

Analogy Events: How Examples are Used During Problem Solving

KURT VANLEHN

University of Pittsburgh

In attempting to fit a model of analogical problem solving to protocol data of students solving physics problems, several unexpected observations were made. Analogies between examples and exercises (a form of case-based reasoning) consisted of two distinct types of events. During an *initialization* event, the solver retrieved an example, set up a mapping between it and the problem, and decided whether the example was useful. During a *transfer* event, the solver inferred something about the problem's solution. Many different types of initialization and transfer events were observed. Poor solvers tended to follow the example verbatim, copying each solution line over to the problem. Good solvers tried to solve the problem themselves, but referred to the example when they got stuck, or wanted to check a step, or wanted to avoid a detailed calculation. Rather than learn from analogies, both Good and Poor solvers tended to repeat analogies at subsequent similar situations. A revised version of the model is proposed (but not yet implemented) that appears to be consistent with all the findings observed in this and other studies of the same subjects.

I. INTRODUCTION

The seminal Chi, Bassok, Lewis, Reimann, and Glasser (1989) study is remembered for having discovered that students who explained physics examples to themselves learned more than students who merely read the examples. However, Chi et al. (1989) also found that learning was correlated with the way students used examples during problem solving. After studying 3 physics examples, the subjects solved 19 physics problems. They could refer to the examples as often as they wished while solving problems. All the students referred to the examples, but not equally often. The more effective learners (called Good solvers by Chi et al., 1989) referred to the examples less frequently during problem solving than the less effective learners (called Poor solvers). That is, learning was inversely correlated with frequency of example reference. When students did refer to an example, the

Good solvers tended to read a single line in the middle of the example, whereas the Poor solvers tended to start at the beginning of the example and read many lines before returning to problem solving. Thus, Chi et al. (1989) actually found a three-way correlation: effective learning correlated with both self-explanation and reference to examples during problem solving.

In developing a computational model of the self-explanation effect, VanLehn, Jones, and Chi (1992) assumed that these references to examples were caused by analogical problem solving, and that differences in the way students studied examples caused differences in their analogical problem solving. These assumptions were built into the computational model, Cascade. In particular, Cascade would refer to an example only when it reached an impasse. That is, when it didn't know what to do, only then would it refer to an example to find out what to do. Thus, if students learned a lot by self-explaining the examples, then they would have fewer impasses, so they would refer to the examples less (VanLehn & Jones, 1993b; VanLehn et al., 1992). This was how Cascade explained the three-way correlation.

However, when Cascade was fit to the protocols of the Chi subjects, (VanLehn & Jones, 1993c), it was discovered that the students often referred to examples when Cascade wouldn't. In particular, they would refer to the examples even when they could not be at an impasse.

The inaccuracy of Cascade's model of analogy motivated the present study, which was designed both to thoroughly evaluate Cascade's model of analogy and to deepen our understanding of analogical problem solving in preparation for building an improved model of it. Since a hypothesis-testing design would not be appropriate for such a goal, the study used a traditional hypothesis-formation design, namely building a descriptive taxonomy bottom-up. Each episode where the student referred to an example was categorized three ways:

- *Purpose*: Why did the student refer to the example?
- *Content*: What information did the student transfer from the example to the problem?
- *Repetition*: Did students repeat this same reference to the example later?

The analogy episodes of each type were counted in order to see if Good solvers tended to use analogy differently than Poor solvers. There did turn out to be significant differences.

This study is one of a series of analyses of the same protocols, namely the ones collected during the Chi et al. (1989) study. It could be that the results are specific to these students or this experimental setting. However, this methodological risk is necessary. The overarching goal of this series of analyses is to understand how students learn a conceptually rich, non-trivially large knowledge base over an extended period of naturalistic study. This goal motivates taking protocols of subjects over a relatively long period (approximately 4.5 hours per subject), which results in a corpus of unusually complex behavior. To understand it requires multiple analyses from multiple perspectives. In short, in order to understand complex, naturally-occurring, lengthy learning, one needs to concentrate on the same corpus of data until it is fully understood. With this study, I believe I have at last understood

what these subjects are doing. The discussion section presents this model, albeit informally.

This paper begins with a brief review of the literature on models of analogical problem solving, followed by a description of Cascade. The descriptive study and its results are presented next. A discussion section presents causal interpretations of the various observed correlations, and weaves these into an integrated (albeit speculative) account that is consistent with all the data known about these rather famous subjects and their learning.

II. MODELS OF ANALOGICAL PROBLEM SOLVING

A Modal Model of Analogical Problem Solving

Most models of analogical problem solving are essentially similar. They assume examples are stored in memory as cases (Kolodner, 1993) or other kinds of mental structures. When the model begins to solve a problem, it retrieves from memory a case whose problem is similar to the problem being solved. It then adapts the example's solution and thereby solves the problem. Adaptation is sometimes done by recursion, namely finding an example of a similar episode of adaptation in memory and applying it to achieve the required adaptation (Kolodner, 1993, chapter 11). Adaptation may also be done by using small, rule-like pieces of domain knowledge (e.g., Falkenhainer, 1990; Kambhampati, 1993; Veloso & Carbonell, 1993). After the problem has been solved, its solution may be abstracted, paired with the problem, and stored in memory as a new case. All solvers do retrieval and adaptation, but only some do the third step of generalization and storage.

In most models, the unit of information transferred from the case to the problem is a whole solution. For instance, if the problem is to design a meal, then JULIA (Hinricks, 1992) transfers a whole menu. If the problem is to explain why a hot brick in a bucket of cold water gradually cools off, PHINEAS (Falkenhainer, 1990) transfers a whole qualitative process model. Only a few models transfer just a piece of a solution. For instance, when CELIA (Redmond, 1990) tries to solve the problem of diagnosing a stalled automobile, it transfers pieces of cases (called *snippets*) such as the hypothesis that the stall could be caused by a loose spark plug wire. ARCHIE-2 (Domeshek & Kolodner, 1992; Domeshek & Kolodner, 1993) helps design buildings by transferring pieces of a plan for a whole building (called *stories*) concerning relationships between building subsystems (e.g., the entryway) and design goals (e.g., security). However, systems like CELIA and ARCHIE-2 are much less common than systems that transfer whole solutions.

For most analogical problem solvers, the purpose of referring to the case is that analogical reasoning is their preferred way to solve problems, or even their only way. In this respect, they are quite different from Cascade, which refers to an example only when it has reached an impasse while doing rule-based problem solving. Some analogical problem solvers do have units of small, abstract knowledge (often represented as rules) that could in principle be combined to solve problems, but they prefer to use them only for adapting a case's solution. For instance, PRIAR (Kambhampati, 1993) solves hierarchical planning problems by transferring a whole plan, then removing parts of the plan that are invalid for

the problem and replacing them with non-primitive tasks. These tasks are then solved (i.e., are replaced by plans that have only primitive tasks) using an ordinary hierarchical planner. The planner could in principle have generated a plan from scratch, but it used analogy instead to create most of the plan, then generated the rest using rule-like knowledge.

As mentioned earlier, this study describes analogy via three classifications: content, purpose and repetition. This brief review of the analogical problem solving literature indicates that along the content dimensions, most models transfer a whole problem solution. Along the purpose dimension, most models refer to a case in order to solve a whole problem because analogical problem solving is their preferred solution method. Along the repetition dimension, some models avoid repetition by storing new cases and some do not. Thus, let us define the *modal* model of analogical problem solving as a solver that prefers to solve a problem by importing a whole problem solution.

Cascade is Not a Modal Model of Analogical Problem Solving

Cascade is different from most models of analogical problem solving in two ways. First, it normally does rule-based problem solving and resorts to analogy only when it reaches an impasse. Second, when it does use analogy, it transfers only a small piece of an example's solution. Thus, in terms of both the purpose and content of its analogical references, it is the opposite of the modal model of analogical problem solving. This section motivates these and other features of Cascade's design. Because the motivation depends on subjects' behavior, brief descriptions of their tasks and their behaviors are presented first.

A Gross Description of the Subjects' Analogical Problem Solving

The 9 Chi et al. (1989) subjects studied 3 physics examples then solved 19 physics problems. The problems were ordered by their similarity to the examples. The first 3 problems were isomorphic to the 3 examples. (The examples and these 3 problems are shown in the Appendix.) Although the given and sought quantities were different, corresponding examples and problems had the same diagrams, forces and equations. The next group of three problems were also isomorphic to the examples, but the value of gravitational acceleration (g) was different, which affected the equations. The third set of three problems replaced objects in the examples with different objects (e.g., blocks became balloons) in different geometric relationships, but the problems were still isomorphic to examples at the level of forces, and the diagrams were visually similar to the example diagrams. The fourth set of three problems used completely different objects and diagrams, although they remained isomorphic at the level of forces. The remaining 7 problems were not isomorphic to the examples.

Analogical retrievals were only partly memory retrievals. The students usually flipped through the textbook searching for an appropriate example. Since both the examples and problems had diagrams, they could use those instead of the text to determine if the example was similar to the problem. The examples' diagrams were quite distinct from one another, so there was only one case of incorrect analogical retrieval. It occurred when a subject searched the textbook for an example to help her solve the three-strings problem (problem

1 in the Appendix) and found an example of a block hanging from a single string which appeared earlier in the textbook than the three target examples.

Analogical application consisted of episodes of reading the example alternating with episodes of working on the problem. For instance, subject S110 began to solve the first pulley problem (number 3 in the Appendix) with the following episodes:

1. He reads the problem statement, then says "Well it shouldn't be too hard to figure this out if I look in the example."
2. He pages through the textbook, finds the pulley example, and says, "Looking at the free body diagram here." (A free body diagram shows the forces acting on the object whose motion is being analyzed. In the Appendix, the free body diagrams for the examples appear to the right of the diagrams.)
3. He then draws a free body diagram for the problem. He apparently is not copying the example's diagram verbatim, because he labels the vectors differently. In fact, he labels the two weight vectors with W , thus incorrectly indicating that they have the same magnitude.
4. He then says, "And there's its weight, which is . . . which is 20 kilograms times 9.8 meters . . . umm . . . which is its weight." This calculation, which he has probably entered on his calculator, is appropriate for the larger object, whose mass is given, but not for the smaller object whose mass is unknown. When he goes to enter his result on the diagram, he seems to notice his mistake, because he pauses for a long time and does not write the result of his calculation down.
5. Eventually he says, "subject is looking at example 8 again." He appears to read it silently for a moment, then copies the first two equations verbatim from the example.
6. He says, "Since the larger mass is going to be accelerating downward, it has a negative acceleration. And then . . . We plug those . . . We solve one of these for T ."

The remainder of the solution is a long struggle with the algebra of solving simultaneous linear equations. The point here is that the subject alternated between referring to the example (the even numbered episodes above) and working on the problem (the odd numbered episodes). This alternation is characteristic of all the subjects whenever they are doing analogical problem solving. Some subjects solved problems without referring to the examples at all. But no subject just referred to the example once, then wrote down a solution, as one would expect if they were using the modal model of analogical problem solving. Whenever an example was referred to in order to solve a problem, it was referred to repeatedly in alternation with ordinary, nonanalogical problem solving.

Cascade's Model of Analogical Problem Solving

When Jones and I first examined the problem solving protocols, it seemed to us that many of the references to examples were triggered by impasses, such as the one in episode 4 above. In that episode, the subject's impasse is apparently caused by noticing a mistake. Other impasses were caused by not knowing what to do next. Referring to the example when one is stuck seemed to us to be both highly plausible and consistent with informal examination of the protocols. Thus, we made the simplifying assumption that *all* analogical references were triggered by impasses. (Since we wanted Cascade to learn during problem

solving, we assumed that it would first try to use its learning techniques to resolve an impasse, then use analogy if the techniques failed.) Assuming all analogy was triggered by impasses meant that the number of references to examples was a function of the number of impasses Cascade had, which in turn was a function of how much it learned during example studying. Thus, making impasses trigger all analogical problem solving would explain the three-way correlation between self-explanation, learning and analogy observed by Chi et al. (1989): The more self-explanation, the more learning during example studying and hence the fewer impasses during problem solving, and the fewer episodes of analogy.

Like all models of analogical problem solving, Cascade's had retrieval and application components. We chose not to model the subjects' retrieval processes because they used a combination of memory search and textbook search to find examples. Thus, we simply gave Cascade a table pairing examples with problems.

The application component began by finding a mapping between the current problem and the examples' problems. There are many algorithms for finding mappings, but most would return the same mapping on the particular problems and examples used in the study. Moreover, our subjects never made any mistakes when finding mappings, so we had no error data to shed light on the nature of their mapping process. Thus, Cascade used a simple heuristic algorithm (VanLehn & Jones, 1993b, p. 273).

Since this algorithm was computationally expensive, Cascade saved the resulting mapping. If a subsequent impasse triggered another episode of analogy for the same problem, then not only would Cascade retrieve the same example, it would also retrieve the same mapping rather than recalculate it.

After a mapping had been calculated or retrieved, Cascade resolved the impasse by searching through the lines of the example in order to find one that addressed the goal that caused the impasse. Cascade's goals are mostly about finding values for quantities, so most of its impasses occurred when it did not know a physics principle (equation) that contained the quantity it was currently seeking. Thus, the analogical application process usually involved searching through the example's lines for an equation that contained a sought quantity. When Cascade found such an equation, it applied the analogical mapping to create a version of the equation specific to the problem. This resolved the impasse, and problem solving continued.

Cascade could in principle have saved the resulting problem-specific equations in memory. This would reduce the number of episodes of analogical problem solving in the future. Cascade would have fewer impasses because it would often find the quantities it was seeking in the saved equations. Although human students could undoubtedly remember such equations for at least a short while, they probably assigned a different epistemological status to them than to equations from the textbook. As one student said, "This is called copying too much from the book. I hate that." Thus, even if students remembered the equations they copied from the examples, they might prefer not to use them. This intuition motivated the design decision to have Cascade not save the results of its analogical application.

In summary, the three components of analogy as implemented in Cascade are:

- *Analogical retrieval*: Not modeled. Cascade used a table pairing problems with examples.
- *Analogical application*: Mapping is done once and saved. Cascade transfers a single line of the example, usually an equation, if it can find one that achieves the goal that caused the impasse that triggered the analogy.
- *Analogical generalization*: Cascade does not store the results of its analogical application.

In addition, Cascade would try analogical problem solving only when it reached an impasse during normal problem solving.

Analogical Search Control

When we began to test Cascade, we discovered that it lacked search control knowledge that the subjects apparently possessed. That is, when a goal could be achieved by multiple rules, Cascade would often choose the wrong rule. This made it impossible for Cascade to learn certain rules that the subjects did learn (see VanLehn, Jones & Chi, 1992, p. 23).

We hypothesized that subjects relied upon the examples to control their search. In particular, when faced with a decision about which of several ways to achieve a goal, they would recall how a similar goal was achieved in solving an analogous example, and use that way to achieve this goal. An earlier system, Eureka (Jones, 1989) was also based on this assumption.

This was modeled in Cascade by a mechanism called *analogical search control*. The older mechanism (discussed above) was named *transformational analogy* (after Carbonell, 1986) in order to distinguish it from analogical search control.

To implement analogical search control, two changes were made to Cascade. First, whenever it successfully achieved a goal during either problem solving or example explaining, it would store in memory a triple consisting of the problem, the goal and the rule used to achieve the goal. These triples were viewed as a crude approximation to an episodic memory. The second change was to have the problem solver do analogy whenever it was faced with a search control decision. In particular, it would retrieve and map an example, using the same mechanisms as transformational analogy, then search the example for a triple that had a goal corresponding (via the map) to the current goal. If it found one, then it would follow the triple's advice and select its rule for achieving the current goal.

Analogical search control was a success from several standpoints. First, it enabled Cascade to learn rules that it could not otherwise learn. Second, it added a second explanation for the self-explanation effect. If Cascade did not self-explain the examples, then it had no triples to refer to, so it would get lost more often during problem solving and learn fewer rules. Third, analogical search control modeled the intuition that students learn more than just physics rules from studying the examples because they also learn "how to" knowledge. According to Cascade, they learn very specific pieces of advice about how to solve certain physics problems. In this respect, the triples stored by Cascade served somewhat the same function as the cases stored by the modal analogical problem solver.

III. DESIGN OF THE CURRENT STUDY

This section describes the reasons for undertaking the study and the analytic methods used during the study. It begins with a description of the results from the initial evaluations of Cascade's model of analogy.

Empirical Difficulties in Matching Cascade to the Protocols

Difficulties with Analogical Search Control

From the beginning, there were empirical difficulties with analogical search control. If Cascade self-explains the examples, then analogical search control succeeds more often because self-explanation causes more triples to be learned. However, the results of Chi et al. (1989) indicated that the subjects who self-explained the examples referred *less* frequently to them during problem solving than subjects who did not self-explain the examples. This suggested that analogical search control was mostly achieved by mental reference to the examples rather than observable reference, so most of the observable uses of analogy would correspond to episodes of transformational analogy, which did occur more frequently when Cascade did not self-explain the examples. In order to fit the Chi et al. (1989) data, between 91.8% and 98.3% of the 551 successful analogical search control applications had to be mental rather than observable (VanLehn, Jones, & Chi, 1992, p. 26). This seemed a bit implausible, especially given that the information transferred during analogical search control was very specific and detailed. Although these figures are based on counting successful retrievals of search control information from the examples, the conclusions would not be affected if all retrievals were counted including those that failed to find useful search control information. With either method of counting, Cascade's modeling of the Good solvers used analogical search control much more than its modeling of the Poor solvers, whereas Chi et al. (1989) observed Good solvers to refer less frequently to examples than the Poor solvers. To resolve the discrepancy requires assuming that most references to the examples for analogical search control are mental rather than observable.

More difficulties were discovered when Cascade was matched to the protocols of the 9 subjects. When modeling the 9 subjects' behavior during problem solving, Cascade made a total of 3,461 search control decisions, and 97.7% of these matched the decisions made by subjects (VanLehn & Jones, 1993c, p. 65). However, if analogical search control had been turned off, Cascade would still match 95.6% of the subjects' decisions. In this case, search control decisions would be made by the order in which the rules appeared in a file. This default ordering was apparently not a bad approximation to the subjects' search control knowledge. It seems plausible that adding only a few heuristics would allow a nearly perfect match to the subject's behavior without using analogical search control. For instance, one heuristic that would make many search control decisions correctly is, "try not to generate the same equation twice."

In summary, the evidence suggests that analogical search control, while successful computationally, was not a good model of the subjects' search control process. It seems more

plausible that subjects use a few heuristics drawn from their prior experience with algebraic problem solving or induced from the physics examples.

Difficulties with Transformational Analogy

Transformational analogy turned out to be more common than predicted. When Cascade was matched to the problem solving protocols, there were 258 episodes of analogical problem solving among the 9 subjects, but Cascade did transformational analogies during only 176 (68%) of the 258 episodes (VanLehn & Jones, 1993c, p. 59). Clearly, transformational analogy (or some kind of analogy) was occurring more frequently than expected among the subjects.

Moreover, in 26 of the cases where Cascade failed to predict an analogy episode, there was clear evidence that the subjects could have used regular problem solving instead, because they had done so on earlier occasions (VanLehn & Jones, 1993c, p. 60). Thus, in those cases, analogy was probably not triggered by impasses. It seemed plausible that subjects had simply decided to use the example rather than their general knowledge.

The way we fit Cascade to the protocols could have hidden many more such cases. If a subject never used a particular rule but always referred to the example instead, then we assumed the rule was missing from their knowledge and that impasses were causing the analogies. However, it could have been that the student actually knew the rule but nonetheless preferred to use analogy. Thus, there could have been many more than 26 cases of "voluntary" analogy.

These results suggest that there is something fundamentally wrong with Cascade's model of analogy. This would not be too surprising, since analogy was not the main focus during the development of Cascade.

Motivations for a New Coding of the Protocol Data

Since we had already collected 258 cases of analogy while matching Cascade to the protocols, it seemed like an easy task to induce a more adequate model of analogy from them. However, it became clear that these 258 cases were not appropriate for several reasons.

First, in many of them, the subject was merely repeating an analogy done on an earlier problem. Since there were four isomorphic problems for each example, many subjects simply repeated the analogies they had done on the first problem during their solution of the other three. This made it difficult to discern whether those subsequent occurrences of analogy were caused by impasses or habit. That is, if a missing rule caused an impasse on the first problem which in turn caused an analogy, it would usually still be missing on the second problem where it could cause another impasse and another episode of analogy. On the other hand, the subject might simply remember that they had used analogy on the first problem at this time, and simply do it again on the second problem. In short, memories of earlier episodes of analogy made interpretation of subsequent episodes difficult.

Second, cases of analogy were only collected when they corresponded to actions that Cascade did, where "action" means either a step achieved by analogical problem solving or by rule-based problem solving. This was a natural way to collect data for evaluating the

match between Cascade's behavior and the protocols. However, it omitted some cases of analogy and sometimes treated consecutive episodes of analogy as a single episode.

Consequently, it was decided to collect a new corpus of analogy events. In order to avoid the difficulties of interpretation caused by memories of earlier episodes of analogy, the new events were collected from only the first 3 problems (those shown in the Appendix). In order to avoid omitting events or merging them, Cascade's analysis of the protocols was ignored and the new corpus was collected by examining only the protocols themselves.

Analytic Methods

Since the main purpose of the study was to induce an improved model of analogical problem solving, a bottom-up taxonomic approach was used, which seemed appropriate for such a hypothesis-formation goal. However, the analysis was structured around the three major differences between Cascade and the modal problem solver:

- *Purpose:* Cascade only did analogy when it could not achieve a goal (i.e., an impasse that causes transformational analogy) or had multiple ways to achieve a goal (which caused an episode of analogical search control). The modal model does analogical problem solving as its primary method for solving the problem.
- *Content:* Cascade transfers a single line (e.g., an equation) during each analogy event. The modal model transfers the whole solution and adapts it.
- *Repetition:* Cascade does not store the results of analogical application, so it has to repeat analogies. Some modal models store results in memory as new cases, and thus avoid such repetitions.

The analogy events were taxonomized once according to their apparent purpose and then again according to the content of the information transferred. This analysis was done on the analogy events that occurred during the first 3 problems, for the reasons described earlier. In order to determine whether analogies were repeated, the protocols of the remaining isomorphic problems were analyzed. At each place where an analogy event occurred during one of the first 3 problems, similar places (if any) in the remaining problems were examined to see if that analogy was done again.

In order to study the correlation of analogy events with effectiveness of learning, this study retained the Chi et al. (1989) classification of the subjects as Good or Poor solvers based on their scores. The frequencies of each taxonomic category were determined in order to see if Good solvers used analogy differently than Poor solvers.

One of the 9 subjects never referred to the examples during the first three problems, so her data was excluded. However, because the student who was eliminated was a Good solver, analyses were conducted with 3 Good and 5 Poor solvers. This uneven distribution was used not only to make the analyses comparable with earlier analyses but also because two subjects (S102 and S105) were nearly tied for best of the Poor solvers. It seemed better to keep both in the Poor classification than to pick one arbitrarily as the better learner. Chi et al. (1989) solved this problem by eliminating S102 from the analyses, but since a Good

solver had already been eliminated here, it was preferable to retain S102, classified as a Poor solver.

In order to contrast Good and Poor solvers, it was necessary to define a unit of analysis that would allow numerical contrasts (Chi, 1997). This was achieved in the first analysis by isolating episodes of the problem solving wherein the subject referred explicitly to an example. Such episodes were called *analogy events*. A reference to an example was coded only when the subjects said they were referring to an example (e.g., "Okay, I'm going to go back and look at that pulley problem.") or when they had already referred to an example recently and seemed to be referring to it again (e.g., "What did they use? They used cosine."). However, because the protocols were transcribed from audio tape recordings, it is likely that subjects occasionally referred to the example without giving any audible indication of it.

Frequency of Analogical Reference

Before presenting the taxonomic results, it is worth a moment discussing the frequency of analogical reference. The 8 subjects solved 3 problems each, and the resulting 24 protocols yielded 148 analogy events. On average, the subjects referred to an example about 6.2 times per problem.

Good and Poor solvers appear to use analogy about the same amount. The Good solvers produced an average of 19.3 analogy events each, while the Poor solvers produced 18.0 analogy events each, which is an insignificant difference (t-test, $p = .68$). This is surprising, given that Chi et al. (1989, p. 173) found that Good solvers produced fewer analogical references than Poor solvers, although the difference was only marginally significant. The discrepancy is probably due to the fact that the two analyses were based on different sets of Good and Poor solvers. This analysis eliminated a Good subject who never used analogy, and Chi et al. eliminated a Poor subject who used analogy 18 times. If no subjects are eliminated, then averages are 14.5 and 18.0, which is still insignificant (t-test, $p = .59$). It appears that differences among the learners may lie primarily in the way they use analogy, and not in how much they use it. The remainder of this section reports various categorizations of the analogy events aimed at characterizing differences in analogy usage among the learners.

IV. THE PURPOSE ANALYSIS

As discussed earlier, analogy is viewed as consisting of several processes including retrieval, mapping, application and generalization. When the same example and problem are used in several analogies, some of the initial processes, such as mapping, can be shared. This motivates looking in the protocols for analogy events that support the "initialize" analogical events but do not actually include any analogical application.

Thus, analogy events were coded as either initialization events or transfer events. A transfer event was defined as one where the student attempted to import some piece of information from an example and incorporate it into a problem's solutions. An initialization event was defined to be an episode where the student did *not* attempt to import a piece

of information to be used in the problem solving. Thus, initialization events are defined to exclude transfer, but transfer events can include some initialization.

Initialization Events

Of the 148 analogy events, 119 were coded as transfer events and 29 were coded as initialization events. In all the initialization events, the subject opened the book to an example if it was not already opened to the correct page, read some of the example, and decided that it was useful for solving the current problem. This is why the category is named "initialization." This section discussed the initialization events. The next section discussed the transfer events.

When Did Initialization Occur?

In 20 of the 29 initialization events, the event occurred soon after the problem was presented. All 3 problems consisted of a diagram and text. Sometimes subjects referred to an example after viewing the diagram but before reading the text. At other times, they read both diagram and text before referring to an example. (Subjects were instructed to read all text aloud. Reference to the diagrams could sometimes be inferred from their comments, and sometimes from the sound of pages being turned.) For the 3 problems analyzed, the information in the diagrams was sufficient for determining whether an example and problem were similar enough that the example could be useful in solving the problem (see Appendix). The Good solvers may have ascertained this. Table 1 shows when each of the 20 initialization events occurred and whether it belonged to a Good or Poor learner. The table indicates that Good solvers referred to the example more frequently after seeing just the diagram than the Poor solvers ($\chi(1) = 9.69, p = 0.02$). Overall, most students preferred to read the whole problem statement before referring to the example.

The other 9 initialization events occurred in the middle of a problem. They all occurred on the inclined plane problem (see the Appendix), which had 3 parts, labeled a, b and c. Parts b and c changed the physical situation slightly from part a (e.g., part b began, "Now suppose that we cut the string. . ."). Some students treated parts b and c as distinct problems, and went back to look at the examples before starting to solve them. Such behavior was coded as an initialization event. Thus, initialization events all occurred either at the very beginning of solving a problem, or at the beginning of a variant of a problem. That analogies began early in the problem solving is consistent with the findings of several earlier studies (Bassok, 1990; Bassok & Holyoak, 1989; Faries & Reiser, 1988; Hinsley, Hayes, & Simon, 1977; Ross, 1989b).

Although analogical problem solving often began with an initialization event, there were exceptions. In 20 of the 24 problems analyzed, the first reference to an example was an initialization event. On the remaining 4 problems, the first reference to the example was a transfer event. For instance, S105 started to solve a problem by drawing a force diagram for it. When he was finished, he decided to check it, so he referred to an example. He studied the example for a moment and discovered that he had omitted a force from his diagram. Thus, his first reference to the example imported some information about the problem's

TABLE 1
When Did Initialization Events Occur?

<i>When Occurred</i>	<i>Good Solvers</i>	<i>Poor Solvers</i>	<i>Total</i>
After the diagram alone	3 (38%)	0 (0%)	3 (15%)
After text and diagram	5 (63%)	12 (100%)	17 (85%)
Total	8	12	20

solution, which makes it a transfer event. Thus, in 20 of the 24 problems, subjects started the process of analogical problem solving before they had even begun to solve the problem, while in 4 of the 24 problems, the subjects started to solve the problem without help, and made their first reference to an example in order to obtain help with part of the solution.

As discussed earlier, models of analogy differ in how they combine the constituent processes of analogy, so it is interesting to consider which processes are included in initialization events. An initialization event certainly includes retrieving an example and deciding if it is similar enough to the problem to warrant using it for help in solving the problem. An initialization event may also include finding a mapping between the elements of the problem and the elements of the example. Although students never spoke about this mapping, it is difficult to imagine how a student could make a decision about using the example without doing some mapping.

The distinction between initialization and transfer events has appeared in no computational model of analogy known to the author. The closest approximation is probably Cascade. Cascade would often refer to the same example several times in the course of solving a problem, but all these references were transfer events. On the first one, it would save the mapping. On subsequent references, it would retrieve that mapping rather than recalculate it, because this saved it considerable computation. However, Cascade did not have initialization events.

Amount of Elaboration During Initialization

According to current models of analogy, students may differ on how much elaboration of the problem and the example they do before concluding that the two are analogous. The issue of deep vs. shallow analogy has dominated the experimental literature on analogy (e.g., Brown & Kane, 1988; Brown, Kane, & Echols, 1986; Clement & Gentner, 1991; Faries & Reiser, 1988; Gentner & Landers, 1985; Gick, 1990; Gick & Holyoak, 1980; Gick & Holyoak, 1983; Holyoak, 1985; Novick, 1988; Reed, 1987; Ross, 1987; Ross, 1989a; Ross & Kennedy, 1990). The difference in amount of elaboration and presumably the depth of the analogy was evident in the behavior of the subjects. During one initialization event, for instance, the subject calculated whether two angles summed to 90 degrees.

In order to determine how much elaboration subjects did and to determine whether Good solvers did more elaboration than Poor solvers, the 20 initialization events that occurred at the beginnings of problems were re-analyzed. Subjects were coded as doing shallow elaboration if they simply referred to the example (e.g., "Ok, here's the pulley example.") or re-read part of it without further comment. Subjects were coded as doing

deep elaboration if they made any comment about the correspondence between the problem and example beyond merely stating that it was acceptable. Most of the initialization events occurred without such deep elaborations (see Table 2). The Good solvers preferred shallow matching just as frequently as Poor solvers ($\chi(1) = 2.10$). The predilection for shallow matching is understandable given that these problems are nearly isomorphic to the examples.

In summary, the main finding is that most students preferred to initialize an analogy right away, as soon as they had read the problem, rather than wait until they needed help in solving the problem. A few subjects retrieved examples on the basis of the diagram alone without even reading the problem. All subjects tended to accept an analogy without a great deal of elaboration, probably because the problems were so similar to the examples.

Transfer Events

Of the 148 analogy events, most (119) were transfer events. That is, the subjects made some kind of analogical inference that helped them solve the problem. Each event was analyzed in order to find out what goal or purpose provoked the analogy event. This analysis is presented in the following section.

What Goal or Purpose Provoked the Analogy Event?

Each transfer event was analyzed to determine why the subject decided to refer to the example at that time. The analysis used not only the protocol of the event itself, but also the context of that event. Categories were invented freely during coding in order to represent the inferred purposes as accurately as possible, then similar categories were aggregated to form the five categories shown in Table 3. Each category will be discussed in turn, along with the prior modeling and empirical work that motivates them.

Repairing an Impasse

Several computational models invoke analogy whenever regular problem solving reaches an impasse (Anderson, 1993; VanLehn et al., 1992). Such models suggest that transfer events should often occur because the subject is stuck. Subjects were coded as repairing an impasse if they had a specific goal, could not achieve it and referred to the example. For instance, Table 4 shows the protocol of one subject who got stuck when he had obtained the mass of an object, measured in kilograms, but the problem asked for its weight, which is measured in Newtons. Around line 235, the subject refers to the example. This transfer

TABLE 2
How Deep Were the Elaborations?

	Good Solvers	Poor Solvers	Total
Deep elaborations	3 (38%)	3 (25%)	6 (30%)
Shallow elaborations	5 (63%)	9 (75%)	14 (70%)
Total	8	12	20

TABLE 3
Purposes for Analogical Transfer Events

Purpose	Good Solvers	Poor Solvers	Total
Repairing an impasse	9 (20%)	5 (7%)	14 (12%)
Following the example	8 (17%)	44 (60%)	52 (44%)
Implementing a step	16 (35%)	10 (14%)	26 (22%)
Checking	10 (22%)	11 (15%)	21 (18%)
Miscellaneous	3 (7%)	3 (4%)	6 (5%)
Total	46	73	119

event is clearly motivated by his impasse, which starts around line 223. As Table 3 indicates, impasses provoked 14 transfer events.

Following the Example

As several investigators have noted (e.g., Pirolli & Anderson, 1985) subjects sometimes follow an example slavishly. Whenever they get finished writing a line, they look at the example, read its next line and copy it. They make small changes as they copy lines in order to adapt them to the problem being solved. In this study, subjects copied the examples' diagrams and equations, but rarely copied the lines consisting of English explanations. Surprisingly, subjects even copied equations that were not directly used in generating the example's answer, but were included in the example for didactic purposes. This indicates that they had turned control totally over to the example. They were no longer solving the

TABLE 4
An Impasse-Driven Transfer Event

222	Use the mass of the block.
223	It's in kilograms.
224	It doesn't matter, mass and weight are the same...aren't
225	they?
226	No.
227	I think they are.
228	Yeah, a Newton is the weight of an object.
229	Um,....
230	[Experimenter: Keep talking.]
231	Okay.
232	Oh shoot.
233	How do I have to figure it out?
234	I might have to figure that out.
235	Um, just looking at this....shoot.
236	What do you think? [Pause.]
237	[Experimenter: What are you looking at?]
238	I'm looking up over here and I'd say this is equal to its
239	weight.
240	[Experimenter: What is "this"?)
241	Also, I'll put this as square.
242	It says $F_c = -F_c$ which is equal to its negative weight.
243	So, that's it weight

problem per se, but were engaged in the task of copying the example's lines instead. As Table 3 shows, 52 transfer events were motivated by following the example closely, and Poor solvers were responsible for most of them.

Implementing a Step

In contrast to subjects who turned control over to the example, some subjects continued to focus on solving the problem. Most applied a form of means-ends analysis, wherein one tries to find a value for a variable by writing an equation that contains that variable. If there are unknowns in the equation, then the strategy recursively seeks values for them as well. While pursuing this strategy, subjects would often obtain equations from the examples instead of from general principles such as Newton's law. This is illustrated by the following excerpt:

73 [Reads: B, with what acceleration are they moving?]

74 What acceleration are they moving?

75 Okay, the acceleration equals mass of two minus mass of one

76 over mass of two plus mass of one times gravity.

The subject had already consulted an example for solving part A of the problem, and the example remained open in front of her. She read the sought quantity from part B of the problem statement, and immediately copied an equation from the example,

$$a = \frac{m_2 - m_1}{m_1 + m_2} g$$

that had exactly that variable as its only unknown. She was not following the example slavishly. In fact, the equation that she copied is not the first one in the example, but third (see example 3 in the Appendix). This illustrates the application of the means-ends strategy with equations drawn from the example.

Some subjects did not use means-ends analysis, but instead followed a plan such as the standard one of drawing forces, assigning coordinate axes and writing equations for solving the problem. In the course of following their plan, they would sometimes refer to the example for equations or other technical details. Thus, all transfer events that were motivated by obtaining technical details from an example in order to pursue means-end analysis of a plan were lumped together and called "Implementing a step." As Table 3 shows, there were 26 transfer events that were motivated by implementing a step.

This category differs from the category of "repairing an impasse" in that there is no sign of the subject being stuck. The subject neither paused nor expressed frustration. In fitting Cascade to the protocol data, we often had to force it to use analogy even when it possessed appropriate general knowledge and used it on other occasions (VanLehn & Jones, 1993c). Clearly, not all analogical references are driven by impasses, even when subjects are solving problems instead of copying examples.

Checking an Action or a Decision

Some subjects would write an equation or vector, then refer to the example to see if the action was done correctly. As Table 3 shows, there were 21 transfer events of this type. No computational model of analogical problem solving includes the capability of checking its reasoning against an example.

Miscellaneous Transfer Events

This residual category contains 6 transfer events where it was difficult to understand why the subject was referring to the example.

Differences Between Good and Poor Solvers

As Table 3 shows, there were significant differences in the distribution of transfer events across Good and Poor solvers ($\chi(4) = 20.2, p = .0002$). Most (60%) of the transfer events of the Poor solvers were caused by following an example verbatim, whereas most (76%) of the transfer events of the Good solvers were caused by following a general problem solving strategy and referring to the example to repair an impasse, or to implement a step or to check a step. That is, the Good solvers preferred to do the problem solving themselves whereas Poor solvers preferred to let the example solve the problem, so to speak.

V. THE CONTENT ANALYSIS

As discussed in the introduction, different models of analogical problem solving make different assumptions about the kinds of information that are transferred during analogy. The term "transfer" really means analogical inference, since the information may be modified as it is imported from the example to the problem. To put it more formally, after subjects have decided to refer to the example, what conclusions do they draw about the problem that are warranted by the existence of a similar conclusion in the example's solution and not by some general principle of physics?

Type of Information Transferred

In order to see the kinds of analogical inferences the subjects made, each of the 119 transfer events was coded into one of the following categories:

- *Line*. The subject transferred a whole equation, vector, or diagram.
- *Part of a line*. The subject transferred a detail from a line, such as whether a projection function was sine or cosine, or whether a vector went up or down.
- *Search control*. Subjects made the decision on what steps to do by consulting the example and seeing what steps it did.
- *Checking*. Subjects decided whether their most recent action or decision was correct by consulting the example.
- *Failure*. The subject failed to find anything useful during this transfer event.

TABLE 5
What Information was Transferred

Type	Good	Poor	Total
Line	24 (52%)	40 (55%)	64 (54%)
Part of a line	2 (4%)	1 (1%)	3 (3%)
Search Control	1 (2%)	10 (14%)	11 (9%)
Checking	9 (20%)	9 (12%)	18 (15%)
Failure	7 (15%)	2 (3%)	9 (8%)
Miscellaneous	3 (7%)	11 (15%)	14 (12%)
Total	46	73	119

- *Miscellaneous.* A variety of analogical inferences that were made only once or twice, such as deciding that the example wasn't analogous to the problem after all.

Table 5 shows the distribution of transfer events across these categories. Although there is a statistically significant difference between the distributions of Good and Poor solvers ($\chi(5) = 10.0, p = .02$), it is neither large nor easily interpreted. The basic result is simply that most students, both Good and Poor, transferred whole lines from the example to the problem.

Many computational models of analogical problem solvers transfer whole solution plans. However, there was only one instance of this among the 119 transfer events. (It was included with the search control events in Table 3.) It seems that this kind of analogy, despite its popularity in AI, occurs infrequently among physics problem solvers.

Does Explanation Precede Transfer?

Most models of analogical problem solving transfer only a solution or a part of the solution, but some also transfer the derivation of the solution as well (e.g., the Carbonell, 1986, model of derivational analogy). A solver that understands how an example's result is derived can adapt it more intelligently to the problem. Thus, one would expect the Good solvers to use derivational analogy more frequently than non-derivational analogy, and Poor solvers should use non-derivational analogy more than derivational analogy. To check this prediction, each of the 64 transfer events where a line was transferred was analyzed to see if the student explained the example's line before transferring it. The resulting counts (see Table 6) show a small trend in the expected direction, but it did not reach significance ($\chi(1) = 5.0, p = .17$).

TABLE 6
Was An Example Line Explained?

	Good	Poor	Total
Explained	7 (29%)	6 (15%)	13 (20%)
Not explained	17 (71%)	34 (85%)	51 (80%)
Total	24	40	64

This result is consistent with those of Chi et al. (1989, p. 175). They coded the number of self-explanations of the example that occurred during the initial 3 isomorphic problems, and found that their Good solvers generated just as many explanations as their Poor solvers. However, they noted that almost all the self-explanations of the Good solvers were generated by one subject who had generated the fewest self-explanations during example-studying. This suggests that self-explaining the example during analogical problem solving is a successful means of learning, although not particularly common.

VI. THE REPETITION ANALYSIS

Several studies have often observed subjects doing analogy once or twice, then relying on their memories thereafter (Anderson, Farrell, & Saurers, 1985; Pirolli & Anderson, 1985). Many models of analogical problem solving (e.g., Elio & Scharf, 1990; Reimann & Schult, 1993; Reimann, Wichmann, & Schult, 1993) assume that repeated analogies lead to generalization (schema or case building) and subsequent use of memory rather than written examples. An analysis was conducted in order to find out if analogies were repeated or if students began to rely on their memories instead.

Was There a Good/Poor Difference in Repetition of Analogical References?

As mentioned earlier, the subjects solved 12 problems that were isomorphic to the examples. Each of the 3 examples had 4 isomorphic problems. In the preceding analysis, all coding was done on protocols from the 3 problems that were presented first. In order to study the subject's learning, it is necessary to examine protocols from the remaining 9 problems. For each of the 119 transfer events (which occurred during the first 3 problems), places in the subsequent protocols were located where that same subject might use analogy in the same way. That is, given a particular transfer event from the first isomorph of an example, places in each of the remaining 3 isomorphic problems' solutions were located where the subject was working on roughly the same goal. These episodes where a transfer event could be repeated were classified using the following categories:

1. The subjects repeated the transfer event; that is, they referred to the example, then took the same action or made the same decision they did before.
2. The subjects took the same action or decision that they would have taken had they repeated the transfer event, but they did so without referring explicitly to the example.
3. Although this episode is parallel to the one where the target transfer event occurred, the differences between the two problems made it unnecessary or impossible to do the transfer again.
4. The protocol made it impossible to determine whether the subject repeated the transfer event or not.

Although each episode was coded as above, for the purposes of analysis the categories were collapsed to indicated whether the analogy caused generalization or not. That is, a transfer event was categorized either as Unlearned, if the subject always referred back to

the example and never used memory, or as Learned, if the subject eventually stopped referring to the example and started to use memory. More specifically, for each of the 119 transfer events, the transfer was coded as Unlearned if on all the subsequent occasions where it could occur, the student either referred to the example again (category 1), avoided the necessity for the transfer (category 3) or did something uninterpretable (category 4). The transfer was coded as Learned if on all subsequent occasions where it could occur, the student either made the same inference as he or she did in the transfer event without referring to the example (category 2), avoided the necessity for the transfer (category 3) or did something we couldn't interpret (category 4). A transfer event was also coded as Learned if there were a few reoccurrences of the analogy (category 1) followed by consistent use of the inference without analogy (category 2). In short, a transfer event was coded as Unlearned if the student kept repeating it, and Learned if the student eventually stopped referring to the example but continued making the same inferences anyway.

Table 7 presents the results for all the transfer events that were coded as Learned or Unlearned. Events that were neither Learned nor Unlearned, according to the definitions above, are not included. There were no significant differences between Good and Poor solvers ($\chi(1) = 0.9, p = .34$). The main finding is that only half of the events were Learned, contrary to the findings of earlier studies. Although most transfer events had 4 occasions where they could occur, one for each of the 4 isomorphic problems, this may not have been enough repetition to adequately strengthen the students' memory.

Were Different Types of Analogical References Repeated More Than Others?

It may be that subjects had no difficulties recalling earlier episodes of analogy, but they decided nonetheless to refer to the example again. If this were the case, then we might see different patterns of Learned/Unlearned depending on the students' purpose in referring to the example. Table 8 breaks down the results by the 4 categories used for coding purposes earlier. There was a significant difference between categories ($\chi(3) = 14.1, p = .003$). When the students referred to the example originally in order to check a step that they made, then they seldom (19%) repeated that check on subsequent problems; otherwise, they often (60%) repeated the analogy. This makes intuitive sense. If a rule has been checked once, and either verified or corrected, then the student will probably have enough confidence in it that it will not be checked again.

If students referred to the example in order to repair an impasse, one would expect them to learn from the experience and avoid having the impasse in the future (Anderson, 1993;

TABLE 7
Was the Transferred Information Learned?

	<i>Good</i>	<i>Poor</i>	<i>Total</i>
Learned	17 (53%)	23 (50%)	40 (51%)
Unlearned	15 (47%)	23 (50%)	38 (49%)
Total	32	46	78

TABLE 8
Learning of Each Type of Transfer Event

<i>Purpose</i>	<i>Learned</i>	<i>Unlearned</i>	<i>Total</i>
Repairing an impasse	2	3	5
Following the example	11	16	27
Implementing a step	6	14	20
Checking a step	21	5	26
Total	40	38	78

VanLehn, 1986). Thus, we should see a higher rate of learning in the Repairing an Impasse category than the overall rate of learning. This did not occur.

If students referred to an example because they were following it verbatim, then one would expect a lower than normal rate of learning, because it would be difficult to memorize a whole example. This too did not occur.

VII. DISCUSSION

In this section, a summary of the findings is presented first, followed by two extensive discussions. The first concerns the causal impact of analogy on learning. The second presents a revised version of Cascade that is consistent with all the known findings.

A Summary of the Findings

One purpose of this study was to present a descriptive account of analogical problem solving strategies. Contrary to the impression fostered by the analogy literature, many types of analogy were observed even within this small sample of subjects.

Table 9 presents a taxonomy of the 148 observed analogy events (references to the example) according to their apparent purpose. They fell clearly into two classes. Initialization events are episodes when the student retrieves an example, compares it to the problem and decides whether it is worth using the example as a guide to solving the problem, whereas during transfer events, students actually made some inferences about the problem's solution on the basis of the analogy between the example and the problem. There were 29 initialization events and 119 transfer events, for an average of 1.2 initialization events and 5.0 transfer events per problem. The average number of initialization events per problem was higher than 1 because one of the problems had multiple parts, and some students treated each part as a separate problem requiring a separate initialization event in order to determine which example to use.

When initialization events occurred, they occurred at the beginnings of problems, whereas transfer events occurred later in the problem solving. The Good students were particularly likely to initialize an analogy after viewing just the diagram alone. However, 17% of the problems did not include an initialization event. In these cases, initialization was combined with the first transfer event.

From the transfer event analysis, there appear to be two major types of problem solving. Some students prefer to follow the example verbatim, copying each line of the example's

TABLE 9
A Taxonomy of Analogy Events According to Purpose

Initializing the analogy: Retrieval, mapping & deciding if the example is useful (29)
Transferring information from the example to the problem
Copying the example's solution verbatim, piece by piece (52)
Solving problems with occasional reference to the example
Referring to the example in order to generate the next action
At an impasse (14)
Not at an impasse (26)
Referring to the example in order to check the action just generated (21)
Miscellaneous (6)

solution and adapting it to the problem. Other students prefer to solve the problem using general knowledge, but refer to the example when they want a technical detail (e.g., equations), when their general knowledge fails them and they reach an impasse, or when they want to check a step after performing it. Both types of analogical problem solving tended to import whole solution lines from the example to the problem. Perhaps the most interesting finding was that Poor solvers usually (60%) prefer to follow the example, whereas Good solvers usually (76%) prefer to try to solve the problem themselves, referring to the example only occasionally. This finding is discussed at length in the next section.

Subjects not only had a diversity of purposes for referring to the example, the type of information they transferred also varied. Most of the time, the subjects transferred a whole equation, a vector, or diagram. It was also common to make a decision, based on the example, about whether an equation, vector, or diagram just written down in the problem was correct. Other smaller pieces of information were occasionally transferred. Although many models of analogical problem solving transfer an abstract solution plan from the example to the problem, it was remarkable that this occurred in only one analogy event. Perhaps because the students were allowed to refer to the examples at any time and the examples' solutions were isomorphic to the problems' solutions, the students may have felt no need to abstract a solution plan from the example.

Having done a transfer event once, there were usually 3 or more opportunities for the student to do it again on later problems. About half the time, students eventually stopped doing the analogy at these points, indicating that they had learned something. However, the other half of the time, students continued to refer to the example at all these opportunities. The type of transfer event did not seem to influence the learning rate much, except that students rarely repeated analogies that were used to check the correctness of their steps. Perhaps the students had no motivation to commit the example's information to memory because they were allowed to refer to the examples at any time. It may also be that the particular items that they transferred from the examples (equations, vectors and diagrams, mostly) were so detailed and easily confused that it would take more than 3 repetitions in order to generate a robust, reliable memory of the item.

Both Good and Poor solvers preferred to initialize the analogy without doing much elaboration of either the example or the problem. During a transfer event, they seldom explained part of the example before importing it to the problem. Although these observa-

tions suggest that most analogy was based on surface features, the problems were nearly isomorphic to the example, so deeper features may not have been needed anyway.

Why are Analogy Strategies Correlated with Learning?

As discussed earlier, the Chi et al. (1989) study found a three-way correlation between self-explanation, example-exercise analogy and learning. The next question to ask, of course, is which correlations indicate an underlying causation. Figure 1 shows the possibilities. Effective learners tend to self-explain the examples and refer less extensively to the examples during problem solving. Since self-explanation occurred before problem solving in this study, the three arrows show the only possible causal pathways.

Only one of the three correlations is currently thought to be based on direct causation. Self-explaining an example consists mostly of rederiving each line in the example (Chi & VanLehn, 1991; VanLehn et al., 1992), and it appears to be directly responsible for learning physics principles (VanLehn, in press). Thus, of the three lines of possible causation shown in Figure 1, only line C represents an empirically verified line of causation. Either line A or line B or both would be sufficient to produce the observed correlation between example using strategies and learning.

An intuitively plausible explanation of the 3-way causation is that if one studies the examples hard, then one doesn't need to refer to them as much when solving problems. In terms of Figure 1, line A is causal and line B is not.

One version of this hypothesis was embedded in Cascade. It assumed that there was only one type of analogical problem solving (transformational analogy), and that students would use it only when they reached an impasse. Thus, the more they learned during example studying, the fewer impasses they had during problem solving, so the fewer analogical references.

Another version of the hypothesis could be based on the assumption that the modal analogical problem solver underlies all references to analogy among both Good and Poor solvers. On this account, self-explanation caused the Good solvers to commit more of the example to memory. Hence, when Good subjects refer to an example, it usually is a mental reference rather than a physical one, whereas the Poor subjects must actually re-read the example because they didn't commit it to memory during example studying. This explains why Good solvers display fewer references to the examples.

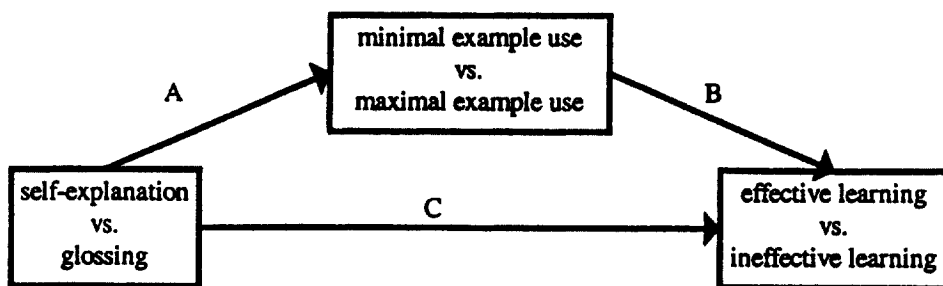


Figure 1. Possible causes of learning.

The common themes in these explanations are (a) there is just one kind of analogical problem solving, and (b) the Good solvers learned more during self-explanations and (c) that affected their analogical problem solving. Referring to Figure 1, let us call this the A-only hypothesis.

There are several problems with the A-only hypothesis. First, it predicts that Good solvers should have fewer references to the examples than Poor solvers. This frequency difference was only marginally significant in the Chi et al. study, and not present at all in this one. There *were* differences in the way Good and Poor solvers used analogy, but they were differences in *manner* rather than frequency. A second problem is that it is unlikely that even the most thorough self-explainers could remember all that much about the examples. The examples simply have too many physical, algebraic and geometric details. It would be difficult for them to be committed to memory during the few minutes that were spent on example studying. The Good solvers studied examples for an average of 13 minutes each, while the Poor solvers studied them for an average of 7.4 minutes each (Chi et al., 1989, p. 159). As shown in this study, students often continued to refer to the examples through 4 opportunities, which supports the idea that the information transferred from the examples was difficult to remember. A third problem is that the analysis of Table 3 indicates that Good solvers appeared to have different reasons for referring to the example than Poor solvers, which makes it unlikely that Good and Poor solvers are pursuing exactly the same basic approach to analogy with the Good solvers retrieving much of the example's information from memory while the Poor solvers read the examples instead.

Contrary to the A-only hypothesis, suppose that self-explanation does not directly cause differences in analogical problem solving, but differences in analogical problem solving do cause differences in learning. In particular, suppose there are at least two different methods for doing analogical problem solving, they cause different amounts of learning, and students can choose which one they want to use regardless of what they did during example studying. In terms of Figure 1, let us call this the B-only hypothesis

The B-only hypothesis predicts that because there are two strategies for studying examples, self-explanation and non-self-explanation, and there are at least two strategies for analogical problem solving, there ought to be some students with each of the possible combinations of the strategies. Moreover, if all combinations occurred, we could perhaps see whether there is any contribution of analogical problem solving strategies to learning that is independent of the contribution due to self-explanation. This can be tested statistically. From the protocol analyses of VanLehn and Jones (1993c), the number of references to the examples during all 19 problems of the study was obtained. From the analyses of Chi et al. (1989), the number of self-explanatory utterances and the scores were obtained. After partialling out the correlation between self-explanations and scores, the partial correlation between analogical references and scores was $-.77$, which was significantly different from zero ($p = .013$).

The logic of this analysis is best illustrated by considering three subjects, who uttered approximately the same number of self-explanations (18, 19 and 18), but had very different scores (8.0, 12.4 and 14.4, respectively). The number of times each student referred to the examples overall was 79, 42 and 18, respectively. The higher the number of analogical refer-

ences, the lower the score despite the fact that the number of self-explanations was approximately the same. This suggests not only that analogical problem solving strategies are independent of self-explanation strategies, but that minimizing use of the example during problem solving helps learning regardless of how much was learned during self-explanation.

The B-only hypothesis is supported by a simple causal mechanism, backed by yet another analysis of the Chi et al. (1989) data. VanLehn (in press) examined all places in the protocols where students could learn a physics principle because the solution demanded an application of that principle at that time. There were 73 learning opportunities during problem solving. Some students encountered the opportunities because they were trying to solve the problems, and often they would learn something from the opportunity. Other students avoided the opportunities because they were copying the example lines instead of deriving them, and thus missed many of the learning opportunities. Of the 73 learning opportunities, 52 were lost because students were copying lines instead of trying to derive them themselves. This analysis explains why students who preferred to copy example solutions learn less physics: They simply missed the opportunities for learning principles.

In short, the evidence suggests that it is not adequate to assume that there is just one type of analogical problem solving which is affected by self-explanation. The evidence suggests instead that there are multiple types of analogical problem solving and that they affect learning directly. In terms of Figure 1, we now have evidence for both B and C. We also have evidence that A and C without B are insufficient. Thus, whether self-explanation directly affects analogical problem solving is still an open question, but analogical problem solving does seem to directly affect learning. Needless to say, most of this evidence is indirect, so further experimentation is certainly warranted.

Toward a More Complete Model of Physics Learning

One purpose of a cognitive model is to organize evidence from diverse sources. Unfortunately, one must often mix well-supported hypotheses with mere conjectures in order to make the model computationally sufficient (VanLehn, Brown, & Greeno, 1984). In order to organize the evidence presented here, a cognitive model of physics learning will be presented even though it has considerably more speculation than I would like. It extends and integrates the two models presented earlier, Cascade and the modal model.

Four Basic Strategies

Let us begin by simplifying the vast complexity of the space of individual strategies. Let us assume that there are two basic strategies for studying examples, self-explanation and glossing, and that there are two basic strategies for solving problems, rule-based reasoning (RBR) and analogical or case-based reasoning (CBR). Students can freely choose which example lines to self-explain or gloss, and what parts of the problem to solve via RBR or CBR. We use Cascade's model of self-explanation, glossing and RBR, and the modal model's account for CBR.

Although the choice of problem solving strategy is computationally independent of the choice of example-studying strategy, it is likely that there is a mild correlation. Dweck

(Dweck, 1986; Dweck & Leggett, 1988) has found that students tend, in a given instructional situation, to adopt either a learning orientation or a performance orientation. Students with a learning orientation focus on finding out the underlying principles of the domain, whereas students with a performance orientation focus on generating error-free solutions. This orientation is not a permanent trait or aptitude of the student, because students will switch orientation when the rewards are changed to favor one strategy over the other. Nonetheless, in the Chi et al. (1989) study, it is likely that students who had a learning orientation during example studying, and thus chose the self-explanation strategy, probably retained a learning orientation during problem solving, and thus chose the RBR strategy. Student with a performance orientation probably chose glossing for example studying and CBR for problem solving. These considerations suggest that there will be a correlation between the choice of example studying and problem solving strategies. In general, strategy choice is a complex issue, especially in this study, where students were not given feedback on their errors and thus must assess their own performance.

Computational Constraints on Problem Solving Success

The task domain of physics is such that RBR always leads to correct solutions given that the solver has flawless knowledge of the principles (rules) of physics, but CBR leads to correct solutions only under certain circumstances. If the problem to be solved is isomorphic to the example (case), then CBR will generate a correct solution. If it is not, then CBR will sometimes lead to incorrect solutions depending on which components of the case are transferred. To discuss this, the components of an example's solution need to be described, and this is most easily done in the context of a specific example.

Consider example 3 in the Appendix, which concerns a pulley with two blocks suspended from it via a string. A plan for the solution of the first goal (finding the tension in the string), as generated by an ordinary goal-based reasoner, is shown in Figure 2. The goals that generate written actions are indicated with an asterisk, and the lines they generate are beneath them. Note that the sequence of writing actions does not correspond to those shown in the example. The example shown in the Appendix drew both free-body diagrams, then generated both applications of Newton's law, then solved both equations simultaneously. In order to generate a plan that would make the actions appear in that order, a partial-order planner (e.g., Russell & Norvig, 1995) would be needed. However, the basic approach in either case is to apply Newton's law to both masses, then solve the resulting equations.

This analysis suggests three types of information about the solution that CBR could transfer from the example to a problem. First, the reasoner could transfer any of the specific lines in the solution. (Let us restrict "line" to mean either an equation or a diagram, since students rarely write text in their solutions.) For instance, the reasoner could transfer the free-body diagram for mass m_1 or perhaps the final equation, $T = 2*m_1*m_2*g / (m_1 + m_2)$. Second, CBR could transfer the specific solution plan. For instance, it could transfer all the information shown in Figure 2 down to the goals marked with asterisks, but not including the written lines themselves (because they are included in the first type of information that

could be transferred by CBR). Third, CBR could transfer the abstract solution plan or basic approach, which for this example is simply to apply Newton's law to both blocks.

If the problem is even moderately similar to the example, and CBR transfers only the basic approach or the solution plan, then it is likely that the problem will be solved correctly by subsequent rule-based reasoning. However, if CBR transfers any of the lines, then the solution will be correct only if the example and the problem are nearly isomorphic. This cannot be proved in general, but it can be illustrated with one of the problems from the study. Figure 3 shows a problem that is not isomorphic to any of the examples in the study, but is superficially similar to both the inclined plane example and the pulley system example. Two subjects used RBR. They got correct answers. As far as one could tell from protocol data, no subject used CBR that transferred a basic approach or solution plan. Seven subjects used CBR that copied lines. They got incorrect answers. For instance, in order to find the tension in the string, 3 subjects copied the last equation of the pulley example, plugged in values for m_1 and m_2 , and thus got an incorrect answer. One subject copied the equation, but substituted $m_1 \sin(\theta)$ for m_1 in a vain attempt to compensate for the support given to the mass by the inclined plane. In order to find the acceleration of the two masses, most of the CBR subjects copied the equation $a = -g \sin(\theta)$ from the inclined plane example for the acceleration of m_1 and the equation $a = (m_1 - m_2)g / (m_1 + m_2)$ from the pulley example for the acceleration of m_2 . One subject averaged the two accelerations in order to find the single acceleration that the problem asks for, but the other subjects apparently

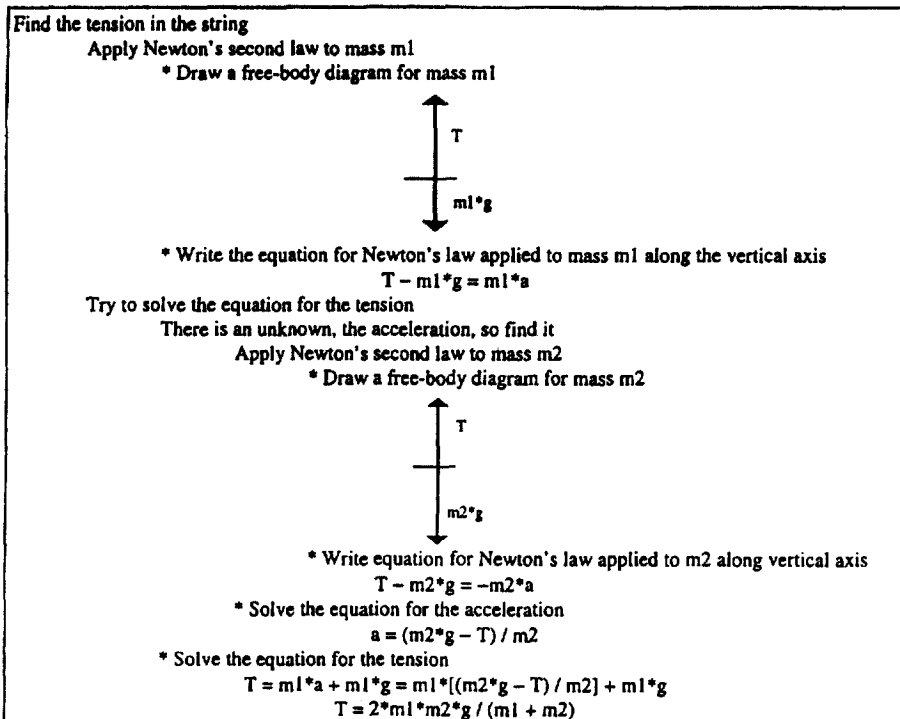


Figure 2. A plan for solving example 3's problem.

assumed that the accelerations of the blocks could be different. All these incorrect solutions were generated by transferring lines from the example to the problem.

There are computational reasons why CBR that transfers lines usually fails with physics problems. Unless the problem is isomorphic to the example, it will have different forces acting on the bodies. This means that the free-body diagrams must be different (not copies) and hence the equations must be different. Normally, CBR handles such differences by characterizing the difference in the problems, then applying that characterization to adapt the example's solution to become a hopefully correct solution to the problem. In order to do this properly for physics problems, an ideal CBR reasoner must store an extremely detailed account of the example's solution in memory, or it must rederive the example's solution in order to refresh its memory of the details. Subjects rarely bother with the details, so their CBR fails. This point is best demonstrated via an illustration.

The problem of Figure 3 differs from the problem of the pulley example in that m_1 is supported by the inclined plane while it is hanging freely in the example. An ideal CBR reasoner could characterize this difference as adding a new contact between bodies, and hence a new force. This characterization would allow it to draw a correct free-body diagram for m_1 (shown in Figure 3), provided that it use correct rules to identify the new force (a normal force, denoted N) and make it perpendicular to the plane. So far so good. Now the ideal CBR reasoner could compare this free-body diagram to the one for the example, and see that in the problem, a component of the gravitational force opposes the pull of tension instead of the whole gravitational force, as in the example's free-body diagram. With the difference thus characterized, the reasoner could adapt equations from the examples by

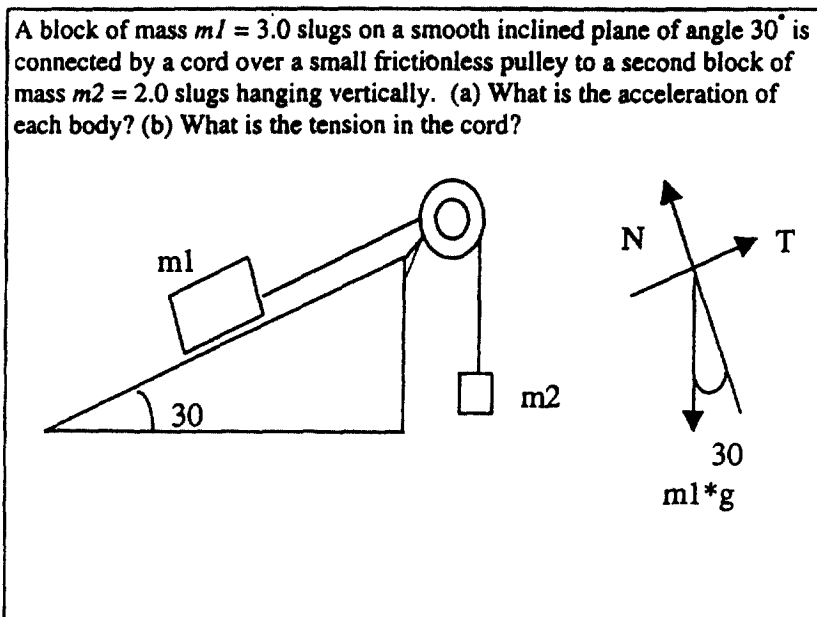


Figure 3. A problem similar but not isomorphic to two examples.

substituting the magnitude of the component ($m_1 * g * \sin(\theta)$) for the magnitude of the gravitational force ($m_1 * g$) in the example's equations. In fact, it could simplify the substitution by leaving g alone and just substituting $m_1 * \sin(\theta)$ for m_1 . This is just what one of the subjects tried to do. However, the equations in the example that mention m_1 have two occurrences of it. Only one comes from the magnitude of the gravitational force, $m_1 * g$. The other comes from the term $m_1 * a$ in Newton's law. If the ideal CBR reasoner knew which of the instances of m_1 to perform the substitution on, then it could successfully adapt any of the equations. To know this, it must either have in memory a whole derivation of the equation, or regenerate that derivation by self-explaining the example's solution in detail at the time it is used for CBR. For human CBR reasoners, remembering so much algebraic detail would be extremely difficult, and regenerating it via self-explanation saves no work compared to just solving the problem via RBR. Thus, the computational and memory requirements for proper use of CBR based on transferring lines are too high for human students, so they use incomplete or heuristic forms of CBR, and thus get the problems wrong.

The bottom line is that from a computational perspective, only 2 of the 3 general strategies for solving physics problems succeed. RBR and CBR with plans yield correct solutions when used with correct rule-based knowledge. Using CBR with lines requires such attention to detail that human students use incomplete or heuristic approximations, and thus generate incorrect solutions.

Instructional Constraints on What Knowledge Students are Likely to Possess

The preceding section argued that RBR and CBR with plans would succeed given that students had correct and complete rule-based knowledge. But how readily available is that knowledge? This section discusses the incompleteness of textbooks and how students can overcome it.

As demonstrated in earlier analyses of the physics textbook (VanLehn, in press), the authors leave out two key types of rule-based knowledge: "minor" physics principles and heuristics for choosing plans. The minor physics principles include rules such as, "A taut string exerts a tension force on the objects that it is tied to," and "If an object sliding along a surface has an acceleration, it is parallel to the surface." Although adapted from common sense and obvious to the expert physicist, students seem unwilling to assume that it is permissible to use such reasoning, or perhaps they are unaware that they even need to make such inferences. Nonetheless, some students do learn the minor physics principles by diligent self-explanation of the examples and other means (VanLehn, in press).

The other kind of rule that was missing from the textbook were heuristics for choosing plans, such as, "If the problem seeks the acceleration of an object, try to apply Newton's second law to it," and "If the problem involves several taut strings tied together at a knot, and then try applying Newton's law to the knot." Although students could learn these rules in the same way that they learned the minor physics rules (indeed, that is how Cascade learns them), it is unlikely that they do. During example studying, only one heuristic seems to have been learned, and only by at most 3 students. Three students were puzzled by example 1's choice of the knot as the body to apply Newton's law to, but ended up accept-

ing it without further explanation (VanLehn, in press). During problem solving, there were no apparent learning events for heuristics. Thus, it seems more likely that students are making most subsequent decisions about plans via CBR.

The bottom line is that textbooks often leave out certain types of rule-based knowledge. Minor physics principles can be learned during self-explanation of examples. Rules for choosing plans could in principle be learned, but it seems that students rarely do that spontaneously. Choice of plans is probably done by CBR rather than RBR even among students who routinely use RBR for everything else.

Behavior During Example Studying

So far, assumptions have been made about the kinds of strategies students can use, the success of those strategies during problem solving, and the availability of the knowledge required by the strategies for success. This section discusses what those assumptions imply for the behavior of students during example studying.

Students enter the example-studying phase with incomplete knowledge of physics rules. All students are missing some rules because they are simply not mentioned in the textbook. In addition, some students have forgotten (or never encoded) some of the rules that are mentioned in the textbook.

Students who choose the self-explanation strategy can and do fill some of the gaps in their knowledge. The learning processes are those used by Cascade, as well as some others observed in subsequent analyses (VanLehn, in press). Students who choose to gloss the examples retain their knowledge gaps.

Regardless of the strategy used to study the example, students acquire a rudimentary case for each example. That is, they can remember that the example exists, what its diagram looks like, and perhaps some other things as well. However, they cannot remember all the details of the solution, even if they self-explained it thoroughly.

Behavior During Problem Solving

When students solve problems, they can choose to use either CBR or RBR, regardless of how they choose to study the examples. We will consider first what happens as they solve problems that are isomorphic to the examples, then consider how their behavior changes as they solve subsequent non-isomorphic problems.

When students choose to use CBR with lines, they may think they need to transfer and adapt the earlier lines of the solution in order to successfully adapt the later lines of the solution. Thus, they would transfer many or all of the lines of the example. While this conservative strategy for CBR is generally a good idea, for isomorphic physics problems it is inefficient. For instance, substituting the problem's value for a mass into an earlier equation doesn't have any impact on the transfer of a later equation because the same substitution is done both times. Consequently, by the time the CBR-with-lines students are doing non-isomorphic problems, they are only transferring the lines they need, and those are typically the final equations in the examples. In the discussion of Figure 3, several illustrations of this kind of behavior were mentioned.

When students choose to use RBR, they rapidly discover that they must use CBR with plans in order to decide what plan to use. This may even occur without much awareness, since the planning decision for the isomorphic problems is a simple one (which body or bodies should Newton's law be applied to?) and the problems are completely isomorphic to the example. Thus, it is probably "obvious" to the students which bodies to apply Newton's law to.

Students' subsequent RBR processing goes smoothly if they have self-explained the examples and thus have relatively few gaps in their knowledge. If they were not so diligent when studying examples, then they will reach impasses as they solve even the isomorphic problems. However, because the example is available, they need not try to discover a rule for each impasse as they would have to do if they were self-explaining the example. They have the option of avoiding the opportunity to learn and simply transferring a line of the example instead. This will remove the impasse but leave the knowledge gap intact. Alternatively, they could take the impasse as evidence that they need to learn something, and use both their own reasoning and the example to ferret out the missing rule. As a result, the group of RBR-using students should gradually split into those who have filled their knowledge gaps during either example studying or problem solving, and those who preserved their knowledge gaps by relying on transferring lines from the example whenever they reached an impasse. However, both groups of students will use CBR for choosing plans.

When these two groups of RBR students reach the non-isomorphic problems, those with the knowledge gaps will have a great deal of trouble solving the problems. They will still rely too heavily on transferring lines from the examples. Indeed, their behavior may be indistinguishable from students who have used CBR-with-lines throughout the problem solving. The RBR students who have avoided using CBR with lines and filled most of their knowledge gaps will have little trouble successfully completing the non-isomorphic problems.

Final Outcomes

By the end of the problem solving, the only students who will have learned physics well enough to successfully solve non-isomorphic problems will be those that avoid using CBR-with-lines. They may or may not have self-explained the examples during example studying. If they did, they filled many of their knowledge gaps then. If they merely glossed the examples, they reacted to impasses during problem solving by filling their knowledge gaps then, often by referring back to the examples and self-explaining the relevant lines.

Students who did not self-explain the examples and did CBR with lines throughout the problem solving simply avoided all chances for discovering their knowledge gaps. Not only were they unable to solve the non-isomorphic problems via CBR, they were also unable to use RBR because they lacked some of the requisite rules.

Even students who self-explained the examples but decided nonetheless to use CBR-with-lines during problem solving did not learn physics well. During problem solving, they did not use the rules that they had acquired while studying the textbook and the example. Consequently, they probably forgot much of their rule-based physics knowledge. Thus, they could not solve the non-isomorphic problems with RBR even if they tried.

Consistency with the Main Findings

In this section, all the main findings from the studies of these subjects are described briefly and their consistency with the above account is discussed.

One central hypothesis is that Good solvers actually learned more than Poor solvers, and that their learning consisted of constructing new physics principles. This hypothesis is reflected in the revised model. The model assumes that although students acquire cases (memories of examples), and these cases are used for selecting a plan or basic approach to solving a problem, the cases are not detailed enough to support correct problem solving during the non-planning phases. In particular, correct adaptation of a case requires not only remembering the equations, but how each term in the equation was derived. According to the model, students do not retain such details in their cases, so when they try to use CBR they make errors. Students who rely on principles instead make fewer errors provided that they have learned those principles. Thus, students who have high scores must have learned many principles. Some findings supporting this hypothesis are:

1. The Good solvers were defined as subjects who scored above the median on the 19 problems (both isomorphic and non-isomorphic), and Poor solvers were defined as those scoring below the median. However, Good and Poor solvers were equivalent on several measures of prior knowledge (Chi & VanLehn, 1991), so the Good solvers learned more than the Poor solvers.
2. The components of knowledge gained by the Good solvers reflected the knowledge emphasized in the examples, so the Good solver's learning probably consisted of constructing new understandings of physics rather than increasing their access to physics knowledge presented in the textbook (Chi & VanLehn, 1991).
3. During problem solving, most of the errors generated by Poor solvers were due to their ignorance of physics principles, whereas most errors of the Good solvers were due to other sources (VanLehn & Jones, 1993a).

The revised model uses Cascade's model of self-explanation during example studying. In particular, self-explanation consists of rederiving example lines. Individuals have free choice on whether to self-explain or not. Some choose to self-explain more lines of the examples than others. Some findings supporting this claim are:

1. The Good solvers uttered more self-explanations than the Poor solvers (Chi et al., 1989; Chi & VanLehn, 1991).
2. The structure and content of the self-explanations of the Good solvers was similar to the structure and content of the self-explanations of the Poor solvers. The Good solvers simply uttered more self-explanations than the Poor solvers (Chi & VanLehn, 1991).
3. More than half the self-explanations mentioned knowledge that could not have been obtained from the textbook but could be obtained from the examples themselves combined with common sense reasoning (Chi & VanLehn, 1991).
4. The content of 63% of the subjects' self-explanations matched that of Cascade's. However, most of the self-explanations that did not match Cascade's concerned algebraic manipulation, which was not modeled by Cascade (VanLehn & Jones, 1993c).

Another hypothesis of the revised model inherited from Cascade is that rederivation reaches impasses whenever required knowledge is missing, and that impasses trigger some kind of reasoning that can result in constructing a principle. Poor solvers can learn the same way as good solvers, but because they choose to gloss the examples instead of rederiving each line from first principles, they lose the opportunity to learn new principles. Some findings supporting this hypothesis are:

1. During example studying, Good solvers mentioned impasses (comprehension failures) more often than Poor solvers, and more frequently followed these with self-explanations (Chi et al., 1989).
2. Of 70 opportunities to learn rules during example studying, students learned rules in 11 cases and lost the opportunity to learn in 59 cases by glossing the relevant lines in the example instead of self-explaining them (VanLehn, in press).
3. Almost all episodes wherein a student constructed a rule during either example studying or problem solving were triggered by impasses (VanLehn, in press).
4. When students constructed a rule, they sometimes (56%) reasoned shallowly from the textbook or the example, and sometimes (44%) used deep reasoning based on their existing physics and common sense knowledge. Deep reasoning produced correct rules more frequently than shallow reasoning. Good solvers more frequently used deep reasoning than Poor solvers (VanLehn, in press).

The revised model proposes that during problem solving, students have free choice between RBR and CBR. Some students may choose to use RBR on most parts of most problems, and some may choose to use mostly CBR. The revised version of Cascade should have a model of RBR and a model of CBR. Its model of RBR should be similar to Cascade's current model. In particular, when a student is reasoning with principles via RBR, the student can reach impasses and learn just as they did when self-explaining. If they choose to use mostly CBR, then they lose the opportunity to do such learning. The only modifications to Cascade required by the new data are that (a) the revised model should use CBR instead of analogical search control to plan the problem solving or at least to choose which bodies to analyze, and (b) it should use heuristics instead of analogical search control to make other search control decisions. Some findings supporting this part of the model are:

1. Approximately 90% of the subjects' inferences during problem solving were matched by Cascade (VanLehn & Jones, 1993c).
2. The search control choices made by the subjects matched those made by Cascade in 95.6% of the cases using just its default rule ordering. Adding analogical search control improved the match by only 2.1% (VanLehn & Jones, 1993c).
3. Of 73 opportunities to learn rules during problem solving, students learned rules in 14 cases and lost the opportunity to learn in 52 cases by copying the relevant lines from the example instead of trying to figure them out themselves (VanLehn, in press).

The revised version of Cascade needs significant changes in its model of analogical problem solving. It needs to have initialization events which are distinct from transfer events. It needs to have voluntary choice about whether to use analogy (CBR) or principles (RBR), instead of using CBR only when RBR reaches an impasse. It should have the abil-

ity to self-explain a line before transferring it, even though this option was not popular among the subjects. It should be able to transfer other kinds of results than lines, even though lines were the most popular type of content transferred by the subjects. It should have the ability to abstract the transferred information and store it in memory where it can be used the next time it is needed, thus removing the necessity of referring again to the example. These items would be linked with the students' understanding of the overall example, thus forming a case, albeit one with fragmentary information at varying levels of detail and generality. Some findings supporting this view are:

1. During solution of the initial 3 isomorphic problems, the Poor solvers tended to start reading the example at the beginning and read many more lines, whereas the Good solvers tended to start reading in the middle of the example and read only one or two lines (Chi et al., 1989).
2. The subjects' analogical retrievals were almost always correct, in that subjects referred to the example that was maximally similar to the problem. There was no statistical difference between the success rates of the Good and Poor solvers (VanLehn & Jones, 1993d).
3. The subjects' analogical mappings were almost always correct, in that subjects correctly paired objects in the example with objects in the problem. There was no statistical difference between the success rates of the Good and Poor solvers (VanLehn & Jones, 1993d).
4. In 14% of the cases where subjects referred to the example, they had earlier used problem solving knowledge that they could have used in this case. That is, in these cases they were clearly not forced to use analogy by their ignorance (VanLehn & Jones, 1993c).
5. When subjects referred to an example, they occasionally (20%) self-explained the line before transferring it. There was no Good/Poor difference in the amount of self-explanation (this paper; Chi et al., 1989).
6. Analogical problem solving consisted mostly of an initialization event followed by a few transfer events (this paper).
7. When solving isomorphic problems, Poor solvers tended to copy the examples verbatim, whereas Good solvers tried to do the problem themselves and referred to the example sparingly (this paper).
8. When students transferred information from the examples, it was generally an equation or a vector (this paper).
9. About half the time, both Good and Poor solvers would repeat an analogical reference at every opportunity instead of committing the transferred information to memory (this paper).

Perhaps the biggest missing piece in Cascade is some account of when and why students choose to use performance-oriented strategies (glossing, CBR with lines) as opposed to using learning-oriented strategies (self-explanation, RBR). In the revised version of Cascade described above, this choice was left as an unexplained parameter. One would test the model by fitting it to a subject's observed strategy choices, then seeing if the model predicts the subject's learning and errors. However, a much stronger, more explanatory model would predict the students' strategy choices. Perhaps the work of Dweck (1986; 1988),

Lepper (1990; 1993) and others could be woven into a motivational model of strategy choice.

Conclusions

This study began as an evaluation of Cascade's model of analogy. The model turned out to be inadequate in several ways, notably in its account of Poor learners' analogical behavior. In the process of evaluating Cascade, several new aspects of analogical problem solving were uncovered, such as the distinction between initialization events and transfer events, and the existence of two strategies for using analogy: as a supplement to RBR or as a replacement. These observations led to a proposal for an enhanced version of Cascade that drops analogical search control, adopts some aspects of CBR, and explicitly represents the Poor solvers' preferred CBR-with-lines strategy.

The next steps in this line of research are of course (a) to implement the revised model and (b) to run experiments to test the results discovered here via indirect methods.

Acknowledgments: The research was supported by the Cognitive Science Division of ONR under grant N00014-92-J-1945. Preparation of this document was completed while the author was a Fellow at the Center for Advanced Study in the Behavioral Sciences supported by Spencer Foundation Grant number 199400132. I am grateful to Rich Catrambone, Clayton Lewis, Peter Pirolli, and Chas Murray for helpful reviews.

REFERENCES

- Anderson, J. R. (1993). *Rules of the Mind*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Anderson, J. R., Farrell, R., & Saurers, R. (1985). Learning to program in Lisp. *Cognitive Science*, 8, 87-129.
- Bassok, M. (1990). Transfer of domain-specific problem-solving procedures. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 16(3), 522-533.
- Bassok, M., & Holyoak, K. J. (1989). Interdomain transfer between isomorphic topics in algebra and physics. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 15, 153-166.
- Brown, A. L., & Kane, M. J. (1988). Preschool children can learn to transfer: Learning to learn and learning from example. *Cognitive Psychology*, 20, 493-523.
- Brown, A. L., Kane, M. J., & Echols, C. H. (1986). Young children's mental models determine analogical transfer across problems with a common goal structure. *Cognitive Development*, 1, 103-121.
- Carbonell, J. G. (1986). Derivational analogy: A theory of reconstructive problem solving and expertise acquisition. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine Learning: An Artificial Intelligence Approach*, Vol. 2 (pp. 371-392). San Mateo, CA: Morgan Kaufman.
- Chi, M. T. H. (1997). Quantifying qualitative analyses of verbal data: A practical guide. *Journal of the Learning Sciences*, 6, 271-316.
- 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., & VanLehn, K. (1991). The content of physics self-explanations. *The Journal of the Learning Sciences*, 1, 69-105.
- Clement, C. A., & Gentner, D. (1991). Systematicity as a selection constraint in analogical mapping. *Cognitive Science*, 15, 89-132.
- Domeshek, E., & Kolodner, J. L. (1992). A case-based design aid for architecture. In J. Gero (Ed.), *Artificial Intelligence in Design 1992*. Boston, MA: Kluwer.
- Domeshek, E., & Kolodner, J. L. (1993). Finding the points of large cases. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 7(2), 87-96.
- Dweck, C. S. (1986). Motivational processes affecting learning. *American Psychologist*, 41, 1040-1048.

- Dweck, C. S., & Leggett, E. L. (1988). A social-cognitive approach to motivation and personality. *Psychological Review*, 95(2), 256-273.
- Elio, R., & Scharf, P. B. (1990). Modeling novice-to-expert shifts in problem-solving strategy and knowledge organization. *Cognitive Science*, 14, 579-639.
- Falkenhainer, B. (1990). A unified approach to explanation and theory formation. In J. Shrager & P. Langley (Eds.), *Computational Models of Scientific Discovery and Theory Formation* (pp. 157-196). San Mateo, CA: Morgan Kaufman.
- Faries, J. M., & Reiser, B. J. (1988). Access and use of previous solutions in a problem solving situation. In V. L. Patel & G. L. Groen (Eds.), *Proceedings of the Tenth Annual Conference of the Cognitive Science Society* (pp. 433-439). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Gentner, D., & Landers, R. (1985). Analogical reminding: A good match is hard to find. *Proceedings of the International Conference on Systems, Man and Cybernetics* (pp. 607-613). Tucson, AZ: IEEE.
- Gick, M. L. (1990). Transfer in insight problems: The effects of different types of similarity. In K. J. Gilhooly, M. T. G. Keane, R. H. Logie, & G. Erdos (Eds.), *Lines of Thinking*, Vol. 1. New York: Wiley.
- Gick, M. L., & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology*, 12, 306-355.
- Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, 15, 1-38.
- Hinricks, T. R. (1992). *Problem Solving in Open Worlds: A Case Study in Design*. Hillsdale, NJ: Erlbaum.
- Hinsley, D. A., Hayes, J. R., & Simon, H. A. (1977). From words to equations: Meaning and representation in algebra word problems. In P. Carpenter & M. A. Just (Eds.), *Cognitive Processes in Comprehension*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Holyoak, K. J. (1985). The pragmatics of analogical transfer. In G. H. Bower (Ed.), *The Psychology of Learning and Motivation*, Vol. 19 (pp. 59-87). New York: Academic Press.
- Jones, R. (1989). *A model of retrieval in problem solving*. Unpublished Doctoral Dissertation, University of California at Irvine, Irvine, CA.
- Kambhampati, S. (1993). Supporting flexible plan reuse. In S. Minton (Ed.), *Machine Learning Methods for Planning* (pp. 391-434). San Mateo, CA: Morgan Kaufmann.
- Kolodner, J. (1993). *Case-Based Reasoning*. San Mateo, CA: Morgan Kaufmann.
- Lepper, M. R., Aspinwall, L., Mumme, D., & Chabay, R. W. (1990). Self-Perception and Social Perception Processes in Tutoring: Subtle Control Strategies of Expert Tutors. In J. M. Olson & M. P. Zanna (Eds.), *Self Inference Processes: The Sixth Ontario Symposium in Social Psychology*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Lepper, M. R., Woolverton, M., Mumme, D. L., & Gurtner, J. L. (1993). Motivational techniques of expert human tutors: Lessons for the design of computer-based tutors. In S. P. LaJoie & S. J. Derry (Eds.), *Computers as Cognitive Tools* (pp. 75-105). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Novick, L. R. (1988). Analogical transfer, problem similarity and expertise. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 14, 510-520.
- Pirolli, P. L., & Anderson, J. R. (1985). The role of learning from examples in the acquisition of recursive programming skills. *Canadian Journal of Psychology*, 39, 240-272.
- Redmond, M. (1990). Distributed cases for case-based reasoning; Facilitating use of multiple cases, *Proceedings of AAAI-90*. Menlo Park, CA: AAAI Press.
- Reed, S. K. (1987). A structure-mapping model for word problems. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 13, 124-139.
- Reimann, P., & Schult, T. J. (1993). Understanding and using worked-out examples: A computational model. In G. Strube & K. F. Wender (Eds.), *The Cognitive Psychology of Knowledge*. Amsterdam: Elsevier.
- Reimann, P., Wichmann, S., & Schult, T. J. (1993). A learning strategy model for worked-out examples. In P. Brna, S. Ohlsson, & H. Pain (Eds.), *Artificial Intelligence in Education: Proceedings of AI-ED93* (pp. 290-297). Charlottesville, VA: Association of Advancement of Computing in Education.
- Ross, B. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 13, 629-639.
- Ross, B. H. (1989a). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 15(3), 456-468.
- Ross, B. H. (1989b). Reminders in learning and instruction. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning*. Cambridge: Cambridge University Press.
- Ross, B. H., & Kennedy, P. T. (1990). Generalizing from the use of earlier examples in problem solving. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 16(1), 42-55.
- Russell, S., & Norvig, P. (1995). *Artificial Intelligence: A Modern Approach*. Los Altos, CA: Morgan-Kaufman.

- VanLehn, K. (1986). Arithmetic procedures are induced from examples. In J. Hiebert (Ed.), *Conceptual and Procedural Knowledge: The case of mathematics*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- VanLehn, K. (in press). Rule learning events in the acquisition of a complex skill: An evaluation of Cascade. *The Journal of the Learning Sciences*.
- VanLehn, K., Brown, J. S., & Greeno, J. G. (1984). Competitive argumentation in computational theories of cognition. In W. Kintsch, J. Miller, & P. Polson (Eds.), *Methods and Tactics in Cognitive Science* (pp. 235–262). Hillsdale, NJ: Erlbaum.
- VanLehn, K., & Jones, R. M. (1993a). Better learners use analogical problem solving sparingly. In P. E. Utgoff (Ed.), *Machine Learning: Proceedings of the Tenth Annual Conference* (pp. 338–345). San Mateo, CA: Morgan Kaufmann.
- VanLehn, K., & Jones, R. M. (1993b). Integration of analogical search control and explanation-based learning of correctness. In S. Minton (Ed.), *Machine Learning Methods for Planning* (pp. 273–315). Los Altos, CA: Morgan Kaufmann.
- VanLehn, K., & Jones, R. M. (1993c). Learning by explaining examples to oneself: A computational model. In S. Chipman & A. Meyrowitz (Eds.), *Cognitive Models of Complex Learning* (pp. 25–82). Boston, MA: Kluwer Academic Publishers.
- VanLehn, K., & Jones, R. M. (1993d). What mediates the self-explanation effect? Knowledge gaps, schemas or analogies? In M. Polson (Ed.), *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society* (pp. 1034–1039). Hillsdale, NJ: Lawrence Erlbaum Associates.
- VanLehn, K., Jones, R. M., & Chi, M. T. H. (1992). A model of the self-explanation effect. *The Journal of the Learning Sciences*, 2(1), 1–59.
- Veloso, M. M., & Carbonell, J. G. (1993). Toward scaling up machine learning: A case study with derivational analogy in Prodigy. In S. Minton (Ed.), *Machine Learning Methods for Planning* (pp. 233–272). San Mateo, CA: Morgan Kaufmann.

APPENDIX: EXAMPLES AND PROBLEMS

Example 1

Problem:

The left side of Figure 4 shows an object of weight W hung by strings. Consider the knot at the junction of the three strings to be “the body.” The body remains at rest under the action of the 3 forces that are shown on the right of Figure 4. Suppose we are given the magnitude of one of the forces. How can we find the magnitude of the other forces?




Figure 4. Magnitude of forces.

Solution:

F_a , F_b and F_c are all the forces acting on the body. Since the body is unaccelerated, $F_a + F_b + F_c = 0$. Choosing the x- and y-axes as shown, we can write this vector equation as three scalar equations:

$$F_{ax} + F_{bx} = 0 \text{ and } F_{ay} + F_{by} + F_{cy} = 0$$

using Eq. 5-2. The third scalar equation for the z-axis is simply:

$$F_{az} = F_{bz} = F_{cz} = 0$$

that is, the vectors all lie in the x-y plane, so that they have no z-components. From the figure, we see that:

$$F_{ax} = -F_a \cos(30^\circ) = -0.866 F_a \text{ and } F_{ay} = F_a \sin(30^\circ) = 0.500 F_a, \text{ and}$$

$$F_{bx} = F_b \cos(45^\circ) = 0.707 F_b \text{ and } F_{by} = F_b \sin(45^\circ) = 0.707 F_b.$$

Also,

$$F_{cy} = -F_c = -W,$$

because the string C merely serves to transmit the force on one end to the junction at its other end. Substituting these results into our original equations, we obtain:

$$-0.866 F_a + 0.707 F_b = 0 \text{ and } 0.500 F_a + 0.707 F_b - W = 0.$$

If we are given the magnitude of any one of these three forces, we can solve these equations for the other two. For example, if $W=100$ N, we obtain $F_a = 73.3$ N and $F_b = 89.7$ N.

Example 2

We wish to analyze the motion of a block on a smooth incline.

(a) *Static case.* The left side of figure 5 shows a block of mass m kept at rest on a smooth plane, inclined at an angle θ with the horizontal, by means of a string attached to the verti-




Figure 5. Analyzing the motion of a block on a smooth incline.

cal wall. The forces acting on the block, which we choose as “the body,” are shown on the right side of the Figure 5. F is the force exerted on the block by the string; mg is the force exerted on the block by the earth, that is, its weight; N is the force exerted on the block by the inclined surface. N , called the normal force, is perpendicular (that is, normal) to the surface of contact because there is no frictional force between the surfaces. If there were a frictional force, N would have a component parallel to the incline. Because we wish to analyze the motion of the block, we choose ALL the forces acting ON the block. Note that the block will exert forces on other bodies in its environment (the string, the earth, the surface of the incline) in accordance with the action-reaction principle. These forces, however, are not needed to determine the motion of the block because they do not act on the block.

Suppose that θ and m are given. How do we find F and N ? Since the block is unaccelerated, we obtain

$$F + N + mg = 0.$$

It is convenient to choose the x -axis of our reference frame to be along the incline and the y -axis to be normal to the incline. With this choice of coordinates, only one force, mg , must be resolved into components in solving the problem. The two scalar equations obtained by resolving mg along the x - and y -axes are:

$$F - mg \sin(\theta) = 0 \text{ and } N - mg \cos(\theta) = 0$$

from which F and N can be obtained if θ and m are given. Note that these equations reduce to expected results for the special cases of $\theta=0^\circ$ and $\theta=90^\circ$.

(b) *Dynamic case.* Now suppose we cut the string. Then the force, F , the pull of the string on the block, will be removed. The resultant force on the block will no longer be zero, and the block will accelerate. What is its acceleration? From Eq. 5-2 we have:

$$F_x = ma_x, \text{ and } F_y = ma_y.$$

Using these relations we obtain:

$$N - mg \cos(\theta) = ma_y = 0 \text{ and } -mg \sin(\theta) = ma_x,$$

which yield:

$$a_y = 0 \text{ and } a_x = -g \sin(\theta).$$

The acceleration is directed down the incline with a magnitude of $g \sin(\theta)$.

(c) *Constant speed.* With what force would we need to pull on the string so that the block could move up the incline at constant speed? In this case, the block has zero acceleration, just as it did in the static situation discussed above. The analysis proceeds exactly as in part a, with the result that the required force has a magnitude:

$$F = mg \sin(\theta).$$

Notice that the same force would be required if the block were to slide down the incline with constant speed.

Example 3

Problem:

Consider two unequal masses connected by a massless string which passes over a frictionless and massless pulley, as shown on the left side of Figure 6. Let m_2 be greater than m_1 . Find the tension in the string and the acceleration of the masses.

Solution:

We consider an upward acceleration positive. If the acceleration of m_1 is a , the acceleration of m_2 must be $-a$. The forces acting on m_1 and m_2 are shown on the right. The equation of motion for m_1 is:

$$T - m_1g = m_1a$$

and for m_2 is

$$T - m_2g = -m_2a.$$

Combining these equations, we obtain:

$$a = \frac{m_2 - m_1}{m_1 + m_2}g$$

and:

$$T = \frac{2m_1m_2}{m_1 + m_2}g.$$




Figure 6. Determining tension and acceleration of masses in a string.




Figure 7. Determining the mass of a block when it is hanging from three strings.




Figure 8. Determining the mass and acceleration of a block resting on a frictionless inclined surface.

Problem 1

In Figure 7 a block is hanging from three strings. If the tension in string 1 is 18N , what is the mass of the block?

Problem 2

The block on Figure 8 is resting on a frictionless inclined ($\theta = 30^\circ$) surface, held still by a string.

- a. If the tension in the string is 5 N , what is the mass of the block?
- b. If the string were to suddenly break, what would the acceleration of the block be?




Figure 9. A massless and frictionless pulley connecting two masses with different weights.

- c. With what force would you need to pull the string so that the block would move down the incline with a constant speed?

Problem 3

The tension in the string pictured on Figure 9 is 50 N. The string runs over a massless and frictionless pulley, and connects two masses that have weights that are different from one another.

- a. If the mass of the larger object is 20 kg, what must the mass of the smaller object be?
b. With what acceleration are they moving?