to be able to respond with the strategy of asking the student "Should the answer be positive or negative?" followed by "Which is larger, the '3' or 'j'?]
T3: Hmm. No. Try again please. What is Tom's age? Hint: pick one of these: j-3, 3-j, j+3, j*3, j/3, or 3/j. [This is a rephrasing as explained in the discussion section]
S3: j-3
T4: Correct. Calculate Mary's age. Use A instead of j-3 for Tom's age.
S4: p-a
T5: Correct. You are correct that p-A is Mary's age. Now, can you state this in terms of the numbers

given in the problem. [Ask the student to algebraically substitute]
S5: p-(j-3)
T6: Correct.

Two More Strategies that Foster Translating Between Multiple Representations

Mathematical communication (Schoenfeld et. al. 1992, Koedinger & Nathan, in press) is increasingly seen as a goal for students, and in particular, translating between algebraic (e.g. "m/s+h"), English (e.g. "the total trip time") and concrete (e.g. "7 hours") representations. So far we have discussed students translating from the English to algebraic or concrete representations. The next two strategies explore asking students to translate in different ways including: 1) translating from the algebraic to the English and 2) within the English representation.

Subgoal Reification Strategy

The *subgoal reification strategy* is inspired by Corbett & Anderson's (1995) lisp tutor that showed learning gains when the students we asked to post (or reify) their subgoals before attempting to write lisp code. We have added to our tutor the ability to choose a similar strategy. At T2 below, the student is asked for the name of the subgoal. At T3 the student is asked to explain in English how to achieve that subgoal. At T4 the tutor asks for the subgoal in algebra. T5 and T6 repeat the pattern from T3 and T4.

T1: Hello. [Again, doing the "bike-trip" problem] Please write an expression for the total time the trip took.
S1: help!
T2: In order to find the total time, what do you first have to find out? [The student is presented with pull-down menu to select a quantity-noun phrase]
S2: [menu:] The time actually on the bikes
T3: Correct. Can you explain how you would find the time actually on the bikes)? [Presented with the following sentence to fill in the blanks: "The time actually on the bikes is equal to (select quantity) (select operator) (select quantity)." The student can select the 2 quantities and the operator from menus.]
S4: The time actually on the bikes is equal to [1st menu:] the distance biked [2nd menu:] divided by [3rd menu:] the speed of the bikes.
T4: Good. Now say what that [note the simple pronoun substitution] in symbols.
S5: m/s
T5: Correct, Can you explain how you would find the total time of the trip using the time on the bikes?
S6: The total time of the trip is equal to [1st menu:] the amount of time for the break [2nd menu:]

17

added to [3rd menu:] the time actually on the bikes.
T6: Can you say that in symbols?
S6: b+m/s
T7: Correct.

Self-Explanation of a Worked Example

The next strategy is inspired by the literature on worked examples and the self-explanation effect. Sweller (1988) proposed giving students worked out examples. Chi (1996) and Conati et. al. (1997) have both proposed asking students to self-explain. This strategy also gets at the goal of practicing translation from the algebraic representation to the verbal representation which Corbett et. al. (2000) have shown to be effective for learning symbolization. In this example, the student is told the answer and the asked to explain it.

T1: Hello. [Again, doing the "bike-trip" problem] Please write an expression for the total time the trip took.
S1: m+s
T2: No. The correct answer is "b+m/s". Let me ask you some questions to make sure you understand that. Can you explain where the m/s came from? [This is similar to T3 from above. The student is presented with a sentence they must complete. The sentence begins "'m/s' is equal to" and is followed by three pull-down menus to compose the answer.]
S2: "m/s" is equal to [1st menu:] the distance ridden [2nd menu:] divided by [3rd menu:] the speed ridden.
T3: Correct. What does m/s represent?
S3: [menu:] The time actually on bikes
T4: Correct. Can you explain where the b+m/s came from?
S4: "b+m/s" is equal to [1st menu:] the time for the break [2nd menu:] plus [3rd menu:] the time actually on bikes.
T5: Correct. What does "b+m/s" represent?
S5: [menu:] The time actually on bikes
T6: Correct. It looks like you understand that problem.

Discussion

How is Ms. Lindquist different from traditional model-tracing tutors? Traditional model-tracing tutors offer feedback primarily through two mechanisms: 1) a "buggy" message that is triggered by the student's input matching with a known buggy rule or 2) in response to the student requesting a hint. A student can usually continue to ask for a hint getting progressively more explicit messages. These "hint chains" sometime terminate in messages that tell the student exactly what to do, but the model-tracing paradigm does not require this.

Ms. Lindquist is different in that she has a greater capacity to carry on a dialog by being able to handle embedded, multi-step dialogs (e.g. "Where did you get the '5'?") Though many model-tracing systems phrase their hints as questions, they are usually just rhetorical questions. When Ms. Lindquist ask you a question, she expects you to answer it. She does this because her strategies suggest that these questions are good ways to "break-down" complicated problems. If Ms. Lindquist has run out of options on how to "break-down" a problem, then she adopts a strategy similar to model-tracing tutors and gives progressively more explicit rephrasings of a question. Most of Ms. Lindquist's rephrasings terminate with multiple choice questions rather then simply telling the student the answer. Through all of this, we hope to make the student more active.

Our experienced tutor was not always satisfied with a student's answer, particularly when it looked like the student might have guessed. Sometimes she would ask reflective follow-up questions to test the student's understanding. Ms. Lindquist does a similar thing: when a student has likely guessed (indicated by having got the answer only after having reached the most explicit rephrasing) the student is asked one of a few types of follow-up questions. A long-term research goal is to learn which types of follow-up work best.

Another way Ms. Lindquist is more similar to human tutors is in being more active. Others have also viewed tutor dialogs as primarily tutor initiated (Graesser et. al. 1999 and Lepper et. al. 1997.) Ms. Lindquist does not wait for a student to ask for help. Our experience tutor made a comment on average once every 20 seconds! In summary, we view Ms. Lindquist as capturing a balance between strategic and tactical responses. She is able to break down problems with different tutorial strategies while at the same time she can make more tactical decisions in response to particular situations (e.g. common errors or student guessing.)

We look forward to measuring the effectiveness of Ms. Lindquist by comparing it to an effective benchmarked 2nd generation ITS (Koedinger et. al. 1997.) We also look forward to comparing the effectiveness of different strategies by allowing the tutor to pick randomly among the tutorial strategies. We can then measure the effectiveness of a strategy by seeing if the student correctly answers an isomorphic problem later in the curriculum. Ms. Lindquist is available on the web at www.AlgebraTutor.org.

References

Anderson, J. R. (1993). *Rules of the Mind.* Hillsdale, NJ: Erlbaum.

Anderson, J. R., Corbett, A. T., Koedinger, K. R., & Pelletier, R. (1995) Cognitive tutors: lessons learned. *The Journal of the Learning Sciences,* 4 (2), 167-207.

Anderson, J. R., Boyle, D. F., & Reiser, B. J. (1985). Intelligent tutoring systems. *Science,* 228, 456-462.

Chi, M. T. H. (1996) Constructing self-explanations and scaffolded explanations in tutoring. *Applied Cognitive Psychology,* 10, S33-S49

Clancey, W. J., (1982) Tutoring rules for guiding a case method dialog. In D. Sleeman & J. S. Brown (Eds.) *Intelligent Tutoring Systems* London: Academic Press. (pp. 201-226.)

Corbett, A. T., and Anderson, J. R., (1995) Knowledge decomposition and subgoal reification in the ACT programming tutor. in *Proceedings of Artificial Intelligence in Education* (pp. 469-476)

Corbett, A. T., McLaughlin, M., Scarpinatto, C., & Hadley, W. (2000) Analyzing and generating mathematical models: an algebra II cognitive tutor design study. To appear in the Proceedings of Intelligent Tutoring Systems Conference.

Conati, C., Larkin, J. and VanLehn, K. (1997) A computer framework to support self-explanation. In : du Bolay, B. and Mizoguchi, R.(Eds.) Proceedings of AI-ED 97 World Conference on Artificial Intelligence in Education. Vol.39, pp. 279-276, Amsterdam: IO Press.

Freedman, R. (2000 to appear) Using a reactive planner as the basis for a dialogue agent. In *Proceedings of the Thirteenth Florida Artificial Intelligence Research Symposium* (FLAIRS '00), Orlando.

Graesser, A.C., Wiemer-Hastings, K., Wiemer-Hastings, P., Kreuz, R., & the TRG (1999). AutoTutor: A simulation of a human tutor. *Journal of Cognitive Systems Research,* 1, 35-51.

Heffernan, N. T., & Koedinger, K. R.(1997) The composition effect in symbolizing: the role of symbol production versus text comprehension. *Proceeding of the Nineteenth Annual Conference of the Cognitive Science Society* 307-312. Hillsdale, NJ: Lawrence Erlbaum Associates.

Heffernan, N. T., & Koedinger, K. R. (1998) The composition effect in symbolizing: the role of symbol production vs. text comprehension. *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society,* 307-312. Hillsdale, NJ: Lawrence Erlbaum Associates.

Koedinger, K. R., Anderson, J.R., Hadley, W.H., & Mark, M. A. (1997). Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education, 8,* 30-43.

Koedinger, K. R., & Anderson, J. R. (1998). Illustrating principled design: The early evolution of a cognitive tutor for algebra symbolization. In *Interactive Learning Environments,* 5, 161-180.

Koedinger, K. R. & Nathan, M. J. (submitted to). The real story behind story problems: Effects of representations on quantitative reasoning. Submitted to *Cognitive Psychology.*

Lepper, M. R., Drake, M. F., & O'Donnell-Johnson, T. (1997) "Scaffolding techniques of expert human tutors. In K. Hogan & M. Pressley (Eds.) *Scaffolding Student Learning: Instructional approaches and Issues.* Cambridge MA: Brookline Books (pp. 108-144.)

McArthur, D., Stasz, C., & Zmuidzinas, M. (1990) Tutoring techniques in algebra. *Cognition and Instruction.* 7 (pp. 197-244.)

Merrill, D. C., Reiser, B. J, Merrill, S. K., & Landes, S. (1995) Tutoring: guided learning by doing. Cognition and Instruction 13(3) (pp. 315-372.)

Moore, J. D. (1993) What makes Human Explanations effective? In *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Erlbaum. (pp. 131-136).

Schoenfeld, A., Gamoran, M., Kessel, C., Leonard, M., Or-Bach, R., & Arcavi, A. (1992) Toward a comprehensive model of human tutoring in complex subject matter domains. *Journal of Mathematical Behavior,* 11, 293-319

Shulman, L. (1986). Those who understand: Knowledge growth in teaching. *Educational Researcher, 15,* 4-14.

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science,* 12, 257-285.

VanLehn, K, Anderson, J., Ashley, K., Chi. M., Corbett, A., . Koedinger, K., Lesgold, A., Levin, L., Moore, M., and Pollack, M., NSF Grant 9720359. *CIRCLE: Center for Interdisciplinary Research on Constructive Learning Environments.* NSF Learning and Intelligent Systems Center. January, 1998 to January, 2003.