Pedagogical Content Knowledge in a Tutorial Dialogue System to Support Self-Explanation

Vincent Aleven, Octav Popescu, and Kenneth R. Koedinger

Human-Computer Interaction Institute

Carnegie Mellon University

aleven@cs.cmu.edu, octav@cmu.edu, koedinger @cs.cmu.edu

Abstract: We are engaged in a research project to create a tutorial dialogue system that helps students learn through self-explanation. Our current prototype is able to analyze students' general explanations of their problem-solving steps, stated in their own words, recognize the types of omissions that we often see in these explanations, and provide feedback. Our approach to architectural tradeoffs is to equip the system with a sophisticated NLU component but to keep dialogue management simple. The system has a knowledge-based NLU component, which performed with 81% accuracy in a preliminary evaluation study. The system's approach to dialogue management can be characterised as “classify and react”. In each dialogue cycle, the system classifies the student input with respect to a hierarchy of explanation categories that represent common ways of stating complete or incomplete explanations of geometry rules. The system then provides feedback based on that classification. We consider what extensions are necessary or desirable in order to make the dialogues more robust.

INTRODUCTION

Self-explanation is an effective metacognitive strategy. Explaining examples or problem-solving steps helps students learn with greater understanding (Chi, et al., 1989; 1994; Berardi-Coletta, et al., 1985; Gagne & Smith, 1962). Yet few students are good self-explainers, even when prompted (Renkl, et al., 1998). So how can we leverage self-explanation to improve learning in actual classrooms? The AI & Education literature provides evidence that self-explanation can be supported effectively by 2nd-generation tutors (Aleven et al 1999; Conati & VanLehn, 2000). However, these systems did not interact in natural language. It is plausible that students will learn even better when they explain in their own words. Natural language allows for flexible expression of partial knowledge: Students can show what they know and the tutor can help them to construct more complete knowledge. Also, articulation forces attention to relevant features. Finally, combining visual and verbal learning modes may create “dual codes” in memory which may facilitate recall (Paivio, 1986). However, it appears that these potential advantages will not be fully obtained without tutoring or giving feedback to students. When students worked with a tutor version that prompted them to explain their steps in their own words, but did not check explanations, they often ignored these prompts and provided almost no good explanations (Aleven & Koedinger, 2000b).

We are preparing to test the hypothesis that students learn better when they explain in their own words and receive feedback on their explanations. To this end we are developing a tutorial dialogue system, the Geometry Explanation Tutor, that assists students as they generate general explanations of their problem-solving steps in their own words. The system engages students in a restricted form of dialogue to help them improve explanations that are not sufficiently precise. We have a working prototype and are starting a phase of pilot testing. The Geometry Explanation

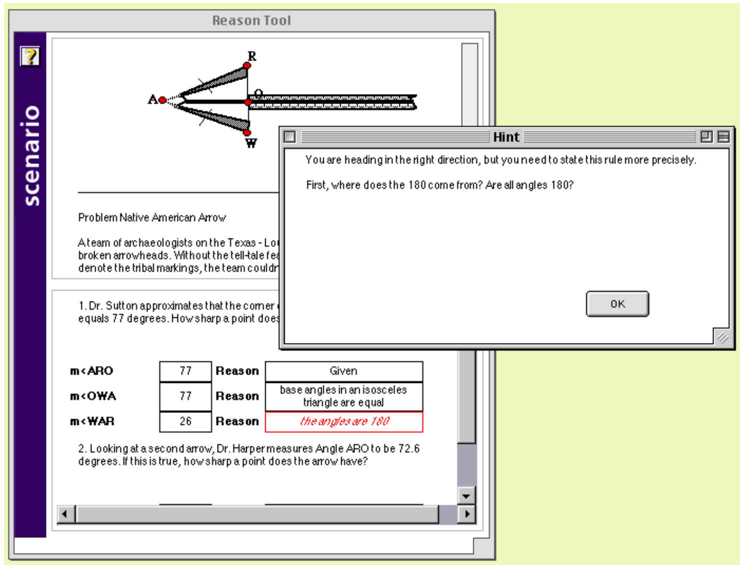

Figure 1: The Geometry Explanation Tutor

Tutor is built on top of an existing 2nd-generation system for geometry problem solving, the PACT Geometry Tutor (Aleven et al., 1999), which is currently in use in about five schools in the Pittsburgh area and elsewhere.

In designing the architecture of the system, we are faced with a number of choices. Thus we find ourselves asking the question, as phrased in the call for papers, "Where is the biggest bang for the buck?" A significant architectural decision has been to equip the system with a fairly sophisticated NLU component (Popescu & Koedinger, 2000). A detailed understanding of the explanations is needed if the system is to provide detailed feedback. A second decision has been to keep the system's dialogue planning and management component as simple as possible, but in Einstein's words, no simpler than that. We follow the approach taken by Heffernan and Koedinger (2000) in developing Ms. Lindquist, an algebra symbolization tutor, and focus on identifying the *pedagogical content knowledge* needed to help students produce accurate and complete explanations. By pedagogical content knowledge they mean domain-specific strategies that experienced human tutors use to help students deal with common difficulties and to scaffold students' problem-solving efforts. Pedagogical content knowledge also includes knowledge about students, their typical errors and typical, often rugged, pathways to learning success.

We foresee that a tutor that helps students to generate accurate geometry explanations needs to have knowledge about (1) how to provide good and detailed comments that help students to improve explanations that are incomplete and (2) how to lead students to good explanations if they have difficulty getting started. So far, we have focused on the first need. The analysis of several small corpora of student explanations indicated that students explanations of geometry tend to be incomplete more often than wrong. The system therefore has a hierarchy of explanation categories that represent common ways of stating full and partial explanations of geometry rules. It decides what feedback to give to student explanations primarily by classifying them into this hierarchy. While it is not difficult to see the limitations of the current system, it is not easy to predict what improvements will give the greatest bang for the buck. Thus, in extending the system, we plan to be guided by results of frequent preliminary evaluation and pilot studies, adding more sophisticated mechanisms or strategies only when the data suggest that they will improve students' learning.

In this paper, we describe the current architecture of the Geometry Explanation Tutor and illustrate its current capabilities by means of dialogue examples. We present results from a

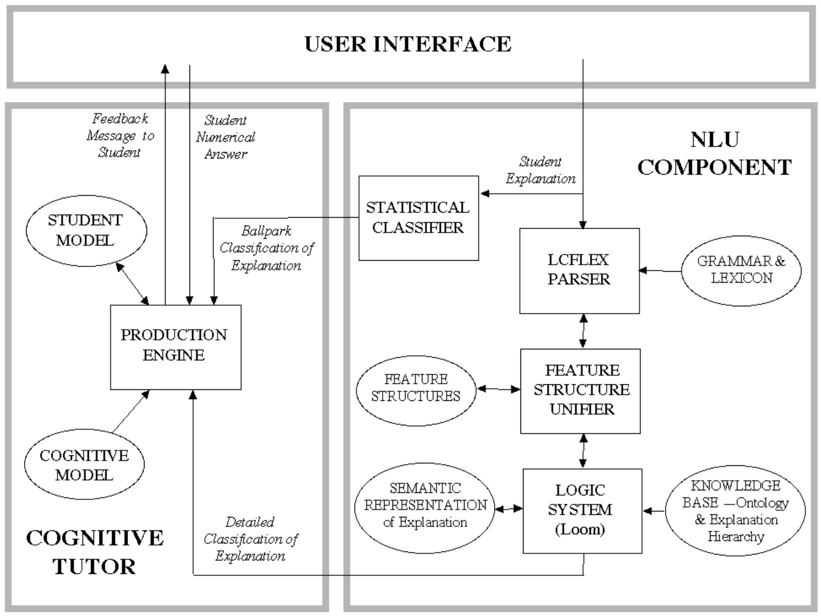

Figure 2: Architecture of the Geometry Explanation Tutor

preliminary evaluation of the accuracy of the system's NLU component. Finally, we discuss what limitations need to be addressed most urgently: What pedagogical content knowledge we will need to add and how far we will have to push the system's dialogue management architecture.

THE GEOMETRY EXPLANATION TUTOR

The Geometry Explanation Tutor covers one of the six units that make up the curriculum of the original PACT Geometry Tutor, namely, the unit that deals with the geometric properties of angles. The Geometry Explanation Tutor provides problem-solving support, just like other Cognitive Tutors (Koedinger, et al., 1997). It monitors students as they work through problems and provides assistance in the form of feedback and context-sensitive hints. Unlike other Cognitive Tutors, the Geometry Explanation Tutor requires that students explain their steps and engages students in a restricted form of dialogue in order to help students state geometry rules accurately (Popescu & Koedinger, 2000; Aleven, et al., in press). The system has been pilot-tested with 20 of our colleagues and staff and with two high-school students.

System Architecture

The Geometry Explanation Tutor is based on the standard Cognitive Tutor architecture (Anderson et al., 1995), augmented with a NLU component (see Figure 2). In each dialogue cycle, the NLU component creates a semantic representation of the student's explanation and *classifies* that representation with respect to the system's hierarchy of explanation categories. The Cognitive Tutor module then checks whether the student's explanation focuses on the right geometry rule and decides how to *react* (i.e., what feedback to give to the student).

An important knowledge source is the hierarchy of explanation categories, which constitutes the system's pedagogical content knowledge. The explanation categories in this hierarchy represent ways of stating each geometry rule correctly, as well as frequently occurring ways of stating rules incorrectly. An excerpt of this hierarchy is shown in Figure 3. Each node represents a class of explanations that have the same meaning, but may have vastly different surface forms. A canonical example of a sentence that falls in each category is shown in each node. Explanation

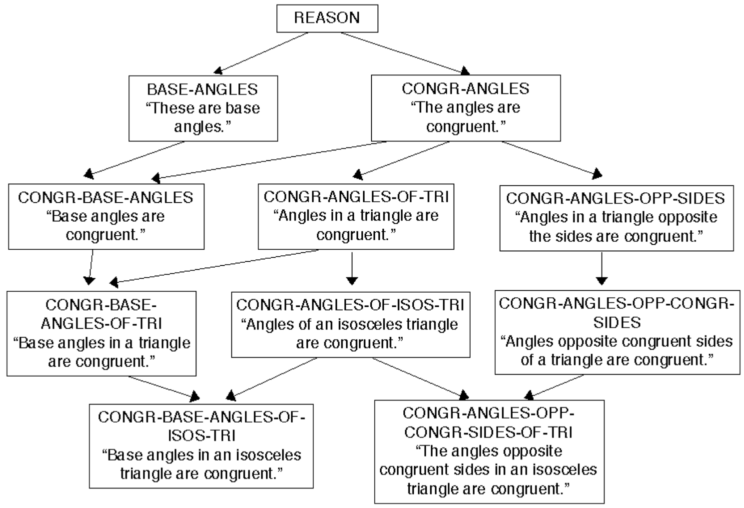

Figure 3: Excerpts from the explanation hierarchy, represented in the system's knowledge base

categories at the bottom of the hierarchy represent correct and complete ways of stating the isosceles triangle rule. Explanations categories higher up in the hierarchy represent progressively more incomplete ways of stating this rule. The hierarchy also includes information about how to respond to student explanations. Attached to each category is a feedback message that is appropriate when an explanation by the student is classified under that category. We have identified about 140 explanation categories, related to the 25 geometry rules that make up the tutor's Angles unit. A key point is that these categories were driven from observations of real student data, contained in several corpora of student explanations. Thus, this rich network captures categories that occur frequently as learners progress towards success.

The hierarchy is implemented as a Loom knowledge base (MacGregor, 1991). This knowledge base also contains an ontology of the domain, which consists of geometry objects such as angles and lines, as well as relations such as congruency, adjacency, etc. The ontology covers the material of the Angles unit of the tutor curriculum. Currently, the knowledge base contains definitions for about 310 concepts and 90 relations.

The NLU component parses student input using a unification-based approach. We employ the LCFLEX parser, an active chart parser (Rose & Lavie, 1999), in combination with a feature structure unifier. We have developed a grammar of about 200 rules. The parser and unifier build a feature structure encoding the syntax of the sentence. They also direct Loom to build a semantic representation. In the process, the Loom classifier tests the coherence of the semantic representation with respect to semantic constraints expressed in the system's domain ontology. When a coherent semantic representation has been constructed, Loom classifies it with respect to the explanation categories.

The Cognitive Tutor module decides how to respond to the student. First, the tutor determines which geometry rule the student should be explaining. This rule must be one that justifies the current problem-solving step. The tutor determines the set of applicable geometry rules by running its cognitive model of geometry problem solving. This model captures the typical and desired knowledge of geometry students, represented in the form of production rules. The tutor then selects an appropriate feedback message as follows. If the student explanation is a complete statement of a relevant geometry rule, the tutor accepts the explanation. If the explanation is only a partial statement, the tutor selects an appropriate feedback message. Usually, this will be the feedback message associated with the category under which the explanation was classified. If the explanation is only the name of a geometry rule, the tutor asks

Table 1: Example dialogue with the Geometry Explanation Tutor. Shown on the left are a student's successive attempts at explaining the given geometry rule (underlining shows what was added since the previous attempt) and the tutor's responses. Shown on the right is the explanation category under which the tutor classified the student's explanation.

1.	**S:** congruent	NIL
2.	**T:** The Geometry Tutor is not able to understand your explanation. Sorry! Please add to your explanation or type something different. Please state a GENERAL geometry rule. Do not mention the names of the angles in the problem.	
3.	**S:** <u>the base angles are congruent</u>	CONGRUENT-BASE-ANGLES
4.	**T:** You are getting there, but you need to be more precise. Are base angles always congruent?	
5.	**S:** the base angles are congruent <u>here</u>	CONGRUENT-BASE-ANGLES
6.	**T:** You are getting there, but you need to be more precise. Are base angles always congruent?	
7.	**S:** Base angles <u>in a triangle</u> are congruent	CONGRUENT-BASE-ANGLES-OF-TRIANGLE
8.	**T:** Well.... Almost! Are the base angles in EVERY triangle congruent?	
9.	**S:** The base angles in an <u>isosceles</u> triangle are congruent	CONGRUENT-BASE-ANGLES-OF-ISOSCELES-TRIANGLE
10.	**T:** [ACCEPTED as correct explanation. Students get non-verbal positive feedback.]	

the student to state the rule. Finally, if the explanation focuses on the wrong geometry rule, the tutor will inform the student of this fact.

We have begun to experiment with the integration of a statistical Naïve Bayes text classifier (Mitchell, 1997, Ch. 6). When the knowledge-based NLU component fails to produce an analysis, the statistical classifier is used to determine whether the student's explanation is in the ballpark (i.e., focuses on the correct geometry rule). If so, the tutor will print a feedback message saying that the student appears to be on the right track. This helps in dealing with unexpected input, as is illustrated below. We are looking for further ways to leverage the statistical text classifier.

EXAMPLE DIALOGUES WITH THE GEOMETRY EXPLANATION TUTOR

The Geometry Explanation Tutor's capabilities are illustrated in two dialogues shown in Tables 1 and 2. These dialogues are adapted from protocols collected during a session in which about 20 of our colleagues and staff worked with the system. In order to make the dialogs reflect the system's current capabilities, we have replaced the system's responses with responses generated by the *current* system version. The downside is that the dialogues are not 100% realistic.

In the first example, the tutor helps a student in explaining the isosceles triangle theorem, which says (stated somewhat informally) that the base angles in an isosceles triangle are congruent. Starting with a minimum effort explanation (step 1) and testing the tutor's sense of humor along the way (steps 3-5), the student incrementally refines an explanation until the explanation is accepted by the tutor (step 10). The tutor produces appropriate responses using its pedagogical content knowledge, captured in the explanation hierarchy. For example, category CONGRUENT-BASE-ANGLES-OF-ISOSCELES-TRIANGLE represents correct and complete statements of the isosceles triangle rule (see Figure 3, bottom left). Student explanations in this category get the thumbs up from the tutor, provided of course that the isosceles triangle rule justifies the problem-solving step being explained (see step 9 and 10). On the other hand, category CONGRUENT-BASE-ANGLES-OF-TRIANGLE represents statements such as "the

Table 2: Example Dialogue with the Geometry Explanation Tutor

1.	**S:** 180 is what you get when you add the interior angles	EQUAL-180
2.	**T:** Can you be more specific? What is 180? And what tells you that this is so?	
3.	**S:** interior	NIL
4.	**T:** It appears that you are getting at the right rule, the triangle sum rule. But the tutor has difficulty understanding your explanation. Please make small changes .	
5.	**S:** triangle sum	REFERENT
6.	**T:** It looks like you have the right idea. Can you state this geometry rule more completely?	
7.	**S:** triangle sum is 180	SUM-OF-TRIANGLE-180
8.	**T:** Very close. The sum of what exactly is equal to 180? Not of the triangle itself but of its... ?	
9.	**S:** add the interior angles is 180	INTERIOR-ANGLES-180
10.	**T:** That is good, but you need to be more specific. Which interior angles do you mean?	
11.	**S:** interior angles add up to 180	INTERIOR-ANGLES-SUM-180
12.	**T:** Very good. The sum of interior angles of WHAT is 180 degrees?	
13.	**S:** interior angles of the triangle add up to 180	ANGLES-SUM-OF-TRIANGLE-180
14.	**T: [**ACCEPTED as correct explanation. Students get non-verbal positive feedback.]	

base angles of a triangle are congruent," which are incomplete (Figure 3, left, one level from bottom). In response to statements in this category, the tutor generates a feedback message that suggests that the statement is an overgeneralization, provided again that the rule that needs to be explained is the isosceles triangle rule (steps 7 and 8). The feedback message is the one associated to the given category. Generally, the dialogue is smooth. Of course one prefers to avoid tutor responses like "the tutor does not understand," as shown in step 2 or instances where the tutor repeats itself in subsequent dialogue turns, (steps 4 and 6). Under the given circumstances however these responses were quite appropriate.

In the second dialogue (shown in Table 2), the student starts off rather well. The student's first explanation attempt could have been completed simply by adding "of a triangle" at the end, so that the sentence reads "180 is what you get when you add the interior angles of a triangle." Unfortunately, the tutor feedback does not make this clear. After a minimalist strategy in step 3, the student quickly gets on track again and gradually improves the explanation until it is complete in step 13. With the exception of the first tutor message, the tutor's feedback seems appropriate and helpful. The reason that the tutor did not produce a more helpful message in step 2 is that the NLU component currently does not handle the construction "is what you get when you add …"

The example illustrates that the statistical text classifier sometimes enables the tutor to produce a more helpful feedback message than it could if it only had the knowledge-based NLU component. In step 3, the student's answer "interior" is not classified under any explanation category by the knowledge-based NLU component. The statistical classifier however returns TRIANGLE-SUM as the most likely category. This enables the tutor to acknowledge (in step 4) that the student is on the right track ("it appears that you are getting at the right rule, the triangle sum rule"). Without the statistical classifier, the tutor could only have said, "the tutor does not understand your explanation."

Table 3: Classification accuracy of the knowledge-based NLU component

Classification result	***N***	**%**
Classified correctly		
Explanation was actually complete	194	29.9
Explanation was actually incomplete	227	35.0
Explanation was actually a reference	102	15.7
Classified under overly general category		
Explanation was actually complete	60	9.3
Explanation was actually incomplete	38	5.9
Classified incorrectly	3	0.5
Not classified	24	3.7
Total	648	100

Further, the example dialogue illustrates that the tutor accepts some common forms of abbreviations that students make. Students very often say "the angles are 180" when they mean strictly speaking that the measures of the angles are 180. This is a form of metonymy, the phenomenon of referring to a concept by means of a related concept [Jurafsky & Martin, 2000]. A prime example is "New York called" where it was the guy or girl from that city who called. The tutor accepts common forms of metonymy without complaint. For example, the tutor responds to the sentence "interior angles add up to 180" (step 11) as if the student said "the measures of interior angles add up to 180". Similarly, the sentence "a linear pair is 180" is treated as if the student had said "the measures of the angles in a linear pair are 180" (double metonymy).

However, the tutor is not so accommodating that it accepts all abbreviations or elliptical expressions. Nor should it be. A tutor whose goal it is to help students learn to "speak mathematics" should be helpful but should also insist on a certain level of precision in language. For example, the tutor does not interpret "triangle sum is 180" (step 7) as meaning "the sum of the measures of the angles of a triangle is 180 degrees," even though one might argue that this is what was meant. The tutor does not assume that the sum was implied.

The challenge in dealing with metonymy is to construct a semantic representation that observes semantic constraints even though the input sentence does not. Some forms of metonymy are dealt with in our system by having the NLU component recover the missing structure (Popescu & Koedinger, 2000). The semantic representations that are constructed in this way are the same as those that would have been built if the same sentence had occurred without abbreviations. The abbreviations that the tutor needs to know about are handled differently. The semantic representation that is constructed reflects the abbreviation, so that the tutor can take notice. For example, for the sentence "a triangle sums to 180" the NLU component builds a semantic representation that does not include angles. This gives the tutor the option of insisting on a more complete explanation (step 8). Exactly on what level of precision the system should insist is not easy to determine. The choice should ultimately be driven by what is best for student learning.

PRELIMINARY EVALUATION OF THE NLU COMPONENT

We conducted a preliminary evaluation of the classification accuracy of the knowledge-based NLU component. As test data we used a corpus of 648 explanations collected during a session in which about 20 of our colleagues and staff worked on the system. The explanations were labeled by hand by two authors of the paper, who assigned each explanation to the most specific category in the explanation hierarchy to which it belongs. New categories were invented as needed. A total of 138 categories were used, 92 of which were represented in the data set. The system's knowledge base was then extended to include definitions for many of the new categories. The system was then run to classify the 648 explanations.

Table 4: Example Dialogue with the Geometry Explanation Tutor

9.	**S:** that two angles in a isosceles triangle are equal	CONGRUENT-ANGLES-OF-ISOSCELES-TRIANGLE
10.	**T:** Correct, except for one thing: are all three angles in an isosceles triangle congruent?	
11.	**S:** that two <u>congruant</u> angles in a isosceles triangle are equal	CONGRUENT-ANGLES-OF-ISOSCELES-TRIANGLE
12.	**T:** Correct, except for one thing: are all three angles in an isosceles triangle congruent?	
15.	**S:** that two congruant <u>base angles</u> are equal	CONGRUENT-BASE-ANGLES
16.	**T:** You are getting there, but you need to be more precise. Are base angles always congruent?	
19.	**S:** base angles <u>in a isosceles triangle</u> are equal	CONGRUENT-BASE-ANGLES-OF-ISOSCELES-TRIANGLE
20.	**T: [**ACCEPTED as correct explanation. Students get non-verbal positive feedback.]	

As shown in Table 3, the system classified 81% of the explanations correctly. Of the correctly classified explanations, 80% (421 out of 523) fall under categories of full or partial explanations. The rest were references, meaning that the student stated only the name of a geometry rule. The system classified a further 15% of the explanations under categories that were too general, although not strictly wrong. The remaining 4% of explanations were either not classified at all or under categories that were unrelated to the correct category. Interestingly, when the system classified an explanation as being complete, it was correct 100% of the time.

An accuracy score of 81% is very encouraging, especially given the fact that we are dealing with a very fine-grained classification task, where small differences between categories are the rule rather than the exception. There are two caveats. First, the accuracy results were obtained with a data set that was used during the development of the system. Second, this data set was obtained with subjects who are more advanced than those in the target population (high-school students). More work is needed before we expect to see the same accuracy score with new data and students from the target population. Nonetheless, the results provide a preliminary indication that knowledge-based NLU is an appropriate choice for analysing geometry explanations.

LIMITATIONS OF CLASSIFY AND REACT

Currently, the system's response in each dialogue turn is based only on the classification of the student's last explanation attempt. No further context is taken into account. This way, the tutor can respond to the types of omissions we often see in students' explanations and can sometimes produce a sense of coherent dialogue, as illustrated in the examples. However, one does not have to look far to see the limitations of the approach. For example, the system has no memory of what went on before in the dialogue. It is therefore not able to detect situations where students stagnate or regress and will not respond adequately. Also, the tutor is not able to engage in multi-turn strategies or to lead students through a directed line of reasoning, as human tutors often do. But which of these limitations is most worth addressing? Which will have the greatest impact on learning? In the next section, we illustrate two multi-turn tutorial strategies for the current domain and discuss how we plan to explore their utility. In the current section, we illustrate a form of stagnation and discuss how the tutor can be made to respond in a more helpful manner.

In contrast to the previous examples, the current dialogue example (see Table 4) involves a student of the same age as students in the target population (10^{th} graders), although the student was definitely better than average. Further, the responses shown are the actual system responses.

We skip the first part of the dialogue and omit two steps from the dialogue that contained spelling errors—at the time, the tutor did not have spelling correction, but currently it does. In step 9, the student is very close to the correct explanation ("two angles in an isosceles triangle congruent"). The explanation is missing only the term "base angles". The tutor's response in step 10, "Are all *three* angles in an isosceles triangle congruent?" was designed to hint at that fact but is not quite appropriate. The message writer had not anticipated that the student might use the word "two" in his explanation. This can be fixed simply by crafting a better message.

Next, the student adds the word "congruant" [sic] to the explanation (step 11). This is not an improvement over the previous explanation attempt (step 9). One might say it is worse, because the purported explanation is now a tautology. This may well be a sign that the student does not fully understand what he typed. The tutor however is oblivious to the problem and simply repeats the feedback message that it gave before, in spite of the fact that this message did not help (step 13). This is unsatisfactory. A likely cause of the problem is that the student does not know the concept of base angles or at least does not think of using the term in this context[1]. The tutor should realize that and provide more helpful feedback. For example, the second time around, the tutor should have said: "WHICH angles in an isosceles triangle are congruent? What is the right term to use here?" If that message again does not help, then the tutor should cut to the chase and simply tell the student to use the term "base angles" and explain what the term means.

This problem needs to be addressed. In the current example, the students quickly gets back on track, but this will not happen as easily with all students. The tutor needs to be able to help students over the hump if they get stuck. In order to be able to detect this kind of stagnation, the tutor needs to keep a history list of the categories under which the student's explanation attempts were classified. Further, it needs to have multiple, increasingly specific messages associated with each explanation category. As a further way to help avoid stagnation, the system needs to provide helpful hints when the student clicks the "Help" button. These messages need to be sensitive to the current state of the student's explanation and need to be coordinated with the feedback messages, so that help and feedback can be interleaved in any order. Quite possibly, this can be achieved simply by using the same sequences of messages for help and feedback. We will have to make sure that students do not abuse the feedback messages and help messages to let the tutor provide the explanation without making a real effort to come up with an explanation themselves. (We have observed similar minimum effort tactics in the patterns of students hint use, see Aleven & Koedinger, 2000a.) One way to discourage such tactics is not to make the tutor too forthcoming with information, even if this may conflict with other maxims of good tutoring.

TUTORIAL STRATEGIES TO HELP STUDENTS EXPLAIN

A category of pedagogical content knowledge that is currently not addressed in the tutor is strategies that can be used to lead students to explanations, if they are struggling to produce an explanation. We look at two examples of such strategies.

Using a "procedural replay" as a bridge

We have observed that students often provide a "procedural replay" when asked to provide an explanation of a geometry problem-solving step. That is, they merely restate the arithmetic that was used to come up with a numeric answer. Such an "explanation" of course does little to justify why that was the correct thing to do. However, the prevalence of procedural replays suggest that one way to help students who are struggling to provide an explanation is (a) to ask how they came

---

[1] One might argue that the student's response in step 11 indicates that he does not know the meaning of the term "congruent", an important term in geometry. However, the use of the term congruent might reflect a superficial repair strategy, a shot in the dark. The student does not know how to improve the explanation in step 9. The tutor has just introduced the term "congruent" (step 10), so why not use that term?

up with their numeric answer (thus inviting them to provide a procedural replay) and then (b) to ask how they knew that this procedure was the right thing to do. This strategy thus grounds the explanation task in an activity that the student has just completed successfully. This strategy is not unlike the "concrete instance" strategy in Ms. Lindquist (Heffernan & Koedinger, 2000). The strategy is illustrated in the following hypothetical dialogue about the triangle sum theorem.

Tutor: How did you find your answer?
Student: subtracted 76 and 61 from 180.
Tutor: So you subtracted two other angles from 180. Which angles?
Student: XZY and YXZ
Tutor: OK, And why did you use the 180? How did you know that this was the right number to use?
Student: Because of the triangle.
Tutor: OK. So there is a triangle and there is the 180. Can you state a general rule about triangles?
Student: The triangle totals 180.
Tutor: …

One of us tried out this tutorial strategy on a few occasions while assisting students during a classroom study with a tutor version that did not provide feedback (Aleven & Koedinger, 2000b). The strategy seemed natural and effective at least for some of the geometry rules. We note that this strategy goes well beyond what the Geometry Explanation Tutor currently can do. First, the dialogue shown above does not conform to the assumption made in the current system that each utterance by the student is an attempt at providing a complete explanation. Further, the dialogue shown above seems to require some form of dialogue planning. Before we add such facilities to the system, we need to know more about how effective this tutorial strategy (and others that human tutors may use) will be in improving students' learning. A Wizard of Oz kind study, in which a human tutor communicates with the student through the computer interface (simulating the smarts of a computer tutor) would be appropriate to explore that question further.

Counterexamples

A second tutorial strategy that is likely to be useful involves the use of counterexamples when students state overgeneralizations. Consider the angle addition theorem, which says: "The measure of an angle formed by two adjacent angles is equal to the sum of the measures of those two angles." This rule appears to be one of the more difficult rules to explain. Let us assume that the student needs to explain the angle addition rule and starts out with the kind of intuitive, informal explanation that students tend to provide with some frequency (shown below). How to get the student to state this geometry rule more formally? The judicious use of counterexamples may be a lot better than just asking "can you state your explanation in more mathematical terms?"

Student: the big angle is equal to the sum of the two small ones.
Tutor: OK - true in this diagram. Now let us look for a general rule. Did you mean that whenever you have a big angle and two small ones, the big one is the sum of the two small ones?
Student: Yeah, like it is in the problem here.
Tutor: Let me change this problem diagram a little bit [in a drawing of the diagram, moves the two adjacent angles apart and draws a third angle]. Here we have a big angle and two small ones. According to your theory, the big angle is the sum of the two small ones. Correct?
Student: no the two small ones need to be next to each other
Tutor: OK, they need to be adjacent. So you are saying that whenever you have two adjacent angles and a third angle, the third angle is equal to the sum of the other two?
Student: Yes.

Tutor: Well, let me change this diagram one more time. [Moves the two angles so that they are adjacent again but keeps a third angle.]

It seems very worthwhile to have the tutor show counterexamples, as this would expose students to a venerable and ubiquitous strategy of mathematical reasoning. On the other hand, it is not clear that this kind of extended dialogue should be carried out in the context of problem-solving exercises. This might distract too much from solving the geometry problem at hand. It may be better to have students engage in activities that focus explicitly on creating and stating definitions based on examples and non-examples of the term to be defined. The (very influential) curricular guidelines of the National Council of Teachers of Mathematics include this kind of mathematical argumentation as an important objective (NCTM, 1989).

The current system is not capable of generating the dialogue shown above, for much the same reasons that it cannot generate the "procedural replay as bridge" dialogue. This is not to say that the current system could not present counterexamples. Certainly, its feedback messages could be modified to do just that. However, within the classify-and-react framework, it may be quite difficult to recover when the student does not understand the counterexample. Also, it may be difficult to stick to the strategy when a first counterexample gets the student to go only halfway (as illustrated in the dialogue shown above). At this point it is not quite clear how important it is to have such capabilities. This question is best explored by means of a Wizard of Oz study.

CONCLUSION

We are involved in a project to develop a tutorial dialogue system that helps students learn through self-explanation. The main purpose is to help students learn geometry problem-solving skills with greater understanding. A secondary purpose is to get students to learn to "speak mathematics", that is, to help students to learn basic math communication skills. With respect to the field of cognitive science, our goal is to test the hypothesis that self-explanation has a greater impact on learning if students explain in their own words, rather than through a structured computer interface, such as a menu.

Our development strategy is to equip the system with a sophisticated NLU component and to keep the dialogue management component simple. Thus, our efforts so far have focused on developing an NLU component that provides detailed analysis of students' explanations. A preliminary evaluation study showed that this component accurately classifies 81 % of student explanations and somewhat reasonable classifications on all but 4% of student explanations. Work on the NLU component continues in order to improve its performance.

Currently, the system's pedagogical content knowledge consists of a hierarchy of explanation categories, which represent common ways of providing complete or partially complete statements of geometry rules. The system uses this knowledge in each dialogue turn to classify the student's explanation and to select appropriate feedback messages. This approach enables the tutor to respond to the types of omissions we often see in students' explanations and produce reasonably effective dialogue. However, some extensions are needed in order to make the dialogue more robust. The tutor must be able it to detect situations where a student stagnates and is not able to improve her explanation even after receiving tutor feedback. The tutor must be able to help students over the hump in such situations. To do so, the system needs to have a dialogue history and multiple levels of feedback messages associated with each explanation category. It will also be necessary to coordinate the hint messages and the feedback messages.

At this point, it is not quite clear that the tutor needs to engage in multi-turn tutorial strategies such as "use procedural replay as bridge" and "counterexamples". To investigate the importance of such strategies, we will follow the 3rd-generation methodology exemplified by many other projects, namely, to study expert human tutors and perform Wizard of Oz studies. More importantly, we will build alternative versions of the tutor and experimentally test whether our changes lead to greater student learning.

Acknowledgements

This research is sponsored by an NSF grant to the Center for Interdisciplinary Research on Constructive Learning (CIRCLE), a joint research center located at the University of Pittsburgh and Carnegie Mellon University.

REFERENCES

Aleven V., Popescu, O. & Koedinger, K. R. (to appear). Towards tutorial dialog to support self-explanation: Adding natural language understanding to a cognitive tutor. To appear in *AI-ED 2001.*

Aleven, V. & Koedinger, K. R. (2000a). Limitations of Student Control: Do Student Know when they need help? In G. Gauthier, C. Frasson, and K. VanLehn (Eds.), *Proceedings ITS 2000* (pp. 292-303). Berlin: Springer Verlag.

Aleven, V. & Koedinger , K. R. (2000b). The Need for Tutorial Dialog to Support Self-Explanation. In C. P. Rose & R. Freedman (Eds.), *Building Dialogue Systems for Tutorial Applications, Papers of the 2000 AAAI Fall Symposium* (pp. 65-73). Menlo Park, CA: AAAI Press.

Aleven, V., K. R. Koedinger, and K. Cross. Tutoring Answer Explanation Fosters Learning with Understanding. In *Proceedings of AIED-99,* edited by S. P. Lajoie and M. Vivet, 199-206. Amsterdam: IOS Press, 1999

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

Berardi-Coletta, B., Buyer, L. S., Dominowsky, R. L., & Rellinger, E. R. (1995). Metacognition and Problem-Solving: A Process-Oriented Approach. Journal of Experimental Psychology: Learning *Memory, and Cognition,* 21 (1) 205-223.

Chi, M. T. H., Bassok, M., Lewis, M. W., Reimann, P., & Glaser, R. (1989). Self-explanations: How students study and use examples in learning to solve problems. *Cognitive Science,* 13, 145-182.

Chi, M. T. H., N. de Leeuw, M. Chiu, and C. Lavancher (1994). Eliciting self-explanations improves understanding. *Cognitive Science, 18*, 439-477.

Conati C. & VanLehn K. (2000). Toward Computer-Based Support of Meta-Cognitive Skills: a Computational Framework to Coach Self-Explanation. *International Journal of Artificial Intelligence in Education,* 11.

Gagné, R. M., & Smith, E. C. (1962). A Study of the Effects of Verbalization on Problem Solving. Journal of Experimental Psychology, 63(1), 12-18.

Heffernan, N. T. & Koedinger, K. R. (2000). Intelligent Tutoring Systems are Missing the Tutor: Building a More Strategic Dialog-Based Tutor. In C. P. Rose & R. Freedman (Eds.), *Building Dialogue Systems for Tutorial Applications, Papers of the 2000 AAAI Fall Symposium* (pp. 14-19). Menlo Park, CA: AAAI Press.

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.

MacGregor, R, (1991). The Evolving Technology of Classification-Based Knowledge Representation Systems. In J. Sowa (ed.), *Principles of Semantic Networks: Explorations in the Representation of Knowledg.* San Mateo, CA: Morgan Kaufmann.

Mitchell, T, 1997. *Machine Learning.* McGraw-Hill.

NCTM, 1989. Curriculum and Evaluation Standards for School Mathematics. National Council of Teachers of Mathematics. Reston, VA: The Council. See also http://standards-e.nctm.org/index.htm.

Paivio, A. (1986). *Mental representations: A dual coding approach.* New York: Oxford University Press.

Popescu, O. & Koedinger , K. R. (2000). Towards Understanding Geometry Explanations. In C. P. Rose & R. Freedman (Eds.), *Building Dialogue Systems for Tutorial Applications, Papers of the 2000 AAAI Fall Symposium* (pp. 80-86). Menlo Park, CA: AAAI Press.

Renkl, A., Stark, R., Gruber, H., & Mandl, H. (1998). Learning from Worked-Out Examples: the Effects of Example Variability and Elicited Self-Explanations. *Contemporary Educational Psychology,* 23, 90-108.

Rose, C. P. & Lavie, A. (1999). LCFlex: An Efficient Robust Left-Corner Parser. User's Guide, Carnegie Mellon University.