Applying Cognitive Theory to Statistics Instruction

Marsha C. Lovett and Joel B. Greenhouse

Carnegie Mellon University

Abstract

This paper presents five principles of learning, derived from cognitive theory and supported by empirical results in cognitive psychology. To bridge the gap between theory and practice, each of these principles is transformed into a practical guideline and exemplified in a real teaching context. It is argued that this approach of putting cognitive theory into practice can offer several benefits to statistics education: a means for explaining and understanding why reform efforts work, a set of guidelines that can help instructors make well-informed design decisions when implementing these reforms, and a framework for generating new and effective instructional innovations.

Key Words: Statistical education, Instructional technique, Pedagogy

1. INTRODUCTION

At the heart of the reform movements in mathematics and science education lies the question: How can students' learning be improved? During the last twenty years, reform movements in math and science education have made new strides to address this question. For example, in elementary and secondary mathematics instruction, there has been a shift toward teaching mathematics in the context of real-world problems so that students can see its usefulness in concrete and familiar situations (NCTM 1989). In calculus instruction, there is a new (and still controversial) interest in offloading the burden of calculation to computational tools so that students can focus more on learning the conceptual issues at hand (Tucker and Leitzel 1995). And similarly, in physics education, qualitative understanding of fundamental principles has risen in importance relative to quantitative equation solving (McDermott 1984). A common goal of these educational reform movements is to proliferate new instructional techniques that will be effective in the classroom.

During the same time period, research in cognitive psychology has also addressed the question of how learning can be improved. Here, learning is defined in similar terms—as a change in knowledge and skills that enables new and different kinds of performance—but the question of how to improve it is addressed from a different perspective. In cognitive psychology, the focus is on understanding the basic phenomena related to learning, such as memory, skill acquisition, and problem solving. The goal is to acquire a body of empirical results which characterize these phenomena and to develop precise theories which make it possible to explain and predict observed data. In particular, one line of inquiry in psychology involves developing and testing *unified theories of cognition* that are implemented as computer models (cf. Newell 1982). One advantage of such theories is that they specify a set of general mechanisms designed to explain learning and performance across a broad range of situations. Another advantage is that their mechanisms are specified in a quantitative form such that predictions can be systematically derived, even for complex situations where the mechanisms may interact in complicated ways. Together, these features have enabled cognitive theories to explain and/or predict patterns of behavior in human learning and performance across a wide variety of tasks (Anderson and Lebiere 1998).

Unfortunately, these two fields—the "educational" and the "psychological"—are not often brought together so that one can inform the other. The former deals more with issues of practice (the "how") and the latter with issues of theory (the "why"). A reasonable characterization is that the fields are making progress on parallel paths. For instance, many of the instructional techniques that have blossomed in recent reform movements (e.g., active learning) are consistent with current psychological research but were not developed on that basis.

In this paper, we try to bring the two fields together in the context of statistics instruction. Specifically, we present five principles of learning that are derived from cognitive theory. These principles, stated briefly, are:

1. Students learn best what they practice and perform on their own.
2. Knowledge tends to be specific to the context in which it is learned.
3. Learning is more efficient when students receive real-time feedback on errors.
4. Learning involves integrating new knowledge with existing knowledge.
5. Learning becomes less efficient as the mental load students must carry increases.

We do not stop with this list of abstract principles, however. For each principle, we document its empirical support in the psychological literature and, based on the details of that research, generate a practical guideline for instructional design. In so doing, we establish a set of practical guidelines that corresponds to the set of theoretical principles. Moreover, we demonstrate how each of these guidelines can be applied to statistics instruction by giving examples from an introductory statistics course designed by one of us. Thus, the goal of this paper is two-fold: (1) to present a general, theoretical framework for understanding and predicting the effects of various instructional methods on student learning and (2) to demonstrate how that framework can be put to practical use in the context of statistics instruction.

The remainder of the paper is organized in five sections. Section 2 sketches current cognitive theory, highlighting several themes that arise in our five principles of learning. Section 3 describes the introductory statistics course at Carnegie Mellon, which serves as an illustrative example throughout the paper. Next, Section 4 discusses our five principles of learning in detail: Each principle is derived in terms of cognitive theory, supported by specific empirical studies, transformed into a practical guideline for instruction, and exemplified in the context of the Carnegie Mellon course. Section 5 then discusses some additional practical issues regarding how to apply the instructional guidelines in designing or re-designing a course. Finally, Section 6 sets our approach in the context of the statistics education reform movement more generally.

2. OVERVIEW OF COGNITIVE THEORY

Here, we provide a brief sketch of current cognitive theory as embodied in the ACT-R cognitive architecture (Anderson and Lebiere 1998). ACT-R is a unified theory of cognition that (1) posits a set of general mechanisms to explain learning and performance in almost any situation and (2) makes definite predictions that can be generated from an existing computer implementation, even for very complex tasks. This theory has shown its usefulness by accurately predicting human learning and performance data in a wide variety of tasks and experiments using the same set of general mechanisms (e.g., Lovett and Anderson 1996; Anderson and Matessa 1997; Anderson and Lebiere 1998). Moreover, the theory can capture individual differences in learning and performance that persist across different tasks (Lovett, Reder, and Lebiere in press).

ACT-R distinguishes between two types of knowledge: declarative knowledge for facts and procedural knowledge for skills. The difference between these two types of knowledge lies in how they are represented. Individual pieces of declarative knowledge are represented as nodes in a network of associations. Hence, the fact that "7 is the sum of 3 and 4" would be represented as a node with links to facts that tend to co-occur with this fact (e.g., nodes for the facts that 7, 3, and 4 are numbers and nodes for other, related arithmetic facts). In contrast, individual pieces of procedural knowledge are represented as "IF <goal> AND <condition> THEN <action>" rules that specify the actions to be taken under certain goals and conditions. Hence, the procedural knowledge for how to use addition facts could be represented by the statement "IF one's goal is to obtain the sum of *a* and *b* AND one can retrieve the fact that *c* is the sum of *a* and *b* THEN say *c*". This simplified example shows how declarative and procedural knowledge must work together to guide performance. For instance, the declarative knowledge that "7 is the sum of 3 and 4" is essentially inert until it is used by a piece of procedural knowledge. Suppose the current goal were to obtain the sum of 3 and 4. When this is the case, the procedural rule above would be instantiated with $a = 3$ and $b = 4$, which would lead to retrieval of the declarative fact that "7 is the sum of 3 and 4." Only if this fact is successfully retrieved can the procedure's action be taken, i.e., "say 7". Generalizing from this example, then, any cognitive ability is taken to be the result of a complex interaction among various pieces of declarative and procedural knowledge.

This example also highlights how the learner's current goal provides important contextual information to constrain which pieces of procedural and declarative knowledge will work together. First, only when the goal involves a sum is the procedural rule mentioned above deemed relevant, i.e., there is a match between the current goal and that rule's goal statement. Second, the "3" and "4" in the current goal help in binding the variables *a* and *b* in that procedural rule, which in turn constrains the to-be-retrieved addition fact to include these same numbers. Third, because the declarative nodes for "3" and "4" are included in the current goal and because these nodes are linked to certain other (related) nodes in declarative memory, this makes related facts (e.g., "7 is the sum of 3 and 4") more accessible. In sum, depending on the contents of the learner's current goal, different pieces of knowledge will tend to be retrieved and applied.

Each piece of knowledge—of either type—has associated with it a continuously valued quantity that reflects the "strength" of that knowledge. For declarative knowledge, strength reflects the overall accessibility of a particular fact; the stronger a declarative fact, the faster and more accurately an individual can access that knowledge. For procedural knowledge, strength reflects the likelihood that a particular skill will be used when its prerequisite goal and condition are satisfied; the stronger a procedural rule is, the more likely its action will be taken—instead of the actions corresponding to competing procedural rules, including inaction. For both types of knowledge, ACT-R specifies a functional form that relates the strength of knowledge to certain performance characteristics (e.g., time to access declarative knowledge, accuracy of declarative retrieval, and probability of applying procedural knowledge). Thus, given a description of an individual's state of knowledge, the theory makes specific quantitative predictions about how that individual will perform. Moreover, these relationships enable the theory to represent various differences among individual learners: learners can differ in terms of the presence or absence of different pieces of knowledge or in terms of the strength of those pieces of knowledge.

Given this representation of knowledge, learning has two parts: (1) updating existing knowledge and (2) acquiring new knowledge. In ACT-R, existing knowledge is updated via mechanisms that change the strength of knowledge based on its use. New knowledge is acquired via mechanisms that transform completed goals into more lasting representations. All of these learning mechanisms are defined using specific functions and algorithms in the computer implementation of the theory, but for this discussion, we describe them at a fairly abstract level.

According to ACT-R, the process of strengthening occurs continually through the natural course of activity. To be specific, any access of a piece of declarative knowledge (for whatever reason) leads to an increase in its declarative strength, thereby improving its subsequent accessibility. Each such increase is added to the declarative fact's current strength level (producing practice effects), but the size of this incremental increase decays as a function of time since the fact was accessed (producing forgetting effects). Thus, repeatedly solving the problem "What is the sum of 3 and 4?" will lead to many retrievals of the fact that "7 is the sum of 3 and 4." Each such retrieval adds to that fact's declarative strength, making it subsequently easier to retrieve. Note, however, that solving this problem multiple times does not affect the strength of other addition facts; the strengthening process is *specific* to the declarative node(s) used.

There is another kind of declarative strengthening process that is relevant here—the strengthening of the associative links between declarative nodes. The more often two nodes are used in the same context, the stronger the link between them becomes. This makes it easier to access one in the context of the other in future situations, i.e., the strength of one node gets propagated to the other as a function of the two nodes' link strength. For example, when the problem "what is the sum of 3 and 4" is repeatedly solved, in addition to strengthening the node "7 is the sum of 3 and 4", there is strengthening of the links between the "3" and "4" nodes and the "7 is the sum of 3 and 4" node. With these links strengthened, more of the strength associated with the "3" and "4" nodes can be propagated to the addition fact "7 is the sum of 3 and 4," making it easier to retrieve in future contexts that involve "3" and "4". (This may be beneficial when the future goal is again "what is the sum of 3 and 4", but it could lead to errors when the future goal is "what is the *product* of 3 and 4".) The theory constrains this propagation of strength along links, however, by positing that strength only propagates from nodes that are part of the goal and that these nodes must share a limited amount of strength. Thus, the more nodes that are part of the current goal (i.e., the more complex the current goal), the less strength each node has and hence the less strength it can propagate to associated nodes.

In the case of procedural knowledge, strengthening also occurs with each use of an individual piece of knowledge. Thus, if a learner performs some task that involves the application of a particular set of procedural rules, only these rules' strengths will be adjusted. The direction of the adjustment for procedural knowledge depends on whether the procedural knowledge led to a success or failure. All the pieces of procedural knowledge whose actions were taken preceding a success have their strength increased, and all the pieces of procedural knowledge whose actions were taken preceding a failure have their strength decreased. Thus, procedural knowledge that tends to lead to success will be stronger than procedural knowledge that tends to lead to failure, and hence the former will be more likely to be used in the future.

Besides updating existing knowledge, there is a need to learn new knowledge. The acquisition of new declarative and procedural knowledge occurs as a byproduct of completing the current goal—either successfully or unsuccessfully. That is, any time a new fact is generated or a new procedure is applied in service of completing the current goal, a corresponding new piece of knowledge is acquired. Moreover, that new piece of knowledge incorporates certain features that were part of the learner's representation of the current goal. For example, suppose a learner has the goal of finding the sum of 3 plus 4 but does not yet know any relevant addition facts. In this case, the procedural rule exemplified above (which involves retrieving the relevant fact) will not be helpful, so a different procedure will have to be used, e.g., counting on fingers. With this counting procedure, the learner can produce the result 7, thereby achieving the goal of finding the sum of 3 plus 4. As a byproduct of doing this, the learner has generated a "new" fact; this corresponds to the acquisition of a new node representing the fact that 7 is the sum of 3 plus 4. Moreover, this new node is linked to nodes for the numbers 3 and 4 because they were all used in the same context—achieving the goal to obtain the sum of 3 plus 4.

In this example, the knowledge generated was a new fact, so declarative knowledge was acquired. A similar process occurs for the acquisition of new procedural knowledge. Whenever the learner cannot use existing procedures to achieve a goal but is able to do so using other, general-purpose knowledge (e.g., reasoning from examples or from instructions), knowledge for the new procedure is acquired in the form of new rules. The new rules are created in such a way that (1) information about the current goal is incorporated in the new rules' goal statements and (2) the conditions under which the new procedure was applied are represented in the new rules' condition statements. These aspects of the new rules will constrain subsequent opportunities for re-use, i.e., the rules will be more likely to be applied in future situations with similar goals and conditions.

This brief discussion of how knowledge is strengthened and acquired has several reoccurring themes. First, there is the idea that cognitive abilities can be decomposed into separate pieces of declarative and procedural knowledge. Learning, then, must take place at the level of these individual pieces of knowledge. Second, existing knowledge is strengthened based on its use: Using a piece of knowledge will tend to make it stronger and hence easier to access in the future, a kind of "rich get richer" effect that does not generalize beyond the pieces of knowledge actually used. Third, the goal provides important contextual information for learning: the contents of the learner's current goal influences what knowledge gets used, acquired, and strengthened. These themes have particular implications for learning, which we will discuss in Section 4 when we show how current cognitive theory leads to the specification of our five principles of learning. Before considering these theoretical implications further, however, we describe a specific statistics course that provides a context for how the principles can be put into practice.

3. CASE STUDY: A STATISTICAL REASONING COURSE AT CARNEGIE MELLON

In this section, we describe an introductory statistical reasoning course at Carnegie Mellon University, designed by one of us in 1991. The design of the course was motivated by the specification of four fundamental course goals: that students learn to (1) apply the techniques of exploratory data analysis for data reduction and summary, (2) understand the concept of sampling variability, (3) understand and critically evaluate the effectiveness of different approaches for producing data, and (4) understand the use and interpretation of inferential statistical procedures. The curriculum focuses on the using and interpreting data-analysis techniques without teaching all the probabilistic or mathematical underpinnings. The idea is that before learning the quantitative aspects of statistical reasoning, students can build useful skills and intuitions and become interested in solving statistical problems. The course gives students the opportunity to engage in authentic statistical reasoning activities and, we hope, experience the excitement of scientific discovery.

These course goals were operationalized by considering the kind of data-analysis problems that students should be able to solve when they complete the course. In order to solve such problems, students must learn many different "pieces" of knowledge and integrate them in some unified whole. Thus, in the course, students learn and practice the component skills of data analysis, at first individual skills on simplified problems and then, later, combinations of skills on realistic problems of larger scope. Indeed, the format of the course was designed to teach these skills using several reform-based instructional techniques—collaborative learning, active learning, and the use of computers to aid students in the practice of statistics. In place of traditional recitation sections, weekly computer labs were established in which students work in pairs at a computer, using a commercially available statistics package (e.g., Data Desk 1992 or Minitab 1994) to complete assigned exercises. These exercises are presented to the students in the lab as a handout that describes a data set, provides detailed instructions that guide the students through the analysis, and asks students to interpret the results of their analysis. Students work with real data sets designed both to engage them in the analysis and to exemplify the general applicability of statistical methods. Students are rewarded for being innovative and for trying to learn how to solve new problems, and they are encouraged to collaborate in order to learn from each other. In addition to solving data-analysis problems during these computer lab sessions, students solve related problems on their own for their homework assignments.

This description of the course is consistent with several reform-movement ideas such as active learning, collaborative learning, and the effective use of computers in instruction, but it still leaves out many potentially important features of the overall course design: How are topics sequenced? How are the computer lab sessions integrated into the course, and how are they run? How are questions structured in lab exercises, homework assignments, and examinations (and how do these three formats overlap)? These “detail” questions are not always addressed in descriptions of innovative course designs. And yet, they are critical to instructors working to develop new materials because the success of an instructional innovation can be directly influenced by *how* such questions are answered (cf. Gage 1991; and Sections 4.1 and 4.5 below). Indeed, we believe that the success of the Carnegie Mellon course stems in large part from the details of its implementation. Thus, in the next section, we answer such “detail” questions by providing concrete examples from the Carnegie Mellon course but we do so only after presenting a relevant principle of learning and its corresponding practical guideline, stated in general terms. In this way, we wish to emphasize that the same basic approach of using principles of learning to guide instructional design can be applied to a variety of courses.

4. PUTTING THE PRINCIPLES INTO PRACTICE

To reiterate, the five principles of learning we discuss in this paper are:

1. Students learn best what they practice and perform on their own.
2. Knowledge tends to be specific to the context in which it is learned.
3. Learning is more efficient when students receive real-time feedback on errors.
4. Learning involves integrating new knowledge with existing knowledge.
5. Learning becomes less efficient as the mental load students must carry increases.

These principles likely seem reasonable and intuitive. Nevertheless, as given, they are still only abstract statements that do not offer a very clear picture of how best to design educational activities that will promote deep learning. Just like the ubiquitous labels "active learning," "collaborative learning," etc., the above statements hide a good deal of complexity and fail to offer concrete suggestions (or a means to discriminate between alternatives) for improving instruction.

Thus, in the subsections below, we elaborate on the five principles of learning in several ways. For each principle, we show how it can be derived from cognitive theory, and we document its empirical support. Then, based on this information, we generate one or more practical guidelines that instructors can use and flexibly apply in a variety of instructional situations. Finally, to concretize each guideline, we describe how a particular aspect of the Carnegie Mellon introductory statistics course exemplifies that guideline.

4.1 Principle 1

At first blush, this principle seems merely to mimic the aphorism "practice makes perfect." Recall, however, that the first two themes of the cognitive theory described above are (1) cognitive abilities can be decomposed into separate pieces of knowledge and (2) these pieces of knowledge are strengthened based on their use. These two ideas imply that the more learners engage in processing that requires them to access certain pieces of knowledge, the more they will learn those pieces of knowledge (and *not* other pieces of un-accessed knowledge). This leads to Principle 1: Students learn best what they practice and perform on their own.

Research on learning has supported this aspect of cognitive theory by showing that the benefits of practice are actually quite circumscribed. For example, take the skills involved in (i) writing new computer programs and (ii) evaluating existing computer code. These two sets of skills seem quite closely related—so much so, in fact, that it is reasonable to expect that practice on one will benefit both greatly. To the contrary, Kessler (1988) found that students who spent their time creating new computer programs did not improve much in their ability to evaluate programs and vice versa. The same asymmetry in learning was found when students practiced either translating calculus word problems into equations or solving the equations themselves; when students practiced one skill and then were tested on both, they showed great improvement on the practiced skill but no improvement on the other, related skill (Singley and Anderson 1989). Instructors may find such results familiar: Students perform well on homework problems but then poorly on test problems that seem (to the instructor) quite closely related. One possible explanation of such poor test performance is that the test problems actually require certain subskills that the students did not actually get to practice while solving the homework problems, i.e., students' lack of practice on these subskills may be impairing their overall performance. The learning process that underlies these results is quite straightforward: students improve greatly on the particular skills and subskills that they actually practice but improve only slightly (or not at all) on related skills that they do not practice. That is, the *content* of practice can greatly affect the degree to which the practice helps students achieve their specified learning goals.

All of this suggests that instructors need to pay special attention to exactly what concepts and skills students are practicing when they complete various assignments. Hence, we generate the following practical guideline corresponding to Principle 1: identify the skills and subskills students are supposed to learn, and then give students opportunities to perform and practice *all* of those skills. Note that this guideline implies that instructors should both identify a set of relevant skills and knowledge for their students to learn and then design instructional opportunities that allow students to practice this set. Garfield (1995) has similarly argued for a consistency in the learning activities students perform as part of coursework and for a specification of the learning outcomes that constitute the goals of the course.

In the Carnegie Mellon statistical reasoning course, students' opportunities for practice are carefully chosen for content, and practice is repeated throughout the semester. As mentioned earlier, in designing the curriculum for this course, a specific set of target skills was identified. Different activities were then designed to give students practice at those component skills. Specifically, during each week of the course, students practice applying a few new skills—first, under supervision in the computer labs and then without supervision on related homework assignments. This offers multiple opportunities to practice the same concepts or skills. Moreover, that practice is not limited to the current week's assignments. Throughout the semester, these basic concepts and skills are repeatedly practiced and integrated with other new skills. This way, students get ample practice on the skills that they will need to achieve the course goals. Moreover, they get practice at combining skills in different ways; this is important, too, because cognitive theory suggests that even more general skills that involve synthesis (e.g., comparing different approaches and managing long solutions) will not be learned unless they too are practiced.

4.2 Principle 2

The third theme of cognitive theory as described above is that the student's current goal provides important contextual information that influences what is learned. For example, when a declarative fact is retrieved in the context of a particular goal, not only is that fact strengthened but the links between that fact and other facts that describe the current goal are strengthened. This link strengthening makes the former fact subsequently easier to retrieve under similar, future goals. The context of the current goal is also incorporated into newly learned procedural knowledge. Here, a description of the current goal gets represented in the new rule's goal statement. Thus, the new rule will only be selected for application in future situations when the current goal is similar to the goal that was current during learning. All of this implies that knowledge will be more easily accessed in contexts that are similar to the student's learning context, which leads to Principle 2: Knowledge tends to be specific to the context in which it is learned.

This principle has received empirical support via the finding that students do not naturally generalize what they have learned to new situations. This oft-cited finding is called a "failure to transfer." It applies both when students fail to transfer what they have learned to new problems (e.g., problems that appear different from what the students have already encountered) and to new situations (e.g., out-of-the-class situations where students may not even consider applying their relevant skills). Many studies in the laboratory and in the classroom have shown that this difficulty of transferring what students have learned on one set of problems to another set of problems is great—even when the second set of problems is very similar to the first. For example, Reed, Dempster, and Ettinger (1985) studied college students in an algebra class. All students initially saw solutions to several different kinds of algebra word problems. Then students were asked to solve new problems that were either equivalent (identical except for different numbers) or similar (requiring slight adaptation of a previous solution). Four experiments showed that students had extreme difficulty solving even the equivalent problems—except when the previous solutions were made available during problem solving—and that students almost never solved the similar problems. These results are likely reminiscent of instructors' experiences with students who manage to solve a homework problem correctly but fail to apply the same skills to solve a closely related test problem. Earlier we suggested that such poor test performance might be attributed to the fact that students did not actually practice all the skills they need for the test problem. The current discussion suggests that even when students have practiced all the relevant skills, there may still be the question of whether students have learned the necessary skills at a sufficient level of generality to be able to apply them appropriately.

Several laboratory experiments have tried to gain insight into this learning problem by exploring what conditions actually do facilitate transfer. The basic approach aims to encourage students to learn new knowledge and skills in a general way such that they can apply the knowledge appropriately in a variety of situations. One intervention that has worked in several different domains involves giving students multiple problems that have related solution structures but that appear different. For example, Paas and Van Merrienboer (1991) gave students sets of geometry problems (with solutions) that were either similar or varied in appearance. (The fact that solutions accompanied the problems in this study relates to the issue of feedback, which we discuss in the next subsection.) Note that both types of problem sets exercised the same set of skills in order to control for the content of practice. After studying these problems, the students were tested on a new set of problems that were unfamiliar to all. The new problems were considered "transfer problems" because they forced students to use the target skills in new or more complicated ways. The students who had practiced with the "varied" problem set performed better on this transfer test than did students who had practiced with the "similar" problem set, even though they did not take significantly longer during the study phase. These and other related results suggest that students tend to learn more generally applicable knowledge and skills when the problems they encounter appear at least somewhat varied (Elio and Anderson 1984; Paas and Van Merrienboer 1991; Ranzijn 1991).

It is also worth noting that students who have been asked to compare and contrast problems that appear different because of their different cover stories also show more generalized learning. In a series of laboratory experiments, Cummins (1992) found that simply instructing algebra students to reflect on the similarities and differences between pairs of problems led them to transfer their algebra skills better than students without these instructions. Cummins's theory is that, by comparing problems, students end up reflecting on the deep (rather than superficial) relationships between problems and develop an a more generalized "schema" for how problems are solved.

Putting this principle into practice leads to the following guideline: give students problems that vary in appearance so their practice will involve applying knowledge and skills in a variety of ways. This not only provides more opportunities for practice but more opportunities of the kind that encourage students to generalize their understanding. This guideline is implemented in our introductory statistics course by virtue of the fact that students get to work on real-world problems that cover a variety of contexts, from the changes in infant mortality rate in nineteenth century Sweden to the efficacy of pharmacological interventions for the prevention of the recurrence of depression. That is, students get to apply the same statistical ideas to problems that are superficially very different. This technique can be used in other courses as well. For example, in a probability course for engineers, a colleague regularly gives students a homework assignment composed of ten problems with very different superficial features (e.g., a problem about solar flares, a problem about highway driving speeds). Unbeknownst to students, these problems all have the same solution structure. The assignment is to solve any three of the ten problems and then comment on the purpose of the assignment. Note that this problem-solving assignment (a) gives students multiple problems to solve, (b) makes those problems appear different even though they are similar in solution structure, and (c) encourages students to reflect on the problems' relationships.

4.3 Principle 3

Most of the learning mechanisms posited by the cognitive theory described above have the common feature that some kind of learning occurs regardless of whether the learner succeeded or failed in achieving the current goal. This suggests that, in many situations, learners will be strengthening incorrect knowledge, acquiring invalid procedures or strengthening inappropriate connections—in essence, "practicing bad habits". Thus, it is very important for learners to avoid errors or, if they cannot avoid errors, to compensate for the strengthening of incorrect knowledge with opportunities to practice the corresponding correct knowledge. This leads to Principle 3: Learning is more efficient when students receive real-time feedback on errors.

Experiments directed at this issue have manipulated the immediacy of feedback given to students as they practice solving problems. For example, in a study reported in Anderson, Conrad, and Corbett (1989), students learning to program in LISP were either (1) given feedback just after each mistake they made or (2) given an opportunity to request feedback at the end of each problem. This experiment and others like it have shown that immediate feedback, relative to delayed feedback, leads to significant reductions in the time taken for students to achieve a desired level of performance. Similarly, the sizable learning gains exhibited by students learning from human tutors is largely attributed to the rich feedback tutors can give (Bloom 1980). One-on-one tutoring, however, is not always possible. Instructors need other ways of giving students real-time feedback during problem solving and learning. This leads to the following somewhat modest guideline: try to "close the loop" as tightly as possible between students' thinking and the instructor's feedback.

This guideline is instantiated in our introductory statistics course's computer laboratories, where students are working on data-analysis exercises in pairs (so they can potentially provide feedback to each other) with teaching assistants circulating throughout the room to check on their work. In fact, each lab assignment includes several "checkpoints" where the students must contact a teaching assistant and demonstrate their understanding up to that point in the exercise. This process was motivated to give students more feedback *during* these supervised practice sessions. Students are now guaranteed to be "on track" at certain key points in the problem. Another approach that offers immediate feedback to students on their understanding is the "peer instruction" technique (Mazur 1997). In peer instruction, the instructor poses a question to the class, students discuss their answers in pairs, and then the instructor continues (e.g., by discussing common misconceptions and/or processes for generating a good answer). This peer instruction technique has been used effectively in classes with large numbers of students.

4.4 Principle 4

One of the important learning mechanisms posited by cognitive theory involves the strengthening of links in the network of declarative knowledge. Here the theory claims that links between pairs of nodes in the network are strengthened based on how often the learner accesses the corresponding pair of facts in the same context. These links are important for learning because the stronger a link between two facts, the more easily one fact can be retrieved in the context of the other, i.e., the more easily a student can make appropriate and useful associations between concepts and/or ideas. This implies that the entire network of associations held by a learner must be considered in making predictions about learning and performance: It is not just important for individual facts to be strengthened but for the appropriate connections between them to be strengthened as well. These issues regarding the importance of making proper links between pieces of knowledge lead to Principle 4: Learning involves integrating new knowledge with existing knowledge. This principle applies in two related situations: (1) integrating what students will learn in a course with what they already know and (2) integrating material that students will learn later in the course with material they learn early on in a course.

Students do not enter the classroom as blank slates; their prior knowledge can have an impact on their learning. Research on classroom learning shows that students often interpret technical terms loosely based on the way those terms are used in daily life. For example, the terms "speed" and "acceleration" are used quite loosely in daily life but require special interpretations in introductory physics class, a difference that can create an obstacle to learning (Reif and Allen 1992). More relevant to statistics instruction are examples involving students' prior conceptions and misconceptions of terms such as "chance" and "probability". Garfield and her colleagues have shown that these misconceptions can make learning certain concepts much more difficult (Garfield and delMas 1991). For example, if students can draw on pre-existing (but incorrect) knowledge when they must interpret and reason about probabilities, they may not see the need to acquire new knowledge. Even if they do eventually learn an appropriate definition for the term “probability”, this new knowledge will tend to be strongly linked to their old, inappropriate definition, making it difficult for them to consistently answer probability questions correctly. Therefore, it can be helpful for instructors to know about students' prior knowledge and conceptions. Then, they can build on the strengths of reasoning that students have in order to shore up their weaknesses. Consider the example of conditional probability problems. Although people’s intuitive reasoning on these problems often leads to errors when the problems are stated in terms of probabilities (Tversky and Kahneman 1982), controlled laboratory experiments have shown that people reason quite well when the same problems are stated in terms of frequencies (e.g., 850 cases out of 1000 instead of 85% probability, etc.; Gigerenzer and Hoffrage 1995). Thus, by using the frequency format as a starting point, students could link their new knowledge about probabilities to correct ideas about frequency.

The idea of helping students create appropriate links between pieces of knowledge appropriately also applies to the problem of presenting new material throughout a course. For example, students often view what they are learning as a set of isolated facts (Hammer 1994; Schoenfeld 1988) when, from the instructor’s point of view, there is a clear structure to the material being taught. Students may not see this structure unless it is explicitly presented. Moreover, controlled laboratory experiments have shown that students learn new verbal material better when it is presented in an organized structure (e.g., hierarchically as compared to merely in a list). Cognitive theory would explain this by positing that, in the hierarchical case, students acquire not only the individual words to be learned but also a set of links associating related words, which makes it easier to retrieve one in the context of the other. This same idea has been applied successfully in classroom situations—both when students are learning a new set of facts and when they are learning new problem-solving procedures. For example, Eylon and Reif (1984) presented different groups of students the same physics material on gravitational acceleration, but the information was organized in either a hierarchical fashion (with higher levels representing information most important to the task and lower levels representing the details) or in a linear fashion (an unorganized list of ideas). When students were asked to recall the material or to use the material to solve new problems, the hierarchical group outperformed the linear group (even when the linear group received more time to study the materials). Follow-up studies aimed at inferring students’ internal knowledge structures from patterns in the way they recalled the material revealed that the external organization of the presented material was reflected in students’ internal knowledge organizations.

Putting these ideas into practice leads to the following guideline: study students' relevant initial conceptions and misconceptions and then organize instruction so that students will build appropriate knowledge structures. This guideline is applied in the statistical reasoning course at several levels. At a global curriculum level, the Carnegie Mellon introductory sequence was revised to first teach students about describing and analyzing data, attempting to build upon their intuitions about describing data before teaching the underlying probability theory. At a more local level, the organization of material within the introductory course was very closely prescribed: students first learn about a new concept or procedure with a few examples; in these examples, the new concept is motivated in terms of other, closely related ideas that have been presented earlier in the course; then after some initial practice, students work on larger, more complex problems that require linking the new idea to more distantly related knowledge. In this way, students are encouraged to use related ideas on a common problem which should help them strengthen appropriate links in their knowledge structures.

4.5 Principle 5

To account for the fact that people have limitations on the amount of information they can attend to at once, the cognitive theory described above posits a constraint on the processing of goal-relevant information: There is a fixed amount of strength attributed to the goal, and this strength must be shared among its different components; thus, the more features required to describe the current goal, the weaker each corresponding fact becomes and hence the less strength that can be passed to linked nodes that may be relevant to completing the current goal. Thus, according to the theory's notion that strength impacts accessibility of information, when there are more elements in the current goal, it will be more difficult to access needed information. This leads to Principle 5: Learning becomes less efficient as the mental load students must carry increases. Here, "mental load" can be interpreted in terms of the amount of information needed to represent the current goal. If working on the current goal requires keeping track of a lot of information, each individual piece will have little strength and will not propagate much strength to goal-related facts that may be needed for its solution. On the other hand, if the current goal is fairly simple in that it requires keeping track of only a few pieces of information, each piece will have a lot of strength to propagate to its linked nodes, making these goal-related facts easier to access. Note that when students are learning to perform a new task, their current goal will tend to be fairly complicated: It must include a certain amount of information to represent the current task at hand (e.g., the features of the current problem being solved) as well as additional information regarding the students' parallel goal of learning something about that task. This suggests that learning should be viewed as a demanding task; it will proceed more effectively when the complexity of activities to be performed *during learning* is reduced.

Several researchers have studied how learning is affected by the combined mental load of performing an assigned activity while learning. For example, Sweller and his colleagues have found that students who are just beginning to learn to solve problems in a particular area can benefit from worked example problems (e.g., Sweller and Cooper 1985; Sweller 1988). The idea here is that seeing worked examples before solving new problems makes the subsequent problem solving an easier task. Note that this "examples-then-problems" activity also reduces the errors students make when solving problems on their own, so students' learning may also benefit from some of the issues related to immediate feedback (Principle 4.3).

Beyond the general advantage of initially giving students some worked examples to study, Ward and Sweller (1990) have revealed that different ways of formatting the worked examples can lead to more or less learning. Again, this difference in learning relates to differences in the mental effort students must expend in the different situations. In particular, Ward and Sweller found that the more information students must integrate on their own while processing the worked examples, the lower their learning outcomes, and the more the example's format places relevant information where it will be needed, the better the learning outcomes. For example, in one study introductory physics students either received worked examples with the solution equations incorporated into relevant parts of the problem statement or worked examples with the solution equations formatted separately from the problem description. Students in the former group (integrated problem text and equations) solved more test problems overall and performed better on transfer test items than did students in the latter group (separated problem text and equations). These results have been replicated with several different kinds of materials.

Putting this principle into practice leads to the following guideline: make the necessary information readily available to students during learning and offload extraneous processing during problem solving so that students can focus their attention on learning the material at hand. This guideline is implemented in our introductory statistics course in two ways. First, students are taught to use the computer to generate summary statistics and graphical displays when they are working on data-analysis exercises. This frees them from having to do detailed calculations (i.e., makes the task of data analysis easier) and hence allows students to focus their efforts on learning the larger task at hand. Second, students are given ample practice at applying new subskills in simpler contexts before they have to solve more complex problems. This way, by the time they encounter more complex and difficult problems, they are already comfortable at applying most of the subskills involved. They no longer need to labor over the basic steps in the problem solution but rather can focus on the larger problem of learning how to put those steps together.

5. BRINGING IT ALL TOGETHER

In the above subsections, we have described each of the five principles of learning in relative isolation (see Table 1 for a summary). However, in designing or re-designing a course, much more must be considered. In the subsections below, we address two such important issues. First, we acknowledge that there is more to instructional design than developing and implementing effective instructional techniques, the focus of this article. Second, even when one is focused on designing effective instructional techniques for a particular course, it is important to apply the five principles of learning *jointly*, so that they can mutually guide this part of the design process.

Table 1. Summary of Principles and Guidelines for Instructional Design

#	Principles of Learning	Guidelines for Instructional Design
1	Students learn best what they practice and perform on their own.	Identify the skills and subskills students are supposed to learn, and then give students opportunities to perform and practice *all* of those skills.
2	Knowledge tends to be specific to the context in which it is learned.	Give students problems with different contexts so they exercise what they have learned in a variety of ways.
3	Learning is more efficient when students receive real-time feedback as they solve problems.	Try to “close the loop” as tightly as possible between students’ thinking and the instructor’s feedback.
4	Learning involves integrating new knowledge with existing knowledge.	Study students' relevant initial conceptions and misconceptions and then sequence instruction to build on what students already know.
5	Students' learning becomes less efficient as the mental load they must carry increases.	Make the necessary information readily available to students during learning and offload extraneous processing during problem solving so that students can focus their attention on learning the material at hand.

5.1 A Schematic for Instructional Design

Educators have described the general process of instructional design in terms of multiple inter-related steps (e.g., Dick 1997):

> 1. Needs assessment: Identify course goals. These goals should depend on what students do not already know and what they need for the future, e.g., in downstream courses.
>
> 2. Operationalizing goals: Establish performance criteria for determining the degree to which instructional goals have been met. This is important for assessment (step 5), and it provides a specification of desired performance required for step 3.
>
> 3. Decomposing desired competency: Decompose the knowledge required for good performance into a set of learnable "pieces" of knowledge (e.g., facts, skills, concepts) that can become the basis of the curriculum and the course materials. (In educational circles, this step is often called "task analysis".)
>
> 4. Determining the instructional method: Consider what instructional method is most appropriate for the current situation. The strengths and weaknesses of candidate methods are weighed, and then a plan for implementation is devised.
>
> 5. Assessment of outcomes: Generate and administer assessments to evaluate the impact of the instructional design.

The first three steps of instructional design can be summarized as focusing on the *content* of what is to be taught. The decisions made here are very important in that they determine the objectives and curriculum of a course. In contrast, the fourth step, like this paper, focuses on the *techniques* of instruction—in what format will the course be taught, in what activities will students engage, etc. Design decisions regarding these issues of technique are also very important because they influence how well the chosen curriculum and course objectives will be executed. The fifth step of instructional design represents assessment of student outcomes. This step is important because it can provide important information regarding not only whether an instructional design was "successful" but what aspects of the design need to be improved. This last point highlights the fact that instructional design is an iterative process, with assessment feeding into a new cycle of change and refinement.

The current paper has focused on design decisions relating to the techniques of instruction (step 4). As the above schematic for instructional design suggests, these decisions are generally made after decisions regarding the content of instruction. Thus, we would argue that current cognitive theory can be a helpful guide in deciding *how* to teach a particular course. Deciding *what* to teach in a particular course, however, is a challenge that rests with the instructor; here, guidance may come from conventions in the field or from information on what students need to know for downstream coursework or on the job. As we will discuss further in Section 6, it is interesting to note that the statistics education reform movement has produced change in both the content and the techniques of statistics instruction.

5.2 Applying the Principles Together

In all of our discussions above, we have considered the five principles of learning—and their implications for instruction—separately. However, all of these principles must come into play simultaneously as students engage in learning. This implies that our principles and guidelines should be *jointly* applied to guide the design and development of effective learning activities. In this section, we offer a more fully described example from the

Carnegie Mellon introductory statistics course to show how multiple constraints can be combined in practice.

A challenging problem for beginning students doing data analysis is the matching of the appropriate methods to the given data. An approach that we have found to be successful is to train students from the outset of the course to (1) identify the variables in a problem as either response or explanatory, and then (2) to classify each variable as either quantitative or categorical. This approach yields a 2×2 table with data analytic methods in the cells of the table (e.g., a problem with a quantitative response variable and a single categorical explanatory variable would suggest using boxplots on the same scale for the exploratory analysis and ANOVA for the confirmatory analysis). In the Carnegie Mellon course, these skills have been broken down into separate subskills. Therefore, students get specific practice (in laboratory and homework assignments) that targets these skills. Consistent with Principle 1, this practice involves students making the identifications and classifications on their own and doing so repeatedly throughout the course. These are not skills students get to practice once or twice; a rough estimate is that students practice these subskills no less than one hundred times throughout the entire semester. Also, by virtue of the repeated opportunities to practice this skill, students get to see the skill applied in a variety of contexts. For example, each week's laboratory covers a new topic applied to a new data analysis problem (i.e., a new context) where these variable-classification skills are practiced. Thus, consistent with Principle 2, students' are more likely to see the general applicability of these skills because they are practicing the skills across a wide range of situations.

These practice opportunities initially arise in simple problems where students are explicitly told to perform these steps for each problem and are given hints regarding what issues to consider. Thus, consistent with Principle 3, students are given more feedback and support during initial learning of these skills. This "scaffolding" is gradually removed as the semester continues, with students eventually performing these skills in the context of much more difficult problems. Here, consistent with Principle 4, the curriculum builds on what students already know and have learned: Students learn these variable classification subskills early so they can later combine them with other skills to solve more complex problems (in the sense that a data set may contain several different response variables reflecting different questions of interest to be addressed). Note that by the time students reach more complex problems, they have already practiced these particular subskills enough to have become comfortable at applying them. In this way, the overall mental load of learning by solving complex problems is reduced because much of the problem solving students must do on complex problems is re-use of skills they have already learned. This is consistent with Principle 5, because each week students are usually only learning to work with one new concept or statistical technique.

6. DISCUSSION

In this paper, we have presented five principles of learning that can guide the design of effective statistics instruction. Each principle rests on psychological theory and research that, we hope, both confirms instructors' good intuitions and also highlights additional important, but perhaps less salient, features of learning. By explicitly stating these principles (and their corresponding practical guidelines), we have tried to offer a new framework for developing innovative instructional designs. We believe this kind of theory-driven approach can be advantageously applied in statistics education research. In this section, then, we provide a brief overview of the statistics education reform movement to place our approach in a more general context and to argue that our approach can complement what has already been done in this area.

6.1 Statistics Education Reform: Changes in Instructional Content

One of the major shifts produced by the statistics education reform movement involves the content of statistics courses—especially introductory level courses. Here, the trend is toward emphasizing students' practical use of statistical reasoning relative to their memorization of statistical formulas and procedures. This is similar to a recent trend in mathematics education more generally, where the content of instruction has come to focus more on real-world problems and the usefulness of mathematics in many familiar contexts.

In the case of statistics education, the emphasis on the "practice of statistics" can be seen through a number of different changes to course curricula. Course goals no longer refer to students' ability to derive particular statistical formulas or to compute certain statistics by hand, but rather they refer to "statistical literacy" and students' ability to reason statistically about real-world problems. For example, the course called CHANCE (Snell 1996) builds its entire curriculum around statistical problems that arise as current issues in the media. In addition, course curricula have been affected more generally by an array of new textbooks dedicated to the analysis and interpretation of real data sets and to the qualitative (vs. quantitative) presentation of statistical concepts (e.g., Cobb 1987; Moore 1997c; Moore and McCabe 1993; Rossman 1996; Utts 1996). Thus, in courses using these texts, problems typically require students to work with a statistical software package in order to answer questions regarding a given data set. In some cases, students are even responsible for collecting their own data. As an aid to instructors wanting to assign such real-world problems, several groups have made available databases of real datasets (e.g., DASL, JSE, EESEE) which are indexed by statistical concept and application area.

6.2 Statistics Education Reform: Changes in Instructional Technique

Beyond the shift towards more practice-oriented content, a major aim of the statistical education reform movement has been to improve the instructional techniques used in statistics courses. Table 2 lists some reform-based techniques with examples of their use. This list is not meant to be comprehensive but rather to demonstrate the variety of innovative instructional techniques being employed in statistics classes today.

Table 2. Annotated Overview of Educational Reforms

Reform-Based Technique	Examples of Use
Collaborative learning	• Students work together to solve problems (e.g., Borresen 1990; Dietz 1993) or discuss concepts, sharing ideas and understanding (Garfield 1993).
Active learning	• Students engage in data collection (e.g., Rossman 1996; Scheaffer et al. 1996; Spurrier et al. 1995), reflect on and explore statistical concepts (Lan et al. 1993), and solve problems on their own.
Target misconceptions	• Instruction is designed so that students will be confronted with their misconceptions and then have the opportunity to reflect and derive a more coherent conceptual understanding (Garfield and delMas 1991).
Use of technology	• Several textbooks (e.g., Rossman 1996) and multi-media resources (e.g., Velleman 1996) are designed to coordinate the presentation of new material with the use of statistical software. Also, simulation programs allow students to explore statistical concepts in discovery-world environments (e.g., Lang, Coyne, and Wackerly 1993; Loomis and Boynton 1992; Velleman 1996).

These methods tend to lead to improvements in students' interest in statistics, their learning outcomes, or both (e.g., Borresen 1990; Lan, Bradley, and Parr 1993; Stedman 1993; Cohen et al. 1996; Garfield 1996; Grabowski and Harkness 1996). To take a few examples, Lan et al. (1993) found that students who were encouraged to reflect on their learning (by recording time spent working on different concepts and estimating their own efficacy at solving problems using those concepts) scored higher on in-class examinations than did two different groups of "control" students. And, on the issue of collaborative learning, Borresen (1990) compared two groups of students in an introductory statistics class: those who worked on their assignments in small groups during class and those who worked individually on the same assignments for homework. In terms of a measure based on total points received in the course, Borresen found that the "small-group" students significantly outperformed the "individual-learning" students.

As the results above suggest, assessments of the outcomes of instructional reforms often show encouraging results. However, there are also caveats. For example, Cohen and his colleagues (e.g., Cohen, Tsai, and Chechile 1995; Cohen, Checile, Burns, and Tsai 1996; Cohen and Checile 1997) have completed several in-depth assessments of a hands-on curriculum in which students use an instructional software package to learn statistics. Although these students exhibited greater learning gains (posttest – pretest) than did "control" students, Cohen and Checile (1997, p. 110) remark that "even those students with adequate basic mathematical skills [who had used the hands-on instructional software] still scored only an average of 57% [correct] on the [post-] test of conceptual understanding." While this is a significant improvement relative to that group's average score of 42% correct at pretest, it shows that students still have a lot to learn.

Garfield and delMas (1991) have similarly found that, while innovative instructional interventions lead to gains in pre- to post-test performance, certain misconceptions held by students do not disappear, leaving absolute levels of performance after instruction unsatisfactory. In our own assessment of our introductory statistical reasoning course at Carnegie Mellon, we found significant higher learning gains among students who took the course relative to a group of control students. Nevertheless, the gain was attributable to large improvements on 75% of the assessment items and little or no improvement on the other 25% of the items. Assessment results such as these, derived from comprehensive study designs, attest to the inherent difficulty of statistical concepts and to the potential need for additional guidance in improving instruction.

6.3 How Can Statistics Education Benefit from a Cognitive Psychology Perspective?

The reform movement in statistics education has made substantial progress in changing the nature of instruction—both in its content and technique. Perhaps just as importantly, it has heightened the awareness among instructors of what statistics education is and what it can be. For example, various terms and ideas such as "active learning vs. passive learning", "hands-on practice", and "misconception-targeted instruction" are becoming part of the statistics educator's standard vocabulary. These ideas have helped instructors begin to analyze the relative merits of different types of instruction and have led to the development of many innovative, reform-based courses. Moreover, the reform movement has opened discussions among instructors and facilitated the sharing of information regarding "what worked" when a particular reform idea was implemented. In sum, there is a vast amount of research on improving statistics instruction, much of it producing very encouraging results.

There is still a call for more improvement, however, especially in helping students to grasp the broader issues of statistical reasoning (e.g., Moore 1997ab). In response to this call, some have argued that the next era of reform needs to focus on using technology more fully and on getting more instructors to embrace the techniques of the current reform movement (e.g., Garfield 1997; Hawkins 1997; Hoerl et al. 1997; Moore 1997ab). It is important to note, however, that both of these approaches require individual instructors to know a good deal about *how* technology should best be incorporated and *how* reform-based instructional techniques should best be implemented. Making all the design decisions required to answer these "how?" questions is far from trivial, even for an instructor familiar with current reform-based techniques.

This is where we believe a cognitive psychology perspective can be helpful. First, much of the existing statistics education work does not document the details of particular instructional reforms, so many of these design issues are left as "exercises for the reader" even though it may not be easy to decide on such implementation details. The principles of learning presented in this paper offer guidance in specifying these details in a way that will likely lead to effective instruction. Second, although there are many cases of "what worked" for a particular instructor implementing an educational reform, cases of "what did not work" are quite scarce. Without the possibility of comparing these two kinds of cases, it can be much more difficult to ascertain what aspects of a particular instructional innovation are critical to its success and therefore most important to replicate in re-implementing the innovation in other contexts. Because the principles of learning discussed in this paper are derived from cognitive theory, they aim to identify what is generally important for learning and thus can guide instructors in identifying the important components of a course design. This can be especially important in designing course activities and materials when there are no similar past cases from which to analogize. Third, new instructors who have not yet acquired a database of their own (or others') good teaching examples may have great difficulty in deciding how to teach a particular course. Such instructors would likely benefit greatly if they learned about a general framework that could guide them in making the many important design decisions inherent in teaching (e.g., what materials to use, what problems to assign, etc.). This framework could also help structure their knowledge as they gain expertise at teaching and gradually acquire their own database of teaching experiences. To take into account all these issues, we advocate a *combined* approach in which instructors use both "bottom-up" reasoning—drawing on their own experiences and other documented cases of good instructional designs—and "top-down" reasoning—using principles of learning derived from cognitive theory—to guide their design decisions in the way we have explicated here. By doing so, our approach can complement the existing research on statistics education and hopefully extend the progress that has been made already.

6.4 Conclusion

Instructors working in the field have a difficult job of implementing new instructional techniques. They often are responsible for all aspects of instructional design, development, and application. Most of the literature on educational reforms, however, describes only the basic properties of reform techniques and not the process of how to design a course. In practice, this leaves instructors with little guidance for implementing a new teaching method in their own courses, especially when their courses are different from the ones to which the method was originally applied. In contrast, the principles we laid out above aim to describe the processes of learning that pertain to all courses. In this way, the principles can guide the process of instructional design for a wide range of courses. Indeed, other colleagues at Carnegie Mellon have applied these ideas to a variety of courses with success.

Applying these principles of learning to course and curriculum design is akin to applying basic physical principles to various engineering designs. Creating designs with these principles in mind can increase the likelihood of developing a course that works. These principles also point to ways for analyzing what features of a course are more or less effective and for predicting how various new learning activities will impact students' learning. The essential idea is that students' processing (e.g., how students learn, what processes they engage while performing various learning activities) is an important part of what makes instruction effective. Focusing on student learning, therefore, can help us understand how best to improve it.

References

Anderson, J. R., Conrad, F. G., and Corbett, A. T. (1989), "Skill acquisition and the LISP tutor," *Cognitive Science,* 13*,* 467-505.

Anderson, J. R., and Lebiere, C. (1998), *Atomic Components of Thought*, Mahwah, NJ: Erlbaum.

Anderson, J. R., and Matessa, M. P. (1997), "A production system theory of serial memory," *Psychological Review,* 104*,* 728-748.

Bloom, B. S. (1984), "The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring," *Educational Researcher,* 13*,* 4-16.

Borresen, C. R. (1990), "Success in introductory statistics with small groups," *College Teaching,* 38, 26-28.

Cobb, G. W. (1987), "Introductory textbooks: A framework for evaluation," *Journal of the American Statistical Association,* 82*,* 321-339.

Cohen, S., and Chechile, R. A. (1997), "Overview of ConStats and the ConStats Assessment," in *Research on the Role of Technology in Teaching and Learning Statistics*, eds. J. B. Garfield and G. Burrill, Voorburg, The Netherlands: International Statistical Institute.

Cohen, S., Smith, G., Chechile, R. A., Burns, G. and Tsai, F. (1996), "Identifying impediments to learning probability and statistics from an assessment of instructional software," *Journal of Educational and Behavioral Statistics,* 21, 35-54.

Cohen, S., Tsai, F., and Chechile, R. (1995), "A model for assessing student interaction with educational software," *Behavior Research Methods, Instruments, and Computers,* 27*,* 251-256.

Cummins, D. D. (1992), "Role of analogical reasoning in induction of problem categories," *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 18, 1103-1138.

Data Desk (1993), Statistical Software (Student Version 4), Ithaca, NY: Data Description.

Dick, W. (1997), "A model for the systematic design of instruction," in *Instructional Design: International Perspectives, Volume 1. Theory, Research, and Models,* eds. R. D. Tennyson, F. Schott, N. Seel, and S. Dijkstra, Mahwah, NJ: Erlbaum, pp. 361-369.

Dietz, E. J. (1993), "A cooperative learning activity on methods of selecting a sample," *The American Statistician,* 47, 104-108.

Elio, R., Anderson, J. R. (1984), "The effects of information order and learning mode on schema abstraction," *Memory & Cognition,* 12, 20-30.

Eylon, B-S., and Reif, F. (1984), "Effects of knowledge organization on task performance," *Cognition and Instruction,* 1, 5-44.

Gage, N. L. (1991), "The obviousness of social and educational research results," *Educational Researcher,* 20, 10-16.

Garfield, J. (1993), "Teaching Statistics Using Small-Group Cooperative Learning," *Journal of Statistics Education,* 1, 1.

Garfield, J. (1995), "How students learn statistics," *International Statistical Review,* 63, 25-34.

Garfield, J. (1996), "Assessing student learning in the context of evaluating a Chance course," *Communications in Statistics—Theory and Method,* 25, 2863-2873.

Garfield, J. (1997), Discussion of "New pedagogy and new content: The case of statistics," by D. Moore, *International Statistical Review,* 65*,* 137-141.

Garfield, J., and delMas, R. (1991), "Students' conceptions of probability," in *Proceedings of the Third International Conference on Teaching Statistics* (Volume 1), ed. D. Vere-Jones, Voorburg, The Netherlands: International Statistical Institute, pp. 340-349.

Gigerenzer, G., and Hoffrage, U. (1995), "How to improve Bayesian reasoning without instructions: Frequency formats," *Psychological Review,* 102*,* 684-704.

Grabowski, B. L., and Harkness, W. L. (1996), "Enhancing Statistics Education with Expert Systems: More than an Advisory System," *Journal of Statistics Education,* 4, 3.

Hammer, D. (1994), "Epistemological beliefs in introductory physics," *Cognition and Instruction,* 12, 151-183.

Hawkins, A. (1997), Discussion of "New pedagogy and new content: The case of statistics," by D. Moore, *International Statistical Review,* 65, 141-146.

Hoerl, R., Hahn, G., and Doganaksoy, N. (1997), Discussion of "New pedagogy and new content: The case of statistics," by D. Moore, *International Statistical Review,* 65, 147-153.

Kessler, C. (1988), *Transfer of programming skills in novice LISP learners.* Doctoral dissertation. Carnegie Mellon.

Lan, W. Y., Bradley, L., and Parr, G. (1993), "The effects of a self-monitoring process on college students' learning in an introductory statistics course," *Journal of Experimental Education,* 62, 26-40.

Lang, J., Coyne, G., and Wackerly, D. (1993), *ExplorStat*, University of Florida.

Loomis, J., and Boynton, G. (1992), *Stat Visuals*, University of California at Santa Barbara.

Lovett, M. C., and Anderson, J. R. (1996), "History of success and current context in problem solving: Combined influences on operator selection," *Cognitive Psychology,* 31, 168-217.

Mazur, E. (1997), *Peer Instruction.* Upper Saddle River, NJ: Prentice Hall.

McDermott, L. C. (1984), "Research on conceptual understanding in mechanics," *Physics Today,* 37, 24-32.

Minitab (1996), Statistical Software (Student Version 9). Reading, MA: Addison-Wesley Publishing.

Moore, D. S., and McCabe, G. P. (1993), *Introduction to the Practice of Statistics*, New York: W. H. Freeman.

Moore, D. S. (1997a), "New pedagogy and new content: The case of statistics," *International Statistical Review,* 65, 123-165.

Moore, D. S. (1997b), Response to "New pedagogy and new content: The case of statistics," by D. Moore, *International Statistical Review,* 65, 162-165.

Moore, D. S. (1997c). *Statistics: Concepts and Controversies.* San Francisco: W. H. Freeman.

National Council of Teachers of Mathematics (1989), *Curriculum and Evaluation Standards For School Mathematics.* Reston, VA: NCTM.

Newell, A. (1982), *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.

Paas, F. G. W. C., and Van Merrienboer, J. J. G. (1991), "Variability of worked examples and transfer of geometrical problem-solving skills: A cognitive-load approach," *Journal of Educational Psychology,* 86, 122-133.

Ranzijn, F. J. A. (1991), "The number of video examples and the dispersion of examples as instructional design variables in teaching concepts," *Journal of Experimental Education,* 59, 320-330.

Reed, S. K., Dempster, A., and Ettinger, M. (1985), "Usefulness of analogous solutions for solving algebra word problems," *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 11, 106-125.

Reif, F. and Allen, S. (1992), "Cognition for interpreting scientific concepts: A study of acceleration," *Cognition and Instruction,* 9, 1-44.

Rossman, A. J. (1996), *Workshop Statistics: Discovery with Data*, New York: Springer.

Schoenfeld, A. H. (1988), "When good teaching leads to bad results: The disasters of 'well-taught' mathematics courses," *Educational Psychologist,* 23, 145-166.

Singley, M. K., and Anderson, J. R. (1989), *The Transfer of Cognitive Skill.* Cambridge, MA: Harvard University Press.

Snell, J. L. (1996), "Teaching a CHANCE course," *Communications in Statistics—Theory and Method,* 25, 2853-2862.

Stedman, M. E., (1993), "Statistical pedagogy: Employing student-generated data sets in introductory statistics," *Psychological Reports,* 72, 1036-1038.

Sweller, J. (1988), "Cognitive load during problem solving: Effects on learning," *Cognitive Science,* 12, 257-285.

Sweller, J., and Cooper, G. (1985), "The use of worked examples as a substitute for problem solving in algebra," *Cognition and Instruction,* 2, 59-89.

Tucker, A. and Leitzel, J. (1995), *Assessing Calculus Reform Efforts.* Washington, DC: Mathematical Association of America.

Tversky, A., and Kahneman, D. (1982), "Evidential impact of base rates, in *Judgment under uncertainty: Heuristics and biases,* eds. D. Kahneman, P. Slovic, and A. Tversky, Cambridge University Press.

Utts, J. M. (1996), *Seeing Through Statistics,* New York: Duxbury Press.

Velleman, P. (1996), *ActivStats*, Ithaca, NY: Data Description.

Ward, M., and Sweller, J. (1990), "Structuring effective worked examples," *Cognition and Instruction,* 7, 1-39.