Conceptual and Meta Learning during Coached Problem Solving

Kurt VanLehn

LRDC, University of Pittsburgh, Pittsburgh PA 15260, USA
http://www.cs.pitt.edu/ vanlehn vanlehn@cs.pitt.edu

Abstract. Coached problem solving is known to be effective for teaching cognitive skills. Simple forms of coached problem solving are used in many ITS. This paper first considers how university physics can be taught via coached problem solving. It then discusses how coached problem solving can be extended to support two other forms of learning: conceptual learning and meta learning.

1 Introduction

Coached problem solving occurs when a tutor and a student solve problems together. Sometimes the student takes the lead and the tutors merely indicates agreement with each step. At other times the tutor walks the student through particularly difficult parts of the solution. However, at all times they are focussed on solving the problem.

Coached problem solving is a kind of cognitive apprenticeship (Collins et al., 1989) wherein the activity is restricted to solving problems. When the coach is taking the lead in solving the problem, the coach is "modeling" the cognitive processes to be learned. A coach who seldom interrupts has "faded" her "scaffolding."

Simple forms of coached problem solving have been used by intelligent tutoring systems. For instance, model-tracing tutors (Anderson et al., 1995) use coached problem solving. As novice Lisp students enter each symbol, the tutor gives immediate negative feedback if the symbol is incorrect. If the student's second attempt to enter a correct symbol also fails, the tutor gives a hint. If the hint does not suffice to get the student to enter the correct symbol, the tutor tells the student what to enter.

By definition, coached problem solving does not include extended discussions of basic principles or concepts of the task domain. This is not to say that such discussion are valueless. Human tutors, even in procedural tasks domains such as Lisp coding, do interrupt coached problem solving with such discussions (Anderson et al., 1985; McArthur et al., 1990; Merrill et al., 1992). [1]

[1] During my talk, data from our current study of human physics tutors will be presented showing exactly how much time is spent on coached problem solving as opposed to other, more general discussions. A brief taxonomy of such "diversions" will be presented as well.

Coached problem solving is an effective form of instruction, especially for procedural knowledge. This is evident in the ubiquity of coached problem solving among human tutors (McArthur et al., 1990; Fox, 1993; Merrill et al., 1992; Lepper et al., 1993) and its success in machine tutors (Anderson et al., 1995; Lesgold et al., 1992; Mark and Greer, 1995; Shute, 1991; Shute and Psotka, 1996).

The success of coached problem solving is consistent with a widely accepted views on cognitive skill acquisition (VanLehn, 1996). Cognitive skill acquisition consists of three phases. (a) An initial phase wherein student acquire preliminary knowledge of the skill, primarily by reading descriptions of it. They do not actually try to apply the skill to solve problems during this relatively short stage. (b) The intermediate phase consists of learning how to solve problems by applying the knowledge first encountered in phase 1. This phase often lasts a long time. Students often need considerable help during at the beginning of this stage, and refer often to text, examples, teachers, peers or other sources of information about the skill. By the end of this phase, they can solve most problems without help. (c) Occasionally, a third phase occurs wherein students who can already perform the task continue to practice it until it becomes nearly automatic. During this phase, their speed and accuracy continues to increase, but the gains decrease according to a power law. Coached problem solving is relevant to the intermediate phase only, so we will confine subsequent discussion of learning theory to that stage.

During the intermediate phase, learning consists of two basic processes (VanLehn, 1996):

- *Practicing the application of existing knowledge.* This causes the knowledge to become easier to recall and apply on subsequent occasions.
- *Acquiring new knowledge.* This often occurs when students try to solve a problem, reach an impasse, realize that they are lacking some domain knowledge, and seek that particular piece of knowledge from a resource such as a textbook, solved example, teacher or peer.

If the instructional material presented during the initial phase were perfectly complete and students' understanding and memory for it were also perfect, then it would not be necessary to acquire new knowledge during the intermediate phase. But such perfection never occurs in practice, so students must often interrupt their problem solving to repair or extend their knowledge.

Coached problem solving encourages both forms of learning.

- The constant focus on solving problems maximizes the amount of practice that students get on applying their existing knowledge.
- The coach facilitates acquisition of new knowledge in two ways. (a) When the tutor provides immediate negative feedback own student errors, students find it easier to locate the incorrect or missing knowledge that caused the error. (b) When students have found an misconception in their thinking, the tutor is immediately available as a resource to aid the students' "debugging" of their knowledge.

Because coached problem solving encourages both forms of learning, one would expect it to accelerate learning, and that is indeed what one finds in the data.

In short, coached problem solving is a well understood, successful method of instruction. In the spirit of building on top of a success, rather than starting anew, this paper considers how coached problem solving can be extended to help students engage in two forms of learning that seem foreign to conventional coached problem solving, namely

- *conceptual learning*, which means acquiring the basic concepts of a task domain, using them precisely and recognizing their applicability in situations far removed from those studied in the classroom, and
- *meta learning*, which means acquiring new strategies for studying (learning) that are more effective than the ones that the student habitually uses.

This paper discusses our plans for encouraging these two kinds of learning in the context of a particular coached problem solving system, Andes, that is being built at LRDC and the US Naval Academy. The next section describes coached problem solving in Andes, and the sections following it describe extensions that will encourage meta learning and conceptual learning. Unfortunately, the implementation of Andes began just recently. The complete system is scheduled for testing in fall of 1997, although pieces of the system will be pilot tested as early as summer 1996. Thus, the following discussion is based on literature review and design. We hope it will be provocative.

2 Coached problem solving in Andes

This section first introduces the subject matter taught by Andes and the kinds of activities that students engage in when using Andes. It then begins to compare Andes to existing instructional systems for physics, indicating how its increased sophistication is needed in order to overcome their limitations.

2.1 An introduction to the task domain

The Andes system helps students learn classical Newtonian mechanics, a major part of university-level physics. This subject is arguably one of the most fundamental of all sciences. It is a prerequisite of virtually all advanced science and engineering courses in college.

A traditional physics problem asks the student to derive values for quantities. For instance, the problem in Figure 1, asks the student to find two quantities, an acceleration and a tension. At minimum, the solver must write down a system of equations and solve them. However, before writing the equations, a good solver will also produce force diagrams, such as those shown in Figure 1B, and will mentally plan the solution of the problem. For instance, the solver might say,

> Okay, let's see what is going on here. The objects can either move clockwise, counterclockwise or not at all. I guess that the heavier object

Problem:

Consider two unequal masses connected by a string that runs over a massless and frictionless pulley, as shown in figure A. Let m2 be greater than m1. Find the tension in the string and the acceleration of the masses.

Solution:

If the acceleration of m1 is a, the acceleration of m2 must be −a. The forces acting on m1 and m2 are shown in figure B.

The equation of motion for m1 is
$$T - m1*g = m1*a.$$

The equation of motion for m2 is
$$T - m2*g = - m2*a.$$

Combining, we obtain

$$a = \frac{m2 - m1}{m2 + m1} g$$

and

$$T = \frac{2*m1*m2}{m1 + m2} g$$

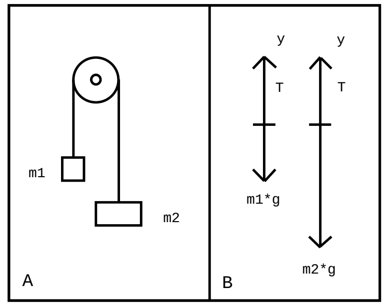

Fig. 1. A traditional physics example

> will be the one going down, and that's m2. They want the acceleration of the two blocks, which suggests applying Newton's law to them. They move in opposite directions, so it won't work to put them together in a single system. I'll have to use two systems, one for each block, and apply Newton's law twice. So let's draw the forces for each system. [Draws the 2 force diagrams shown in pane B of Figure 1.]

This type of problem solving involves no algebraic manipulation. Therefore it is called *qualitative analysis* in the physics education literature. It consists of 3 stages:

1. Determining how the objects will move over time. In the illustration just given, the student decided that one block would fall as the other block rose.
2. Planning a solution to the problem by choosing systems[2] and deciding which major principles, such as Newton's law or Conservation of Energy, to apply to each system. In the illustration, the student chose two systems, one for block 1 and one for block 2, and decided to apply Newton's law to each.

---

[2] A system consists of (1) a set of objects that are lumped together and treated as a single body, (2) a period of time, and (3) a reference frame, which is used for measuring velocity and other quantities. In the illustration, the student mentioned the bodies of the two systems. Each body consisted of a single object, a block. The student did not mention a time period or reference frame, as those details can be added later if necessary.

3. Delineating the relevant mechanical quantities, such as forces, accelerations, energies, momenta, etc.

For complex problems, holding all this information in working memory may be difficult, so solvers may write some of it down. In the illustration above, for instance, the subject drew a force diagram.

Students often prefer to solve problems without first doing a qualitative analysis (Van Heuvelen, 1991a). They often do not even draw force diagrams (only 8% did on one of van Heuvelen's (1991a) tests). Students instead plunge right into writing and solving algebraic equations. The students' algebraic focus of attention often leads them to overlook important qualitative features of the problem (e.g., certain forces) and to make unwarranted but algebraically convenient assumptions.

Qualitative analysis can not only be used for planning the solution to problems, it can also be used alone to solve problems such as the ones shown in Figure 2. Such problems require no algebraic manipulation for their solution, so they are called *qualitative* problems.

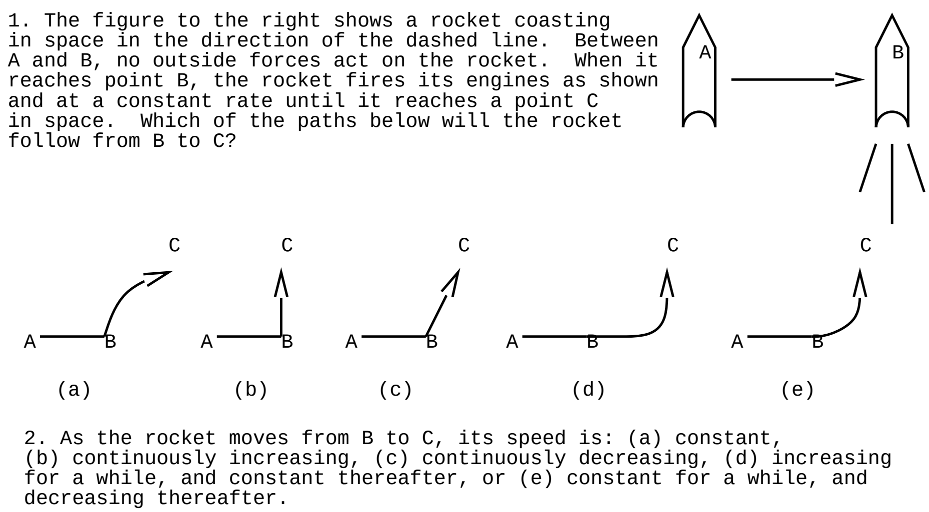

Fig. 2. Some qualitative physics problems

Physics educators have adopted qualitative analysis as a primary instructional objective. There are several reasons for this:

– Physicists reason qualitatively not only on the simple problems used in physics courses (Chi et al., 1981; Larkin, 1983) but also when discussing physics with colleagues or presenting their results at conferences (Van Heuvelen, 1991a).

- After students leave their physics courses, those who do not become physicists will probably never solve another quantitative physics problem. Indeed, they probably couldn't even if they had to, because they will have forgotten many of the algebraic details of principles they had once mastered. However, they will have to do qualitative analyses, or at least partial analyses. They should be able to delineate forces, energies and momenta, and relate them to particle trajectories. Qualitative analysis is needed in many more scientific, technical and practical applications than quantitative analysis.
- Many students have misconceptions prior to physics courses, and these misconceptions survive intact when the instruction allows them to solve problems via algebraic reasoning (Halloun and Hestenes, 1985; Pfundt and Duit, 1991). When the instruction stresses qualitative analysis, fewer misconceptions escape remediation (Van Heuvelen, 1991b; Heller and Reif, 1984).

These are just some of the reasons that qualitative analysis has been adopted as an explicit instructional objective.

2.2 Student activities while using Andes

Andes must have a user interface that allows students to conveniently express their qualitative analyses. This interface serves two purposes. It reifies the abstract notion of a qualitative analysis, and that in itself can cause significant learning (Collins and Brown, 1990; Koedinger and Anderson, 1993; Singley, 1990; Merrill et al., 1992). Second, it allows students to show their qualitative analyses to the tutor so it can check them and offer advice. We are not yet sure what the user interface will look like, but one idea is to use textual forms. For instance, the following form might be generated by Andes for the qualitative analysis of the balloon-pulley problem of figure 1 (the italicized text is entered by the student):

- What happens?
  - Balloon A moves *straight up*
  - Mass A moves *straight up*
  - Balloon B moves *straight down*
  - Mass B moves *straight down*
  - Moving at the same speed are: *balloon A, mass A, balloon B, mass B*
- What is your solution plan?
  - Apply *Newton's law* to the system consisting of *mass A*.
  - Apply *Newton's law* to the system consisting of *mass B*.

Part of qualitative analysis is to delineate forces and other quantities. Forces are conventionally depicted on a diagram such as the ones in panes A and B of figure 1. Andes will have students display forces with such diagrams. Similar diagrams or textual forms will be invented for displaying other mechanical quantities, such as energies, momenta and motion vectors.

The students' activity when using Andes will consist of simply solving problems or studying problems that have already been solved (examples). When solving problems, students typically will enter a qualitative analysis using a form such as the one above and one or more vector diagrams, and then enter equations. They (or Andes) would then solve the equations algebraically. Some problems, such as the ones in figure 2, would require only a qualitative analysis. Other activities might provide a qualitative analysis but ask students to generate equations. Such activities serve partially as examples and partially as problems. Many such mixtures are possible.

2.3 Immediate feedback

The Andes system, as presented so far, does nothing that could not be done with pencil and paper. Physics educators have often used form-based scaffolding of qualitative problems solving, sometimes with remarkably good effects (Van Heuvelen, 1991b). Andes will go further in several ways, the first of which is providing a capability for immediate feedback.

Andes will know every possible correct solution path to the problems it presents. If a student's entry is not along any of the correct solution paths, the student can be given immediate negative feedback.

Sometimes the feedback will consists of simply highlighting the entry and saying that it is incorrect. However, Andes cannot use such "flag" tutoring with all errors, because many entries can be corrected by guessing (e.g., if the entry $sin(30)$ is flagged as wrong, some student immediately try $cos(30)$). When guessing seems likely to lead to a correct entry, the feedback will take the form of a question about the line of reasoning leading up to the correct entry (e.g., "Is the side adjacent or opposite?"). Although other physics tutoring systems have provided immediate feedback (Kane and Sherwood, 1980), they have not tried to prevent guess-based corrections to errors.

2.4 Fading

Fading is another feature of Andes that differentiates it from conventional physics instruction. Fading, which is the last phase in a cognitive apprenticeship (Collins et al., 1989), is the gradual removal of explicit supports for a cognitive process. Fading must be done in such a way that students continue to perform the target cognitive processes.

In the case of Andes, there are two forms of scaffolding: the user interface for qualitative analysis and the immediate feedback. Both should be faded. In particular, Andes will initially *require* that students enter a qualitative analysis before producing the equations, and it will gradually stops requiring such an explicit entry.

If a tutor does not successfully fade its scaffolding, then students may simply revert to their old bad habits as soon as they exit the tutor. This seems to have occurred with a physics tutoring system implemented by Bruce Sherwood and his colleagues as part of the Plato system (Kane and Sherwood, 1980). It was similar to Andes except that the solutions to problems were authored by humans instead of a model. In particular, it explicitly taught qualitative analysis and insisted that students do it before writing equations. When the final exam scores were compared, Plato students scored significantly higher than students taught conventionally, but the size of the effect was disappointing (approximate 35% of sigma) (Jones et al., 1983). Although the lackluster result is open to many explanations, the one that Sherwood endorsed during a recent talk (September, 1996) is that students simply did not use qualitative analysis on the final exam. Some students thought they didn't need to use it because they could correctly solve the problems without it and it would just slow them down. Other students thought that qualitative analysis did not apply to exam problems! Sherwood's interpretation corresponds with the experience of physics instructors who have explicitly and forcefully taught qualitative analysis in their classes and recitations, only to find students using purely algebraic problem solving on the homework and exams (Van Heuvelen, 1991a). In short, it is one thing to teach students how to do qualitative analysis, and it is quite another thing to get them to actually do it routinely. It is crucial that Andes fade its support for qualitative analysis in such a way that students continue to perform it either on paper or mentally.

Immediate feedback should also be faded, but for a different reason. Fading prevents students from becoming too dependent on feedback (Collins et al., 1989).

2.5 Procedural help

The features of Andes that have that has been described so far could be implemented with minor extensions to Olae, a student modeling and assessment system for physics (Martin and VanLehn, 1995b; Martin and VanLehn, 1995a; Martin and VanLehn, 1993). Olae could determine whether a students entry was correct and it could assess the student's probability of mastery for each physics rule. With the addition of a few heuristics, it could use rule-mastery to control fading of the scaffolding and immediate feedback. Although such a system would probably be better than its predecessors, it would occasionally be extremely frustrating to students because it would have no capability of providing helpful hints when they get stuck. As it turns out, providing such help requires a significant extensions.

In order to find out what kind of help students would need, we conducted a pilot study with a human tutor who played the role of Andes. The tutor sat in a different room from the student, but could see a copy of the student's screen. When the student wanted help, the student would type "help," and the tutor would pick up the phone and talk to the student. Notes were taken during the sessions with 9 students, and 73 help requests were observed.

Of the 73 help requests, 18 were simply requests for immediate feedback. The students often just asked, "Is that ok?" referring to their most recent entry. This kind of help could be provided by an Olae-based tutor.

The most common help request (27 occurrences) was the students' indication, in various ways, that they were lost and needed a hint about what to do next. Providing such a hint requires knowing which solution path the student has been pursuing so far and what would be the next reasonable step along it. This is a form a plan recognition. The Olae framework was extended to do this form of plan recognition and a new student modelling engine, Pola, was implemented (Conati and VanLehn, 1996; Conati and VanLehn, 1995). Pola will become the basis for the student modelling module of Andes. Once it has selected an appropriate next step for the student, Andes will construct a series of increasingly specific hints leading up to that step.

The next most common help request (14 occurrences) consisted of asking for help in achieving a specific goal. Some examples, paraphrased to make them easier to understand, are:

- How do I find the acceleration of block2?
- How can I find out which direction that force acts?
- How do I convert kilograms to Newtons?
- How to I eliminate mass from that equation?
- How do I solve that equation for acceleration?

Before giving help in achieving the goal, Andes should check that the goal is appropriate. Sometimes the proposed goal is not on a solution path. Sometimes the goal reveals a misunderstanding (e.g., one cannot convert kilograms to Newtons because kilograms measures mass and Newtons measure force). If the goal is appropriate, then a succession of increasingly more specific hints can be devised based on the goal tree for achieving the stated goal, as is done in many tutors (e.g., (Anderson et al., 1995)).

The remaining help requests (14) all revealed a conceptual misunderstanding of some kind. They will be discussed in the next section.

In most of the help requests ($27 + 14 = 41$ of 73), the students wanted to know what to do next. Andes will tell them, albeit obliquely via a series of hints. The hints will be based on Andes' data base of correct solution paths. Thus, this type of help is called "procedural."

The capabilities discussed so far make Andes similar to a model-tracing tutor (Anderson et al., 1995). Like the model-tracing tutors, Andes can give immediate feedback and procedural help. Unlike the model-tracing tutors, Andes intermingles example studying and problem solving, and it fades its scaffolding. Most importantly, Andes does not confine the students to only a small set of solution paths, as most of the model-tracing tutors do. Enforcing a rigid problem solving style on students is possible in physics, but it would be completely unacceptable to many physics instructors. Andes allows students to pursue any correct solution path, even though causes extraordinary combinatorial problems during student modeling (Conati and VanLehn, 1996; Conati and VanLehn, 1995).

Nonetheless, despite the high technology required to implement feedback, procedural help and the other Andes' features discussed so far, the students will see approximately the same level of support and the same kinds of activities as those afforded by a model tracing tutor. Thus, one would expect approximately the same pedagogical improvement over classroom instruction (Shute and Psotka, 1996). In order to obtain qualitatively greater benefits, we need to encourage new kinds of learning. The next sections describe them.

3 Conceptual help

In the discussion of types of learning, it was mentioned that coached problem solving, and model tracing tutors in particular, encouraged both acquiring of new knowledge and practicing the application of existing knowledge. How can we improve on that?

One answer arises from considering exactly how procedural help supports the acquisition of new knowledge. Suppose a student is lacking a particular rule (call it the target rule) and that hinting does not cause the student to invent it. Eventually, the procedural help system will simply tell the student what to do next. From this, the some students might induce a very specific version of the target rule. Other students might realize the need to generalize their experience, and ask the tutor to explain why the action it suggested is correct. The tutor prints the line of reasoning leading up to the action. Suppose our student is perceptive enough to determine which inference in the line of reasoning is due to the unknown (target) rule. The student asks the tutor to explain that inference. Now the tutor is at a loss for words. Its only representation of knowledge is the rule itself. It cannot explain or justify any of its rules simply because it has no representation of knowledge other than the rules themselves. In this respect, the tutor's knowledge is fundamentally shallow. The rules contain just enough information to allow them to solve problems correctly; they do not need to represent their intellectual heritage, so they do not.

Better students might be stimulated by the tutor to wonder about the veracity of its rules, seek information from a teacher or text, and end up learning just would one would like them to learn. But such good students probably do not need a tutor in the first place. Less capable or motivated students will simply accept the tutor's "rote" version of the rules. Perhaps the worst outcome may occur with students who seek a deeper explanation for rules but cannot find one; they may accept the tutor's rule in order to move forward on the problem solving, but doubt its veracity. Thus, if the tutor cannot explain individual rules, students may acquire "rote" rules whose veracity they doubt.

This assumes that individual rules actually have deep justifications that are worth knowing. This is not always the case. Some rules are just arbitrary. For instance, the Lisp tutor has many rules of the form "if you want to code a conditional statement, set the goal of entering COND." This rule reflects an arbitrary choice by John McCarthy, the inventor of Lisp, to start conditional statements with the word "COND." This rule cannot be proved via scientific experiments or mathematical logic. In order to learn Lisp, students must accept it "without proof" because it is not the sort of rule that needs proof.

Other rules are true by assumptions. For instance, an equation solving tutor has rules expressing the distributive law, the associative law and other fundamental postulates of arithmetic. Although a student could debate the truth of these rules, the debate would swiftly lead into the deepest parts of the foundations of mathematics where there are no easy answers to satisfy the curious student. For most students, it is better simply to treat the laws as extraordinarily useful assumptions and not to try to justify them further.

Most the model tracing tutors to date have been in mathematics or programming, where most of the rules are either arbitrary or true by assumptions. In physics, many rules are open to debate. For instance, many students find it difficult to believe that when an object rests on a surface, the surface exerts a force on it (called the normal force). Physics problem solvers typically represent this assertion as a single rule. The problem solver does not "know" why the normal force rule is true. If asked, it could not justify the rule to the student. Yet such a justification is exactly what some students need to hear in order to deeply understand the normal forces.

There is a converse problem as well. Students often have specific beliefs that are not true. Many of these have been observed. The Pfundt and Duit (1991) bibliography of the misconceptions literature is in its fourth edition and has thousands of entries. As Ploetzner and VanLehn (in press) demonstrated, many mechanics misconceptions can be easily represented as rules. For instance, one rule is, "If an object is moving in a certain direction, then there is a force on the object in that direction." As many investigators have found, simply telling the students that there is no such force will not convince them to abandon belief in it. Even producing experimental evidence against the belief is seldom convincing (Smith et al., 1993). Nonetheless, the tutor must do something to undermine belief in incorrect rules.

In short, we need to extend the model tracing paradigm in two ways. The tutor needs to be able to justify correct rules. It also must explicitly represent the most popular incorrect beliefs as rules and provide arguments against them.

Physics educators have tried many techniques to alter students' beliefs. There is a vast literature full of good ideas that we have only begun to tap. Our plan is to construct a library of "minilessons," where a minilesson is a short multimedia lesson on a particular concept.

Similar themes seem to occur in multiple minilessons. For instance, one theme is viewing the situation at the atomic level. Clement and his colleagues showed that students could be convinced to believe in the normal force rule showing the students that even the most rigid surface is springy at the atomic level (Murray et al., 1990). Placing an object on the surface causes it to indent ever so slightly. Once the students believe this (and they are generally more willing to accept statements about atomic phenomena than observable phenomena), then they can be easily convinced that normal forces exist.

Another common theme involves debugging confusions about the notion of differentiation with respect to time. Differentiation underlies several pairs of concepts that are often confused by students: Acceleration is the differential of velocity. Force is the differential of momentum. Work is the differential of energy. Students often confuse quantities related by a differential. For instance, many students believe that any object in motion has a force acting on it in the direction of motion, and that the faster the object moves, the stronger the force. Although there is no such force, any object in motion does have a momentum that is in the direction of motion and proportional to the object's speed. Thus, confusing force and momentum is one source for this misconception. (There are others.) Similarly, students often confuse acceleration and velocity (Reif, 1987). They think that the acceleration at the apex of a vertically thrown object is zero, when in fact it is the velocity that is zero. The minilessons that teach students to distinguish acceleration from velocity should appeal to the same differentiation sub-text as the minilesson that teaches students to distinguish force from momentum. We hope that this will encourage learning of the underlying abstraction. Similar comments apply to the underlying abstraction of constraint-based interactions (Chi, 1992; Chi et al., 1994b; Slotta et al., 1995), which seems to be involved in many difficult-to-learn concepts.

Minilessons will be initiated under a variety of circumstances. When the procedural help system's hints have finally gotten the student to make a correct entry, and the student model indicates that it is likely that the student did not know one of the rules involved in generating that action, then the conceptual help system will try to determine how ell the student learned the target rule. (Natural language interaction would be ideal for this interrogation, but well beyond the state of the art.) If the help system or the student feel that the rule is not adequately understood, either may initiate the minilesson.

Minilessons are also initiated during "terminology help," where terminology help is a simple hypertext technique based on Moore and Swartout (1990). Whenever the system prints anything substantive—including the statement of the physics problem, feedback messages, procedural help, etc.—most of the technical phrases in the text are underlined. Clicking on a underlined phrase causes a short explanation of its meaning in this context to be printed. This explanation will also have some underlined phrases. The need for terminology help became apparent in analyzing the 14 non-procedural help requests that we observed in the pilot study. Of these, 5 were requests for clarification of phrases in the problem statement. For instance, on one problem, which involved a parachute slowing the descent of a woman, mentioned "the retarding force on the parachutist." Students often did not realize that the parachute was pulling upward on the woman and thus was exerting a force on her, so they did not understand what the phrase "retarding force" was referring to. In 6 cases, students asked general questions about physics or algebra, such as "what is a system of equations?" Thus, in 11 of the 14 cases, some kind of terminology help would have been appropriate.

Some of the terminology help messages will have a minilesson button attached to them. Clicking on the button initiates a minilesson Thus, minilessons can be entered from both the procedural and terminology help systems.

3.1 Meta help

So far, we have assumed that the only goal of the tutor is to teach the students the task domain, which is physic in this case. However, if we expand the instructional goals, new kinds of learning will need to occur. This section proposes a new instructional goal for Andes, and discusses how it can be achieved.

Some student employ better strategies for studying than others. For instance, Chi et al. (1989) found that some student studied examples by explaining each line of the solution to themselves, whereas other students merely read the examples thorough in a cursory manner knowing that they could come back to them later if necessary. The students who self-explained the examples learned more not only during the example studying phase but also during a subsequent problem solving phase as well (VanLehn et al., 1992; VanLehn, 1995b). Thus, there is a correlation between a studying strategy, self-explanation, and the students' learning. This same correlation has been found in many other situations (Pirolli and Bielaczyc, 1989; Ferguson-Hessler and de Jong, 1990; Pressley et al., 1992; Chi et al., 1994a; Recker and Pirolli, 1995; Lovett, 1992). Protocol analysis and cognitive modeling have resulted in a thorough understanding of what the self-explanation strategy is and why it increases learning (VanLehn et al., 1992; VanLehn and Jones, 1995; Pirolli and Recker, 1994; Recker and Pirolli, 1995).

Recently, another naturally occurring studying strategy was discovered (VanLehn, 1995a). When solving problems, some students notice that the problem is similar to an already solved problem or example, but they do not immediately refer to it. Instead, they try to solve their problem without help, and refer to the example only when they reach an impasse and need help. This studying strategy is called Min analogy. Students using Min analogy learned more during problem solving than students who used a strategy called Max analogy, wherein they refer to the example as soon as they notice its relevance and continually refer to it as they work on the problem.

Other effective studying strategies have been observed, albeit less formally. For instance, Collins and Brown (1990) suggested that reflecting on ones' solution to a problem can speed learning. VanLehn (1991,1995b) suggested that adopting a wary, reflective stance while solving problems helps students notice opportunities for improvement in their methods.

There is evidence that studying strategies can be taught. Bielaczyc, Recker and Brown (1994) taught students how to self-explain in the classic cognitive apprenticeship manner. That is, they first demonstrated self-explanation (using a videotape of an actor), then helped the student practice self-explanation as they studied Lisp, then finally faded their help leaving the student to self-explain without prompting or help. The students who were taught to self-explain learned more than students who received an equivalent amount of time on the Lisp-learning task. Chi et al. (Chi et al., 1994a) showed that merely prompting a student to self-explain after each line caused them to learn the skill, and that in turn caused them to learn more of the target material (cardiophysiology). Shauble, Raghavan and Glaser (1993) taught students to reflect on their problem solving with a special tutoring system, and this caused increases in learning (Raghavan, personal communication, May 1995).

Andes will teach effective studying strategies not only because they should help the students learning physics, but also because they should help the students in their other courses as well.

For instance, in order to teach self-explanation, Andes will have students study examples using the "poor man's eyetracker," a software module developed for Olae. The example appears on the screen, but all its text, equations and diagrams are hidden by boxes that are the same shape as the hidden material. When the student clicks on a box, the box is removed and the material is exposed for the student to read. When the student clicks on another box, the first box is restored and the new box is removed. In this fashion, the student can read the example in any way desired but the machine can determine which lines the student has read, for how long the student spent contemplating the line, and whether the student went back to earlier lines while studying this line. On the basis of these data, the tutor decides whether the student has self-explained the example. If the student has not, then the tutor will not let them leave the example, but instead says, "You went through that example rather fast. There are some interesting things to learn from it. Perhaps you should seek them out." If the student goes back and studies the lines where, according to the tutor's analysis, rules unfamiliar to the student have been applied, then the tutor will decide that the student is self-explaining the example and learning from it. On the other hand, if the student again glosses the example, the tutor will make an even stronger suggestion, and perhaps even point out the line(s) that require attention. We are not yet sure of the details of the interaction (we need pilot subjects to find this out), but the intent is to use the eye-tracker to monitor the student's studying strategy and use blunt suggestions to modify it if necessary.

Because the sort of studying strategies that Andes will teach are sometimes called meta strategies (in order to distinguish them from problem solving strategies, such as means-ends analysis), we call this kind of tutoring *meta help*.

4 Conclusions

This paper has presented a plan for a tutoring system that goes several steps beyond the state of the art in model-tracing tutoring systems. The following are some of the less exciting but still important advances:

1. Andes will support a more flexible form of problem solving that allows students to pursue any correct solution path, rather than restricting them to ones that the tutor is prepared to recognize. This causes combinatorial difficulties in the student modelling module, which we have begun to conquer (Conati and VanLehn, 1996).
2. Andes will intermingle example studying and problem solving, which is intended to more closely approximate the kinds of mixed initiative problem solving that occurs during human-tutor coached problem solving.
3. Andes will fade its scaffolding by gradually removing both feedback and the requirement that students fill out a qualitative analysis form before doing algebraic manipulations.
4. Andes will provide terminological help by using hypertext for all substantive messages from the tutor; clicking on a confusing phrase will cause an explanation of it to be printed.
5. Andes will replace its customary minimal feedback with more focussed feedback whenever it is likely that receiving minimal feedback would allow students to guess the correct entry rather than figuring it out.

It is fairly clear how to implement most of these features. We expect them to provide minor enhancements either to students' learning or to the acceptance of the tutor by physics instructors.

However, our hope for major enhancements lies in two rather novel and risky forms of help: conceptual help and meta help. Conceptual help is based on the observation that in many task domains, problem solving rules only represent the minimal knowledge required to solve problems. In such domains, many rules have behind them a rich explanation and justification that students need to know. Such knowledge is often called "conceptual" knowledge. Andes' conceptual help system will consist of a library of minilessons and a variety of heuristics for invoking them. The most difficult part of the conceptual help system is coming up with minilessons that actually work. We hope to find several explanatory themes or techniques that can be used to design multiple minilessons.

Meta help is designed to get students to change their studying techniques. This is not simply another procedural skill that needs to be taught. Many students have a motivational set that emphasizes moving as quickly as possible thorough the material while getting as many problems correct as possible. In Carol Dweck's seminal work (Dweck, 1986), such students are called "performance oriented." Other students, whom Dweck calls "learning oriented," care less about getting the right answers than in learning the material. The overall goal of the meta help system is to not only teach students effective studying techniques but to change their orientation from performance to learning.

5 Acknowledgements

This research is supported by ONR's Cognitive Science Division under grant N00014-96-1-0260, and by DARPA's Computer Aided Education and Training Initiative under grant N66001-95-C-8367. The LRDC Andes team currently comprises Patricia Albacete, Cristina Conati, Marek Druzdzel, Abigail Gertner, Zhendong Niu and Anders Weinstein. The US Naval Academy Andes team comprises Dave Correll, Kay Schulze, Robert Shelby and Mary Wintersgill. This paper could not have been written without their crucial input during the long process of pulling the Andes design together.

References

Anderson, J., Farrell, R., and Saurers, R. (1985). Learning to program in Lisp. *Cognitive Science*, 8:87–129.

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

Bielaczyc, K., Pirolli, P., and Brown, A. L. (1995). Training in self-explanation and self-regulation strategies: Investigating the effects of knowledge acquisition activities on problem-solving. *Cognition and Instruction*, 13(2):221–252.

Chi, M., Bassok, M., Lewis, M., Reimann, P., and Glaser, R. (1989). Self-explanations: How students study and use examples in learning to solve problems. *Cognitive Science*, 15:145–182.

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

Chi, M., Feltovich, P., and Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5:121–152.

Chi, M. T. H. (1992). Conceptual change within and across ontological categories. In Giere, R., editor, *Cognitive models of science: Minnesota studies in the philosophy of science*. University of Minnesota Press, Minneapolis, MN.

Chi, M. T. H., Slotta, J. D., and de Leeuw, N. (1994b). From things to processes: A theory of conceptual change for learning science concepts. *Learning and Instruction*, 4:27–43.

Collins, A. and Brown, J. S. (1990). The computer as a tool for learning through reflection. In Mandl, H. and Lesgold, A., editors, *Learning issues for intelligent tutoring systems*. Springer, New York.

Collins, A., Brown, J. S., and Newman, S. E. (1989). Cognitive apprenticeship: Teaching the craft of reading, writing and mathematics. In Resnick, L. B., editor, *Knowing, learning and instruction: Essays in honor of Robert Glaser*, pages 543–494. Lawrence Erlbaum Associates, Hillsdale, NJ.

Conati, C. and VanLehn, K. (1995). A student modeling technique for problem solving in domains with large solution spaces. In Greer, J., editor, *Proceedings of the 1995 Artificial Intelligence and Education Conference*. Association for the Advancement of Computers in Education, Charlotsville, NC. Abstract only.

Conati, C. and VanLehn, K. (1996). Pola: A student modeling framework for probabilistic on-line assessment of problem solving performance. In Chin, D. N., Crosby, M., Carberry, S., and Zukerman, I., editors, *Proceedings of UM-96, the Fifth International Conference on User Modeling*, pages 75–82. User Modeling, Inc., Kailua-Kona, Hawaii.

Dweck, C. S. (1986). Motivational processes affecting learning. *American Psychologist*, 41:1040–1048.

Ferguson-Hessler, M. and de Jong, T. (1990). Studying physics texts: Differences in study processes between good and poor solvers. *Cognition and Instruction*, 7:41–54.

Fox, B. A. (1993). *The Human Tutorial Dialogue Project: Issues in the Design of Instructional Systems*. Lawrence Erlbaum Associates, Hillsdale, NJ.

Halloun, I. A. and Hestenes, D. (1985). Common sense concepts about motion. *American Journal of Physics*, 53(11):1056–1065.

Heller, J. I. and Reif, F. (1984). Prescribing effective human problem-solving processes: Problem descriptions in physics. *Cognition and Instruction*, 1(2):177–216.

Jones, L. M., Kane, D., Sherwood, B. A., and Avner, R. A. (1983). A final-exam comparison involving computer-based instruction. *American Journal of Physics*, 51(6):533–8.

Kane, D. and Sherwood, B. (1980). A computer-based course in classical mechanics. *Computers and Education*, 4:15–36.

Koedinger, K. and Anderson, J. R. (1993). Reifying implicit planning in geometry: Guidelines for model-based intelligent tutoring system design. In Lajoie, S. P. and Derry, S. J., editors, *Computers as cognitive tools*. Lawrence Erlbaum Associates, Hillsdale, NJ.

Larkin, J. (1983). The role of problem representation in physics. In Gentner, D. and Stevens, A., editors, *Mental Models*. Lawrence Erlbaum Associates, Hillsdale, NJ.

Lepper, M. R., Woolverton, M., Mumme, D. L., and Gurtner, J. (1993). Motivational techniques of expert human tutors: Lessons for the design of computer-based tutors. In Lajoie, S. P. and Derry, S. J., editors, *Computers as Cognitive Tools*, pages 75–105. Lawrence Erlbaum Associates, Hillsdale, NJ.

Lesgold, A., Lajoie, S., Bunzo, M., and Eggan, G. (1992). Sherlock: A coached practice environment for an electronics troubleshooting job. In Larkin, J. and Chabay, R., editors, *Computer Assisted Instruction and Intelligent Tutoring Systems: Shared Goals and Complementary Approaches*, pages 201–238. Lawrence Erlbaum Associates, Hillsdale, NJ.

Lovett, M. C. (1992). Learning by problem solving versus by examples: The benefits of generating and receiving information. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*. Lawrence Erlbaum Associates, Hillsdale, NJ.

Mark, M. A. and Greer, J. E. (1995). The vcr tutor: Effective instruction for device operation. *The Journal of the Learning Sciences*, 4(2):209–246.

Martin, J. and VanLehn, K. (1993). OLAE: Progress toward a multi-activity, Bayesian student modeler. In Brna, S. P., Ohlsson, S., and Pain, H., editors, *Artificial Intelligence in Education, 1993: Proceedings of AI-ED 93*, Charlottesville, VA. Association for the Advancement of Computing in Education.

Martin, J. and VanLehn, K. (1995a). A Bayesian approach to cognitive assessment. In Nichols, P., Chipman, S., and Brennan, S., editors, *Cognitively diagnostic assessment*. Lawrence Erlbaum Associates, Hillsdale, NJ.

Martin, J. and VanLehn, K. (1995b). Student assessment using Bayesian nets. *International Journal of Human-Computer Studies*, 42:575–591.

McArthur, D., Stasz, C., and Zmuidzinas, M. (1990). Tutoring techniques in algebra. *Cognition and Instruction*, 7(3):197–244.

Merrill, D., Reiser, B. J., Ranney, M., and Trafton, J. (1992). Effective tutoring techniques: A comparison of human tutors and intelligent tutoring systems. *The Journal of the Learning Sciences*, 2(3):277–306.

Moore, J. D. and Swartout, W. R. (1990). Pointing: A way toward explanation dialogue. In *Proceedings of the American Association of Artificial Intelligence Conference.* AAAI Press, Menlo Park, CA.

Murray, T., Schultz, K., Brown, D., and Clement, J. (1990). An analogy-based computer tutor for remediating physics misconceptions. *Interactive Learning Environments*, 1(2):79–101.

Pfundt, H. and Duit, R. (1991). *Bibliography: Students' alternative frameworks and science education.* Institute for Science Education, Kiel, FRG.

Pirolli, P. and Bielaczyc, K. (1989). Empirical analyses of self-explanation and transfer in learning to program. In *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society*, pages 459–457, Hillsdale, NJ. Lawrence Erlbaum Associates.

Pirolli, P. and Recker, M. (1994). Learning strategies and transfer in the domain of programming. *Cognition and Instruction*, 12(3):235–275.

Ploetzner, R. and VanLehn, K. (1996). The acquisition of informal physics knowledge during formal physics training. *Cognition and Instruction.* In press.

Pressley, M., Wood, E., Woloshyn, V., Martin, V., King, A., and Menke, D. (1992). Encouraging mindful use of prior knowledge: Attempting to construct explanatory answers facilitates learning. *Educational Psychologist*, 27:91–109.

Recker, M. M. and Pirolli, P. (1995). Modeling individual differences in students' learning strategies. *The Journal of the Learning Sciences*, 4(1):1–38.

Reif, F. (1987). Interpretation of scientific or mathematical concepts: Cognitive issues and instructional implications. *Cognitive Science*, 11(4):395–416.

Schauble, L., Raghavan, K., and Glaser, R. (1993). The discovery and reflection notation: A graphical trace for supporting self-regulation in computer-based laboratories. In Lajoie, S. P. and Derry, S. J., editors, *Computers as Cognitive Tools*, pages 319–337. Lawrence Erlbaum Associates, Hillsdale, NJ.

Shute, V. J. (1991). Who is likely to learn programming skills? *Journal of Educational Comuting Research*, 6:1–24.

Shute, V. J. and Psotka, J. (1996). Intelligent tutoring systems: Past, present and future. In Jonassen, D., editor, *Handbook of Research on Educational Communications and Technology.* Scholastic Publications. In press.

Singley, M. (1990). The reification of goal structures in a calculus tutor: Effects on problem solving performance. *Interactive Learning Environments*, 1:102–123.

Slotta, J. D., Chi, M. T. H., and Joram, E. (1995). Assessing students' misclassifications of physics concepts: An ontological basis for conceptual change. *Cognition and Instruction*, 13(3):373–400.

Smith, J. P., diSessa, A. A., and Roschelle, J. (1993). Misconceptions reconceived: A constructivist analysis of knowledge in transition. *The Journal of the Learning Sciences*, 2(2):115–164.

Van Heuvelen, A. (1991a). Learning to think like a physicist: A review of research-based instructional strategies. *American Journal of Physics*, 59(10):891–897.

Van Heuvelen, A. (1991b). Overview, case study physics. *American Journal of Physics*, 59(10):898–907.

VanLehn, K. (1991). Rule acquisition events in the discovery of problem solving strategies. *Cognitive Science*, 15(1):1–47.

VanLehn, K. (1995a). Looking in the book: The effects of example-exercise analogy on learning. Submitted for publication.

VanLehn, K. (1995b). Rule learning events in the acquisition of a complex skill. Submitted for publication.

VanLehn, K. (1996). Cognitive skill acquisition. In Spence, J., Darly, J., and Foss, D. J., editors, *Annual Review of Psychology, Vol. 47*, pages 513–539. Annual Reviews, Palo Alto, CA.

VanLehn, K. and Jones, R. M. (1995). Is the self-explanation effect caused by learning rules, schemas or examples? Submitted for publication.

VanLehn, K., Jones, R. M., and Chi, M. T. H. (1992). A model of the self-explanation effect. *The Journal of the Learning Sciences*, 2(1):1–59.