Applying Cognitive Theory to Statistics Instruction

Marsha C. LOVETT and Joel B. GREENHOUSE

This article 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: Instructional technique; Pedagogy; Statistical education.

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 20 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 off-loading the burden of calculation to technology so that students can focus more on learning the conceptual issues at hand (e.g., Tucker and Leitzel 1995). Similarly, in physics education, qualitative understanding of fundamental principles has risen in importance relative to quantitative equation solving (e.g., McDermott 1984). A common goal of these educational reform movements is to promulgate 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—for example, memory, skill acquisition, and problem solving. The goal is to acquire a body of empirical results that characterizes these phenomena and to develop precise theories that explain and predict observed data. In particular, one line of inquiry involves developing and testing *unified theories of cognition* that are implemented as computer models (see Newell 1990). These theories specify a fixed set of general mechanisms designed to explain learning and performance across a broad range of situations. Moreover, they specify their mechanisms quantitatively so 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 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 primarily with issues of practice and the latter with issues of theory. 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., collaborative learning) are consistent with current cognitive psychological research but tend not to be developed directly from those results. Perhaps more importantly, the degree of separation between the fields has made it even more difficult for instructors to reap the benefits of the combined research. Educational research tends to emphasize what instructional methods work in a particular situation but not how instructors can generalize from that to make effective instructional decisions in their own classes. Psychological research tends to emphasize how learning proceeds in the abstract but not how specific instructional methods influence learning.

This article tries to bring the two fields together within the area of statistics instruction in a way that is of practical use to instructors. Specifically, we present five principles of learning, derived from cognitive theory and applied to education. 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.

Marsha C. Lovett is Research Scientist, Center for Innovation in Learning, Carnegie Mellon University, Pittsburgh, PA 15213 (E-mail: lovett@cmu.edu). Joel B. Greenhouse is Professor of Statistics and Associate Dean in the College of Humanities and Social Sciences, Carnegie Mellon University, Pittsburgh, PA 15213 (E-mail: joel@stat.cmu.edu). This work was partially supported by grants 97-20354 and 98-19950 from the National Science Foundation and by grant MH-30915 from the National Institutes of Health. The authors thank the associate editor, two anonymous reviewers, and David Moore for helpful suggestions and comments on this article.

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 only abstract statements that do not offer a clear picture of how best to design educational activities that will promote deep learning. To combat this, for each principle we document its empirical support in the psychological literature and, based on that research, generate a corresponding practical guideline. 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 the authors. Hence, the goal of this article 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 statistics instruction. The remainder of the article includes the following: a description of the introductory statistics course at Carnegie Mellon to which the above principles have been applied; a sketch of current cognitive theory and an elaboration of our five principles of learning; and a discussion of how our approach complements the achievements of the statistics education reform movement.

2. CASE STUDY: THE STATISTICAL REASONING COURSE AT CARNEGIE MELLON

In this section, we describe an introductory, one-semester statistical reasoning course for students in the humanities and social sciences 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 use and interpretation of 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 (Cobb 1992; Moore 1992).

The above course goals were operationalized by considering the kind of data-analysis problems that students should be able to solve when they complete the course. To solve such problems, students must learn many different “pieces” of knowledge and integrate them in some unified whole. 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. The format of the course was designed to teach these skills using several reform-based instructional techniques—collaborative learning, ac-

tive learning, and the use of computers to aid students in the practice of statistics. During weekly lab sessions, 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 a lab handout that describes a dataset, provides detailed instructions that guide the students through the analysis, and asks students to interpret the results of their analysis. Students work with real datasets designed both to engage them in the analysis and to exemplify the general applicability of statistical methods. Students are rewarded for trying to learn how to solve new problems, and they are encouraged to learn from each other. In addition to solving data-analysis problems in the computer labs, 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 problems chosen for lab exercises, homework assignments, and examinations? These “detail” questions are not always addressed in descriptions of innovative course designs. And yet, they are critical to instructors because the success of an instructional innovation can be directly influenced by exactly how such questions are answered (see Gage 1991; and Sections 3.1 and 3.5). Indeed, we believe that the success of the Carnegie Mellon course stems in large part from the details of its implementation. In the next section, then, we present our five learning principles in the context of this course to demonstrate how a theory-driven approach can help in answering the above kinds of questions. More generally, these principles offer a general framework that instructors can apply to a variety of courses when they need to make instructional decisions.

3. PRINCIPLES OF LEARNING

As mentioned in the introduction, current cognitive theories aim to explain and predict human learning and performance data by positing a fixed set of representations of knowledge and a fixed set of mechanisms for acquiring and using those knowledge representations (e.g., Anderson and Lebiere 1998; Holland, Holyoak, Nisbett, and Thagard 1986; Newell 1990). Although these theories are not often applied to classroom learning (Anderson, Conrad, and Corbett 1989 is one exception), their success at capturing data collected in a variety of laboratory tasks (Anderson and Matessa 1997; Lovett and Anderson 1996; Singley and Anderson 1989) suggests that they can offer important insights for understanding and predicting students’ learning outcomes.

Indeed, these theories share several claims about learning and information processing. The first such claim is that cognitive abilities can be decomposed into separate pieces of knowledge, each of which can be categorized as either

procedural (representing skills and procedures) or declarative (representing facts and ideas). For example, to answer the question “What is $13 + 48$?” one must access and use several pieces of knowledge. Assuming one used the strategy of long addition, these would include both procedural rules for adding the numbers in the one’s column, writing the sum mod 10, handling the carry, moving to the next column, and so on, and declarative facts for the various sums required by this problem. The second theoretical claim is that these knowledge pieces are acquired and strengthened according to their use. Although the different theories specify somewhat different mechanisms for knowledge acquisition and strengthening, they tend to share the notion that using a given piece of knowledge will make it stronger and hence easier to access in the future. Note that this “rich get richer” effect does not generalize beyond the pieces of knowledge actually used, and it applies equally to uses of knowledge that turn out to be “correct” or “incorrect.” The third theoretical claim is that goals set the context for learning. That is, depending on the learner’s current goal, different pieces of knowledge will be used, acquired, and strengthened. This dependence arises in part because procedural knowledge is represented in the form “IF my goal is $\langle x \rangle$ and $\langle \text{other conditions hold} \rangle$, then take action $\langle y \rangle$.” In addition, in some theories, the learner’s goal serves as a link among the various pieces of knowledge relevant to the task at hand. This linkage promotes associations among related pieces of knowledge so that when one piece is currently being used, other related pieces are made more available. The fourth theoretical claim states that an individual’s cognitive capacity is limited. Different theories represent this limitation in different ways, but the basic idea is the same: there is only so much information that a person can process at once. This limitation, in turn, influences how well people can learn and perform complex tasks. These shared theoretical claims form the foundation of current cognitive theories. They also lead to our five principles of learning, which will be elaborated in the following subsections.

3.1 Principle 1: Students Learn Best What They Practice and Perform on Their Own

At first blush, this principle seems merely to mimic the aphorism “practice makes perfect.” The first two themes of cognitive theory, as described above—(1) cognitive abilities can be decomposed into separate pieces of knowledge and (2) these pieces of knowledge are strengthened based on their use—suggest a more specific relationship between practice and learning. They imply 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.

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 (1) writing new computer programs and (2) 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 yet were pre/posttested 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 get to practice while solving the homework problems; that is, students’ lack of practice on these subskills may be impairing their overall performance. For example, in homework problems that involve computing and interpreting inferential statistics, students may either be explicitly told which statistics to compute or they may implicitly be cued by virtue of the current week’s topic. In this situation, the students will not have had opportunities to practice the skill of selecting the appropriate test statistics and therefore may have difficulty on an exam which requires this skill as part of a larger problem. 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) 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. For example, a commonly “missed” skill (or set of skills) is selecting the appropriate display or analysis. It is likely missed because this step is so automatic for experts;

it is little practiced because students are usually learning about only one analysis type at a time, so their choice of analysis can be made on trivial (nonconceptual) grounds. For these reasons, there are “synthesis” labs at regular intervals during the course where students work on a problem that requires several types of analyses drawn from the preceding weeks’ material. Another advantage of these particular labs is that they give students practice at combining skills in different ways; this is important, too, because cognitive theory suggests that the skills required for synthesis (e.g., comparing different approaches and managing long solutions) will not be learned unless they too are practiced.

3.2 Principle 2: Knowledge Tends to be Specific to the Context in Which it is Learned

The third theme of cognitive theory as described earlier 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, enabling the new rule to be used only in future situations where 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. Here, learning context typically refers to the type of problem the learner is trying to solve (i.e., learning tends to be tied to particular problems or problem types), but it may also refer to the physical environment.

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 in-

structors’ 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 Merriënboer (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 Merriënboer 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, and the instructor explicitly discusses these commonalities with students. 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 (1) gives students multiple problems to solve; (2) makes those problems appear different even though they are similar in solution structure; and (3) encourages students to reflect on the problems' relationships.

3.3 Principle 3: Learning is More Efficient When Students Receive Real-Time Feedback on Errors

Most of the learning mechanisms posited by cognitive theories 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.

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 1984). 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 cir-

culating 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.

3.4 Principle 4: Learning Involves Integrating New Knowledge With Existing Knowledge

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; that is, 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. Note that the importance of making proper links between pieces of knowledge 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 students' everyday interpretations of statistical terms such as “chance,” “probability,” “hypothesis,” and “variability.” Garfield and her colleagues have shown that students' misconceptions in these areas can make learning certain concepts much more difficult (Garfield and delMas 1991). For example, if students can make use of their pre-existing (but incorrect) knowledge about probabilities in statistics class, 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 interpret probabilities 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 1,000 instead of 85% probability, and so on; Gigerenzer and Hoffrage 1995). This suggests that, by using the frequency format as a starting point, students could link their new knowledge about probabilities to pre-existing *correct* intuitions 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).

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 Carnegie Mellon 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 exam-

ples, 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.

3.5 Principle 5: Learning Becomes Less Efficient as the Mental Load Students Must Carry Increases

To account for the fact that people have limitations on the amount of information they can attend to at once, cognitive theories posit a constraint on people's cognitive capacity. This can be specified in the following terms: the more complex the current goal (i.e., the more information simultaneously needed to solve it), the more difficult it will be to access that needed information. Here, "mental load" can be interpreted in terms of the amount of information simultaneously needed for solving the current goal. Note that when students are learning to perform a new task (e.g., solve addition problems), their current goal will tend to be fairly complicated: It must include a certain amount of information to represent the details of that new task (e.g., that the current problem is $34 + 81$ and the one's column has been summed) as well as additional information regarding the students' parallel goal of learning something about that task (e.g., that "carries" require a special procedure that should be remembered for the future). This suggests that learning should be viewed as a mentally 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 3, p. 5).

Beyond the general advantage of initially giving students some worked examples to study, Ward and Sweller (1990) 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

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. Give students repeated practice at applying certain concepts or skills and time this practice so that it is spread out in time.
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.

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.

4. BRINGING IT ALL TOGETHER

In the preceding subsections, we have described each of the five principles of learning in relative isolation (see Table 1 for a summary). However, in designing or redesigning a course, much more must be considered. First, we note that it is important to apply the five principles of learning *jointly*, so that they can mutually guide the design process. Second, we acknowledge that there is more

to instructional design than developing and implementing effective instructional techniques, the focus of this article. As mentioned in the description of our course design (Section 2), several steps preceded the decisions regarding instructional technique—identifying course goals, establishing performance criteria for those goals, and decomposing the goals into learnable “pieces” of knowledge—and other steps followed—assessment and re-design [see Dick (1997) for a review of the steps of instructional design]. The reform movement in statistics education has confronted these larger issues while still directing considerable effort towards improving instructional techniques. In the following subsections, 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 the existing work in this area. Note that because of space limitations the following only represents a small sample of the work and progress in statistics education.

4.1 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 problem solving 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

Table 2. *Annotated Overview of Educational Reforms*

<i>Reform-based technique</i>	<i>Examples of Use</i>
Collaborative learning	<ul style="list-style-type: none"> • Students work together to solve problems (e.g., Borresen 1990; Dietz 1993) or discuss concepts, sharing ideas and understanding (Garfield 1993).
Active learning	<ul style="list-style-type: none"> • Students are engaged in data collection (e.g., Rossman 1996; Scheaffer et al. 1996; Spurrier et al. 1995), reflection on and exploration of statistical concepts (Lan et al. 1993), and solving problems on their own (cf. Use of technology).
Target misconceptions	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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. • 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).

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 that are dedicated to the analysis and interpretation of real datasets and to the qualitative (instead of quantitative) understanding of statistical concepts (e.g., Cobb 1987; Glenberg 1996; Moore 1997c; Moore and McCabe 1993; Rossman 1996; Scheaffer, Gnanadesikan, Watkins, and Witmer 1996; Utts 1996). In courses using these texts, problems typically require students to work with a statistical software package in order to answer questions regarding a given dataset (e.g., performing exploratory or inferential data analyses and interpreting the results). As an aid to instructors wanting to assign such problems, several groups have made available databases of real datasets (e.g., DASL, JSE, EESEE) which are indexed by statistical concept and application area. In some cases, students are even responsible for collecting their own data (e.g., considering issues of experimental design and data collection).

4.2 Statistics Education Reform: Changes in Instructional Technique

Beyond the shift in the content of statistics courses, a major focus of the statistics education reform movement has been on improving the instructional techniques used in these courses. Table 2 provides a list of some reform-based techniques and brief descriptions of how they are used. This list is not meant to be comprehensive but rather to demonstrate the variety of innovative instructional techniques that are being employed in a variety of statistics classes today.

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; Giraud 1997; Gnanade-

sikan, Scheaffer, Watkins, and Witmer 1997; Grabowski and Harkness 1996; Keeler and Steinhorst 1995; Magel 1998; Smith 1998). 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. Other approaches to getting students actively engaged in learning statistics have also been used (e.g., Gnanadesikan et al. 1997; Smith 1998). For example, Smith (1998) found that incorporating a sequence of projects in a semester-long introductory course led to positive responses from students and improved exam scores relative to the previous semester’s students.

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. Similar results have been shown by other researchers as well (e.g., Dietz 1993; Garfield 1993; Giraud 1997; Keeler and Steinhorst 1995; Magel 1998).

Technology has also played a major role in instructional innovations. Some software packages offer students simulation systems for exploring statistical concepts (e.g., Alper and Raymond 1998; delMas, Garfield, and Chance in press; Finch and Cumming 1998; Finzer and Erickson 1998; Loomis and Boynton 1992; Velleman 1996) while others provide educationally enhanced statistical software packages (e.g., Cohen, Tsai, and Checile 1995; Cohen et al. 1996; Cohen and Checile 1997; Schuyten and Dekeyser 1998; Shaughnessy 1998). These technological aids to learning generally lead to improvements in students’ understanding

and problem-solving skills, most likely because the technology give students more opportunities to consider conceptual implications and work through problems on their own. For example, delMas et al. (in press) demonstrate that students showed better statistical reasoning after working with their simulation world, especially when students explored the system fully.

As the above results 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 Checile 1995; Cohen et al. 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 (post-test–pre-test) than did “control” students, Cohen and Checile (1997, p. 110) remarked 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) similarly found that, although 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 significantly greater learning gains from pre-test to post-test among students who took the course relative to a group of control students. Nevertheless, these gains were attributable to large improvements on 75% of the test items and little or no improvement on the other 25% of the items (see Lovett, Greenhouse, Johnson, and Gluck in press for more details). 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.

4.3 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. It has increased instructors’ awareness of and dialogue about new instructional ideas. For example, terms such as “active learning” and “hands-on practice” 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. 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 1997a,b). 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 1997a,b). 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, the principles of learning presented in this article offer guidance in specifying these design decisions in a way that will likely lead to effective instruction. Second, new instructors could greatly benefit from a theoretical framework that offers general guidance in making instructional decisions and helps structure their knowledge about teaching as they acquire their own database of experiences, read the statistics education literature, and gradually gain expertise. Taking into account these issues as well as the strengths of individual instructors and the statistics education reform movement more generally, we advocate a *combined* approach in which instructors both (1) draw on their own experiences and other documented cases of good instructional designs and (2) apply the principles of learning discussed above to guide their instructional design decisions. In this way, our framework can complement both instructors’ expertise and the existing research on statistics education, extending the progress that has been made already.

5. 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 implementation. Most of the resources available to them, however, fail to emphasize the process of how to design a course. In contrast, the principles we laid out in this article describe the processes of learning that apply to students in general and lead to a set of practical guidelines that can guide the process of instructional design for a wide range of courses.

Applying the 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

- Alper, P., and Raymond, R. (1998), "Some Experience Comparing Simulation and Conventional Methods in an Elementary Statistics Course," in *Proceedings of the Fifth International Conference on the Teaching of Statistics*, Singapore: ISIASE/ISI.
- 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.
- (1992), "Teaching Statistics," in *Heeding the Call for Change: Suggestions for Curricular Action*, MAA Notes Number 22, ed. L. Steen, pp. 3–43.
- Cohen, S., and Checile, 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., Checile, 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 Checile, 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.
- DelMas, R., Garfield, J., and Chance, B. (in press), "Assessing the Effects of a Computer Microworld on Statistical Reasoning," to appear in *Journal of Statistics Education*, <http://www.amstat.org/publications/jse>.
- Dick, W. (1997), "A Model for the Systematic Design of Instruction," in *Instructional Design: International Perspectives: Theory, Research, and Models* (vol. 1), 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., and Anderson, J. R. (1984), "The Effects of Information Order and Learning Mode on Schema Abstraction," *Memory and Cognition*, 12, 20–30.
- Eylon, B.-S., and Reif, F. (1984), "Effects of Knowledge Organization on Task Performance," *Cognition and Instruction*, 1, 5–44.
- Finch, S., and Cumming, G. (1998), "Assessing Conceptual Change in Learning Statistics," in *Proceedings of the Fifth International Conference on the Teaching of Statistics*, Singapore: ISIASE/ISI, pp. TKK–TKKK.
- Finzer, W. F., and Erickson, T. E. (1998), "Dataspace—A Computer Learning Environment for Data Analysis and Statistics Based on Dynamic Dragging, Visualization, Simulation and Networked Collaboration," in *Proceedings of the Fifth International Conference on the Teaching of Statistics*, Singapore: ISIASE/ISI, pp. TKKK–TKKK.
- 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, <http://www.amstat.org/publications/jse>.
- (1995), "How Students Learn Statistics," *International Statistical Review*, 63, 25–34.
- (1996), "Assessing Student Learning in the Context of Evaluating a Chance Course," *Communications in Statistics—Theory and Method*, 25, 2863–2873.
- (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* (vol. 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.
- Giraud, G. (1997), "Cooperative Learning and Statistics Instruction," *Journal of Statistics Education*, 5, <http://www.amstat.org/publications/jse>.
- Glenberg, A. M. (1996), *Learning From Data: An Introduction to Statistical Reasoning* (2nd ed.), Mahwah, NJ: Erlbaum.
- Gnanadesikan, M., Scheaffer, R. L., Watkins, A. E., and Witmer, J. A. (1997), "An Activity-Based Statistics Course," *Journal of Statistics Education*, 5, <http://www.amstat.org/publications/jse>.
- 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, <http://www.amstat.org/publications/jse>.
- 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.
- Holland, J. H., Holyoak, K. J., Nisbett, R. E., and Thagard, P. R. (1987), *Induction: Processes of Inference, Learning, and Discovery*, Cambridge, MA: MIT Press.
- Keeler, C. M., and Steinhurst, R. K. (1995), "Using Small Groups to Promote Active Learning in the Introductory Statistics Course: A Report From the Field," *Journal of Statistics Education*, 3, <http://www.amstat.org/publications/jse>.
- Kessler, C. (1988), "Transfer of Programming Skills in Novice LISP Learners," unpublished 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*, Gainesville, FL: University of Florida.
- Loomis, J., and Boynton, G. (1992), *Stat Visuals*, Santa Barbara, CA: University of California.
- 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.
- Lovett, M. C., Greenhouse, J. B., Johnson, M., and Gluck, K. (in press) "Assessing an Introductory Statistics Course Using the Rasch Model," in preparation.
- Magel, R. C. (1995), "Using Cooperative Learning in a Large Introductory Statistics Class," *Journal of Statistics Education*, 6, <http://www.amstat.org/publications/jse>.
- 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. (1992), "Teaching Statistics as a Respectable Subject," in *Statistics for the Twenty-First Century*, MAA Notes Number 26, eds. F. and S. Gordon, 14–25.
- (1997a), "New Pedagogy and New Content: The Case of Statistics," *International Statistical Review*, 65, 123–165.
- (1997b), Response to "New Pedagogy and New Content: The Case of Statistics," by D. Moore, *International Statistical Review*, 65, 162–165.
- (1997c), *Statistics: Concepts and Controversies*, San Francisco: W. H. Freeman.
- Moore, D. S., and McCabe, G. P. (1993), *Introduction to the Practice of*

- Statistics, New York: W. H. Freeman.
- National Council of Teachers of Mathematics (1989), *Curriculum and Evaluation Standards For School Mathematics*, Reston, VA: NCTM.
- Newell, A. (1990), *Unified Theories of Cognition*, Cambridge, MA: Harvard University Press.
- Paas, F. G. W. C., and Van Merriënboer, 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.
- Scheaffer, R. L., Gnanadesikan, M., Watkins, A., and Witmer, J. (1996), *Activity-Based Statistics*, New York: Springer-Verlag.
- Schoenfeld, A. H. (1988), "When Good Teaching Leads to Bad Results: The Disasters of 'Well-Taught' Mathematics Courses," *Educational Psychologist*, 23, 145–166.
- Schuyten, G., and Dekeyser, H. (1998), "An Electronic Learning Environment for Applied Statistics: Quality Care and Statistic Education in Higher Education," in *Proceedings of the Fifth International Conference on the Teaching of Statistics*, Singapore: ISIASE/ISI, pp. TKKK–TKKK.
- Shughnessy, M. (1998), "Immersion in Data Handling: Using the Chance-Plus Software With Introductory Statistics Students," in *Proceedings of the Fifth International Conference on the Teaching of Statistics*, Singapore: ISIASE/ISI, pp. TKKK–TKKK.
- Singley, M. K., and Anderson, J. R. (1989), *The Transfer of Cognitive Skill*, Cambridge, MA: Harvard University Press.
- Smith, G. (1998), "Learning Statistics by Doing Statistics," *Journal of Statistics Education*, 6, <http://www.amstat.org/publications/jse>.
- 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, New York: 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.