Educational Implications of CELIA: Learning by Observing and Explaining

Michael Redmond
Computer Science
Rutgers University
Camden, NJ 08102
(609) 225-6122
E-mail: redmond@crab.rutgers.edu
keywords: learning; education; examples/cases

Abstract

CELIA is a computational model of how a novice student can quickly become competent at a procedural task through observing and understanding an expert's problem solving. This model was inspired by protocol studies, and implemented in a computer program. This model of a student's effective learning suggests some implications for teaching novices in a new domain. These may be relevant for both human teaching and intelligent tutoring. The implications include: encourage the student to predict, interactive step-by-step presentation of example steps, encourage self-explanation by the student, order example steps to match their logical order, give a variety of examples in early instruction, allow flexible interaction with the student, and present basic background concepts prior to examples. These implications represent hypotheses that follow from the learning model; they suggest further research.

Introduction

Redmond (1992, 1989) presented a computational model, CELIA, (Cases and Explanations in Learning; an Integrated Approach) showing how a student can profit from observing an expert solve example problems. CELIA was inspired and influenced by protocol studies of students in a trade school learning to be automobile mechanics (Lancaster and Kolodner 1987,1988). By explaining the significance of each of the instructor's actions and the relationships between the actions, a student is able to detect parts of his knowledge which are lacking and also retain the example as an understood *case* (Kolodner and Simpson 1984; Hammond 1986; Riesbeck and Schank 1989) for future use. CELIA was implemented as a computer program, showing how, through this type of learning experience, a student can rapidly becomes reasonably competent at the observed task. This working model of a student's effective learning suggests some prescriptive implications for teaching novices. These teaching implications follow from the completed model of learning, and represent hypotheses that can be tested with future research.

Overview of CELIA

An overview of CELIA is necessary for understanding its implications for teaching. For CELIA to get something out of observed problem solving, CELIA must make an active effort to understand. Then it can make use of what it understands in its later problem solving. If CELIA is unable to explain part of the problem solving it can ask the instructor questions during the presentation. The instructor can interrupt his presentation of his problem solving and respond with a hint. Learning can be even more effective if CELIA takes a more active role; it should do more than just explain the instructor's actions. As part of understanding, a student can set up expectations of what the instructor will do. As long as those expectations are met, the student has an adequate understanding. When the student's expectations fail, this failure indicates to the student that he must learn something. CELIA makes use of this strategy as well.

In addition, an effective student tries to explain the instructor's actions. This not only helps the student identify the need to learn, it gives the student an understanding of the expert's problem solving. Thus, CELIA tries to explain the instructor's actions.

Several failures can occur in understanding an expert's actions; these indicate the need to learn. Once the need to learn has been identified, CELIA must be able to recognize a useful method to try, and carry out that learning method. Given that a great number of things can be learned and that experience provides a rich set of methods and materials to use, the learning process must be controlled and biased. Having the understanding process drive learning is a first step in that direction.

Retaining Example. Explaining the instructor's problem solving also helps CELIA remember the example for future use. This remembered episode, or *case*, is very useful to a novice. It shows how a problem can be successfully solved. Case-based Reasoning (CBR) (Kolodner and Simpson 1984) is a method of using previous episodes to suggest solutions to new problems. CBR allows a reasoner to solve problems efficiently when previous similar experiences are available. A student can make use of remembered cases to become functional long before he has a deep understanding of the domain and task.

One of the things discussed in Redmond (1992) is that early in apprenticeship, acquisition of cases is a most powerful learning method. This holds for at least two reasons. First, the knowledge that a case provides is operational — it shows how successful problem solving can be done (rather than being disconnected expertise).

Second, a single new case provides a wealth of information that can provide much useful guidance, whereas a typical single piece of knowledge about a relevant device has limited and infrequent impact on problem solving. As discussed in Redmond (1990), CELIA stores cases in pieces, or *snippets* (Kolodner 1988). This allows the reasoner direct access to small fragments of cases rather than always having to wade through large monolithic cases.

CELIA's Model of a Student

This discussion sets the stage for a brief overview of the model (presented in greater detail in Redmond (1992)). Figure 1 shows our model of explanatory apprenticeship. The boxes show processes and subprocesses, with thin arrows showing their decomposition. The bold arrows show the sequence in which subprocesses are carried out. The learner carries out two processes — understanding the instructor's problem solving, and learning. The focus here is on the process involved in understanding a teacher's actions. Space considerations preclude discussion of the learning methods used once the student has identified the need to learn. The learning process is presented in Redmond (1992).

As can be seen in Figure 1, the understanding process is broken down into three subprocesses:

1. Predict - CELIA predicts the instructor's next reasoning goal, and how it will be carried out. This helps focus CELIA's understanding process.
2. Observe - CELIA observes the expert's actions, comparing to the prediction.
3. Explain - CELIA explains the expert's actions — including recognizing what type of goal the instructor is pursuing, organizing actions within the goal, recognizing what previous goal it follows from (*drawing connections*), and how it affects the problem solving context. These steps also combine to enable the student to create a new case.

Through this process, and with the instructor's help when necessary, a novice can come to understand the expert's problem solving — what the goals are, what goals follow from each other, etc.

Explanation There are several steps that need to be taken in order to understand the instructor's actions. We will discuss three:

1. clustering actions into goals.
2. drawing connections between goals.
3. updating understanding of problem solving context

Failures in these or other parts of understanding can drive CELIA's learning.

Clustering Actions. When the instructor is solving a problem, his actions combine to serve different goals. A cluster of actions may serve one goal. In order for CELIA to truly explain the instructor's actions, CELIA needs to be able to cluster those actions together and determine what goal they serve. It also needs to be able to explain the role of each action with respect to the pursuit of its goal. Failure to make inferences of these types helps CELIA recognize the need to learn.

Drawing Connections. During the course of observing an instructor solving a problem, the most important step for the student is drawing connections between the instructor's goals. A diagnosis has structure, and part of understanding is determining that structure. A major part of the structure is the connections between goals - what goal instance follows from what other goal instance. Although the expert must perform his problem solving steps in some sequential order, that order does not necessarily reflect the relationships between the actions and goals. The instructor's goals usually will not have a purely linear set of relationships. One goal can suggest more than one later goal. The instructor may return to follow up a goal after a different path has proven unfruitful. Thus, many circumstances exist in which the sequence of actions taken by the instructor does not match the underlying relationships between goals. Therefore, the student cannot merely infer the connections between goals by their temporal proximity.

Updating Context. Another type of inference needed for CELIA to explain the instructor's actions is updating CELIA's understanding of the problem solving context. As the instructor takes actions, the results of those actions affect what the instructor will do next. Therefore, inferring the effects of the instructor's actions helps CELIA explain the instructor's *later* actions.

For example, suppose CELIA knows that the instructor has tightened the spark plugs prior to the test of whether the car still stalls. That knowledge helps CELIA explain the instructor's interpretation that the spark plugs being loose was not the problem. CELIA, in following along with the instructor, can keep track of the problem solving context. To keep up and understand, CELIA needs to infer the effects of the instructor's actions on the problem solving context.

This understanding process drives learning. Through the understanding process, CELIA both comes to understand the expert's actions, and also detects the need to learn. An early learner will have many failures when trying to understand the expert's problem solving. When CELIA makes an incorrect prediction, cannot recognize the instructor's reasoning goal, cannot connect it to some previous goal, or is unfamiliar with something mentioned by the instructor, then it recognizes the need to learn. [1]

It would be meaningless to look at the implications for teaching of a model that does not achieve successful learning. To demonstrate that the model embodied in CELIA shows a powerful approach to learning, Figure 2 is repeated from Redmond (1992). CELIA was presented a sequence of 24 examples of expert problem solving involving 8 distinct faults, mostly in the auto fuel system. Ten random orders of the examples were pre-

[1] Figure 1 also shows the different kinds of knowledge used in understanding the expert's actions, including device knowledge, *process knowledge* (knowledge of goals involved in the task), and other cases. Through the learning process, the same types of knowledge can be acquired.

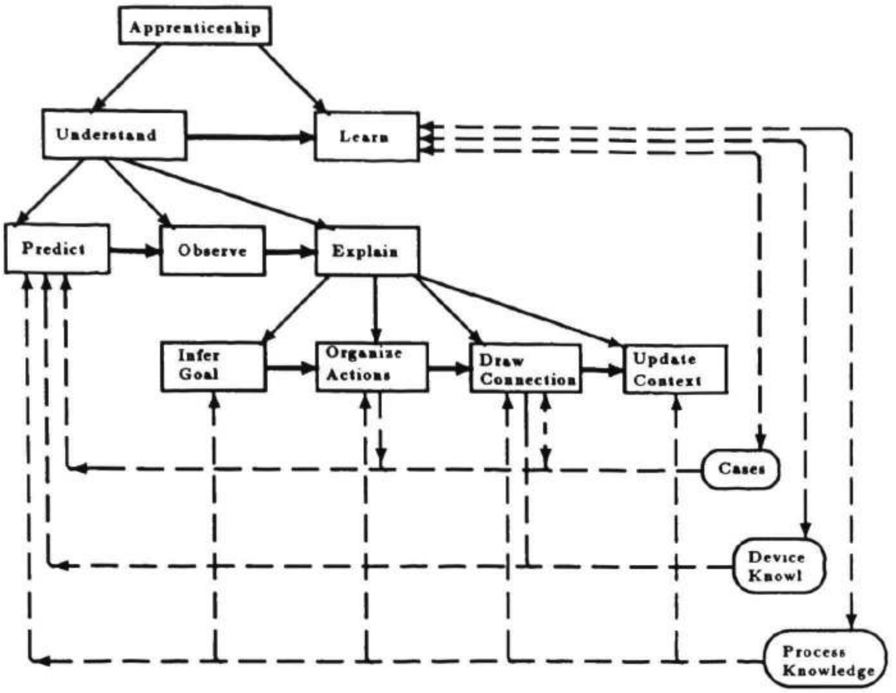

Shows the decomposition of the tasks involved in early apprenticeship. The rectangular boxes represent processes. The bold lines represent the sequence of tasks. Ovals represent knowledge structures. The dashed lines represent knowledge flow to and from tasks. The thin lines indicate the subtasks that a task is decomposed into.

Figure 1: Process model.

sented. The performance measure was the accuracy of the system's predictions of the expert's actions. The improvement comes through acquisition of new cases and improvements in the access of cases based on the experience. As can be seen, CELIA dramatically improves its performance after only a few examples have been seen.

Educational Implications

CELIA was developed as a computational model of how a novice student could effectively learn from observed examples. Besides modeling characteristics of learning, CELIA also has implications for teaching. Since the model focuses on how a student can learn from instruction (in the form of worked out examples), what it has to say about when learning is easy or hard can be equally revealing of what aspects will make teaching effective or ineffective. By analyzing the reasons for the success of CELIA, we have generated some prescriptive guidelines for tutoring.

CELIA is built around the observation that examples help students learn. This is not a unique observation. For instance, Reder, Charney and Morgan (1986) found that instruction that included examples was more effective, and LeFevre and Dixon (1986) found that students prefer examples to written text in learning a procedural task. Thus, examples should be presented to students. CELIA's implications for instruction mainly have to do with how the examples are presented:

1. The student should be encouraged to make predictions of the instructor's actions.
2. The examples should be presented in a slow-paced, interactive, step-by-step manner.
3. The student should be encouraged to try to explain how the example's steps follow from each other (*self-explanation* (Chi, Bassok, Lewis, Reimann, and Glaser 1989)).
4. Steps of examples should be presented such that the order of steps as closely as possible matches the logical order.
5. When possible, present an initial set of examples such that they cover a wide range of the problem-solving to be learned.
6. The instructor should be available for unplanned interaction with the student.
7. The student should receive some background instruction prior to presentation of examples.

Encourage Prediction. The model suggests that the student benefits from actively processing the examples. This follows from a tradition of failure-driven learning in AI (Schank 1982; Kolodner 1987; Hammond 1986). The model suggests that when tutoring, the student should be encouraged to set up predictions. When the prediction fails then the student knows that he needs to learn.

Predictions of an expert's actions helps a student identify their weaknesses and gaps in their knowledge. When a student incorrectly predicts the expert's actions, that

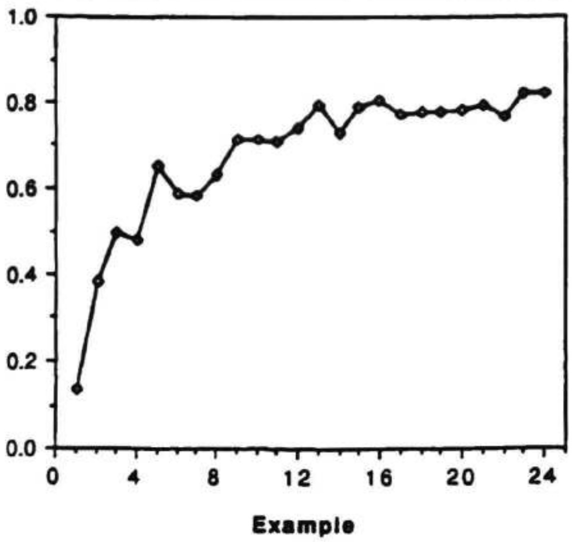

Improvement in accuracy for reasoner predicting observed expert problem solving actions.

Figure 2: CELIA: Results of exposure to sequence of examples.

very directly shows the student that he would have done something incorrect if he were on his own.

Incorrect predictions of an expert's actions help a student recognize the need to learn. Prediction is useful in detecting failures even when the student is doing self-explanation; it provides an extra level of checking of one's knowledge. It is one thing to be able to understand why an observed action was done; it takes greater understanding to actually come up with the correct action on one's own. Prediction gives the student a test of their operational or procedural knowledge.

Thus, when tutoring, our model suggests that the student should be encouraged to actively process the example, including setting up predictions. This helps the student recognize the need to learn.

Interactive Step-by-Step Presentation. It is for the above reason that the examples should be presented in a slow-paced, interactive, step-by-step manner. This gives the student time to make predictions about what will happen next, and to recognize reasoning failures as they are happening, when they can be isolated to a single step. By presenting the solution steps one at a time, the student's failures are isolated. This helps blame assignment.

This suggestion is not without controversy. Anderson's group's work with their lisp tutor has suggested that immediate feedback may not provide a benefit (Schooler and Anderson 1990). By a method called model tracing (Anderson, Conrad, and Corbett 1989), the tutor determines immediately when a student has made an error (and why). Anderson *et al* (1989) found that students who did not receive immediate feedback took longer to solve problems, but did as well on a post-test. Corbett and Anderson (1990) likewise found no impact on a post test from immediate feedback. They believe that the students can learn equally from the successful completed solution regardless of the path they took to get there. This finding may not be applicable in all domains/tasks. In particular, in CELIA's domain, part of the solution and thus part of what needs to be learned is the most appropriate path to the diagnosis.

In addition, Corbett and Anderson (1990) found that the student's who didn't receive immediate feedback incorrectly believed that the material was easier and they had learned the material better than did students who received immediate feedback. Thus, besides having a more difficult time doing blame assignment (if they choose to do it), the students who do not receive immediate feedback will be less motivated to do self-explanation.

Encourage Self-Explanation. The student should also be encouraged to try to explain the steps taken by the expert to himself. The student should be asked to explain how the example's steps follow from each other. One difference between good students and poor students is that good students more frequently attempt to explain worked out examples to themself (Chi *et al* 1989). The model embodied in CELIA predicts this as well. By encouraging students to actively process examples it is more likely that they will recognize when they don't understand something. Thus, they will know to try to learn something.

Self-explanation (Chi, *et. al.* 1989) involves the student trying to explain how the current expert's action fits into the overall problem solving. Self-explanation helps a student in at least two ways. First, if the student is unable to explain the expert's action, that, like an incorrect prediction, shows the student that he doesn't know something. Chi *et al* (1989) found in studies of people that good students identified a lack of knowledge *more frequently* than did poor students, who remained unaware of what they didn't know. Prompting the student to do self-explanation increases the possibility that the student will recognize the gaps in their knowledge, and try to learn something. Second, *successful* explanation of the expert's actions allows the student to come to a better understanding of the expert's problem solving. This in turn may do three things. First, it may help the student remember this example for future use, through developing more or better cues or indices to the example. Also, if the student does recall the example, it is a more fully understood example being recalled. In addition, this successful explanation process may help the student add to or correct previous knowledge about the domain and the task.

Thus, by encouraging students to actively process examples, carrying out self-explanation, the students can enhance their learning in two ways. It is more likely that they will recognize when they don't understand something. Thus, they will know to try to learn something. Also, self-explanation can help the student better understand the expert's actions.

Order Example Steps to Match Logical Order. CELIA has other implications for presentation of examples. In some examples, the order of problem solving steps may be somewhat flexible. One possible order may be easier for the student to understand. Our model suggests one factor that makes a solution easier or harder to understand – whether the temporal order of steps presented is the same as the logical order.

The logical order is the order in which steps are suggested to the problem solver. If step A suggests step C, then step C logically follows from step A. If the steps are partially ordered (the problem solving can be successful with more than one order of steps), than an instructor *may* present step B between steps A and C. In that instance, the logical and temporal order differ.

To ease the process of drawing the connections between goals, when possible the teacher should make the temporal order of the actions pursued correspond to the logical order of problem solving actions. (e.g. when proving a geometry theorem, in some instances it may be better to not do it from beginning to end, since the logical order may be to work both from the given and the conclusion to be proven.)

Presenting problem solving in its natural order preserves the underlying structure of the reasoning needed to solve the problem. When steps follow from recent steps, and from steps taken in the current direction of problem solving, the example is more coherent than otherwise. Such an ordering of steps makes the observed problem solving easier to explain. This implication follows from the process of explaining observed examples specified in the model. CELIA first tries to see if an action follows from the most recent previous action. If that isn't the case, then CELIA considers other possibilities. Thus some orders of steps in an example will be easier for CELIA to explain. When steps of possible solutions are partially ordered, the instructor should present the steps in the order that is the most coherent and the easiest to explain.

Variety of Early Examples. CELIA also has implications for the *order* of examples that should be presented to the student. The instructor may not always have a lot of control over this. However, when he does, the model suggests some advice. Early experience with a variety of different problems provides a good foundation for future learning. These provide a base of examples that can be remembered to make predictions, explain new examples, and to help distinguish what features are particularly salient indicators of the appropriateness of a case. The student will make errors during this early time but will have previous cases that would have suggested the correct action available that can aid learning.

This is a direct contrast from the prediction made by VanLehn's SIERRA (1987). In SIERRA, the student tries to generate a correct general procedure for solving a problem. This requires that each example differ from the procedures followed for previous examples by just a single difference (a single disjunct in the procedure). In CELIA, the learner is accumulating examples or cases for future use rather than inducing a general procedure. Thus, he wants a set of examples that cover a wider space, giving a starting point for more different future problems.

Interaction with Student. In teaching through apprenticeship, the teacher can do more than just present example solutions. The teacher should do other things to aid the student as well. The teacher should gives hints to help the student see the relationships between steps in the examples. Also, when an opportunity arises that could help the student learn something that would have avoided an earlier failure, the teacher should point out the opportunity. In addition, the teacher needs to be able to answer questions when the student makes an incorrect prediction or is unable to draw the connections between problem solving steps.

The model suggests that these are important learning methods. They particularly help CELIA acquire device knowledge, which in turn helps future prediction, drawing connections, learning distinguished indices and censors, even carrying out plausible inference to learn additional device knowledge. These are relatively easy learning methods, so they should be taken advantage of whenever possible.

Present Some Background Knowledge First. Lastly, the model has something to say about background knowledge needed. Though we suggest that exposure to examples allows a student to start becoming proficient without a complete knowledge of the relevant device(s), CELIA depends on the learner having *some* knowledge of the device. A student without any knowledge will be hopeless in trying to draw the connections between the instructor's goals and will be unable to ask focussed questions. Thus, the student needs to have some background knowledge. This contrasts with purely empirical models of learning. The instructor does not need to ensure that the student understands the domain completely before presenting examples. However, the most common basic concepts should be understood. Thus, the model suggests presenting the student with background information prior to working through examples. This can come from assigned reading or knowledge presented by a tutorial. If a domain is simple enough, the background knowledge may be common sense knowledge that the student already has.

Summary

CELIA is a computational model which shows how a novice can quickly become fairly competent by observing and understanding examples. In the model, the learner predicts the expert's actions, observes them, and then explains the actions and how they relate to each other. The explanation includes inferring what type of goal the instructor is pursuing, organizing his actions to determine which aspect of the goal they serve, drawing connections between the instructor's goals, and keeping track of the effects of the actions on the problem solving context.

This model suggests several implications for teaching novice students. The student should be encouraged to make predictions of the instructor's actions. The examples should be presented in a slow-paced, interactive, step-by-step manner. The student should be encouraged to try to explain how the example's steps follow from each other. Steps of examples should be presented such that the order of steps as closely as possible matches the logical order. When possible, present an initial set of examples such that they cover a wide range of the problem-solving to be learned. The instructor should be available for unplanned interaction with the student. The student should receive some background instruction prior to presentation of examples. These implications should be considered hypotheses that can be tested with further research.

Acknowledgements

This research has been supported by the Army Research Institute for the Behavioral and Social Sciences under Contract No. MDA-903-86-C-173, and Contract No. MDA-903-90-K-0112 and by DARPA contract F49620-88-C-0058 monitored by AFOSR. Thanks to Janet Kolodner for her advice and guidance and to Justin Peterson for helpful discussion. Thanks to Joel Martin for helpful comments on an earlier version of this paper.

References

[Anderson, Conrad, & Corbett1989] Anderson, J. R.; Conrad, F. G.; and Corbett, A. T. (1989). Skill acquisition and the lisp tutor. *Cognitive Science* 13:467–505.

[Chi *et al.*1989] Chi, M.; Bassok, M.; Lewis, M.; Reimann, P.; and Glaser, R. (1989). Self-explanations: How students study and use examples to solve problems. *Cognitive Science* 13:145–182.

[Corbett & Anderson1990] Corbett, A. T., and Anderson, J. R. (1990). The effect of feedback control on learning to program with the lisp tutor. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, 796–806. Cambridge, MA: Lawrence Erlbaum Associates.

[Hammond1986] Hammond, K. J. (1986). Learning to anticipate and avoid planning problems through the explanation of failures. In *Proceedings of the National Conference on Artificial Intelligence (AAAI-86)*. Philadelphia, PA: Morgan Kaufmann.

[Kolodner & Simpson1984] Kolodner, J., and Simpson, R. J. (1984). A case for case-based reasoning. Hillsdale, NJ: Lawrence Erlbaum Associates.

[Kolodner1987] Kolodner, J. (1987). Capitalizing on failure through case-based inference. In *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*. Seattle, WA: Lawrence Erlbaum Associates.

[Kolodner1988] Kolodner, J. (1988). Retrieving events from a case memory: a parallel implementation. In *Proceedings of a Workshop on Case-Based Reasoning*. Clearwater, FL: Morgan Kaufmann.

[Lancaster & Kolodner1987] Lancaster, J., and Kolodner, J. (1987). Problem solving in a natural task as a function of experience. In *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*. Seattle, WA: Lawrence Erlbaum Associates.

[Lancaster & Kolodner1988] Lancaster, J., and Kolodner, J. (1988). Varieties of learning from problem solving experience. In *Proceedings of the Tenth Annual Conference of the Cognitive Science Society*. Montreal, Canada: Lawrence Erlbaum Associates.

[LeFevre & Dixon1986] LeFevre, J., and Dixon, P. (1986). Do written instructions need examples?. *Cognition and Instruction* 3:1–30.

[Reder, Charney, & Morgan1986] Reder, L.; Charney, D.; and Morgan, K. (1986). The role of elaborations in learning a skill from an instructional text. *Memory and Cognition* 14:64–78.

[Redmond1989] Redmond, M. (1989). Combining explanation types for learning by understanding instructional examples. In *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society*. Ann Arbor, MI: Lawrence Erlbaum Associates.

[Redmond1990] Redmond, M. (1990). Distributed cases for case-based reasoning; facilitating use of multiple cases. In *Proceedings of the National Conference on Artificial Intelligence (AAAI-90)*. Boston, MA: Morgan Kaufmann.

[Redmond1992] Redmond, M. A. (1992). *Learning by Observing and Explaining Expert Problem Solving*. Ph.D. Dissertation, Georgia Institute of Technology, Atlanta, GA.

[Riesbeck & Schank1989] Riesbeck, C., and Schank, R. (1989). Inside case-based reasoning. Hillsdale, NJ: Lawrence Erlbaum.

[Schank1982] Schank, R. (1982). Dynamic memory: a theory of learning in computers and people. New York, NY: Cambridge University Press.

[Schooler & Anderson1990] Schooler, L. J., and Anderson, J. R. (1990). The disruptive potential of immediate feedback. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, 702–708. Cambridge, MA: Lawrence Erlbaum Associates.

[VanLehn1987] VanLehn, K. (1987). Learning one subprocedure per lesson. *Artificial Intelligence* 31:1–40.