

Applications and Misapplications of Cognitive Psychology to Mathematics Education

**John R. Anderson
Lynne M. Reder
Herbert A. Simon**

**Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213**

**ja@cmu.edu
reder@cmu.edu
has@cs.cmu.edu**

Abstract

There is a frequent misperception that the move from behaviorism to cognitivism implied an abandonment of the possibilities of decomposing knowledge into its elements for purposes of study and decontextualizing these elements for purposes of instruction. We show that cognitivism does not imply outright rejection of decomposition and decontextualization. We critically analyze two movements which are based in part on this rejection--situated learning and constructivism. Situated learning commonly advocates practices that lead to overly specific learning outcomes while constructivism advocates very inefficient learning and assessment procedures. The modern information-processing approach in cognitive psychology would recommend careful analysis of the goals of instruction and thorough empirical study of the efficacy of instructional approaches.

Following on the so-called "cognitive revolution" in psychology that began in the 1960s, education, and particularly mathematics and science education, has been acquiring new insights from psychology, and new approaches and instructional techniques based on these insights. At the same time, cognitive psychologists have been paying increasing attention to education as an area of application of psychological knowledge and as a source of important research problems. There is every reason to believe that as research in cognitive psychology progresses and increasingly addresses itself to educational issues, even closer and more productive links can be formed between psychology and mathematics education.

However, there is a tendency now to present all manner of educational opinion as bearing a stamp of approval from cognitive psychology. For instance, Lamon and Lesh (1992) write in the introduction to a recent book they edited:

"Behavioral psychology (based on factual and procedural rules) has given way to cognitive psychology (based on models for making sense of real-life experiences), and technology-based tools have radically expanded the kinds of situations in which mathematics is useful, while simultaneously increasing

the kinds of mathematics that are useful and the kinds of people who use mathematics on a daily basis. In response to these trends, professional and governmental organizations have reached an unprecedented, theoretically sound, and future-oriented new consensus about the foundations of mathematics in an age of information." (p. 18-19)

In fact, as in many recent publications in mathematics education, much of what is described in that book reflects two movements, "situated learning" and "constructivism", which have been gaining influence on thinking about education and educational research. In our view, some of the central educational recommendations of these movements have questionable psychological foundations. We wish to compare these recommendations with current empirical knowledge about effective and ineffective ways to facilitate learning in mathematics and to reach some conclusions about what are the effective ways. A number of the claims that have been advanced as insights from cognitive psychology are at best highly controversial and at worst directly contradict known research findings. As a consequence, some of the prescriptions for educational reform based on these claims are bound to lead to inferior educational outcomes and to block alternative methods for improvement that are superior.

These two schools, of situated learning and constructivism, are not identical: situated learning emphasizes that knowledge is maintained in the external, social world; constructivism argues that knowledge resides in an individual's internal state, perhaps unknowable to anyone else. However, both schools share the general philosophical positions that knowledge cannot be decomposed or "decontextualized" for purposes of either research or instruction, and each group often appeals to the writings of the other for support. Since rejection of decomposition and decontextualization seems to be the core common ground of this "new look" in mathematics education, we will first examine the degree to which modern cognitive psychology lends support to that rejection.

Decomposition and Decontextualization

In an influential educational paper, Resnick and Resnick (1992) provide a succinct statement of a common theoretical understanding in cognitive psychology called the information-processing approach:

"Information-processing theories of cognition (Anderson, 1983; Newell and Simon 1972), for example, analyze cognitive performances into complexes of rules, but performances critically depend on interactions among those rules. Each rule can be thought of as a component of the total skill, but the rules are not defined independently of one another. The 'competence' of a problem-solving system thus depends on how the complex of rules acts together." (p. 43)

A number of educational researchers (e.g., Shepard, 1991) have cited Resnick and Resnick as reporting that cognitive psychology has shown that cognition cannot be analyzed into components. On the contrary, what the above quote states (and what the cognitive literature they allude to says) is quite the opposite. This literature, incorporating extensive empirical evidence, deals both with the "rules" (components or processes) to which Resnick and Resnick refer, and also, emphatically, with the interactions among these processes: the interaction between these processes and sensory stimuli (the organism's awareness of its current environment), and the interaction of processes with information (other components of knowledge) that has been assembled in memory through previous engagement with the environment. The whole purpose of modeling cognition with computer programs--a central tool in information-processing approaches--is to develop a full picture of these interactions among components of knowledge.

Unlike earlier behaviorist theories, information-processing theories do not posit a simple one-to-one mapping between individual rules or knowledge components and individual bits of behavior. They deny this precisely because continual interaction can be observed among components of knowledge and behavior. Information-processing psychology has advanced rapidly by developing methods both for identifying the components and for studying them in their interactions with their entire contexts. This is the

meaning of the "unified theories of cognition" (e.g., Newell, 1991) which has guided so much of the recent research and theory-building.

Thus, componential analysis is very much alive and well in modern cognitive psychology. The information-processing approach tries both to deepen our understanding of the components and to understand the relations among them and with their environments. Examples of these methods of componential analysis are the use of think-aloud protocols as data (Ericsson and Simon, 1993) and the use of models that simulate the interactions of perceptual, memory, learning and thinking processes over a wide range of cognitive tasks (e.g., Anderson, 1993; Feigenbaum and Simon, 1984; Newell, 1991).

With respect to decomposition, the correct principle is:

Assessing learning and improving learning methods requires careful task analysis at the level of component skills, intimately combined with study of the interaction of these skills in the context of broader tasks and environments.

o much for decomposition; what about decontextualization? Because components interact with one another, it might prove impossible to invoke and study them outside certain contexts. To cite a simple example, processes for carrying out multi-column addition will only be evoked in the context of a problem large enough to require carrying; they cannot be studied by posing problems of adding $3+4$ or $5+2$.

While some context will often be required to assess a component, there are always bounds on how complex such a context need be. It is a well-documented fact of human cognition that large tasks decompose into nearly independent subtasks (Simon, 1981, Chapter 7; Card, Moran & Newell, 1983), so that only the context of the appropriate subtask is needed to study its components. For instance, there is no need to teach or assess the ability to perform multi-column addition in the context of calculating income taxes. The process of adding tax deduction items is the same as the process of taking sums in other tasks. And whether one does the sum by hand or by calculator is unlikely to affect the rest of the tax calculation procedures. Thus, the larger procedure is independent of the summing procedure, just as the summing procedure is independent of the larger procedure.

The addition procedures might become tied into the tax calculation procedures--for example, ignoring cents in calculating the sums. Such specialized subprocedures are especially frequent at high levels of expertise. However, this just means that the expert's procedure involves a structure of different subtasks than the novice's, not that it cannot be analyzed into components nor that these components cannot still be assessed in subtasks of the original task. Thus, with respect to decontextualization, while it may be difficult to get behavioral measures of individual components; these components organize themselves into subtasks to achieve subgoals, and these subgoals can have independent, assessable, behavioral realizations. It does not require recondite research to demonstrate the near-decomposability of human tasks. Every page of a good cookbook contains examples of assumed component procedures (e.g., *sauté*, *parboil*) as do the how-to books in domains like carpentry, plumbing or car repair. Moreover, one can apply these procedures in new contexts such as when a chemistry lab requires us to boil water. Fortunately for us human beings, with our very limited short-term memories, the workings of each component can be understood without simultaneous awareness of the details of all the other components.

With respect to decontextualization, the correct principle is:

Assessing learning and improving learning methods requires research and instruction in contexts that are consistent with the scopes of the skills currently under investigation. Component skills can be viewed within narrower contexts than broad skills. Relating context to task is essential in order to meet the limits of human attention and short-term memory capacity.

This false rejection of decomposition and decontextualization runs deep in modern mathematics education. So, for instance, in the 1993 draft of the NCTM assessment standard for school mathematics, we

find condemnation of the "essentialist view of mathematical knowledge" which assumes "mathematics consists of an accumulation of mathematical concepts and skills" (p.12). We can only say we find frightening the prospect of mathematics education based on such a misconceived rejection of componential analysis.

The major agenda of this paper is to focus on the claims of situated learning and constructivism, discussing them separately and focusing in each case on a small number of central claims that we believe are unwarranted. There are other issues beyond those we discuss, but these are perhaps the most important for choosing among research directions and pedagogical strategies.

Situated Learning

Two of us have been involved in past reviews relevant to situated learning--Simon in support of the mutual compatibility of modern information processing theory and situated cognition (Vera & Simon, 1993) and Reder in an assessment of the effectiveness for training of techniques located at various points along the scale of "situatedness" (Reder & Klatzky in a report of the National Research Council, 1994). We will focus on the four claims of situated learning discussed in the NRC report.

Claim 1: Action is grounded in the concrete situation in which it occurs

That action is situationally grounded is surely the central claim of situated cognition. It means that the potentialities for action cannot be fully described independently of the specific situation, a statement with which we fully concur. But the claim is sometimes exaggerated to assert that all knowledge is specific to the situation in which the task is performed, and that more general knowledge cannot and will not transfer to real-world situations. Supposed examples of this are Lave's (1988) description of Orange County homemakers who did very well at making supermarket best-buy calculations but who did much worse on arithmetically equivalent school-like paper-and-pencil mathematics problems. Another frequently cited example is Carraher, Carraher and Schliemann's (1985) account of Brazilian street children who could perform mathematics when making sales in the street but were unable to answer similar problems presented in a school context.

Even if these claims are valid and generalizable beyond the specific anecdotes that have been cited, they demonstrate *at most* that particular skills practiced in real-life situations do not generalize to school situations. They assuredly do not demonstrate that arithmetic procedures taught in the classroom cannot be applied to enable a shopper to make price comparisons or a street vendor to make change. What such observations call for is closer analyses of the task demands and the use of such analyses to devise teachable procedures that will achieve a balance between the advantages of generality and the advantages of incorporating enough situational context to make transfer likely. What they also call for is research on the feasibility of increasing the application and transfer of knowledge by including ability to transfer as a specific goal in instruction--a skill that is given little attention in most current instruction.

At one level there is nothing new in this claim about the contextualization of learning. There have been numerous demonstrations in experimental psychology that learning can be contextualized (e.g., Godden & Baddeley, 1975; Smith, Glenberg, & Bjork, 1978). For instance, Godden and Baddeley found that divers had difficulty remembering under water what they learned on land or vice versa. However, it is not the case that learning is totally tied to a specific context. For instance, Godden and Baddeley's divers could remember some of what they learned in the other context. In fact, there are many demonstrations of learning that transfers across contexts and of failures to find any context specificity in the learning (e.g., Fernandez & Glenberg, 1985; Saufley, Olaka, & Baversco, 1985) -- a fact that has often frustrated researchers who were looking for context sensitivity.

How tightly learning will be bound to context depends on the kind of knowledge being acquired. Sometimes knowledge is necessarily bound to a specific context by the nature of instruction. Thus, to return to an earlier example, one would not be surprised (and only a little upset) to learn that carrying is

bound to the context of doing base-ten addition and would not generalize to another base system. In other cases, how contextualized the learning is depends on the way the material is studied. If the learner elaborates the knowledge with material from a specific context, it becomes easier to retrieve the knowledge in that same context (Eich, 1985), but perhaps harder in other contexts. One general result is that knowledge is more context bound when it is just taught in a single context (Bjork & Richardson-Klavhen, 1989).

Clearly, some skills, like reading, transfer from one context to another. For instance, the very fact that we can engage in a discussion of the context-dependence of knowledge is itself evidence for the context independence of reading and writing competence. Many of the demonstrations of contextual-binding from the situated camp involve mathematics, but clearly, mathematical competence is not always contextually bound either. Although the issue has seldom been addressed directly, the psychological research literature is full of cases where mathematical competence has transferred from the classroom to all sorts of laboratory situations (sometimes bizarre--the intention was never to show transfer of mathematical skills--e.g., Bassok & Holyoak, 1987; Elio, 1986; Reder & Ritter, 1992). It is not easy to locate the many published demonstrations of mathematical competence generalizing to novel contexts; these results are not indexed under "context-independence of mathematical knowledge" because, until recently, this did not seem to be an issue.

The literature on situation-specificity of learning often comes with a value judgment about the merits of knowledge tied to a nonschool context relative to school-taught knowledge, and an implied or expressed claim that school knowledge is not legitimate. Lave (1986, 1988 p. 195) goes so far as to suggest that school-taught mathematics serves only to justify an arbitrary and unfair class structure. The implication is that school-taught competences do not contribute to on-the-job performance. However, numerous studies show modest to large correlations between school achievement and work performance (e.g., Hunter & Hunter, 1984; Brossiere, Knight, & Sabol, 1985) even after partialling out the effects of general ability measures (which are sometimes larger).

We conclude that action is indeed grounded in the situation where it occurs. We dissent strongly from some of the supposed implications that have been attached to this claim by proponents of situated action, and we have shown that our dissent has strong empirical support. Instead, the evidence shows that:

How contextualized learning is depends on the way the material is studied. Knowledge is more context bound when it is just taught in a single context.

Knowledge does not have to be taught in the precise context in which it will be used, and grave inefficiencies in transfer can result from tying knowledge too tightly to specific, narrow contexts.

We need closer analyses of the task demands to devise teachable procedures that will balance the advantages of generality with the advantages of incorporating enough situational context to make transfer likely.

We also need to study how to increase the application and transfer of knowledge by including ability to transfer as a specific goal in instruction.

In particular, knowledge does not have to be taught in the precise context in which it will be used, and grave inefficiencies in transfer can result from tying knowledge too tightly to specific, narrow contexts.

Claim 2: Knowledge does not transfer between tasks

This second claim, of the failure of knowledge to transfer, can be seen as a corollary of the first. If knowledge is wholly tied to the context of its acquisition, it is not going to transfer to other contexts. Even without strong contextual dependence, one could still claim that there is relatively little transfer, beyond

nearly identical tasks, to different physical contexts. For instance, while one might be able to do fractional math in any context, it might not transfer to learning algebra. There is a long tradition of research on transfer in psychology, going back at least to Weber in 1844 and Fechner in 1858 (Woodworth, 1938, Chapter 8), demonstrating that, depending very much upon the experimental situation and the relation of the material originally learned to the transfer material, there can be either large amounts of transfer, a modest amount of transfer, no transfer at all, or even negative transfer.

The more recent psychological literature is also full of failures to achieve transfer (e.g., Gick & Holyoak, 1980; Hayes & Simon, 1977; Reed, Ernst, & Banerji, 1974; Weisberg, DiCamillo, & Phillips, 1985), but it is also full of successful demonstrations of transfer (e.g., Brown, 1990; Brown & Campione, 1993; Kotovsky & Fallside, 1989; Schoenfeld, 1985; Singley & Anderson, 1989; Smith, 1986). Indeed, in the same domain (Tower of Hanoi isomorphs) quite different amounts of transfer occur depending on the amount of practice with the target task and on the representation of the transfer task (Kotovsky & Fallside, 1989). In general, representation and degree of practice are critical for determining the transfer from one task to another.

Singley and Anderson (1989) argued that transfer between tasks is a function of the degree to which the tasks share cognitive elements. This hypothesis had also been put forth very early in the development of research on transfer (Thorndike & Woodworth, 1901; Woodworth, 1938), but was hard to test experimentally until we acquired our modern capability for identifying task components. Singley and Anderson taught subjects several text editors, one after another and sought to predict transfer (savings in learning a new editor when it was not taught first). They found that subjects learned subsequent text editors more rapidly and that the number of procedural elements shared by two text editors predicted the amount of this transfer. In fact, they obtained large transfer across editors that were very different in surface structure but that had common abstract structures. Singley and Anderson also found that similar principles govern transfer of mathematical competence across multiple domains, although here they had to consider transfer of declarative as well as procedural knowledge. As a general statement of the research reported by Singley and Anderson, transfer varied from one domain to another as a function of the number of symbolic components that were shared. If anything, Singley and Anderson found empirically slightly more transfer than was predicted by their theory.

What about the situations where subjects have shown relatively little transfer? In one famous series of studies (Gick & Holyoak, 1980, 1983), subjects were presented with Duncker's (1945) classic radiation problem: "Suppose you are a doctor faced with a patient who has an inoperable stomach tumor. You have at your disposal rays that can destroy human tissue when directed with sufficient intensity. How can you use these rays to destroy the tumor without destroying the surrounding healthy tissue?" (adapted from Gick & Holyoak, 1983). Prior to their exposure to the target problem, subjects read a story about an analogous military problem and its solution. In the story, a general wishes to capture an enemy fortress. Radiating outward from the fortress are many roads, each mined in such a way that the passing of any large force will cause an explosion. This precludes a full-scale direct attack. The general's plan is to divide his army, send a small group down each road, and converge on the fortress. The common strategy in both problems is to divide the force, attack from different sides, and converge on the target. After reading this story, however, only about 30 percent of the subjects could solve the radiation problem, which is only a limited improvement (although an improvement by a factor of three) over the 10 percent baseline solution rate (Gick & Holyoak, 1980).

One of the striking characteristics of such failures of transfer is how relatively transient they are. Gick and Holyoak were able to increase transfer greatly just by suggesting to subjects that they try to use the problem about the general. Exposing subjects to two such analogs also greatly increased transfer. The amount of transfer appeared to depend in large part on where the attention of subjects was directed during the experiment, which suggests that instruction and training on the cues that signal the relevance of an available skill might well deserve more emphasis than they now typically receive--a promising topic for cognitive research with very important educational implications.

As a methodological comment, we think that there is a tendency to look for transfer in situations where one is least likely to find it. That is, research tends to look for transfer from little practice in one domain to initial performance in another domain. Superficial differences between the two domains will have their largest negative effect when the domains are unfamiliar. We do not require that students show the benefit of one day of calculus on the first day of physics. Rather, we expect that they will be better physics students at year's end for having had a year's study of calculus.

Contrary to the claim that knowledge does not transfer between tasks, the evidence we have reviewed supports the following principles for securing transfer of learning:

Depending upon the experimental situation and the relation of the material originally learned to the transfer material, there can be either large amounts of transfer, a modest amount, no transfer at all, or even negative transfer.

Representation and degree of practice are critical for determining the transfer from one task to another, and transfer varies from one domain to another as a function of the number of symbolic components that are shared.

The amount of transfer depends on where attention is directed during learning. Training on the cues that signal the relevance of an available skill may deserve much more emphasis than they now typically receive in instruction.

Claim 3: Training by abstraction is of little use; real learning occurs in "authentic" situations.

Like Claim 2, the claim that training by abstraction is of little use is a corollary of the earlier claims. Nonetheless, one might argue for it even if one dismisses the others. Claim 3 has been extended into an advocacy for apprenticeship training (Brown, Collins, & Duguid, 1989; Collins, Brown, & Newman, 1989). It is argued that, because current performance will be facilitated to the degree that the context closely matches prior experience, the most effective training is an apprenticeship to others in the performance situation. This claim is used more than any other to challenge the legitimacy of school-based instruction.

Abstract instruction can be ineffective if what is taught in the classroom is not what is required in the job situation. Often this is an indictment of the design of the classroom instruction rather than of the idea of abstract instruction in itself. However, sometimes it is an indictment of the job situation. For instance, Los Angeles police after leaving the police academy are frequently told by more experienced officers "now forget everything you learned" (Independent Commission on the Los Angeles Police Department, 1991: 125). The consequence is that police officers are produced who, ignoring their classroom training in the face of contrary influences during apprenticeship, may violate civil rights and make searches without warrants. Clearly, one needs to create a better correspondence between job performance and abstract classroom instruction and sometimes this means changing the nature of the job performance (including the structure of motivations and rewards) and fighting unwanted and deleterious effects of apprenticeship learning.

Abstract instruction can be quite effective. In unpublished research, Singley found that abstract instruction leads to successful transfer while concrete instruction can lead to failure of transfer. He taught subjects to solve algebra word problems involving mixtures. Some subjects were trained with pictures of the mixtures while other subjects were trained with abstract tabular representations that highlighted the underlying mathematical relationships. It was the abstract training group that was able to transfer better to other kinds of problems that involved analogous mathematical relationships. Perhaps the most striking demonstration of the benefit of abstract instruction comes from Biederman and Shiffrar (1987). They looked at the very difficult task of sexing day-old chicks--something that people spend years learning in an apprentice-like role. They found that 20 minutes of abstract instruction brought novices up to the levels of experts who had years of practice.

The issue of choosing between abstract and very specific instruction can be viewed in the following way. If abstract training is given, learners must also absorb the money and time costs of obtaining supplemental training for each distinct application. But if very specific training is given, they must completely retrain for each application. Which is to be preferred, and to what extent, depends on the balance among (a) the cost of the more general abstract training, (b) the cost of the specific training, (c) the cost of the supplemental training for application of abstract training, and (d) the range of jobs over which the learner is likely to have occasion to apply what was learned. Someone who will spend years performing a single set of very specific tasks might be well advised to focus on specific training. But if the cost of supplemental training is not large (i.e., if there is substantial transfer over the range of tasks), or if technological or other changes are likely to alter tasks substantially over the years, or if the range of tasks the learner is likely to address over time is substantial, then abstract training with supplemental applications training is clearly preferable. It is easy to work out an exercise of this kind by assigning numbers to the various costs and to the variability of the tasks encountered, and thereby to show that there is no solution that is optimal for all cases.

Most modern information-processing theories are "learning-by-doing" theories which imply that learning would occur best with a combination of abstract instruction and concrete illustrations of the lessons of this instruction. Numerous experiments show combining abstract instruction with specific concrete examples (e.g., Cheng, Holyoak, Nisbett, & Oliver, 1986; Fong, Krantz, & Nisbett, 1986; Reed & Actor, 1991) is better than either one alone. One of the most famous studies demonstrating this was performed by Scholckow & Judd (described in Judd, 1908; a conceptual replication by Hendrickson & Schroeder, 1941). They had children practice throwing darts at a target underwater. One group of subjects received an explanation of refraction of light which causes the apparent location of the target to be deceptive. The other group only practiced, receiving no abstract instruction. Both groups did equally well on the practice task which involved a target 12 inches under water, but the group with abstract instruction did much better when asked to transfer to a situation where the target was now under only 4 inches of water.

A variation on the emphasis on apprenticeship training is the emphasis that has been given to using only "authentic" problems (e.g., Lesh & Lamon, 1992). What is authentic is typically ill-defined but there seems to be a strong emphasis on having problems be like the problems students might encounter in everyday life. A focus on underlying cognitive process would suggest that this is a superficial requirement. Rather, we would argue as have others (e.g., Hiebert, Hearnar, Carpenter, Fennema, Fuson, 1994) that the real goal should be to get students motivated and engage in cognitive processes that will transfer. What is important is what cognitive processes a problem evokes and not what real-world trappings it might have.

Abstract instruction can be ineffective if what is taught in the classroom is not what is required in the job situation, but under other conditions, it can be quite effective.

Whether abstract or specific instruction is to be preferred, and to what extent, depends on the balance among (a) the cost of the more general abstract training, (b) the cost of the specific training, (c) the cost of the supplemental training for application of abstract training, and (d) the range of jobs over which the learner is likely to have occasion to apply what was learned.

Most modern information-processing "learning-by-doing" theories imply that learning would occur best with a combination of abstract instruction and concrete illustrations of the lessons of this instruction that get students motivated and engaged in cognitive processes that will transfer. What is important is what cognitive processes a problem evokes and not its real-world trappings.

Claim 4: Instruction needs to be done in a highly social environment

The claim that instruction is only effective in a highly social environment is based on the ideas that (a) virtually all jobs are highly social in nature and (b) learning is closely associated with its context. As we have shown, the second claim is overstated. We suspect that the first claim is also somewhat

overstated, although we are not acquainted with any analyses of existing job surveys that show how much social interaction, and what kind, is involved in various jobs. Clearly, there are jobs that are not social in character and for which this claim does not hold. Likewise, it is clear that there are jobs where performance is highly social. Obviously it is important that people with such jobs learn (within or outside the specific job context) to deal effectively with the social nature of their jobs.

While one must learn to deal with the social aspects of jobs, this is no reason why **all** skills required for these jobs should be trained in a social context. Consider the skills necessary to become a successful tax accountant. While the accountant must learn how to deal with clients, it is not necessary to learn the tax code or how to use a calculator while interacting with a client. It is better to train independent parts of a task separately (see the earlier discussion of nearly independent subtasks under decontextualization) because fewer cognitive resources will then be required for performance, thereby reserving adequate capacity for learning. Thus, it is better to learn the tax code without having to simultaneously interact with the client and better to learn how to deal with a client when the tax code is no longer a burden.

In fact, a large history of research in psychology shows that part training is often more effective when the part component is independent, or nearly so, of the larger task (e.g., Knerr, Morrison, Muman, Stein, Sticha, Hoffman, Buede, & Holding, 1987; Patrick, 1992). Indeed in team training, it is standard to do some part-task training of individuals outside of the team just because it would be expensive and futile to get the whole team together when a single member needs training on a new piece of equipment (Salas, Dickinson, Converse, & Tannenbaum, 1993). In team sports, where a great deal of attention is given to the efficiency of training, the time available is always divided between individual skill training and team training. We will have more to say about the issue of part versus whole training when we discuss the constructivist advocacy of carrying on all instruction in complex learning situations.

Another facet of the claim that instruction is best in a highly social environment comes not from those advocating situated learning, *per se*, but from those advocating the advantages of co-operative learning (e.g., Johnson & Johnson, 1989) as an instructional tool. Co-operative learning, also known as "communities of practice" and "group learning", refers to learning environments where people of equal status work together to enhance their individual acquisition of knowledge and skills. This environment or structure is to be contrasted with tutoring (where the tutor and tutee are of unequal knowledge and status) and team training (where the desired outcome is concerned with team or group performance). In a review by the Committee on Techniques for the Enhancement of Human Performance (National Research Council, 1994), it was noted that research on cooperative learning has frequently not been well controlled (e.g., nonrandom assignments to treatments, uncontrolled "teacher" and treatment effects), that relatively few studies "have successfully demonstrated advantages for cooperative versus individual learning," and that "a number of detrimental effects arising from cooperative learning have been identified--the "free rider," the "sucker," the "status differential," and "ganging up" effects (see e.g., Salomon and Globerson, 1989, pp. 94-95).

As the NRC review of cooperative learning notes, there have been a substantial number of reports of no-differences (e.g., Slavin, 1990), but unfortunately, there have also been a huge number of practitioner-oriented articles about cooperative learning that tend to gloss over difficulties with this approach, and treat it as an academic panacea. Indeed, the approach is applied too liberally without the requisite structuring or *scripting* to make it effective. Cooperative learning needs to be structured with incentives (for children at least) that motivate cooperation and a sharing of the goal structure. Because of this uncritical application it seems likely that the costs of this type of instruction may outweigh the intended benefits. In colleges we find group projects increasingly popular among instructors but some of the difficulties encountered show that group learning can become counterproductive. Students sometimes complain that the difficulty of finding times to meet to work on assignments together make the practice frustrating and that some students exploit the system and assume that other partners in the group will do all the work (and hence acquire all the knowledge and skills). A reported practice among some students is to divide the labor across classes so that one member of a group does all of the work for a project in one programming class, while another carries the burden for a different class. Clearly these situations are not

the intended outcomes of cooperative learning, but are the sorts of things that will occur if there is not thoughtful implementation and scripting of the learning situation.

Our point is not to say that cooperative learning can not be successful nor sometimes better than individual learning. Rather, it is not a panacea that always provides outcomes superior or even equivalent to those of individual training.

Summary: Situated Learning

In general, situated learning focuses on some well-documented phenomena in cognitive psychology and ignores many others: While cognition is partly context-dependent, it is also partly context-independent; while there are dramatic failures of transfer, there are also dramatic successes; while concrete instruction helps, abstract instruction also helps; while some performances benefit from training in a social context, others do not. The development from behaviorism to cognitivism was an awakening to the complexity of human cognition. The analysis offered by situated learning seems a regressive move. What is needed to improve learning and teaching is to continue to deepen our research into the circumstances that determine when narrower or broader contexts are required and when attention to narrower or broader skills are optimal for effective and efficient learning.

In our discussion, we have focused, as do the proponents of situated learning, on cognitive issues. There are, of course, also very important questions about the circumstances under which people are most strongly motivated to learn. Motivational questions lie outside our present discussion, but are at least as complex as the cognitive issues. In particular, there is no simple relation between level of motivation, on the one hand, and the complexity or realism of the context in which the learning takes place, on the other. To cite a simple example, learning by doing in the real-life domain of application is sometimes claimed to be the optimum procedure. Certainly, this is not true, when the tasks are life-threatening for novices (e.g., firefighting), when relevant learning opportunities are infrequent and unpredictable (e.g., learning to fly a plane in bad weather), or when the novice suffers social embarrassment from using inadequate skills in a real-life context (e.g., using a foreign language at a low level of skill). The interaction of motivation with cognition has been described in information-processing terms by Simon (1967, 1994). But an adequate discussion of these issues would call for a separate paper as long as this one.

Constructivism

Constructivism has a less unified position than situated learning. Indeed, under some interpretations, we are constructivists and have been called so by mathematics educators (e.g., Silver, 1987). However, there is a rising interpretation of constructivism that rejects the information-processing approach (Cobb, 1990) which is the subject of discussion here. Such views are often espoused by those claiming to practice "radical constructivism". Even among radical constructivists, positions vary and some theorists seem to be making philosophical claims about the nature of knowledge rather than empirical claims. Indeed, in the extreme, constructivism denies the relevance of empirical data to educational decisions. However, some of the claims also have clear psychological implications that are not always supported.

Claim 1: Knowledge cannot be instructed (transmitted) by a teacher, it can only be constructed by the learner

The constructivist vision of learning is nicely captured by the following quote:

"learning would be viewed as an active, constructive process in which students attempt to resolve problems that arise as they participate in the mathematical practices of the classroom. Such a view emphasizes that the learning-teaching process is interactive in nature and involves the implicit and explicit negotiation of mathematical meanings. In the course of these negotiations, the teacher and students elaborate the taken-as-

shared mathematical reality that constitutes the basis for their ongoing communication" (Cobb, Yackel, & Wood, 1992).

As an example of this, Cobb, Wood, Yackel, Nicholls, Wheatley, Trigatti, & Pertwitz (1991) describe an effort to teach second graders to count by tens. Rather than telling the students the principle directly, they assigned groups of students the task of counting objects bundled in sets of ten. Invariably, the groups discover that counting by tens is more efficient than counting by ones. Building a whole second-grade curriculum around such techniques, they found their students doing as well on traditional skills as students from traditional classrooms, transferring more, and expressing better attitudes about mathematics.

One can readily agree with one part of the constructivist claim: that learning must be an active process. Learning requires a change in the learner, which can only be brought about by what the learner does--what he or she attends to, what activities he or she engages in. The activity of a teacher is relevant to the extent that it causes students to engage in activities they would not otherwise engage in--including, but not limited to, acquiring knowledge provided by the teacher or by books. A teacher may also engage students in tasks, some of which may involve acquisition of skills by working examples. Other tasks include practicing skills to bring them to effective levels, interacting with their fellow students and with the teacher, and so on.

The problem posed to psychology and education is to design a series of experiences for students that will enable them to learn effectively and to motivate them to engage in the corresponding activities. On all of these points, it would be hard to find grounds for disagreement between constructivists and other cognitive psychologists. The more difficult problem, and the one that often leads to different prescriptions, is determining the desirable learning goals and the experiences that, if incorporated in the instructional design, will best enable students to achieve these goals. Of course, arriving at good designs is not a matter for philosophical debate; it requires empirical evidence about how people, and children in particular, actually learn, and what they learn from different educational experiences.

One finds frequent reference to Jean Piaget as providing a scientific basis for constructivism. Piaget has had enormous influence on our understanding of cognitive development and indeed was one of the major figures responsible for the emergence of cognitivism from the earlier behaviorist era in psychology. While it is fair to say that many of his specific claims have been seriously questioned, the general influence of his theoretical perspective remains. Key to constructivism is Piaget's distinction between assimilation and accommodation as mechanisms of learning and development. Assimilation is a relatively passive incorporation of experience into a representation already available to the child. However, when the discrepancies between task demands and the child's cognitive structure become too great, the child will reorganize his or her thoughts. This is called accommodation (and often nowadays, "re-representation").

Piaget emphasized how the child internalizes by making changes in mental structure. The constructivists make frequent reference to this analysis, particularly the non-passive accommodation process. (In this respect, constructivism is quite different from situated learning which emphasizes the external bases of cognition.) A more careful understanding of Piaget would have shown that assimilation of knowledge also plays a critical role in setting the stage for accommodation--that the accommodation cannot proceed without assimilation.

Some constructivists (e.g., Cobb, 1990) have mistakenly implied that modern information-processing theories deal only with assimilation and do not incorporate the more constructive accommodation. Far from this, the learning-by-doing theories that are widely employed in cognitive science are in fact analyses of how cognitive structure accommodates to experience. We will briefly describe here two such analyses, both to correct the misrepresentation of information-processing theory and to establish a more precise framework for discussing the effects of instruction.

In Anderson's (1993) ACT-R, one principal learning mechanism is knowledge compilation. When learners come upon problems they do not know how to solve, they can look at an example of how a similar problem is solved (retrieved either from memory or some external source) and try to solve the problem by analogy to this example. Knowledge compilation is the accommodation process by which new procedures (rules) are created to produce more directly the computation that this retrieve-and-analogize process requires.

In Feigenbaum and Simon's (1984) EPAM, learning involves gradually building up a discrimination net for recognizing objects and taking appropriate actions. A discrimination net is a sequence of tests that are applied to various features of an object. New tests are added as experience indicates that previous tests were inadequate. Gradually, the system develops a complex sensitivity to the situations and stimuli in its environment in a continuing process of re-representation, or accommodation.

These theories provide concrete realizations of what it means for a system to construct knowledge. As such they provide a basis for examining the constructivist's claim that knowledge cannot be instructed. If passive recording is what one means by "instruct" these learning mechanisms cannot be instructed. However, it is quite wrong to claim that what is learned is not influenced by explicit instruction. For instance, in ACT-R's learning by analogy, instruction serves to determine the representation of the examples from which one "constructs" one's understanding, and Pirolli and Anderson (1985) showed in the domain of recursive programming that what one learns from an example is strongly influenced by the instruction that accompanied the example. In EPAM, which has had extensive success in modeling human learning in a variety of perceptual and verbal learning tasks (e.g., Simon & Feigenbaum, 1964), learning is strongly influenced by the sequence of stimuli and the feedback that tells the system when responses are correct, and when they are wrong.

There is a great deal of research showing that, under some circumstances, people are better at remembering information that they create for themselves than information they receive passively (Bobrow & Bower, 1969; Slamecka & Graf, 1972). However, this does not imply that people do not remember what they are told. Indeed, in other cases people remember as well or even better information that is provided than information they create (Slamecka & Katsaiti, 1987; Stern & Bransford, 1979).

When, for whatever reason, students cannot construct the knowledge for themselves, they need some instruction. The argument that knowledge must be constructed is very similar to the earlier arguments that discovery learning is superior to direct instruction. In point of fact, there is very little positive evidence for discovery learning and it is often inferior (e.g., Charney, Reder & Kusbit, 1990). Discovery learning, even when successful in acquiring the desired construct, may take a great deal of valuable time that could have been spent practicing this construct if it had been instructed. Because most of the learning in discovery learning only takes place after the construct has been found, when the search is lengthy or unsuccessful, motivation commonly flags. As Ausubel (1968) wrote, summarizing the findings from the research on discovery learning twenty-five years ago:

"actual examination of the research literature allegedly supportive of learning by discovery reveals that valid evidence of this nature is virtually nonexistent. It appears that the various enthusiasts of the discovery method have been supporting each other research-wise by taking in each other's laundry, so to speak, that is, by citing each other's opinions and assertions as evidence and by generalizing wildly from equivocal and even negative findings." (p. 497-498)

It is sometimes argued that direct instruction leads to "routinization" of knowledge and drives out understanding:

"the more explicit I am about the behavior I wish my students to display, the more likely it is that they will display the behavior without recourse to the understanding which the behavior is meant to indicate; that is, the more likely they will take the form for the substance." Brousseau (1984)

An extension of this argument is that excessive practice will also drive out understanding. This criticism of practice (called "drill and kill," as if this phrase constituted empirical evaluation) is prominent in constructivist writings. Nothing flies more in the face of the last 20 years of research than the assertion that practice is bad. All evidence, from the laboratory and from extensive case studies of professionals, indicates that real competence only comes with extensive practice (e.g., Hayes, 1985; Ericsson, Krampe, Tesche-Romer, 1993). In denying the critical role of practice one is denying children the very thing they need to achieve real competence. The instructional task is not to "kill" motivation by demanding drill, but to find tasks that provide practice while at the same time sustaining interest. Substantial evidence shows that there are a number of ways to do this; "learning-from-examples," a method we will discuss presently, is one such procedure that has been extensively and successfully tested in school situations.

The evidence, then, leads us to the following conclusions about the role of student and teacher in learning:

Learning requires a change in the learner, which can only be brought about by what the learner does. The activity of a teacher is relevant to the extent that it causes students to engage in activities they would not otherwise engage in.

The task is to design a series of experiences for students that will enable them to learn effectively and to motivate them to engage in the corresponding activities.

The learning-by-doing theories that are widely employed in cognitive science are analyses of how cognitive structure accommodates to experience.

When students cannot construct the knowledge for themselves, they need some instruction. There is very little positive evidence for discovery learning and it is often inferior. In particular, it may be costly in time, and when the search is lengthy or unsuccessful, motivation commonly flags.

People are sometimes better at remembering information that they create for themselves than information they receive passively, but in other cases they remember as well or better information that is provided than information they create.

Real competence only comes with extensive practice. The instructional task is not to "kill" motivation by demanding drill, but to find tasks that provide practice while at the same time sustaining interest. There are a number of ways to do this, for instance, by "learning-from-examples."

Claim 2: Knowledge cannot be represented symbolically

The claim of the situated school that knowledge cannot be represented symbolically is more an epistemological claim in the constructivist's hands than a psychological claim. The claim is that there are subtleties in human understanding that defy representation in terms of a set of rules or other symbol structures (e.g., Cobb, 1990). The argument is not really about whether the knowledge is actually so represented in the human head, but whether knowledge, by its very nature, can be represented symbolically. Searle's well-known attempt to show that, in principle, a symbolic system cannot understand language (the "Chinese Room" metaphor, Searle, 1980) is an extension of this claim.

Among the misconceptions underlying the claim that knowledge is non-symbolic is the faulty notion that "symbolic" means "expressed in words and sentences, or in equivalent formal structures." Symbols are much more than formal expressions. Any kind of pattern that can be stored and can refer to some other pattern, say, one in the external world, is a symbol, capable of being processed by an information-processing system. Thus, an EPAM-like system can learn, when a stimulus satisfies certain tests (has certain features), to create an internal symbol that designates the kind of object we know as a cat. EPAM can then also learn and store in memory the name spelled "c-a-t" (also a symbol), and associate it

with the symbol (pattern) that allows it to recognize a cat. But of course the English name and the object (the cat) are denoted by quite different symbols--a cat is not a verbal structure but a furry creature that can sometimes be seen in the environment.

A substantial number of symbolic systems have been built that can store symbol structures representing mental images of external events and can reason about the events pictorially with the help of these structures (Larkin, 1981). Careful comparison with the behavior of human subjects reasoning about pictures or diagrams shows that these systems capture many of the basic properties and processes of human imagery. Searle's Chinese Room story fails because the inhabitants of his postulated room, unlike humans and other symbolic systems, do not have a sensory window on the world: cannot associate a pattern in memory with the external object that can be seen and denoted by that pattern.

Cobb, Yackel, and Wood (1992) present constructivism as a rejection of the "representational view of mind." We and other cognitive psychologists, who do subscribe to a representational view, find little that we can recognize in their characterization of that view. Cobb et al. quote Rorty's mischaracterization of it:

"To know is to represent accurately what is outside the mind; so to understand the possibility and nature of knowledge is to understand the way in which the mind is able to construct such representations" (Cobb, Yackel, and Wood, 1992, p. 3 from Rorty, 1979).

The representational view of mind, as practiced in cognitive psychology, certainly makes no claims that the mind represents the world accurately or completely^[2], and no strong claims about the nature of knowledge as a philosophical issue. The true representational position is compatible with a broad range of notions about the relation of the mind to the world, and about the accuracy or inaccuracy and completeness or incompleteness of our internal representations of the world's features. Its claim simply:

Cognitive competence (in this case mathematical competence) depends on the availability of symbolic structures (e.g., mental patterns or mental images) that are created in response to experience.

In constructivist writings, criticisms of the straw-man position typified by the quotation from Rorty are used to discredit the actual representational view of the mind employed in cognitive psychology. As we have already pointed out in discussing the constructivist's first claim, modern cognitive theories emphatically do not assume that learning is a passive recording of experience.

The misinterpretation of the representational view leads to much confusion about external mathematical representations (e.g., equations, graphs, rules, Dienes blocks, etc.) versus internal representations (e.g., production rules, discrimination nets, mental images). Believing that the representational version of learning records these external representations passively and without transformation into distinct internal representations, constructivists take inadequacies of the external representations as inadequacies of the notion of internal representation. For instance, if a set of rules in a textbook is inadequate this is taken as an inability of production rules to capture the concepts. However, cognitive theories postulate (and provide evidence for) complex processes for transforming (assimilating and accommodating) these external representations to produce internal structures that are not at all isomorphic to the external representations.

While it is true that education has proceeded for centuries without a theory of internal representation, this is no reason to ignore the theories that are now coming from cognitive psychology. Consider the analogy of medicine. For thousands of years before there was any real knowledge of human physiology, remedies for some pathological conditions were known and used, sometimes effectively, by both doctors and others. But the far more powerful methods of modern medicine were developed concurrently with the development of modern physiology and biochemistry, and are squarely based on the latter developments. To acquire powerful interventions in disease, we had to deepen our understanding of the mechanisms of disease--of what was going on in the diseased body.

In the same way, human beings have been learning, and have been teaching their offspring, since the dawn of our species. We have a reasonably powerful "folk medicine," based on lecturing and reading and apprenticeship and tutoring, aided by such technology as paper and the blackboard--a folk medicine that does not demand much knowledge about what goes on in the human head during learning and that has not changed radically since schools first emerged. To go beyond these traditional techniques, we must follow the example of medicine and build (as we have been doing for the past thirty or forty years) a theory of the information processes that underlie skilled performance and skill acquisition: that is to say, we must have a theory of the ways in which knowledge is represented internally, and the ways in which such internal representations are acquired. In fact, cognitive psychology has now progressed a long way toward such a theory, and, as we have seen, a great deal is already known that can be applied, and is beginning to be applied, to improve learning processes.

In summary, contrary to the claim that knowledge cannot be represented symbolically, the evidence indicates the following actual state of affairs:

Symbols are much more than formal expressions.

Any kind of pattern that can be stored and can refer to some other pattern, say, one in the external world, is a symbol, capable of being processed by an information-processing system.

Cognitive competence (in this case mathematical competence) depends on the availability of symbolic structures (e.g., mental patterns or mental images) that are created in response to experience.

Cognitive theories postulate (and provide evidence for) complex processes for transforming (assimilating and accommodating) these external representations to produce internal structures that are quite different from the external representations.

Today instruction is based in large part on "folk psychology." To go beyond these traditional techniques, we must continue to build a theory of the ways in which knowledge is represented internally, and the ways in which such internal representations are acquired.

Claim 3: Knowledge can only be communicated in complex learning situations

Part of the "magical" property of knowledge asserted in the second claim, that there is something in the nature of knowledge that cannot be represented symbolically, is that no simple instructional situation suffices to convey the knowledge, whatever it may be. This assertion is the final consequence of rejecting decontextualization. Thus, constructivists recommend, for example, that children learn all or nearly all of their mathematics in the context of complex problems (e.g., Lesh & Zawojeski, 1992). This recommendation is put forward without any evidence as to its educational effectiveness.

There are two serious problems with this approach, both related to the fact that a complex task will call upon a large number of competences. First, as we noted earlier with respect to part training, a learner who is having difficulty with many of the components can easily be overwhelmed by the processing demands of the complex task. Second, to the extent that many components are well mastered, the student will waste a great deal of time repeating these mastered components to get an opportunity to practice the few components that need additional effort.

There are, of course, reasons sometimes to practice skills in their complex setting. Some of the reasons are motivational and some reflect the special skills that are unique to the complex situation. The student who wishes to play violin in an orchestra would have a hard time making progress if all practice were attempted in the orchestra context. On the other hand, if the student never practiced as a member of an orchestra, critical skills unique to the orchestra would not be acquired. The same arguments can be made in the sports context, and motivational arguments can also be made for complex practice in both contexts. A

child may not see the point of isolated exercises, but will when they are embedded in the real-world task. Children are motivated to practice sports skills because of the prospect of playing in full-scale games. However, they often spend much more time practicing component skills than full-scale games. It seems important both to motivation and to learning to practice one's skills from time to time in full context, but this is not a reason to make this the principal mechanism of learning.

While there may be motivational merit to embedding mathematical practice in complex situations, Geary (1995) notes that there is a lot of reason to doubt how intrinsically motivating complex mathematics is to most students in any context. The kind of sustained practice required to develop excellence in an advanced domain is not inherently motivating to most individuals and requires substantial family and cultural support (Ericsson, Krampe, & Tesch-Romer, 1993). Geary argues, as have others (e.g., Bahrack & Hall, 1991; Stevenson & Stigler, 1992), that it is this difference in cultural support that accounts for the large difference in mathematics achievement between Asian and American children.

Contrary to the contention that knowledge can always be communicated best in complex learning situations, the evidence shows that:

A learner who is having difficulty with components can easily be overwhelmed by the processing demands of a complex task. Further, to the extent that many components are well mastered, the student wastes much time repeating these mastered components to get an opportunity to practice the few components that need additional effort.

There are reasons sometimes to practice skills in their complex setting. Some of the reasons are motivational and some reflect the skills that are unique to the complex situation. While it seems important both to motivation and to learning to practice skills from time to time in full context, this is not a reason to make this the principal mechanism of learning.

Claim 4: It is not possible to apply standard evaluations to assess learning

The denial of the possibility of objective evaluation could be the most radical and far-reaching of the constructivist claims. We put it last because it is not clear how radically this principle is interpreted by all constructivists. Certainly, some constructivists have engaged in rather standard evaluations of constructivist learning interventions (e.g., Cobb, Wood, Yackel, Nicholls, Wheatley, Trigaitti, & Perlwitz, 1992). However, others are very uncomfortable with the idea of evaluation. As Jonassen (1992) writes:

"If you believe, as radical constructivists do, that no objective reality is uniformly interpretable by all learners, then assessing the acquisition of such a reality is not possible. A less radical view suggests that learners will interpret perspectives differently, so evaluation processes should accommodate a wider variety of response options." (p. 144).

In the hands of the most radical constructivists, Claim 4 implies that it is impossible to evaluate any educational hypothesis empirically because any such test necessarily requires a commitment to some arbitrary, culturally-determined, set of values. In the hands of the more moderate constructivists, the claim manifests itself in advocacy of focusing evaluation on the process of learning more than the product, in what are considered "authentic" tasks, and by involving multiple perspectives in the evaluation.

This milder perspective calls for emphasis on more subjective and less precisely defined instruments of evaluation. While we share with most educators their instinctive distaste of four-alternative forced-choice questions and we agree that mathematics assessment should go beyond merely testing computational skills, we question whether the very open-ended assessment being advocated as the proper alternative will lead to either more accurate or more culture-free assessment. The fundamental problem is a failure to specify precisely the competence being tested for and a reliance on subjective judgment instead. We examined a number of recent papers in Wirzup and Streit (1992) addressing this issue. In one paper, Resnick, Briars, and Lesgold (1992) present two examples of answers that are objectively equivalent (and

receive equal scores in their objective assessment scheme). However, they are uncomfortable with this equal assessment and feel a subjective component should be added so one answer would receive a higher score because it displayed greater "communication proficiency." Although the "better" answer had neater handwriting, one might well judge it as just more long-winded than the "worse" answer. "Communication proficiency" is very much in the eyes of the beholder. In another paper, Dossey (1992), in explaining the new NAEP open-ended scoring, states that a student will be given 50% (2 points) for the right answer if the justification for the answer is "not understandable" but will be given 100% (4 points) for the wrong answer if it "does not reflect misunderstanding of either the problem or how to implement the strategy, but rather **seems** to be a copying error or computational error." While we are sympathetic with the sentiments behind such ideas, such subjective judgments will open the door to a great deal of cultural bias in assessment (Rist, 1970). Anytime the word "seems" appears in an assessment, it should be a red flag that the assessors do not know what they are looking for. The information-processing approach would advocate precisely specifying what one is looking for in terms of a cognitive model and then precisely testing for that.

Another sign of the constructivist's discomfort with evaluation manifests itself in the motto that the teacher is the novice and the student the expert (e.g., see papers in von Glasersfeld, 1991). The idea is that every student gathers equal value from every learning experience. The teacher's task is to come to understand and value what the student has learned. As Confrey (1991) writes:

"seldom are students' responses careless or capricious. We must seek out their systematic qualities which are typically grounded in the conceptions of the student...frequently when students' responses deviate from our expectations, they possess the seeds of alternative approaches which can be compelling, historically supported and legitimate if we are willing to challenge our own assumptions." (p. 122)

or as Cobb, Wood, and Yackel (1991) write:

"The approach respects that students are the best judges of what they find problematical and encourages them to construct solutions that they find acceptable given their current ways of knowing." (p. 158).

If the student is supposed to move, in the course of the learning experiences, from a lower to a higher level of competence, we wonder why the student's judgments of the acceptability of solutions are particularly valid. While we value the teacher who can appreciate children's individuality, see their insights and motivate them to do their best and to value learning, there must be definite educational goals. More generally, if the "student as judge" attitude were to dominate education, it would no longer be clear when instruction had failed and when it had succeeded, when it was moving forward and when backward. It is one thing to understand why the student, at a particular stage in understanding, is doing what he or she is doing. It is quite another matter to help the student understand how to move from processes that are "satisfactory" in a limited range of tasks to processes that are more effective over a wider range. As Resnick (1994) argues, many concepts which children naturally come to (e.g., that motion implies force) are not what the culture expects of education and that in these cases "education must follow a different path: still constructivist in the sense that simple telling will not work, but much less dependent on untutored discovery and exploration (p. 489)."

Again, we find important empirical reasons for proceeding in assessment in somewhat different ways from those recommended by constructivists, and particularly, the more radical among them:

We all share an instinctive distaste for four-alternative forced-choice questions, but these are not required to attain validity or reliability in assessment. Accurate and culture-free assessment does require, however that the competence being tested for to be specified precisely without undue reliance on subjective judgment. Subjective judgments open the door to cultural bias in assessment

It cannot be assumed that students' judgments of the acceptability of solutions are particularly valid. If the "student as judge" view were adopted, it would no longer be clear when instruction had failed and when it had succeeded

Summary: Constructivism

To argue for radical constructivism seems to us to engender deep contradictions. Radical constructivists cannot argue for any particular agenda if they deny a consensus as to values. The very act of arguing for a position is to engage in a value-loaded instructional behavior. It would seem that radical constructivists should present us with data about the consequences of various educational alternatives and allow us to construct our own interpretations. (But data beyond anecdotes are rare in such constructivist writings.)

It is not clear how many of those who describe themselves as constructivists really subscribe to an outright rejection of evaluation and instruction. A less radical constructivism may contain no contradictions and may bear some truth. However, to repeat our conclusion with respect to situated learning, such a moderate constructivism contains little that is new and ignores a lot that is already known.

What is to be Done?

In the preceding pages of this paper, we have questioned a number of the basic claims of situated learning and constructivism, but our own recommendations for educational research and practice have mainly been left implicit. In this final section we set forth briefly a program of research and action that is based on the information-processing approach to of cognitive. We will address research first, then instructional practice.

Recommendations for Research

Educational research needs to understand both the component processes that are involved in intellectual tasks and the ways in which these processes must interact for good performance on complex tasks. Of course, as most complex skills are hierarchical in structure, with component skills within component skills, and so on, the inquiries must be carried out at several levels. At the highest level, we should study the structure of real-world skills in both laboratory and real-world settings. Such study requires clear statements of the educational goals--the knowledge and skills aimed at--and careful design of procedures for assessing the degree to which the goals have been achieved. The research will need to study performance at various skill levels, from novice to expert, and to employ a variety of observational methods, including the analysis of verbal and video protocols and the computer modeling of processes--methods that have only recently been refined as a part of the psychological research armory. These methods can yield a specification of the cognitive structures which we want students to acquire. With this cognitive specification in hand, we can use recent learning-by-doing theories in psychology to guide instruction.

It is also important that we do careful empirical study of the instructional programs developed under the information-processing approach and evaluate them carefully in comparison with alternatives. Evaluation should include not only (and perhaps not mainly) the immediate learning effects of instruction for tasks like those used in training, but particularly (1) the retention of knowledge and skills after a substantial time has elapsed from the completion of training (months or even years), and (2) the transferability of the knowledge and skills to a broader range of tasks than those used in the instruction. To take an obvious example from mathematics, research on calculus instruction should be evaluated in large measure (except, possibly, for mathematics majors) by assessing the ability and propensity of students to use the calculus successfully when it is relevant in their work in physics or economics.

There is unanimous agreement that what is desired is not rote learning but learning with understanding. We need research that will tell us how to assess better than we do now when a student is performing by rote, and when and to what degree understanding has been achieved. For a long time there has been evidence (Katon, 1940) that knowledge and skill acquired with understanding is retained better and transferred better than that which is acquired by rote. If this relation can be further validated, tests of retention and transfer can be used to assess understanding and, conversely, achievement of understanding can be used as a predictor of retention and transfer. It would be highly desirable, also, to devise procedures that would help students to assess their own levels of understanding.

Among the processes that have been shown by recent research to have considerable power in speeding the learning process and encouraging the learner to achieve deeper levels of understanding are learning from examples and learning by doing. Computer tutors, using these and other methods, are beginning to show impressive effectiveness, and methods of these sorts can also be implemented with paper and pencil.

There is almost universal consensus that only the active learner is a successful learner. Proponents of situated learning and constructivism have proposed a number of modes of instruction that are aimed at encouraging initiative from students and interaction among them. While we have criticized some of the assumptions underlying current proposals for "child-centered" procedures as both implausible and lacking empirical evidence, we fully agree that the social structure of the environment in which education takes place is of utmost importance from a cognitive, and especially from a motivational, standpoint.

Recommendations for Instruction

We need to be more tentative in our recommendations for instructional methods than in our recommendations for research. Nevertheless, there is already considerable empirical support for the superiority, relative to mainstream classroom methods, of a number of procedures (like the learning-from-examples and learning-by-doing methods already mentioned) that are ready for classroom testing on a large scale.

The use with children of experimental methods, that is, methods that have not been finally assessed and found effective, might seem difficult to justify. Yet the traditional methods we use in the classroom every day have exactly this characteristic--they are highly experimental in that we know very little about their educational efficacy in comparison with alternative methods. There is widespread cynicism among students and even among practiced teachers about the effectiveness of lecturing or repetitive drill (which we would distinguish from carefully designed practice), yet these methods are in widespread use. Equally troublesome, new "theories" of education are introduced into schools every day (without labeling them as experiments) on the basis of their philosophical or common-sense plausibility but without genuine empirical support. We should make a larger place for responsible experimentation that draws on the available knowledge--it deserves at least as large a place as we now provide for faddish, unsystematic and unassessed informal "experiments" or educational "reforms." We would advocate the creation of a "FEA" on analogy to the FDA which would require well designed clinical trials for every educational "drug" that is introduced into the market place.

Overall Conclusions

Given that so much educational reform is presented as a response to the excesses of behaviorism, it is interesting to read the conception of good education from one of the foremost proponents of behaviorism, B. F. Skinner. In his classic Novel, *Walden II*, intended to innovate behaviorism, Skinner's hero Frazier says:

Since our children remain happy, energetic, and curious, we don't need to teach "subjects" at all. We teach only the techniques of learning and thinking. As for geography, literature, the sciences--we give our children opportunity and guidance, and they learn for themselves. In that we dispense with half the

teachers required under the old system, and the education is incomparably better. Our children are not neglected, but they're seldom, if ever, taught anything.

Education in Walden Two is part of the life of the community. We don't need to resort to trumped-up life experiences. Our children begin to work at a very early age. It's no hardship; it's accepted as readily as sport or play. A good share of our education goes on in workshops, laboratories, and fields. It's part of the Walden Two code to encourage children in all the arts and crafts. (Skinner, 1948, p. 119-120).

Cognitive psychology rose up in response to the simplistic conception of human exemplified by the behaviorist views of Skinner, which he represented in Frazier's views. We see that influential schools have arisen, claiming a basis in cognitive psychology, that are advocating Frazier's program but which have almost no grounding in cognitive theory and at least as little grounding in empirical fact. This is particularly grievous because we think information-processing psychology has a lot to offer to mathematics education.

Information-processing psychology would propose that any effective educational practice should begin with detailed, precise cognitive task analysis. This requires first identifying what competences mathematics education seeks to foster. Having these specified, one then has to engage in the labor-intensive process of developing cognitive models that embodied these skills. With these in hand, one can bring to bear well-established principles of learning to facilitate students' acquisition of the cognitive components.

There are a number of ways to implement this agenda. One of the authors (Anderson, Corbett, Koedinger & Pelletier, 1995) has been involved in an effort to follow this program in designing computer tutors in America. Another of the authors (Zhu & Simon, 1988) has been involved in an effort to achieve this in China with paper-and-pencil technology. Both of these efforts have resulted in significant achievement gains--students have learned more and faster than they did by traditional methods. While there have been such local successes, we must conclude that this kind of effort will fail to have any meaningful impact on mathematical competence in America until there is some consensus on the goals of mathematics education and a careful and detailed cognitive analysis has been launched of how to achieve these goals. Current situated and constructivist trends in mathematics education are preventing this from happening because they refuse to focus on details and precise specifications, believing that this would amount to accepting the supposedly discredited tenets of decomposition and decontextualization.

The evidence for such information-processing approaches to education, however incomplete, is enormously stronger than the evidence for the opposite approaches, supposedly based in cognitive psychology, that are currently dominating discussions of mathematics education. And this is our main message: A program of educational reform is being adopted with weak empirical and theoretical bases while a better, and better validated, program stands ready for further development and application, and that is a situation that should be and can be altered.

References

- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Anderson, J. R. (1993). *Rules of the Mind*. Hillsdale, NJ: Erlbaum.
- Anderson, J. R., Corbett, A. T., Koedinger, K., & Pelletier, R. (1995). Cognitive tutors: Lessons learned. *The Journal of Learning Sciences*, 4, 167-207.
- Ausubel, D. P. (1968). *Educational psychology: A cognitive view*. New York: Holt, Rinehart and Winston.
- Bahrick, H. P. & Hall, L. K. (1991). Lifetime maintenance of high school mathematics content. *Journal of Experimental Psychology: General*, 120, 22-33.

- Bassok, M. & Holyoak, K. J. (1989). Transfer of domain-specific problem solving procedures. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 522-533.
- Biederman, I., & Shiffrar, M. (1987). Sexing day-old chicks: A case study and expert systems analysis of a difficult perceptual learning task. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 13, 640-645.
- Bjork, R. A. & Richardson-Klavehn, A. (1989). On the puzzling relationship between environment context and human memory. In C. Izawa (Ed.) *Current issues in cognitive processes: The Tulane Flowerree Symposium on Cognition*. Hillsdale, NJ: Erlbaum.
- Bobrow, D. G. & Bower, G. H. (1969). Comprehension and recall of sentences. *Journal of Experimental Psychology*, 80, 455-461.
- Brousseau, G. (1984). The crucial role of the didactical contract in the analysis and construction of situations in teaching and learning mathematics. In H.-G. Steiner (Ed.), *Theory of mathematics education* (pp. 110-119). Occasional paper 54. Bielefeld, Germany: University of Bielefeld, Institut für Didaktik de Mathematik.
- Brown, A. L. (1994). The advancement of learning. *Educational Researcher*, 23, 4-12.
- Brown, A. L. & Campione, J. C. (1994). Guided discovery in a community of learners. In K. McGilly (ed.) *Classroom lessons: Integrating cognitive theory and classroom practice*, Chapter 9, 229-270. Cambridge, MA: MIT Press.
- Brown, J. S., Collins, A. & Digiuid, P. (1988). Situated cognition and the culture of learning. Technical report form the Institute for Research on Learning, No. IRL88-0008.
- Carraher, T. N., Carraher, D. W. & Schliemann, A. D. (1985). Mathematics in the streets and in the schools. *British Journal of Developmental Psychology*, 3, 21-29.
- Card, S. K., Moran, T. P. & Newell, A. (1983). *The psychology of human-computer interaction*. Hillsdale, NJ: Erlbaum.
- Charney, D.H., Reder, L. M., & Kusbit, G.W. (1990). Goal setting and procedure selection in acquiring computer skills: A Comparison of tutorials, problem-solving, and learner exploration. *Cognition and Instruction*, 7(4), 323-342.
- Cheng, P. W., Holyoak, K. J., Nisbett, R. E., & Oliver, L. M. (1986). Pragmatic versus syntactic approaches to training deductive reasoning. *Cognitive Psychology*, 18, 293 - 328.
- Cobb, P. (1990). A constructivist perspective on information-processing theories of mathematical activity. *International Journal of Educational Research*, 14, 67-92.
- Cobb, P., Wood, T., & Yackel, E. (1991). A constructivist approach to second grade mathematics. In E. von Glasersfeld (Ed.) *Radical constructivism in mathematics education*. Dordrecht, The Netherlands: Kluwer Academic Publishers.
- Cobb, P., Wood, T., Yackel, E., Nicholls, J., Wheatley, G., Trigatti, B., & Perlwitz, M. (1991). Assessment of a problem-centered second grade mathematics project. *Journal for Research in Mathematics Education*, 22, 3-29.

Cobb, P., Yackel, E. & Wood, T. (1992). A constructivist alternative to the representational view of mind in mathematics education. *Journal for Research in Mathematics Education*, 23, 2-33.

Collins, A., Brown, J. S., & Newman, S. (1989). Cognitive apprenticeship: Teaching students the craft of reading, writing, and mathematics. In L. B. Resnick (Ed.), *Knowing, learning, and instruction: Essays in honor of Robert Glaser*. Hillsdale, NJ: Erlbaum.

Confrey, J. (1991). Learning to listen: A student's understanding of powers of ten. In E. von Glasersfeld (Ed.), *Radical constructivism in mathematics education.*, p. 111-138. Dordrecht, The Netherlands: Kluwer Academic Publishers.

Dossey, J. A. (1992). Assessing mathematics: enhancing understanding. In I. Wirszup & R. Streit (Eds.). *Developments in school mathematics education around the world, Volume 3.*, p. 208-222. Proceedings of the Third UCSMP International Conference on Mathematics Education. Reston, VA: NCTM.

Duncker, K. (1945). On problem-solving (translated by L. S. Lees). *Psychological Monographs*, 58, No. 270.

Eich, E. (1985). Context, memory, and integrated item/context imagery. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11, 764-770.

Elio, R. (1986). Representation of similar well-learned cognitive procedures. *Cognitive Science*, 10, 41-74.

Ericsson, K. A., Krampe, R. T. & Tesche-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, 100, 363-406.

Ericsson, K. A. & Simon, H. A. (1993). *Protocol analysis: Verbal reports as data*, Revised edition. Cambridge, MA: MIT Press.

Feigenbaum, E. A. & Simon, H. A. (1984). EPAM-like models of recognition and learning. *Cognitive Science*, 8, 305-336.

Fernandez, A. & Glenberg, A. M. (1985). Changing environmental context does not reliably affect memory. *Memory & Cognition*, 13, 333-345.

Fong, G. T., Krantz, D. H., & Nisbett, R. E. (1986). The effects of statistical training on thinking about everyday problems. *Cognitive Psychology*, 18, 253-292.

Geary, D. C. (1995). Reflections of evolution and culture in children's cognition. *American Psychologist*, 50, 24-37.

Gick, M. L. & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology*, 12, 306-355.

Gick, M. L. & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, 15, 1-38.

Godden, D. R. & Baddeley, A. D. (1975). Context-dependent memory in two natural environments: On land and under water. *British Journal of Psychology*, 66, 325-331.

Hayes, J. R. (1985). The problems in teaching general skills. In J. Segal, S. Chipman & R. Glaser (Eds.), *Thinking and learning*, Vol. 2. Hillsdale, NJ: Erlbaum.

Hayes, J. R. & Simon, H. A. (1977). Psychological differences among problem isomorphs. In J. Castellan, D. B. Pisoni, and G. Potts (Eds.) *Cognitive theory*, volume 2. Hillsdale, NJ: Erlbaum.

Hendrickson, G. & Schroeder, W. H. (1941). Transfer training in learning to hit a submerged target. *Journal of Educational Psychology*, 32, 205-213.

Hiebert, J., Wearne, D., Carpenter, T. P., Fennema, E., Fuson, K., Human, P., Olivier, A., & Murray, H. (1994). Authentic problem solving in mathematics. Presented at the annual meeting of the American Educational Research Association (AERA), New Orleans.

Hunter, J. E. & Hunter, R. F. (1984). Validity and utility of alternative predictors of job performance. *Psychological Bulletin*, 96, 72-98.

Jonassen, D. H. (1992). Evaluating constructivist learning. In T. M. Duffy & D. H. Jonassen (Eds.). *Constructivism and the technology of instruction*, Chapter 12. Hillsdale, NJ: Erlbaum.

Johnson, D. W. & Johnson, R. T. (1989). *Cooperation and competition: Theory and research*. Edina, MN: Interaction Book Company.

Judd, C. H. (1908). The relation of special training to general intelligence. *Educational Review*, 36, 28-42.

Katona, G. *Organizing and memorizing*. (1940) New York, NY: Columbia University Press.

Knerr, C. M., Morrison, J. E., Mumaw, R. J. Stein, D. J., Sticha, P. J., Hoffman, R. G., Buede, D. M. & Holding, D. M. (1987). Simulation-based research in part-task training. AF HRL-TR-86-12, AD-B107293. Brooks Air Force Base, TX: Air Force Human Resources Laboratory.

Kotovsky, K. & Fallside, D. (1989). Representation and transfer in problem solving. In D. Klahr, D. & K. Kotovsky (Eds.) *Complex information processing: The impact of Herbert A. Simon* (pp. 69-108). Hillsdale, NJ: Erlbaum.

Larkin, J. H. (1981). Enriching formal knowledge: A model for learning to solve textbook physics problems. In J. R. Anderson (Ed.) *Cognitive Skills and their Acquisition*. Hillsdale, NJ: Erlbaum, 1981.

Lave, J. (1986). Experiments, tests, jobs and chores: How we learn to do what we do. In K. Borman and J. Reisman (Eds.), *Becoming a worker*. Norwood, NJ: Albex.

Lave, J. (1988). *Cognition in practice: Mind, mathematics, and culture in everyday life*. New York: Cambridge University Press.

Lesh, R. & Lamon, S. J. (1992). *Assessment of authentic performance in school mathematics*. Washington, DC: AAAS Press.

Lesh, R. & Zawojewski, J. S. (1992). Problem solving. In T. R. Post (Ed), *Teaching mathematics in grades K-8: Research-based methods*. p. 49-88. Needham Heights, MA: Allyn and Bacon.

National Research Council (1994). *Learning, remembering, believing: Enhancing human performance*. In D. Druckman & R. A. Bjork (Eds.). Washington, DC: National Academy Press.

Newell, A. & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.

Patrick, J. (1992). *Training: Research and practice*. San Diego, CA: Academic Press.

- Pirolli, P. L., & Anderson, J. R. (1985). The role of practice in fact retrieval. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *11*, 136-153.
- Reder, L. M. & Klatzky, R. (1994) Transfer: Training for Performance. In D. Druckman & R. A. Bjork (Eds.) *Learning, Remembering, Believing: Enhancing team and individual performance* (pp. 25-56). Washington, DC: National Academy Press.
- Reder, L.M. & Ritter, F. (1992) What determines initial feeling of knowing? Familiarity with question terms, not with the answer. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 435-451.
- Reed, S. K., & Actor, C.A. (1991) Use of examples and procedures in problem solving. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *17*, 753-766.
- Reed, S. K., Ernst, G. W., & Banjeri, R. (1974). The role of analogy in transfer between similar problem states. *Cognitive Psychology*, *6*, 436-450
- Resnick, L. B. (1994). Situated rationalism: biological and social preparation for learning. Mapping the mind: domain specificity in cognition and culture (474-493). Cambridge, NY: Cambridge University Press.
- Resnick, L. B., Briars, D. & Lesgold, S. (1992). Certifying accomplishments in mathematics: The new standards examining system. In I. Wirszup & R. Streit (Eds.). *Developments in school mathematics education around the world, Volume 3.*, p. 186-207. Proceedings of the Third UCSMP International Conference on Mathematics Education. Reston, VA: NCTM.
- Resnick, L. B. & Resnick, D. P. (1992). Assessing the thinking curriculum: New tools for educational reform. In B. R. Gifford & M.C. O'Connor (Eds.) *Changing assessments: Alternative views of aptitude, achievement and instruction*. Boston, MA: Kluwer.
- Rist, R.C. (1970). Student social class and teacher expectations: the self-fulfilling prophecy in ghetto education. *Harvard Educational Review*, *40*, 411-451.
- Rorty, R. (1979). *Philosophy and the mirror of nature*. Princeton, NJ: Princeton University Press.
- Salas, E., Dickinson, T. L., Converse, S. A. & Tannenbaum, S. I. (1993) Toward an understanding of team performance and training. In R. W. Swezey and E. Salas (Eds.), *Teams: Their training and performance*. Norwood, NJ: Ablex.
- Salomon, G., & Globerson, T. (1989). When teams do not function the way they ought to. *International Journal of Educational Research*, *13*, 89-98.
- Saufley, W. H., Otaka, S. R., & Bavaresco, J. L. (1985). Context effects: Classroom tests and context independence. *Memory & Cognition*, *13*, 522-528.
- Schoenfeld, A. H. (1985). *Mathematical problem solving*. New York: Academic Press.
- Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, *3*, 414-424.
- Shepard, L. A. (1991). Psychometricians' beliefs about learning. *Educational Researcher*, *20*, 2-16.
- Silver, E. A. (1987). Foundations of cognitive theory and research for mathematics problem-solving. In A. H. Schoenfeld (Ed.) *Cognitive science and mathematics education*. Hillsdale, NJ: LEA.

- Simon, H. A. (1967). Motivational and emotional controls of cognition. *Psychological Review*, 74, 29-39.
- Simon, H. A. (1981). *The sciences of the artificial* Second edition. Cambridge, MA: MIT Press.
- Simon, H. A. (1994). The bottleneck of attention: Connecting thought with motivation. In W. E. Spaulding (Ed.). *Integrative views of motivation, cognition, and emotion*. Lincoln, NE: University of Nebraska Press.
- Simon, H. A. & Feigenbaum, E. A. (1964). In information-processing theory of some effects of similarity, familiarization, meaningfulness in verbal learning. *Journal of Verbal Learning and Verbal Behavior*, 3, 385-396.
- Singley, M. K. & Anderson, J. R. (1989). *Transfer of Cognitive Skill*. Cambridge, MA: Harvard University Press.
- Skinner, B. F. (1948). *Walden two*. New York: Macmillan, Co.
- Slamecka, N. J. & Graf, P. (1978). The generation effect: Delineation of a phenomenon. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 4, 592-604.
- Slamecka, N. J. & Katsaiti, L. T. (1987). The generation effect as an artifact of selective displaced rehearsal. *Journal of Memory and Language*, 26, 589-607.
- Slavin, R. E. (1990). *Cooperative learning: Theory, research and practice*. Englewood Cliffs, NJ: Prentice-Hall.
- Smith, S. B. (1986). *An analysis of transfer between Tower of Hanoi isomorphs*. Unpublished doctoral dissertation, Carnegie-Mellon University, Pittsburgh.
- Smith, S. M., Glenberg, A., & Bjork, R. A. (1978). Environmental context and human memory. *Memory & Cognition*, 6, 342-353.
- Stein, B. S. & Bransford, J. D. (1979). Constraints on effective elaboration: Effects of precision and subject generation. *Journal of Verbal Learning and Verbal Behavior*, 18, 769-777.
- Stevenson, H. W., & Stigler, J. W. (1992). *The learning gap: Why our schools are failing and what we can learn from Japanese and Chinese education*. New York: Summit Books.
- Thorndike, E. L., & Woodworth, R. S. (1901). The influence of improvement in one mental function upon the efficiency of other functions. *Psychological Review*, 9, 374-382.
- Vera, A. H. & Simon, H. A. (1993). Situated action: A symbolic interpretation. *Cognitive Science*, 17, 7-48.
- von Glasersfeld, E. (1991). *Radical constructivism in mathematics education*. Dordrecht, The Netherlands: Kluwer Academic Publishers.
- Weisberg, R. DiCamillo, M. & Phillips, D. (1985). Transferring old associations to new situations: a non-automatic process. *Journal of Verbal Learning and Verbal Behavior*, 17, 219-228.
- Wirszup, I. & Streit, R. (1992). *Developments in school mathematics education around the world, Volume 3*. Proceedings of the Third UCSMP International Conference on Mathematics Education. Reston, VA: NCTM.

Woodworth, R. S. (1938). *Experimental psychology*. New York, NY: Henry Holt & Company.

Zhu, X. & Simon, H. A. (1988). Learning mathematics from examples and by doing. *Cognition and Instruction*, 4, 137-166.