

STEPS: A Preliminary Model of Learning from a Tutor

Sigalit Ur

Intelligent Systems Program
University of Pittsburgh
Pittsburgh, PA 15260
(412) 624-8847

sigalit@pogo.isp.pitt.edu

Kurt VanLehn

Learning Research and Development Center
University of Pittsburgh
Pittsburgh, PA 15260
(412) 624-7458

vanlehn@cs.pitt.edu

Abstract

This paper describes a prototype of a simulated physics student that learns by interacting with a human tutor. The system solves physics problems while showing its work on a workstation screen, and the tutor can intervene at certain points during problem-solving to advise the simulated student. This prototype constitutes an initial cognitive task analysis of the skill of learning from a tutor, which prescribes several tutoring practices that appear to be plausible for both human and computer tutors.

Introduction

STEPS is a simulated, tutable physics student. That is, it is a machine learning program that learns from feedback and hints provided by a human tutor as it solves physics problems. The main motivation for STEPS is to see if it is technically feasible to build a simulation-based tutor training system, and a preliminary evaluation of its potential for training human tutors has been conducted (Ur and VanLehn, 1994). However, this paper focuses on what STEPS has taught us about the cognitive task of learning from a tutor.

Although computational theories of human skill acquisition exist (e.g. Anderson, 1990; Newell, 1991; VanLehn, Jones and Chi, 1992) they do not adequately address *interactive learning*, wherein a learner and a tutor can converse. They apply most readily to instructional situations where the learner works alone or with minimal supervision. STEPS is an extension of one such theory (CASCADE—see VanLehn et al., 1992) to interactive learning, but there are as yet so few empirical studies of human-human tutoring, especially from the tuttee's point of view (see Merrill, Reiser, Ranney and Trafton, 1992) that many of the design decisions in STEPS are unconstrained by empirical evidence. At this point, we cannot claim that STEPS is a theory of interactive learning. However, it does provide an initial cognitive task analysis of the process of learning from a tutor.

The benefit of having a computational cognitive task analysis of learning (i.e., a simulated student) is that one can readily see how to organize instruction so that it will be learned most effectively. In particular, we will indicate a number of tutoring practices that optimize STEPS's learning and appear to be plausible prescriptions for both human and computer tutors. Moreover, we understand completely *why* the prescribed policies optimize learning (or at least, STEPS's learning), which is not the case with the heuristics that are currently used for designing the pedagogical components of intelligent tutoring systems.

First the system is described, and its operation is illustrated with several short sessions. Then we present the tutoring tactics suggested by STEPS's design and discuss their suitability as prescriptions for human and computer tutoring.

The System

STEPS utilizes the simulated physics student embodied in the CASCADE system (see VanLehn, Jones and Chi, 1992 for a description of an earlier version) as its basic problem-solver, and the OLAE interface (Martin and VanLehn, 1993), on which it shows its work and accepts input from the tutor. Both systems have been modified, extended and linked to form STEPS. The next sections will describe the system's interface, problem-solver and learner.

User Interface

The STEPS screen is divided into several windows (see figure 1): the icon window, which displays icons representing physics problems (top bar), the problem window, which displays the text and diagram of the current problem (upper right), the diagram window, which will contain the free body diagram that STEPS constructs (lower right), the solution window, where STEPS can write equations (upper left), and the dialog window, where STEPS comments about its progress, describes new variables it will be using, requests help from the tutor, etc. (lower left).

The tutor chooses a problem to pose to STEPS by clicking on its icon. The problem description is then presented in the problem window, and STEPS starts to solve the problem by drawing vectors (representing forces, velocities and accelerations) and axes in the diagram window and writing equations in the solution window.

After each problem solving action is displayed on the screen, the tutor is given an opportunity to intervene. The tutor can cross out an action that STEPS has taken by pointing to its representation on the screen. The tutor may also draw an arrow in the diagram window or enter an equation in the solution window. In the latter case, the system allows only input that it can parse, enforcing variable names that are consistent with the problem-solver's representation. When the tutor has said all he wishes to say, he clicks the "Go Ahead" button, which passes control back to the problem-solver.

Problem Solving

Currently, the problem solver's knowledge covers a subset of the material in Chapters 2-5 of a college physics textbook (Halliday and Resnick, 1988), including kinematics and




Figure 1: The OLAE screen

Newtonian Mechanics. This knowledge is represented by a set of rules (e.g. If spring A touches object B then A exerts pressure force F on B) and equations (e.g. If X is a massive body, close to earth, then $W=mg$ where W is X's weight, m is X's mass and g is the gravitational constant). Using this knowledge, STEPS can solve quantitative physics problems of the kind described in Figures 2–4. These problems are available to STEPS as lists of predicates which describe the situation, given quantities and sought quantities.

The declarative rules are used by an agenda-based top-level control structure. At each iteration of the top-level loop, possible tasks are proposed. When all possible tasks have been queued, one of them is selected for execution, executed, and marked as having been tried. This process continues until the problem is solved. In the physics domain, the tasks are creating variables for given and sought quantities in the problem, choosing bodies, choosing methods (e.g., forces, energies), drawing the free body diagram, drawing the axes, and writing an equation. Once a task has been chosen for execution, the rules that implement it are enabled, and these fire until quiescence, which signals the next iteration of the top-level loop.

STEPS' progress is visible on the screen as it draws arrows on the free body diagram and writes equations. Each time that the problem solver displays another step in the solution, it pauses and asks the tutor, "Go ahead?" If the tutor answers yes, the problem solver continues from exactly where it left off. However, the tutor can also cross out equations or arrows, or add new equations or arrows before answering yes. When this occurs, STEPS (a) deletes all the goals of the problem solver, (b) handles the tutor's inputs, (c) reconstructs the problem solving goals, then (d) resumes problem solving. This same four-step process also occurs after the problem solver has gotten stuck and asked the tutor for help.

The tutor's input often changes the state of the problem solving so much that it is impossible to continue problem solving from where it left off. That is why STEPS deletes the

Table 1: STEPS's top-level algorithm

```

Get problem from tutor
Repeat until problem solved
  Generate next step (if known)
  Get tutor's advice
  If (case 1) tutor suggested step
    Try to derive tutor's step
  else if Tutor crossed out step
    if (case 2) Tutor also suggested alternative step
      Derive tutor's step
      Try to find divergence between derivations
      Delete rule at divergence
    else if (case 3) Can modify crossed-out step
      Generate modified step
  if Tutor approves it
    Try to derive it
  else (case 4)
    Mark step for later learning
If any steps marked for later learning
  Find first previously unused rule in derivation and delete it
  
```

goals before handling the tutor's input and reconstructs goals afterwards. Goal reconstruction (VanLehn and Ball, 1991) is simple in STEPS' agenda-based control structure. The system checks if new tasks have become possible given the tutor's input, then chooses which task to execute in the usual way. Sometimes a tutor may leave the most recently written equation alone and modify an arrow or equation written earlier. In this case, STEPS assumes that all the reasoning taken since the modified action is now suspect, so it mentally crosses those results out.

The tutor may take several actions before giving control back to STEPS, and how STEPS learns from this input depends on what the specific actions are. Table 1 summarizes STEPS' algorithm, and the next section describes in more detail STEPS' methods for learning new rules and deleting old ones.

Learning

When the tutor demonstrates what the next action should be (case 1 in Table 1), and the system realizes that it cannot generate this action using its current set of rules, it tries to learn new rules that would enable generation of the action in future using Explanation-Based Learning of Correctness (VanLehn et al., 1992). The system tries to derive the correct action using a set of *overly general rules* that are part of its knowledge base. If it is successful, it creates (using Explanation-Based Generalization (Mitchell et al., 1986)) new rules that are more specific than the overly general rules but general enough to apply to more than this specific problem. These new rules become part of the system's knowledge base. There is fairly good evidence that this is how human students learn new rules in this task domain (VanLehn, Jones and Chi, 1992; VanLehn and Jones, 1993).

When the tutor supplies negative feedback on an action, the system assumes the incorrect action was a result of an incorrect rule. What happens next depends on whether the tutor supplied an alternative, correct action.

If the tutor has supplied a correct alternative action (case 2 in Table 1), the system attempts to derive it using its rules. If the attempt is successful, the system traces through the derivation of the correct action and the derivation of the incorrect action, then deletes the rule that caused the derivations to diverge. This is a standard machine learning technique for handling negative feedback (Sleeman et al., 1982). However, there is as yet no evidence that human

students use it because no fine-grained (e.g., protocol) studies of learning from a tutor have been conducted. Nonetheless, it is clear that human students do learn from negative feedback (e.g., Anderson, Conrad and Corbett, 1989) and this is a relatively simple, parsimonious mechanism that does the job.

When no additional feedback is given, STEPS first considers generating its own action to replace the crossed-out one. If it decides it should (case 3 in Table 1) it syntactically modifies its old action and then asks the tutor for confirmation. If the tutor confirms that STEPS' new action is correct, then learning continues just as if the tutor had supplied the correct action.

On the other hand, STEPS can decide that the tutor supplied no correct action for the crossed-out action because there is no such correct action (case 4 in Table 1). In this case, the standard machine learning assignment-of-blame technique will not work, so STEPS must use a riskier approach. It marks the incorrect action as something it needs to think about once problem solving has been done, and proceeds. When the problem has been solved, the system isolates the derivation of the incorrect action and deletes the first rule it finds that has not been used successfully in a previous problem. Obviously, the success of this tactic depends heavily upon the set of problems the system has been exposed to.

Illustrations

In this section, we illustrate STEPS' operation by presenting problems to it in pairs. The first problem of each pair gives it the opportunity to learn some new rules, and the second allows it to demonstrate what it has learned.

The first pair of problems demonstrates the learning of new rules via EBL. In the first problem, the system is required to solve a simple problem involving a block resting on a spring (see figure 2(a)). The system draws the free body diagram for the problem (see figure 2(b)). The free body diagram is incomplete, since the force exerted by the spring on the block is missing because the system does not know about this type of force. The system then derives an erroneous instantiation of Newton's First Law, stating that $mg = 0$. The tutor intervenes at this point, and draws an arrow originating at the block, pointing upwards. He then indicates to the system that it should continue. The system, faced with a hint by the tutor, decides to try to explain the hint to itself. It does so by activating its overly-general rules and trying to derive the action of drawing an arrow like the one drawn by the tutor. It succeeds in doing so by employing rules that suggest that a spring is an instance of a pusher, that pushers push, and that pushes can be forces. Once the system has succeeded in deriving the tutor's suggested action, it adds new rules to its knowledge base. These new rules (shown in figure 2(c)) are the result of applying EBL to the derivation.

STEPS assumes that the reasoning it engaged in after leaving out the spring force is suspect, so it crosses out the equation it generated, and attempts to generate another equation. This time it produces a correct instantiation of the First Law.

The third problem describes a situation where a force was applied for a short period of time (see figure 3(a)). This situation triggers a common misconception in human students — the belief that an agent that moves a body imparts *impetus* to the body (Halloun and Hestenes, 1985). The system draws the free body diagram shown in figure 3(b), which contains

the force that the man exerts on the block. The tutor promptly crosses out the arrow representing this force, and clicks the "Go Ahead" button. Because the tutor has not supplied a corrected version of the crossed-out force and STEPS cannot see how to do so, the system decides that the action just should not have generated at all. Thus, it continues solving the problem, but when it has finished, it searches for the rule that should be blamed for deriving the action of drawing the crossed-out force. The system walks backward along the derivation tree of the action, looking for a rule that has never been used in a problem that was correctly solved. The first such rule is indeed the rule embodying the misconception, and this rule is marked as incorrect. The next problem, involving a rocket that fires its engines for a short period, is solved correctly—the system does not draw the "impetus" force.

In the fifth problem, a rocket that is accelerating in a direction opposite to the direction of its motion (see figure 4(a)). The system's knowledge base contains an incorrect rule, which states that if a body is moving in direction X , its acceleration is also in direction X . This is one manifestation of the misconceptions caused by lack of differentiation between velocity and acceleration (Halloun and Hestenes, 1985; Reif, 1987). The system starts drawing the free body diagram and produces the acceleration arrow shown in figure 4(b). The tutor intervenes, and crosses out the incorrect (downward pointing) acceleration arrow. The system tries to fix the arrow using a simple repair heuristic: Try reversing the direction of an arrow. It draws the acceleration arrow, this time pointing upwards, and asks the tutor "Should it go this way?" The tutor confirms, and the system decides to try to explain to itself why this is true by deriving the correct action — drawing the acceleration arrow pointing upwards. Using the correct rule, which states that when the body is slowing down, the direction of its acceleration is the opposite of the direction of its motion, the system derives the correct action. It then tries to find which rule was responsible for the incorrect action. This is done by walking backward along both derivation trees (the incorrect action's and the correct action's) simultaneously, and finding the point where they diverge. The bad rule is assumed to be the first rule at the divergence point in the incorrect action's derivation, and it is marked as incorrect. The next problem, which holds the possibility of making the same mistake, is solved correctly.

Tutoring Tactics Indicated by STEPS

As we have demonstrated, STEPS learns to solve physics problems while receiving much the same information as a human student would. We cannot yet show that everything it does is matched by human behavior, however many of its learning mechanisms are plausible and have some independent support in the learning literature. Since this is the case, we may assume that tutoring tactics that help STEPS learn more effectively should also be employed by human and machine tutors. Moreover, we can now explain *why* these tactics are more effective. In this section, we will discuss these tutoring tactics.

Immediate Feedback. It is much easier for STEPS to learn from immediate negative feedback than from delayed negative feedback. When faced with delayed negative feedback, which




Figure 2: Problem 1




Figure 3: Problem 3




Figure 4: Problem 5

refers to an action that is not the last action, STEPS mentally crosses out the actions it took since the incorrect action. However, this is not always sufficient. In many cases an incorrect action results in erroneous conclusions that are not manifested in actions but reside in working memory. STEPS does not try to comb through its memory searching for conclusions that are rendered suspect by the tutorial input, so subsequent errors are possible.

Several studies indicate that immediate feedback is better for human learners than delayed feedback (Anderson, Conrad and Corbett 1989; Lewis and Anderson, 1985). Anderson's initial explanation for this effect was that delaying feedback makes it harder for the student to recall the reasoning that led to the incorrect action, thus making it more difficult to locate and repair the incorrect knowledge. However, in view of the fact that students can often reconstruct their reasoning from the scratchwork visible on the paper or screen, Anderson's current explanation is simply that immediate feedback prevents students from wasting time going down unproductive paths (Anderson et al., 1994), so delay does not affect the probability of learning from the feedback, only the efficiency.

STEPS is consistent with Anderson's latest explanation, in that it has no more trouble locating the incorrect knowledge when feedback is delayed than when it is immediate. Thus, delaying feedback wastes STEPS's time but does not affect whether it will learn from the feedback.

However, the development of STEPS suggests that there is a subtle penalty for using delayed feedback. When a tutor points out an "old" error to a human student, the student must retract the conclusions of his reasoning from that point onwards (which is what STEPS does), or engage in a complicated process of dependency-directed backtracking that would leave both his internal memory and his external memory (the page) disordered, and may confuse him in future reasoning. Either way, subsequent errors that the student makes may not be due to knowledge flaws, but to incomplete retraction of obsolete inferences. The tutor should either take this into account when responding to any subsequent errors that the student makes, or start the problem over from the beginning, or move to a new, similar problem.

In short, STEPS "teaches" the tutor that delayed feedback is just as likely to be successful as immediate feedback (although it does waste some time), but it may complicate the interpretation of any subsequent errors while solving the problem. This is a rather non-obvious tutoring tactic, so tutors who are not using the tactic already may find it a difficult discover even while practicing tutoring with STEPS. This is why we believe that STEPS and other simulated students cannot stand alone as tutor training devices. They must be accompanied by instruction on tutoring that explicitly mentions considerations like the ones discussed here.

Low Content Negative Feedback. Whenever STEPS takes an action that is basically correct but has some details wrong, the tutor can either cross out the action and give a correct version (high content negative feedback), or just cross out the action (low content negative feedback).

Low content negative feedback is often used in tutoring. Human tutors often initiate an episode of negative feedback

with a low content indication, such as an overly long pause (Fox, 1993). Many intelligent tutoring systems' first level of negative feedback consists of beeping, highlighting the incorrect action, or issuing some other low content indicator.

In STEPS, if the tutor uses low content feedback, then the learner has to solve two problems that it would not otherwise have to solve. First it must decide whether the action was crossed out because its details are wrong or because the whole thing is wrong (e.g., a non-existent force). STEPS uses problem-specific heuristics to make this decision. If it selects the details-are-wrong interpretation, then it faces the second problem, deciding which details are wrong. STEPS uses syntactic heuristics to guess a corrected version of the action, and then asks the tutor if it is correct. If it solves this second problem correctly, it finally has a corrected action, so learning proceeds just as if the tutor had entered that action in the first place. Thus, low content feedback has no advantage for STEPS over high-content feedback. It only increases the chance that STEPS will misinterpret the feedback.

STEPS's behavior is consistent with human data. McKendree (1990) showed that low content feedback caused less learning than feedback that provided some indication about what was wrong.

STEPS does not "teach" tutors to avoid low content negative feedback, because STEPS can sometimes learn from it. However, it does "teach" tutors to follow up low content negative feedback carefully in order to make sure that the student has learned from it. For instance, if the tutor intended a details-are-wrong interpretation and STEPS does not propose a correction, then the tutor should make the correction at the first opportunity.

Using the Student's Vocabulary. In human-human tutorial interaction, the usual complexities of understanding dialogues exist, though somewhat mitigated by the use of technical words and symbols. STEPS sidesteps these issues by restricting tutorial input in several ways. The tutor can only refer to entities such as forces, accelerations and equations by pointing to them, thereby precluding misidentification of the intended referent, a common source of misunderstanding in natural language dialogues. In addition, STEPS announces the meaning of each new variable as it is created, and expects the tutor to use these variable names. Without these conventions, the STEPS interface would be complicated immensely. A tutor working with a human student could impose these conventions on herself by pointing to the graphic representations of entities she is referring to and using the same variable names as the student whenever possible. Such restrictions in human-human interactions could help students avoid some of the problems of disambiguating tutorial input, and free students' cognitive resources for the important learning task.

Avoiding Shortcuts. A tutor may be tempted to collapse several reasoning steps into one resulting action. Consider the problem presented in Figure 2(a). The problem-solver has not yet learned about pressure forces, so it omits the force exerted by the spring on the block and generates an erroneous equation. The tutor could decide to intervene by proposing the correct equation: $P - mg = 0$. (Indeed, this is exactly what the two pilot subjects did - see (Ur and VanLehn, 1994).)

STEPS cannot understand this intervention, because it does not know what P stands for. In order to be able to understand the tutor's comment, STEPS would have had to guess that the equation is an instantiation of Newton's Law, therefore addends in the equation represent forces, so there must be a force missing from the diagram. It would then have to search for the missing force. This chain of reasoning would be extremely hard to produce. In the first problem described above, the tutor chose to draw in the missing force on the free body diagram, and this gave STEPS sufficient information to learn about pressure forces and to generate a correct equation on its own.

Human students may have similar difficulties in understanding tutors when they engages in "reasoning leaps." Catrambone (1993a, 1993b) has shown that human students also learn better from examples where shortcuts are avoided.

The tutoring tactic suggested by STEPS's solution is difficult to state precisely because it depends on the details of STEPS explanation algorithm. However, the gist of it is that the tutor should always take the same small steps that the student does. This is arguably a reasonable convention for tutors to learn.

Summary. We have discussed four tutoring tactics that STEPS encourages tutors to practice: (a) Delaying feedback can cause subsequent errors on the problem, so consider restarting the problem or going to a new problem when feedback has been delayed significantly. (b) Follow up low content feedback carefully to make sure that the student did not misinterpret it. (c) In order to make it easier for the student to understand which objects in the problem you are referring to, use the student's names for variables and other problem-specific terms, and use pointing whenever possible. (d) When demonstrating a line of reasoning, avoid skipping steps and use steps as small as the student uses. These four tutoring tactics could easily be employed as design policies for the pedagogical component of an intelligent tutoring system. Unlike other prescriptive work, these pieces of advice are each based on a computational model that indicates exactly why each tactic improves learning.

Acknowledgements

The implementation of STEPS was simplified immeasurably by the use of code from CASCADE (by Ted Rees) and OLAE (by Joel Martin and Jonathan Rubin). We would like to thank Stellan Ohlsson and Alan Lesgold for their valuable comments on drafts of this paper. This research was supported by the Cognitive Sciences Division of the Office of Naval Research under grant number N00014-93-1-1161.

References

- Anderson, J. R. (1993). *Rules of the Mind*. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Anderson, J. R., Conrad, F. G., and Corbett, A. T. (1989). Skill acquisition and the LISP tutor. *Cognitive Science*, 14(4):467-505.
- Anderson, J. R., Corbett, A. T., Koedinger, K. R., and Pelletier, R. (1994). Cognitive tutors: Lessons learned. Submitted for publication.
- Catrambone, R. (1991). Helping learners to acquire subgoals to improve transfer. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*, pages 352-357, Hillsdale, NJ. Lawrence Erlbaum Associates.
- Catrambone, R. (1993). The effects of labels in training examples on transfer to novel problems. In preparation.
- Fox, B. A. (1993). *The Human Tutorial Dialogue Project: Issues in the Design of Instructional Systems*. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Halliday, D. and Resnick, R. (1988). *Fundamentals of Physics*. Wiley and Sons, New York, NY.
- Halloun, I. A. and Hestenes, D. (1985). Common sense concepts about motion. *American Journal of Physics*, 53(11):1056-1065.
- Lewis, M. W. and Anderson, J. R. (1985). Discrimination of operator schemata in problem solving: Learning from examples. *Cognitive Psychology*, 17(1):26-65.
- 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*, pages 410-417, Charlottesville, VA. Association for the Advancement of Computing in Education.
- McKendree, J. (1990). Effective feedback content for tutoring complex skills. *Human-Computer Interaction*, 5(4):381-413.
- Merril, D. C., Reiser, B. J., Ranney, M., and Trafton, J. G. (1992). Effective tutoring techniques: A comparison of human tutors and intelligent tutoring systems. *Journal of the Learning Sciences*, 2(3):277-305.
- Mitchell, T. M., Keller, R. M., and Kedar-Cabelli, S. T. (1986). Explanation-based generalization: A unifying view. *Machine Learning*, 1(1):47-80.
- Newell, A. (1990). *Unified Theories of Cognition*. Harvard University Press, Cambridge, MA.
- Reif, F. (1987). Interpretation of scientific or mathematical concepts: Cognitive issues and instructional implications. *Cognitive Science*, 11(4):395-416.
- Sleeman, D., Langley, P., and Mitchell, T. M. (1982). Learning from solution paths: An approach to the credit assignment problem. *AI Magazine*, 3(2).
- Ur, S. and VanLehn, K. (1994). STEPS: A simulated, tutorable physics student. Submitted for publication.
- VanLehn, K. and Ball, W. (1991). Goal reconstruction: How Teton blends situated action and planned action. In VanLehn, K., editor, *Architectures for Intelligence*, pages 147-188. Lawrence Erlbaum Associates, Hillsdale, NJ.
- VanLehn, K. and Jones, R. (1993). What mediates the self-explanation effect? Knowledge gaps, schemas or analogies? In Polson, M. C., editor, *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, pages 1034-1039, Hillsdale, NJ. Lawrence Erlbaum Associates.
- VanLehn, K., Jones, R., and Chi, M. T. H. (1992). A model of the self-explanation effect. *The Journal of the Learning Sciences*, 2(1):1-59.