



I L L I N O I S

---

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

PRODUCTION NOTE

University of Illinois at  
Urbana-Champaign Library  
Large-scale Digitization Project, 2007.



THE LIBRARY OF THE  
DEC 3 1987  
UNIVERSITY OF ILLINOIS  
URBANA-CHAMPAIGN

**Technical Report No. 403**

**COGNITIVE APPRENTICESHIP:  
TEACHING THE CRAFT OF READING,  
WRITING, AND MATHEMATICS**

**Allan Collins  
BBN Laboratories  
John Seely Brown  
Susan E. Newman  
Xerox Palo Alto Research Center**

**January 1987**

# **Center for the Study of Reading**

## **TECHNICAL REPORTS**

**UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN  
174 Children's Research Center  
51 Gerty Drive  
Champaign, Illinois 61820**



# CENTER FOR THE STUDY OF READING

Technical Report No. 403

## COGNITIVE APPRENTICESHIP: TEACHING THE CRAFT OF READING, WRITING, AND MATHEMATICS

Allan Collins  
BBN Laboratories  
John Seely Brown  
Susan E. Newman  
Xerox Palo Alto Research Center

January 1987

University of Illinois  
at Urbana-Champaign  
51 Gerty Drive  
Champaign, Illinois 61820

Bolt Beranek and Newman Inc.  
10 Moulton Street  
Cambridge, Massachusetts 02238

**This research was supported by the National Institute of Education under Contract No. US-NIE-C-400-81-0030 and the Office of Naval Research under Contract No. N00014-85-C-0026. We thank Marlene Scardamalia, Frank Fischer, Lauren Resnick, Alan Schoenfeld, Tom Malone, and Andee Rubin for comments and editing help on this manuscript. In L. B. Resnick (Ed.), Cognition and instruction: Issues and agendas. Hillsdale, NJ: Lawrence Erlbaum Associates, in press. This is a substantially revised version of an earlier draft entitled The New Apprenticeship: Teaching Students the Craft of Reading, Writing, and Mathematics.**

**Abstract**

It is only in the last century, and only in industrialized nations, that formal schooling has emerged as a widespread method of educating the young. Before schools, apprenticeship was the most common means of learning, used to transmit the knowledge required for expert practice in fields from painting and sculpting to medicine and law. Even today, many complex and important skills, such as those required for language use and social interaction, are learned informally through apprenticeshiplike methods--i.e., methods involving not didactic teaching, but observation, coaching, and successive approximation while carrying out a variety of tasks and activities.

The differences between formal schooling and apprenticeship methods are many, but for our purposes, one is most important. Perhaps as a by-product of the specialization of learning in schools, skills and knowledge taught in schools have become abstracted from their uses in the world. In apprenticeship learning, on the other hand, target skills are not only continually in use by skilled practitioners, but are instrumental to the accomplishment of meaningful tasks. Said differently, apprenticeship embeds the learning of skills and knowledge in the social and functional context of their use. This difference is not academic, but has serious implications for the nature of the knowledge that students acquire. This paper attempts to elucidate some of those implications through a proposal for the retooling of apprenticeship methods for the teaching and learning of cognitive skills. Specifically, we propose the development of a new cognitive apprenticeship to teach students the thinking and problem-solving skills involved in school subjects such as reading, writing, and mathematics.

The organization of the paper is as follows: In the first section, we discuss briefly what we believe to be key shortcomings in current curricular and pedagogical practices. We then present some of the structural features of traditional apprenticeship and discuss, in general, what would be required to adapt these characteristics to the teaching and learning of cognitive skills.

In the second section we consider in detail three recently developed pedagogical "success models," which we believe exemplify aspects of apprenticeship methods in teaching the thinking and reasoning skills involved in reading, writing and mathematics. We attempt to show how and why these methods are successful with regard to the development of not only the cognitive, but also the metacognitive, skills required for true expertise.

In the final section, we organize our ideas on the purposes and characteristics of successful teaching into a general framework for the design of learning environments, where "environment" includes the content being taught, the pedagogical methods employed, the sequencing of learning activities, and the sociology of learning. This framework emphasizes how cognitive apprenticeship goes beyond the techniques of traditional apprenticeship. We hope it will be useful to the field in designing, evaluating, and doing research on pedagogical methods, materials, and technologies.

## **COGNITIVE APPRENTICESHIP: TEACHING THE CRAFT OF READING, WRITING, AND MATHEMATICS**

### **1. Schooling and Apprenticeship**

**Schooling and the acquisition of expert practice.** While schools have been relatively successful in organizing and conveying large bodies of conceptual and factual knowledge, standard pedagogical practices render key aspects of expertise invisible to students. In particular, too little attention is paid to the processes that experts engage in to use or acquire knowledge in carrying out complex or realistic tasks. Where processes are addressed, the emphasis is on formulaic methods for solving "textbook" problems, or on the development of low-level subskills in relative isolation. Few resources are devoted to higher-order problem-solving activities that require students to actively integrate and appropriately apply subskills and conceptual knowledge.

As a result, conceptual and problem-solving knowledge acquired in school remains largely unintegrated or inert for many students. In some cases, knowledge remains bound to surface features of problems as they appear in textbooks and class presentations. For example, Schoenfeld (1985) has found that students rely on their knowledge of standard textbook patterns of problem presentation, rather than on their knowledge of problem-solving strategies or intrinsic properties of the problems themselves, for help in solving mathematics problems. Problems that fall outside these patterns do not invoke the appropriate problem-solving methods and relevant conceptual knowledge. In other cases, students fail to use resources available to them to improve their skills because they lack models of the processes required for doing so. For example, in the domain of writing, students are unable to make use of potential models of good writing acquired through reading because they have no understanding of the strategies and processes required to produce such text. Stuck with what Bereiter and Scardamalia (in press) call "knowledge-telling strategies," they are unaware that expert writing involves organizing one's ideas about a topic, elaborating goals to be achieved in the writing, thinking about what the audience is likely to know or believe about the subject, and so on.

In order to make real differences in students' skill, we need both to understand the nature of expert practice and to devise methods that are appropriate to learning that practice. Thus, we must first recognize that cognitive and metacognitive strategies and processes, more centrally than low-level subskills or abstract conceptual and factual knowledge, are the organizing principles of expertise, particularly in domains such as reading, writing, and basic mathematics. Further, because expert practice in these domains rests crucially on the integration of cognitive and metacognitive processes, we believe that it can best be taught through methods that emphasize what Lave (in preparation) calls successive approximation of mature practice, methods that have traditionally been employed in apprenticeship to transmit complex physical processes and skills. We propose that these methods of apprenticeship be adapted to the teaching and learning of complex cognitive skills.

**Traditional apprenticeship.** In order to get an idea of what these methods may look like and why they are likely to be effective, let us first consider some of the crucial features of traditional apprenticeship. We have relied on Lave's (in preparation) careful description of apprenticeship as practiced in a West African tailoring shop for many of our insights into the nature of apprenticeship.

First and foremost, apprenticeship highlights methods for carrying out tasks in a domain. Apprentices learn these methods through a combination of what Lave calls observation, coaching, and practice, or what we, from the teacher's point of view, call modelling, coaching, and fading. In this sequence of activities, the apprentice repeatedly observes the master executing (or modelling) the target process, which usually involves a number of different but interrelated subskills. The apprentice then attempts to execute the process with guidance and help from the master (coaching). A key aspect of coaching is the provision of scaffolding, which is the support, in the form of reminders and

help, that the apprentice requires to approximate the execution of the entire composite of skills. Once the learner has a grasp of the target skill, the master reduces his participation (fades), providing only limited hints, refinements, and feedback to the learner, who practices by successively approximating smooth execution of the whole skill.

Several points are worth emphasizing here. The interplay between observation, scaffolding, and increasingly independent practice aids apprentices both in developing self-monitoring and -correction skills, and in integrating the skills and conceptual knowledge needed to advance toward expertise. Observation plays a surprisingly key role. Lave hypothesizes that it aids learners in developing a conceptual model of the target task or process prior to attempting to execute it. Having a conceptual model is an important factor in apprenticeship's success in teaching complex skills without resorting to lengthy practice of isolated subskills, for three related reasons. First, it provides learners with an advanced organizer for their initial attempts to execute a complex skill, thus allowing them to concentrate more of their attention on execution than would otherwise be possible. Second, a conceptual model provides an interpretative structure for making sense of the feedback, hints, and corrections from the master during interactive coaching sessions. And third, it provides an internalized guide for the period of relatively independent practice by successive approximation. Moreover, development of a conceptual model, which can be continually updated through further observation and feedback, encourages autonomy in what we call reflection (Collins & Brown, in press). Reflection is the process that underlies the ability of learners to compare their own performance, at both micro and macro levels, to the performance of an expert. Such comparisons aid learners in diagnosing difficulties and incrementally adjusting their performance until they reach competence. A conceptual model serves as an internal model of expert performance, and, thus, as a basis for development of self-monitoring and -correction skills.

A second key observation about apprenticeship in general concerns the embedding social context in which learning takes place. Apprenticeship derives many (cognitively important) characteristics from its embedding in a subculture in which most, if not all, members are visible participants in the target skills. As a result, learners have continual access to models of expertise-in-use against which to refine their understanding of complex skills. Moreover, it is not uncommon for apprentices to have access to several masters and, thus, to a variety of models of expertise. Such richness and variety helps apprentices to understand that there may be multiple ways of carrying out a task and to recognize that no one individual embodies all knowledge or expertise. And finally, in the tailoring shop described by Lave, learners have the opportunity to observe other learners at varying degrees of skill; among other things, this encourages them to view learning as an incrementally staged process, while providing them with concrete benchmarks for their own progress.

**From traditional to cognitive apprenticeship.** This paper proposes a rethinking of these aspects of apprenticeship for the teaching and learning of subjects such as reading, writing, and mathematics. We call this rethinking of teaching and learning in school "cognitive apprenticeship" to emphasize two things. First, these methods are aimed primarily at teaching the processes that experts use to handle complex tasks. Where conceptual and factual knowledge is addressed, cognitive apprenticeship emphasizes its uses in solving problems and carrying out tasks. That is, in cognitive apprenticeship, conceptual and factual knowledge is exemplified and situated in the contexts of its use. Conceptual knowledge thus becomes known in terms of its uses in a variety of contexts, encouraging both a deeper understanding of the meaning of the concepts themselves and a rich web of memorable associations between important concepts and problem-solving contexts. It is this dual focus on expert processes and situated learning that we expect to help solve the educational problems of brittle skills and inert knowledge.

Second, the term cognitive apprenticeship refers to the fact that the focus of the learning-through-guided-experience is on cognitive and metacognitive, rather than on physical, skills and processes. While we do not wish to draw a major theoretical distinction between the learning of physical and cognitive skills, there are differences that have practical implications for the organization of teaching



and learning activities and teacher-learner interactions. Most importantly, traditional apprenticeship has evolved to teach domains in which the process of carrying out target skills (1) is external and thus readily available to both student and teacher for observation, comment, refinement, and correction and (2) bears a relatively transparent relationship to concrete products that are the outcome of the skill. The externalization of relevant processes and methods makes possible such characteristics of apprenticeship as its reliance on observation as a primary means of building a conceptual model of a complex target skill. And the relatively transparent relationship, at all stages of production, between process and product facilitates the learner's recognition and diagnosis of errors, upon which the early development of self-correction skills depends.

Applying apprenticeship methods to largely cognitive skills requires the externalization of processes that are usually carried out internally. At least as most subjects are taught and learned in school, teachers cannot make fine adjustments in students' application of skill and knowledge to problems and tasks, because they have no access to the relevant cognitive processes. By the same token, students do not usually have access to the cognitive problem-solving processes of instructors, as a basis for learning through observation and mimicry. Cognitive research, through such methods as protocol analysis, has begun to delineate the cognitive and metacognitive processes that heretofore have tacitly comprised expertise. Cognitive apprenticeship teaching methods are designed, among other things, to bring these tacit processes into the open, where students can observe, enact, and practice them with help from the teacher and from other students.

Cognitive apprenticeship also requires extended techniques to encourage the development of self-correction and -monitoring skills, as we cannot rely on the transparent relationship between process and product that characterizes the learning of such physical skills as tailoring. We have identified two basic means of fostering these crucial metacognitive skills. First, cognitive apprenticeship encourages reflection on differences between novice and expert performance by alternation between expert and novice efforts and by techniques that we have elsewhere called "abstracted replay" (Collins & Brown, in press). Alternation between expert and novice efforts in a shared problem-solving context sensitizes students to the details of expert performance as the basis for incremental adjustments in their own performance. Abstracted replay attempts to focus students' observations and comparisons directly on the determining features of both their own and an expert's performance by highlighting those features in a skillful verbal description, or, in some domains, through use of recording technologies such as computers or videotapes.

A second means of encouraging the development of self-monitoring and -correction skills is based on the insight that these skills require the problem solver to alternate among different cognitive activities while carrying out a complex task. Most notably, complex cognitive activities involve some version of both generative and evaluative processes. However, both types of processes are complex and can be difficult to learn in tandem. Thus, cognitive apprenticeship involves the development and externalization of a producer-critic dialogue that students can gradually internalize. This development and externalization is accomplished through discussion, alternation of teacher and learner roles, and group problem-solving.

**Some caveats.** Obviously, apprenticeship is intended as a suggestive, rather than an exact, model for teaching and learning in the future. In addition to the emphasis on cognitive and metacognitive skills, there are two major differences between cognitive apprenticeship and traditional apprenticeship. First, because traditional apprenticeship is set in the workplace, the problems and tasks that are given to learners arise not from pedagogical concerns, but from the demands of the workplace. Cognitive apprenticeship as we envision it differs from traditional apprenticeship in that the tasks and problems are chosen to illustrate the power of certain techniques or methods, to give students practice in applying these methods in diverse settings, and to slowly increase the complexity of tasks so that component skills and models can be integrated. In short, tasks are sequenced to reflect the changing demands of learning. Letting the job demands select the tasks for students to practice is one of the great inefficiencies of traditional apprenticeship.

On the other hand, the economic bias in traditional apprenticeship has useful, as well as less-than-ideal, effects. For example, apprentices are encouraged to quickly learn skills that are useful, and therefore meaningful within the social context of the workplace. Moreover, apprentices have natural opportunities to realize the value, in concrete economic terms, of their developing skill: Well-executed skills result in saleable products. Cognitive apprenticeship must find a way to create a culture of expert practice for students to participate in, and aspire to, as well as devise meaningful benchmarks and incentives for progress.

A second difference between cognitive apprenticeship and traditional apprenticeship is the emphasis in cognitive apprenticeship on decontextualizing knowledge so that it can be used in many different settings. Traditional apprenticeship emphasizes teaching skills in the context of their use. We propose that cognitive apprenticeship should extend situated learning to diverse settings so that students learn how to apply their skills in different contexts. Moreover, the abstract principles underlying the application of knowledge and skills in different settings should be articulated as fully as possible by the teacher whenever they arise in different contexts.

We do not want to argue that cognitive apprenticeship is the only way to learn. Reading a book or listening to a lecture are important ways to learn, particularly in domains where conceptual and factual knowledge are central. Active listeners or readers, who test their understanding and pursue the issues that are raised in their minds, learn things that apprenticeship can never teach. However, to the degree the reader or listener is passive, they will not learn as much as they would by apprenticeship, because apprenticeship forces them to use their knowledge. Moreover, few people learn to be active readers and listeners on their own, and that is where cognitive apprenticeship is critical--observing the processes by which an expert listener or reader thinks, and practicing these skills under the guidance of the expert, can teach students to learn on their own more skillfully.

Even in domains that rest on elaborate conceptual and factual underpinnings, students must learn the practice or art of solving problems and carrying out tasks. And to achieve expert practice, some version of apprenticeship remains the method of choice. Thus, apprenticeshiplike methods are widely used in graduate education in most domains. Students are expected to learn how to solve problems that arise in the context of carrying out complex tasks, and to extend and make use of their textbook knowledge by undertaking significant projects guided by an expert in the field.

We would argue that the development of expert practice through situated learning and the acquisition of cognitive and metacognitive skills is equally, if not more, important in more elementary domains. This is nowhere more evident than in the foundational domains of reading, writing, and mathematics. These domains are foundational not only because they provide the basis for learning and communication in other school subjects, but also because they engage cognitive and metacognitive processes that are basic to learning and thinking more generally. Unlike school subjects such as chemistry or history, these domains rest on relatively sparse conceptual and factual underpinnings, turning instead on students' robust and efficient execution of a set of cognitive and metacognitive skills. Given effective analyses and externalizable prompts for these skills, we believe that these domains are particularly well suited to teaching methods modelled on cognitive apprenticeship. In the next section of this paper we discuss a set of recently developed, and highly successful, models for teaching the cognitive and metacognitive skills involved in reading, writing, and mathematics in terms of the key notions underlying our cognitive apprenticeship model.

## **2. Three Success Models for Cognitive Apprenticeship**

**Palincsar** and Brown's reciprocal teaching of reading. Palincsar and Brown's (1984) method of teaching reading comprehension, which exemplifies many of the features of cognitive apprenticeship, has proved remarkably effective in raising students' scores on reading comprehension tests, especially those of poor readers. The basic method centers on modelling and coaching students in four

strategic skills: formulating questions based on the text, summarizing the text, making predictions about what will come next, and clarifying difficulties with the text. The method has been used with groups of two to five students, as well as individual students. It is called *Reciprocal Teaching* because the teacher and students take turns playing the role of teacher.

The procedure is as follows: Both the teacher and students read a paragraph silently to themselves. Whoever is playing the role of teacher formulates a question based on the paragraph, constructs a summary, and makes a prediction or clarification if any come to mind. Initially, the teacher models this process, eventually turning it over to the students. When students first undertake the process, the teacher coaches them extensively on how to construct good questions and summaries, offering prompts and critiquing their efforts. In this way, the teacher provides scaffolding for the students, enabling them to take on whatever portion of the task they can. As the students become more proficient, the teacher fades, assuming the role of monitor and providing occasional hints or feedback. Table 1 shows a sequence of dialogues illustrating how scaffolding is used and adjusted over time to help a student formulate questions about a series of texts.

[Insert Table 1 about here.]

Reciprocal Teaching is extremely effective. In a pilot study with individual students who were poor readers, the method raised subjects' reading comprehension test scores from 15% to 85% accuracy after about 20 training sessions. Six months later the students were still at 60% accuracy, recovering to 85% after only one session. In a subsequent study with groups of two students, the scores increased from about 20% to 60% accuracy, with very little change eight weeks later. In classroom studies with groups of five students, test scores increased from about 40% to 75% correct, again with only a slight decline eight weeks later. These are very dramatic effects for any instructional intervention.

Why is Reciprocal Teaching so effective? In our analysis, which reflects in part the views of Palincsar and Brown (Brown & Palincsar, in press, this volume; Palincsar & Brown, 1984) its effectiveness depends upon the co-occurrence of a number of factors.

First, the method engages students in a set of activities that help them form a new conceptual model of the task of reading. In traditional schooling, students learn to identify reading with the subskills of recognizing and pronouncing words and with the activities of scanning text and saying it aloud. Under the new conception, students recognize that reading requires constructive activities such as formulating questions and making summaries and predictions, as well as evaluative ones such as analyzing and clarifying the points of difficulty in the text. Moreover, carrying out these activities by repeatedly reviewing the text helps students realize that reading for understanding is often more than a one-pass operation; it provides them with a more realistic expectation about what will be required of them as they go on to read increasingly difficult texts.

Second, these activities involve the student in using the reading strategies and metacognitive skills necessary for expert reading. In particular:

1. *Formulating questions* is an important strategic activity for understanding difficult texts (Collins, Brown, & Larkin, 1980) because it provides the basis for checking if the text makes sense (self-monitoring). As we can see in Table 1, formulating questions that capture the main ideas of the text sometimes leads to questions that the text raises, but does not answer, as the basis for further inquiry.

2. *Summarizing*, like formulating questions, provides a general test of comprehension and so forms the basis for comprehension monitoring: It is a preliminary phase of self-diagnosis. Students learn that if they cannot form a good summary, then they do not understand the text and had better either reread the text or try to clarify their difficulties (Collins & Smith, 1982).
3. *Clarification* is a key activity in comprehension monitoring that involves detailed self-diagnosis, in which students attempt to isolate and formulate their particular difficulties in understanding a text. While summarizing is a fairly global test of comprehension, usually applied at the paragraph level, clarification attempts to narrow points of difficulty by focusing on word and phrase levels of meaning. Skill at clarifying difficulties provides students with the basis for using evidence from subsequent text to disambiguate the meaning of problematic words or phrases, a key strategy employed by expert readers.
4. *Prediction* involves formulating guesses or hypotheses about what the author of a text is likely to say next, and as such, promotes an overall reading strategy of hypothesis formation and testing. The inclusion of prediction as an explicit strategic activity for beginning readers reflects the fact that skilled reading involves developing expectations and evaluating them as evidence accumulates from the text (Collins & Smith, 1982).

The third factor we think is critical for the success of Reciprocal Teaching is that the teacher models expert strategies in a problem context shared directly and immediately with the students (Brown & Palincsar, in press). This organization of teacher-learner interaction encourages students first to focus their observations and then to reflect on their own performance relative to that of the teacher during subsequent modelling. Here's how it works: Both teacher and students read a paragraph. The teacher then performs the four activities: She articulates the questions she would ask about the paragraph, summarizes it, makes predictions about what would be next, and explains what part of the paragraph gave her difficulty. She may try to explain why she generated a particular question or made a particular prediction. What is crucial here is that the students listen in the context of knowing that they will soon undertake the same task, using that expectation to focus their observations on how those activities are related to the paragraph. After they have tried to do it themselves, and perhaps had difficulties, they listen to the teacher with new knowledge about the task. As they read subsequent passages, they may try to generate a question or summary to themselves, noticing later what she does differently. That is, they can compare their own question or summaries with the questions or summaries she generates. They can then reflect on any differences, trying to understand what led to those differences. We have argued elsewhere that this kind of reflection is critical to learning (Collins & Brown, in press).

Fourth, the technique of providing scaffolding is a crucial factor in the success of Reciprocal Teaching for several reasons. Most importantly, it decomposes the task as necessary for the students to carry it out, thereby helping them to see how, in detail, to go about the task. For example, in formulating questions, the teacher might first want to see if the student can generate a question on his or her own; if not, she might suggest starting a question with "Why" or "How." If the student still can't generate a question, she might suggest formulating a simple "Why" question about the agent in the story. If that fails, she might generate one herself and ask the student to reformulate it in his or her own words. In this way, it gets students started in the new skills, giving them a "feel" for the skills and helping them develop confidence that they can do them. Scaffolding is designed to help students when they are at an *impasse* (Brown & VanLehn, 1980). With successful scaffolding techniques, students get as much support as they need to carry out the task, but no more. Hints and modelling are then gradually faded out, with students taking on more and more of the task as they become more skillful. These techniques of scaffolding and fading slowly build students' confidence that they can master the skills required.

The final aspect of Reciprocal Teaching that we think is critical is having students assume the dual roles of *producer* and *critic*. That is, they must not only be able to produce good questions and summaries, but they also learn to evaluate the summaries or questions of others. By becoming critics as well as producers, students are forced to articulate their knowledge about what makes a good question, prediction, or summary. This knowledge then becomes more readily available for application to their own summaries and questions, thus improving a crucial aspect of their metacognitive skills. Moreover, once articulated, this knowledge can no longer simply reside in tacit form. It becomes more available for performing a variety of tasks; that is, it is freed from its contextual binding, so that it can be used in many different contexts.

**Scardamalia and Bereiter's procedural facilitation of writing.** Scardamalia and Bereiter (1985; Scardamalia, Bereiter, & Steinbach, 1984) have developed an approach to the teaching of writing that relies on elements of cognitive apprenticeship. Based on contrasting models of novice and expert writing strategies, the approach provides explicit procedural supports, in the form of prompts, that are aimed at helping students adopt more sophisticated writing strategies. Like other exemplars of cognitive apprenticeship, their approach is designed to give students a grasp of the complex activities involved in expertise by explicit modelling of expert processes, gradually reduced support or scaffolding for students attempting to engage in the process, and opportunities for reflection on their own and others' efforts.

According to Bereiter and Scardamalia's (in press) analysis of expert-novice differences, children who are novices in writing use a "knowledge-telling" strategy. When given a topic to write on, they immediately produce text by writing the first idea they think of, then the next idea, and so on, until they run out of ideas, at which point they are done. This is a very simple control strategy that finesses most of the difficulties in composing. In contrast, experts spend time not only writing, but also in planning what they are going to write and revising what they have written (Hayes & Flower, 1980). As a result, they engage in a process that Scardamalia and Bereiter call "knowledge transforming," which incorporates the linear generation of text, but is organized around a more complex structure of goal setting and problem solving. Scardamalia and Bereiter (1985) argue that for experts, writing is a "compositional" task in which goals are emergent, i.e., "your knowledge of what you are after grows and changes as part of the process." Emergent goals are products of the fact that "there is a wealth of potentially applicable knowledge and potential routes to the goals."

In order to encourage students to adopt a more sophisticated writing strategy, Scardamalia, Bereiter and colleagues have developed a detailed cognitive analysis of the activities of expert writers. This analysis provides the basis for a set of prompts that they call *procedural facilitation*, designed to reduce students' information-processing burden in trying to carry out complex tasks by allowing them to select from a limited number of diagnostic statements. For example, in their analysis, planning is broken down into five general processes or goals: (1) generating a new idea, (2) improving an idea, (3) elaborating an idea, (4) identifying goals, and (5) putting ideas into a cohesive whole. For each process, they have developed a number of specific prompts, designed to aid students in their planning, as shown in Table 2. These prompts, which are akin to the suggestions made by the teacher in Reciprocal Teaching, serve to simplify the complex process of elaborating and reconsidering one's plans by suggesting specific lines of thinking for students to follow. A comparable analysis and set of prompts has been developed for the revision process as well (Scardamalia & Bereiter, 1983b, 1985).

[Insert Table 2 about here.]

Scardamalia and Bereiter's teaching method, like Reciprocal Teaching, proceeds through a combination of modelling, coaching, scaffolding, and fading. First the teacher models how to use the prompts, which are written on cue cards, in generating ideas about a topic she is going to write on. Table 3 illustrates the kind of modelling done by a teacher during an early phase of instruction. Then the students each try to plan an essay on a new topic using the cue cards, a process the students call "soloing." As in Reciprocal Teaching, students have the opportunity to assume both producer and

critic roles. While each student practices soloing, the teacher, as well as other students, assume the role of evaluating the soloist's performance, by, for example, noticing discrepancies between the soloist's stated goals (e.g., to get readers to appreciate the difficulties of modern dance) and their proposed plans (to describe different kinds of dance). Students also become involved in discussing how to resolve problems that the soloist could not solve. As in the Reciprocal Teaching method, assumption of the role either of critic or producer is incremental, with students taking over more and more of the monitoring and problem-solving process from the teacher, as their skills improve. Moreover, as the students internalize the processes invoked by the prompts, the cue cards are gradually faded out as well.

[Insert Table 3 about here.]

In addition, they have developed specific techniques, called *coinvestigation* (Scardamalia & Bereiter, 1983a), aimed at encouraging students to reflect on both their existing strategies and the new ones they are acquiring. In *coinvestigation*, Scardamalia and Bereiter try to have students think aloud as they carry out some task, such as writing a paragraph linking two sentences together. They propose to the students that together they will jointly try to find out what the students are thinking when they carry out such a task. This motivates the students to consider their reflections as data from an experiment to find out what they think. When students have learned how to reflect on their own thinking, Scardamalia and Bereiter can push them into reflecting on the way experts do the same task. One way they do this is to provide the procedural supports shown in Table 2, so that children can carry out writing tasks in more expert ways. Then they can reflect on how their normal writing methods differ from these more expert methods. The scaffolding provided by the cue cards thus enables students to compare two different writing processes.

Scardamalia and Bereiter have tested the effects of their approach on both the initial planning and the revision of student compositions. In a series of studies (Bereiter & Scardamalia, in press), procedural facilitations were developed to help elementary school students evaluate, diagnose, and decide on revisions for their compositions. Results showed that each type of support was effective independent of the other supports. And when all the facilitations were combined, along with modelling and *coinvestigation*, they resulted in superior revisions for nearly every student and a ten-fold increase in the frequency of idea-level revisions, without any decrease in stylistic revisions. Another study (Scardamalia, et al., 1984) investigated the use of procedural cues to facilitate planning. In this study, students gave the teacher assignments, often chosen to be difficult for her. She used cues like those shown in Table 2 to facilitate planning, modelling the process of using the cues to stimulate her thinking about the assignment (Table 3). Pre- and post-comparisons of think-aloud protocols of a randomly selected portion of the subjects showed significantly more reflective activity on the part of experimental-group students even when prompts were no longer available to them. Time spent in planning increased ten fold. And when students were given unrestricted time to plan, the texts of experimental-group students were judged to be significantly superior in thought content.

Obviously, Scardamalia and Bereiter's methods for teaching writing are bringing about significant changes in the nature and quality of student writing. In addition to the methods and effects already discussed, we believe that there are two key reasons for their success. First of all, as in the Reciprocal Teaching method for reading, their methods help students build a new conception of what the writing process is. Students clearly consider writing to be a linear process of knowledge telling. By explicitly modelling and scaffolding expert processes, they are providing students with a new model of writing that involves planning before they write and revising what they have written. Most children found the view of writing implicit in this analysis to be an entirely new view of the writing process, as shown in their comments during *coinvestigation* ("I don't usually ask myself those questions," "I never thought closely about what I wrote," and "They helped me look over the sentence, which I don't usually do"). Moreover, since students rarely, if ever, see writers at work, they tend to hold naive beliefs about the nature of expert writing, thinking that writing is a smooth and easy process for "good" writers. Live

modelling helps to convey that this is not the case. The model demonstrates struggles, false starts, discouragement, and the like. Modelling also demonstrates for students that in evolving and decomposing a complex set of goals for their writing, expert writers often treat their own thoughts as objects of reflection and inquiry. These sorts of reflective operations underlie the fact that writing is not a linear, but an iterative, process--another new idea for students. Thus, a key effect of this sort of teaching is to radically alter students' understanding of the process.

Second, because writing is a complex compositional task, a key component of expertise is the control structure by which the writer organizes the numerous subactivities or lines of thinking involved in producing high quality text. A clear need of student writers, therefore, is to develop a more useful control structure and related processes than the ones evidenced in "knowledge telling." Their methods encourage this development in an interesting way: The cue cards act to externalize not only the basic cognitive processes involved in planning, but also help students to keep track of the higher-order intentions (such as generating an idea, elaborating or improving an idea, and so on) that organize these basic processes. This externalization aids students in monitoring their own (and others') ongoing progress in the writing task, so that they can determine what kind of general activity is required before moving on to specific prompts. This explicit hierarchical decomposition of general goals and process into more locally useful subprocesses aids students in building an explicit internal model of what might otherwise seem a confusing or random process.

**Schoenfeld's method for teaching mathematical problem solving.** Our third example is Schoenfeld's (1983, 1985) method for teaching mathematical problem solving to college students. Like the other two, this method is based on a new analysis of the knowledge and processes required for expertise, where expertise is understood as the ability to carry out complex problem-solving tasks in a domain. And like the other two, this method incorporates the basic elements of a cognitive apprenticeship, using the methods of modelling, coaching, and fading and of encouraging student reflection on their own problem-solving processes. In addition, Schoenfeld's work introduces some new concerns into our discussion, leading the way toward articulation of a more general framework for the development and evaluation of ideal learning environments in the next section.

One distinction between novices and experts in mathematics is that experts employ heuristic methods, usually acquired tacitly through long experience, to facilitate their problem solving. In order to teach these methods directly, Schoenfeld formulated a set of *heuristic* strategies, derived from the problem-solving heuristics of Polya (1945). These heuristic strategies consist of rules of thumb for how to approach a given problem. One such heuristic specifies how to distinguish *special cases* in solving math problems: For example, for series problems in which there is an integer parameter in the problem statement, one should try the cases  $n = 1, 2, 3, 4$  and try to make an induction on those cases; for geometry problems, one should first examine cases with minimal complexity, such as regular polygons and right triangles. Schoenfeld taught a number of these heuristics and how to apply them in different kinds of math problems. In the experiments he ran, Schoenfeld (1985) found that learning these strategies significantly increased students' problem-solving abilities.

But as he studied students' problem solving further, he became aware of other critical factors affecting their skill, in particular what he calls control strategies and belief systems. In Schoenfeld's analysis, control strategies are concerned with executive decisions, such as generating alternative courses of action, evaluating which will get you closer to a solution, evaluating which you are most likely to be able to carry out, considering what heuristics might apply, evaluating whether you are making progress toward a solution, and so on. Schoenfeld's notion of belief systems includes beliefs about oneself (e.g., math phobia), about the world (e.g., "physical phenomena have physical causes, not psychic causes") and about the domain (e.g., "mathematical proof is of no use in geometry construction problems"). Schoenfeld found that it was critical to teach control strategies and productive beliefs, as well as heuristics.

As with the previous two examples, explicit teaching of these elements of expert practice yields a fundamentally new understanding of mathematics for students. Previously to students, learning mathematics had meant learning a set of mathematical operations and methods, what Schoenfeld calls "resources." Schoenfeld's method is teaching students that doing mathematics consists not only of applying problem-solving procedures, but of reasoning about and managing problems using heuristics, control strategies, and beliefs.

Schoenfeld's teaching (1983, 1985) employs the elements of modelling, coaching, scaffolding, and fading in a variety of activities designed to highlight different aspects of the cognitive processes and knowledge structures required for expertise. For example, as a way of introducing new heuristics, he models their selection and use in solving problems for which they are particularly relevant. In this way, he exhibits the thinking processes (heuristics and control strategies) that go on in expert problem solving, but focuses student observation on the use and management of specific heuristics. Table 4 provides a protocol from one such modelling.

[Insert Table 4 about here.]

Next he gives the class problems to solve that lend themselves to the use of the heuristics he has introduced. During this collective problem solving, he acts as a moderator, soliciting heuristics and solution techniques from the students, while modelling the various control strategies for making judgments about how best to proceed. This division of labor has several effects. First, he turns over some of the problem-solving process to students by having them generate alternative courses of action, but provides major support or scaffolding by managing the decisions about which course to pursue, when to change course, etc. Second, it is significant that he is no longer modelling the entire expert problem-solving process, but a portion of it. In this way, he shifts the focus of student observation during modelling from the application or use of specific heuristics to the application or use of control strategies in managing those heuristics.

Like Scardamalia and Bereiter, Schoenfeld employs a third kind of modelling that is designed to change students' assumptions about the nature of expert problem-solving. He challenges students to find difficult problems, and at the beginning of each class offers to try to solve one of their problems. Occasionally the problems are hard enough that the students see him flounder in the face of real difficulties. During these sessions, he models for students not only the use of heuristics and control strategies, but the fact that one's strategies sometimes fail. In contrast, textbook solutions and classroom demonstrations generally illustrate only the successful solution path, not the search space which contains all of the dead-end attempts. Such solutions reveal neither the exploration one must do in searching for a good method nor the necessary evaluation of the exploration. Seeing how experts deal with problems that are difficult for them is critical to students' developing a belief in their own capabilities. Even experts stumble, flounder, and abandon their search for a solution until another time. Witnessing these struggles helps students realize that thrashing is neither unique to them nor a sign of incompetence.

In addition to class demonstrations and collective problem solving, Schoenfeld has students participate in small-group problem-solving sessions. During these sessions, Schoenfeld acts as a "consultant" to make sure that the groups are proceeding in a reasonable fashion. Typically he asks three questions: (1) what are they doing, (2) why are they doing it, and (3) how success in what they are doing will help them find a solution to the problem. Asking these questions serves two purposes: first, it encourages the students to reflect on their activities, thus promoting the development of general self-monitoring and -diagnosis skills; second, it encourages them to articulate the reasoning behind their choices as they exercise control strategies. Gradually the students, in anticipating his questions, come to ask the questions of themselves, thus gaining control over reflective and metacognitive processes in their problem solving. In these sessions, then, he is fading relative to both helping students generate heuristics and, ultimately, to exercising control over the process. In this way they gradually gain control over the entire problem-solving process.



Schoenfeld (1983) advocates small-group problem solving for several reasons. First, it gives the teacher a chance to coach students while they are engaged in semi-independent problem solving; he cannot really coach them effectively on homework problems or class problems. Second, the necessity for group decision making in choosing among alternative solution methods provokes articulation, through discussion and argumentation, of the issues involved in exercising control processes. Such discussion encourages the development of the metacognitive skills involved in, for example, monitoring and evaluating one's progress. Third, students get little opportunity in school to engage in collaborative efforts; group problem solving gives them practice in the kind of collaboration prevalent in real-world problem solving. Fourth, students are often insecure about their abilities, especially if they have difficulties with the problems. Seeing other students struggle alleviates some of this insecurity as students realize that difficulties in understanding are not unique to them, thus contributing to an enhancement of their beliefs about self relative to others.

We believe that there is another important reason that small-group problem solving is useful for learning: the differentiation and externalization of the roles and activities involved in solving complex problems. Successful problem solving requires that one assume at least three different, though interrelated, roles at different points in the problem-solving process: that of moderator or executive, that of generator of alternative paths, and that of critic of alternatives. Small-group problem solving differentiates and externalizes these roles: Different people naturally take on different roles, and problem solving proceeds along these lines. Thus, group discussion and decision making itself models the interplay among processes that an individual must internalize to be a successful problem solver. And here, as in Reciprocal Teaching, students may play different roles, so that they gain practice in all the activities they need to internalize.

In its use of the techniques of modelling, coaching, and fading, and its promotion of a new understanding of the nature of expertise, Schoenfeld's methods bear important similarities to our other two "success models." However, perhaps because of the requirements both of the domain and of the stage of learning that his students have achieved, Schoenfeld's work introduces some new issues into our discussion of pedagogical methods. First, Schoenfeld places a unique emphasis on the careful sequencing of problems. He has designed problem sequences to achieve four pedagogical goals: motivation, exemplification, practice, and integration. He first tries to show students the power of the heuristics he is teaching by giving them problems they will fail to solve without the heuristics. He then presents a few heuristics that enable students to solve the problems. The change in their ability to solve problems convinces the students that the heuristics are worth learning.

As he introduces each new heuristic, he tries to exemplify it with problems that are particularly "interesting," by which he presumably means problems in which the heuristic is especially effective in helping to solve the problem. Over the next week, he assigns extensive practice problems for which the new heuristic is helpful: He estimates that perhaps one-third of the week's problems involve use of the new heuristic. Finally, after the heuristic has been introduced and practiced, problems involving that heuristic continue to be assigned, but less frequently. As the course progresses, the problems involve use of multiple heuristics, so that students are learning to integrate the use of different heuristics to solve complex problems.

By selection and sequencing of examples and problem sets, Schoenfeld is trying to ensure that students will learn *when* to apply the heuristics, as well as how to apply them. Initially, instruction focuses on how to apply each heuristic; thus, the first problems all involve the heuristic. What varies is the problem context: A given problem might be a series problem or a geometry problem or an algebra problem, but the same heuristic always applies. Once the students know how to apply the heuristic, they must learn to recognize those situations in which the heuristic applies. Therefore, it is important to include problems for which the heuristic does not apply, forcing students to differentiate problems for which the heuristic applies from problems for which it does not. This problem-differentiation ability is critical to transfer of skills. The final phase, during which problems requiring

the heuristic applies are assigned occasionally, is aimed at preventing students from learning to apply the heuristic only to those problems assigned while the heuristic is being taught. (This is typical of the strategies that students derive from school courses.) Unless the need for the heuristic recurs, it will drop out of their repertoire.

There is one final aspect of Schoenfeld's method that we think is critical and that is different from the other methods we have discussed: what he calls "post-mortem" analysis. As with other aspects of Schoenfeld's method, students alternate with the teacher in producing post-mortem analyses. First, after modelling the problem-solving process for a given problem, Schoenfeld recounts the solution method, highlighting the generalizable features of the process (see Table 4). For example, he might note the heuristics that were employed, the points in the solution process where he or the class engaged in generating alternatives, the reasons for the decision to pursue one alternative before another, and so on. In short, he provides what we (Collins & Brown, in press) have labeled an "abstracted replay," that is a recapitulation of some process designed to focus students' attention on the critical decisions or actions. Post-mortem analysis also occurs when individual students explain the process by which they solved their homework problems. Here students are required to generate an abstracted replay of their own problem-solving process, as the basis for a class critique of their methods. The alternation between expert and student post-mortem analyses enables the class to compare student problem-solving processes and strategies with those of the expert; such comparisons provide the basis for diagnosing student difficulties and for making incremental adjustments in student performance. Moreover, generating abstracted replays involves focussing on the strategic as well as the tactical, levels of problem solving; this aids students in developing a hierarchical model of the problem-solving process as the basis for self-monitoring and -correction, and in seeing how to organize local (tactical) processes to accomplish high-level (strategic) goals.

### 3. A Framework for Designing Learning Environments

In our discussion so far, we have described an apprenticeshiplike approach to teaching the skills necessary for expert practice in cognitive domains and considered in detail three recently developed teaching methods, viewed as "success models" of cognitive apprenticeship. Our discussion of these teaching methods has introduced numerous pedagogical and theoretical issues that we believe are important to the design of learning environments generally. To facilitate consideration of these issues, we have developed a framework, outlined in Table 5. The framework describes four dimensions that constitute any learning environment: content, method, sequence, and sociology. Relevant to each of these dimensions are a set of characteristics that we believe should be considered in constructing or evaluating learning environments. We consider these characteristics in detail below, giving examples from reading, writing, and mathematics.

[Insert Table 5 about here.]

**Content.** Recent cognitive research has begun to differentiate the types of knowledge required for expertise in a domain. In particular, researchers have begun to distinguish between the explicit conceptual, factual, and procedural knowledge associated with expertise, and various types of *strategic* knowledge. We use the term strategic knowledge to refer to the usually tacit knowledge that underlies an expert's ability to make use of concepts, facts, and procedures as necessary to solve problems and carry out tasks. This sort of expert problem-solving knowledge involves problem-solving strategies and heuristics, and the strategies that control the problem-solving process at its various levels of decomposition. Another type of strategic knowledge, often overlooked, includes the learning strategies that experts have about how to acquire new concepts, facts, and procedures in their own or another field.

Within our framework, the appropriate target knowledge for an ideal learning environment is likely to include all four categories of expert knowledge, only one of which is often the current focus in schools.

1. *Domain knowledge* includes the conceptual and factual knowledge and procedures explicitly identified with a particular subject matter; these are generally explicated in school textbooks, class lectures, and demonstrations. As we argued in the Introduction, this kind of knowledge, while certainly important, provides insufficient clues for many students about how to actually go about solving problems and carrying out tasks in a domain. Moreover, when it is learned in isolation from realistic problem contexts and expert problem-solving practices, domain knowledge tends to remain inert in situations for which it is appropriate, even for successful students. And finally, while at least some concepts can be formally described, many of the crucial subtleties of their meaning are best acquired through the work of applying them in a variety of problem situations. Indeed, it is only through encountering them in real problem solving that most students will learn the boundary conditions and entailments of much of their domain knowledge.

Examples of domain knowledge in reading are vocabulary, syntax, and phonic rules; the standard procedure for reading is scanning text, either silently or aloud, and constructing an interpretation. For writing, domain knowledge includes much of the same vocabulary and syntactic knowledge, but, in addition, knowledge about rhetorical forms and genres, and about writing drafts and revising. In mathematics, most of the domain knowledge, other than number facts and definitions, consists of procedures for solving different kinds of problems, from addition algorithms to procedures for solving problems in algebra and constructing proofs in geometry.

2. *Problem-solving strategies and heuristics* are generally effective techniques and approaches for accomplishing tasks that might be regarded as "tricks of the trade"; they don't always work, but when they do, they are quite helpful. Most heuristics are tacitly acquired by experts through the practice of solving problems; however, there have been noteworthy attempts to address heuristic learning explicitly. The literature is replete with examples of heuristics for mathematical problem-solving, beginning with Polya (1945); though less widely formalized, useful problem-solving heuristics and strategies can also be identified for more open-ended task domains, such as reading and writing.

For example, a standard heuristic for writing is to plan to rewrite the introduction to a text (and therefore to spend relatively little time crafting it); this heuristic is based on the recognition that a writer's initial plan for a text is likely to undergo radical refinement and revision through the process of writing, and therefore that the beginning of a text often needs to be rewritten to "fit" the emergent organization and arguments of the main body and conclusion. Another strategy, designed to help a writer maintain momentum and "flow of ideas," is to avoid getting bogged down in syntax or other presentational details while getting one's ideas down. In reading, a general strategy for facilitating both comprehension and critical reading is to develop an overview and set of expectations and questions about a text before reading line by line; one can achieve this by looking through tables of contents and reading section headings in chapters to get a sense of the overall organization of the text. Certain kinds of texts, for example, experimental psychology articles, have a standard format corresponding to a paradigmatic argument structure; one can read the introduction and conclusions to understand the major claims being made before attempting to assess whether they are supported by evidence presented in other sections.

3. *Control strategies*, as the name suggests, control the process of carrying out a task. As students acquire more and more heuristics and strategies for solving problems, they encounter a new management or control problem: how to select among the various possible problem-solving strategies, how to decide when to change strategies, and so on. The knowledge that experts have about managing problem solving can be formulated as control strategies. Control strategies require reflection on the problem-solving process in order to

determine how to proceed. Control strategies operate at many different levels. Some are aimed at managing problem solving at a global level and are probably useful across domains; for example, a simple control strategy for solving a complex problem might be to switch to a new part of a problem if one is stuck on another part. Other strategies control selection of domain-specific problem-solving heuristics and strategies for carrying out parts of the task at hand.

Control strategies have monitoring, diagnostic and remedial components; decisions about how to proceed in a task generally depends on an assessment of the current state relative to one's goals, on an analysis of current difficulties, and on what strategies are available for dealing with difficulties. Monitoring strategies can be represented as activities that help students to evaluate their progress in a general way by providing a simple criterion for determining whether or not a given goal is being achieved. For reading, these strategies are called "comprehension monitoring" strategies (Baker & Brown, 1980; Collins & Smith, 1982). For example, a comprehension monitoring strategy might be to try to state the main point of a paragraph one has just read; if one cannot do so, then one has not understood the text. Monitoring strategies lead either to diagnosis or directly to remedial actions. For example, if one does not understand a given paragraph, one may proceed to analyze the source of one's difficulties or simply re-read the text. Diagnosis refers to those processes whereby the problem solver arrives at a useful analysis of the nature or cause of his difficulties. The level of diagnostic analysis required depends on a number of factors, for example, how important understanding the current difficulty is to achieving the overall goals of the activity, or what level of diagnosis is necessary to determine corrective action. A diagnostic activity for reading is what Palincsar and Brown call "clarifying difficulties" with the text, in which students attempt to isolate the particular word or phrase that they don't understand. In order to be useful, diagnoses must point to remedial strategies, that is, to problem-solving or learning activities that will lead out of the difficulty by introducing new knowledge or providing an alternate tack on the problem. Having recognized that their difficulties in understanding a passage lie with a particular word or phrase, readers can employ various strategies, such as looking up words, or continuing to read with the plan of coming back to the difficult passage to see if subsequent evidence from the text resolves the difficulty (Collins & Smith, 1982).

4. *Learning strategies* are strategies for learning any of the other kinds of content described above. Like the other types of process knowledge we have described, knowledge about how to learn ranges from general strategies for exploring a new domain to more local strategies for extending or reconfiguring knowledge as the need arises in solving problems or carrying out a complex task.

For example, if students want to learn to read better on their own, they have to know how to pick texts that expand their vocabulary, but are not too demanding. They also have to know how to check their understanding against other people's, by reading critical reviews of the texts they have read or by discussing the text with someone. If students want to learn to write better, they need to find people to read their writing who can give helpful critiques and explain the reasoning underlying the critiques (most people cannot). They also need to learn to analyze the texts of others in terms of the ways that they are well and badly written. To learn to solve math problems better, it helps to try to solve the example problems presented in the text before reading the solution, to provide a basis for comparing one's own solution method to the solution method in the book. These are just a few of the more general strategies that expert learners acquire. Just as it is possible to teach heuristic and monitoring strategies by apprenticeship, it is possible to teach such learning strategies by apprenticeship.

**Method.** As we have discussed, a key goal in the design of teaching methods should be to help students acquire and integrate cognitive and metacognitive strategies for using, managing, and discovering knowledge. However, it is our belief that the way in which these strategies are acquired and, once acquired, brought to play in problem solving, is both subtle and poorly understood. In general, it seems clear that both acquisition and use of these strategies depend crucially on interactions between the individual's current knowledge and beliefs, the social and physical environment in which the problem solving takes place, and the local details of the problem solving itself as it unfolds. A major direction in current cognitive research is to attempt to formulate explicitly the strategies and skills underlying expert practice, in order to make them a legitimate focus of teaching in schools and other learning environments. Indeed, all three success models we have discussed are based on explicit formulations of cognitive and metacognitive strategies and center their teaching around activities designed to explicitly convey these to students. However, we believe it is also important to consider the possibility that, because of the nature of the relationship between these strategies and the overall problem context, not all of the necessary--and certainly not all of the possible--strategies involved in complex cognitive activities can be captured and made explicit. In this regard, it is worth noting that these strategies and skills have tended to remain tacit and thus to be lost to formal education precisely because they arise from the practice of solving problems, in situ, in the domain. Moreover, we would argue that, even given explicit formulation of strategies, understanding how to use them depends crucially on understanding the way in which they are embedded in the context of actual problem solving.

For these reasons, we believe that teaching methods should be designed to give students the opportunity to observe, engage in, and invent or discover expert strategies in context. Such an approach will enable students to see how these strategies fit together with their factual and conceptual knowledge, and how they cue off and make use of a variety of resources in the social and physical environment. This is the essence of what we mean by situated learning (see Sociology), and the reason why the cognitive apprenticeship method, with its modelling-coaching-fading paradigm, is successful and perhaps indispensable.

The following six teaching methods fall roughly into three groups: The first three (modelling, coaching, and scaffolding) are the core of cognitive apprenticeship, designed to help students acquire an integrated set of cognitive and metacognitive skills through processes of observation and of guided and supported practice. The next two (articulation and reflection) are methods designed to help students both to focus their observations of expert problem solving and to gain conscious access to (and control of) their own problem-solving strategies. The final method (exploration) is aimed at encouraging learner autonomy not only in carrying out expert problem solving processes, but also in defining or formulating the problems to be solved.

1. *Modelling* involves showing an expert carrying out a task so that students can observe and build a conceptual model of the processes that are required to accomplish the task. In cognitive domains, this requires the externalization of usually internal (cognitive) processes and activities--specifically, the heuristics and control processes by which experts make use of basic conceptual and procedural knowledge. For example, a teacher might model the reading process by reading aloud in one voice, while verbalizing her thought processes (e.g., the making and testing of hypotheses about what the text means, what the author intends, what he or she thinks will happen next, and so on) in another voice (Collins & Smith, 1982). Tables 3 and 4 give examples of teacher modelling of expert processes in the domains of writing and mathematics.
2. *Coaching* consists of observing students while they carry out a task and offering hints, scaffolding, feedback, modelling, reminders, and new tasks aimed at bringing their performance closer to expert performance. Coaching may serve to direct students' attention to a previously unnoticed aspect of the task or simply to remind the student of some aspect of the task that is known but has been temporarily overlooked. Coaching focusses on the

enactment and integration of skills in the service of a well-understood goal through highly interactive and highly situated feedback and suggestions. That is, the content of the coaching interaction is immediately related to specific events or problems that arise as the student attempts to carry out the target task. In reading, coaching might consist of having students attempt to give summaries of different texts. The teacher, in the role of coach, might choose tests with interesting difficulties, might remind the student that a summary needs to integrate the whole text into a sentence or two, might suggest how to start constructing a summary, might evaluate the summary a student produces in terms of how it could be improved, or ask another student to evaluate it. Similarly, the description of Scardamalia and Bereiter's classes, and of Schoenfeld's classes provide examples of how the teacher can function as a coach while students try to carry out tasks in writing and mathematics.

3. *Scaffolding* refers to the supports the teacher provides to help the student carry out a task. These supports can either take the forms of suggestions or help, as in Palincsar and Brown's (1984) Reciprocal Teaching, or they can take the form of physical supports, as with the cue cards in Scardamalia, et al.'s (1984) procedural facilitation of writing or the short skis used to teach downhill skiing (Burton, Brown, & Fischer, 1984). When scaffolding is provided by a teacher, it involves the teacher in carrying out parts of the overall task that the student cannot yet manage. As such, it involves a kind of cooperative problem-solving effort by teacher and student in which the express intention is for the student to assume as much of the task on his own as possible, as soon as possible. A requisite of such scaffolding is accurate diagnosis of the student's current skill level or difficulty and the availability of an intermediate step at the appropriate level of difficulty in carrying out the target activity. *Fading* consists of the gradual removal of supports until students are on their own. The three models described employed scaffolding in a variety of ways.
4. *Articulation* includes any method of getting students to articulate their knowledge, reasoning, or problem-solving processes in a domain. We have identified several different methods of articulation. First, *inquiry teaching* (Collins & Stevens, 1982, 1983) is a strategy of questioning students to lead them to articulate and refine "proto-theories" about the four kinds of knowledge enumerated above. For example, an inquiry teacher in reading might systematically question students about why one summary of the text is a good one while another is poor, in order to get the students to formulate an explicit model of what makes a good summary. Second, teachers might encourage students to articulate their thoughts as they carry out their problem solving, as do Scardamalia, et al. (1984). Third, having students assume the critic or monitor role in cooperative activities, as do all three models we discussed, leads students to formulate and articulate their knowledge of problem-solving and control processes.
5. *Reflection* (Brown, 1985a, 1985b; Collins & Brown, in press) involves enabling students to compare their own problem-solving processes with that of an expert, other students, and ultimately, an internal cognitive model of expertise. Reflection is enhanced by the use of techniques for reproducing or "replaying" the performances of both expert and novice for comparison. This can be done through a variety of methods. For example, an expert's skillful post mortem of the problem-solving process, as Schoenfeld showed, can serve as a target for reflective comparison, as can the students' post mortems of their own problem-solving process. Alternately, various recording technologies, such as video or audio recorders and computers, can be employed to reproduce student and expert performance. The levels of detail at which a replay should be done may vary depending on the student's stage of learning, but often some form of "abstracted replay," in which the determining features of expert and student performance are highlighted, is desirable. For reading or writing, methods to encourage reflection might consist of recording students as they think out loud and then replaying the tape for comparison with the thinking of experts and other students.

6. *Exploration* involves pushing students into a mode of problem solving on their own. Forcing students to do exploration is critical for students to learn how to frame questions or problems that are interesting and that they can solve. Exploration is the natural culmination of the fading of supports. It involves not only fading in problem solving, but fading in problem setting as well. But students do not know a priori how to explore a domain productively. So exploration strategies need to be taught as part of learning strategies more generally.

Exploration as a method of teaching involves setting general goals for students, but encouraging them to focus on particular subgoals of interest to them or even to revise the general goals as they come upon something more interesting to pursue. For example, in reading the teacher might send the students to the library to find out what president died in office as a result of a trip to Alaska, or to investigate theories about why the stock market crashed in 1929. In writing, students might be encouraged to write an essay defending the most outrageous thesis they can devise, or to keep a diary of their best ideas or their most traumatic experiences. In mathematics, students might be given a data base on teenagers detailing their backgrounds and how they spend their time and money; the students' task might be to come up with hypotheses about what determines how different groups of teenagers spend their time or money that they test out by analyzing the data base they have been given. The goal is to find general tasks that students will find interesting and turn them loose on them, after they have acquired some basic exploration skills.

**Sequencing.** Lave (in preparation) has suggested that research emphasis on early skill acquisition has resulted in a failure to recognize the changing learning needs of students at different stages of skill acquisition and, consequently, to sequence and structure materials and activities appropriately for those stages. In particular, designers need to understand how to support the phases of both integration and generalization of knowledge and complex skills. We have identified some dimensions or principles that should guide the sequencing of learning activities in order to facilitate the development of robust problem-solving skills.

1. *Increasing complexity* refers to the construction of a sequence of tasks and task environments or microworlds such that more and more of the skills and concepts necessary for expert performance are required (VanLehn & Brown, 1980; Burton, Brown & Fischer, 1984; White, 1984; White & Frederiksen, in press). We doubt that it is possible to sequence skills and tasks such that they undergo a monotonic increase in complexity. Instead, there are more likely to be jumps in complexity as learners are required not only to learn and integrate the interrelated set of skills or activities necessary to carry out an interesting task (even a relatively simple one), but also to manage and direct these activities. For example, in the tailoring apprenticeship described by Lave, apprentices jump from practicing very simple rudimentary skills, such as wielding scissors and needle and sewing scraps, to actually putting together a garment, which requires the integration of sewing skill with a conceptual understanding of the structure of the garment over a series of ordered steps.

There are two mechanisms for helping students manage increasing complexity. First, efforts should be made to control task complexity. As an example, in the tailoring apprenticeship described by Lave (in preparation), apprentices first learn to construct drawers, which have straight lines, few pieces, and no "special features," such as waistbands or pockets. They then learn to construct blouses, which require curved lines, patch pockets, and the integration of a complex subpiece, the collar. The second key mechanism for helping students manage complexity is the use of scaffolding, which enables students to handle at the outset, with the support of the teacher or other helper, the complex set of activities needed to carry out any interesting task.

Presumably in most domains task complexity can vary along a variety of dimensions. For example, in reading, texts can vary in complexity (at least) along the dimensions of syntax, vocabulary, conceptual abstractness, and argumentation. Increasing task complexity might consist of progressing from relatively short texts employing straightforward syntax and concrete description to texts in which complexly interrelated ideas and the use of abstractions make interpretation difficult.

2. *Increasing diversity* refers to the construction of a sequence of tasks in which a wider and wider variety of strategies or skills are required. While it is important to practice a new strategy or skill repeatedly in a sequence of (increasingly complex) tasks, as the skill becomes well learned it becomes increasingly important that tasks requiring a diversity of skills and strategies be introduced so that the student learns to distinguish the conditions under which they do (and do not) apply. Moreover, as students learn to apply skills to more diverse problems and problem situations, their strategies become freed from their contextual bindings (or perhaps more accurately, acquire a richer net of contextual associations) and thus are more readily available for use with unfamiliar or novel problems. For reading, task diversity might be attained by intermixing reading for pleasure, reading for memory (studying), and reading to find out some particular information in the context of some other task. Varying task diversity in writing might be achieved by posing different rhetorical problems, such as writing to persuade an audience of some point of view versus writing descriptive or instructional text, or by introducing specific constraints, such as writing for a particular audience (say the school board) or under different time constraints. We described earlier how Schoenfeld systematically increases diversity in teaching mathematics.
3. *Global before local skills*. In the tailoring apprenticeship described by Lave, apprentices invariably learn to put together a garment from pre-cut pieces before learning to draw and cut out the pieces themselves. This sequencing of activities provides learners with the opportunity to build a conceptual model of how all the pieces of a garment fit together before attempting to produce the pieces. For cognitive domains, this implies sequencing of lessons such that students have a chance to apply a set of skills in constructing an interesting problem solution before they are required to generate or remember those skills. This requires some form of scaffolding (see Methods section). Scaffolding can be applied to different aspects of a problem-solving process, for example, to management and control of the problem solving or to the subprocesses that are required to carry out the task. Global before local skills means that in the sequencing of lessons there is a bias toward supporting the lower-level or composite skills that students must put together in order to carry out a complex task. In algebra, for example, students may be relieved of having to carry out low-level computations in which they lack skill in order to concentrate on the higher-order reasoning and strategies required to solve an interesting problem (Brown, 1985b).

The chief effect of this sequencing principle is to allow students to build a conceptual map, so to speak, before attending to the details of the terrain. In general, having students build a conceptual model of the target skill or process (which is also encouraged by expert modelling) accomplishes two things: First, even when the learner is able to carry out only a portion of a task, having a clear conceptual model of the overall activity both helps him make sense of the pieces that he is carrying out and provides a clear goal toward which to strive as he takes on and integrates more and more of the pieces. Second, the presence of a clear conceptual model of the target task acts as a guide for the learner's performance, thus improving his ability to monitor his own progress and to develop attendant self-correction skills. We also suspect that having such a model helps crucially to prevent students from developing bugs in the acquisition of individual composite skills; having an understanding of the purpose of various skills can help clarify the conditions under which they are applicable, their entailments, their relationships to other processes, and so on.



**Sociology.** The final dimension in our framework concerns the sociology of the learning environment, a critical dimension that is often ignored in decisions about curriculum and pedagogical practice. In her analysis of tailoring apprenticeship, Lave (in preparation) discusses some of the determining features of the embedding social context and the ways that they affect learning. For example, she notes that apprentices learn tailoring not in a special, segregated learning environment, but in a busy tailoring shop. They are surrounded by both masters and other apprentices, all engaged in the target skills at varying levels of expertise. And they are expected, from the beginning, to engage in activities that contribute directly to the production of actual garments, advancing quickly toward independent skilled production. As a result, apprentices learn skills in the context of their application to realistic problems, within a culture focussed on, and defined by, expert practice. They continually see the skills they are learning being used in a way that clearly conveys both how they are integrated into patterns of expertise and their efficacy and value within the subculture. And by advancing in skill, apprentices are increasing their participation in the community, becoming expert practitioners in their own right. These characteristics--the ready availability of models of expertise-in-use, the presence of clear expectations and learning goals, and the integration of skill improvement and social reward--help motivate and ground learning.

Furthermore, we believe that certain aspects of the social organization of apprenticeship encourage productive beliefs about the nature of learning and of expertise that are important to learners' motivation, confidence, and, most importantly, their orientation toward problems that they encounter as they learn. For example, the presence of other learners provides apprentices with calibrations for their own progress, helping them to identify their strengths and weaknesses and thus to focus their efforts for improvement. Moreover, the availability of multiple masters may help learners realize that even experts have different styles and ways of doing things and different special aptitudes. Such a belief encourages learners to understand learning as, in part, using multiple resources in the social context to obtain scaffolding and feedback.

We believe that structuring the social context so as to encourage the development of these productive beliefs sets the stage for the development of cooperative learning styles, such as those found by Levin (1982) in contemporary computer clubs, and of collaborative skill generally. In his study, Levin found that nonexperts were able to successfully bootstrap their knowledge about computers without regular access to high-level expertise by pooling their fragments of knowledge and using other learners as a source of scaffolding for carrying out their tasks. This sort of decoupling of the experience of learning from the availability of an "authority" encourages independent and self-directed learning. Moreover, awareness of the distributed nature of expertise and insight is at the foundation of successful collaboration in all domains. Partly because of this key belief--that knowledge is not concentrated in any single person--skilled collaborators are more likely to be open to, and seek out, help and input from others. As a result, they are better able to take advantage of interactions with others in order to construct better and more satisfactory solutions to complex problems.

From our consideration of these general issues, we have abstracted five critical characteristics affecting the sociology of learning.

1. ***Situated learning.*** A critical element for learning is that students are carrying out tasks and solving problems in an environment that reflects the multiple uses to which their knowledge will be put in the future. This goal serves several different purposes. First students will come to understand the purposes or uses of the knowledge they are learning. Second, they will learn by actively using knowledge, rather than passively receiving it. Third, they will learn the different conditions under which their knowledge can be applied. As we pointed out in the discussion of Schoenfeld's work, students have to learn when to use a particular strategy and when not to use it (i.e., the *application conditions* of their knowledge). Fourth, learning in multiple contexts induces the abstraction of knowledge, so that students acquire knowledge in a dual form, both tied to the contexts of its uses and independent of any particular

context. This unbinding of knowledge from a specific context fosters its transfer to new problems and new domains.

In addition, the reason that Dewey (see Cuban, 1984), Papert (1980), and others have advocated learning from projects rather than isolated problems is, in part, so that students can face the task of formulating their own problems, guided on the one hand by the general goals they set, and on the other hand by the "interesting" phenomena and difficulties they discover through their interaction with the environment. Recognizing and delineating emergent problems, that is, problems that arise while carrying out complex tasks in a rich problem-solving context, is a crucial skill. Emergent problems encountered in projects are ones for which one cannot use knowledge about the instructional designer's goals to help solve the problem as students do in working textbook problems (Schoenfeld, 1985). Instead, problems emerge from interactions between the overall goals and the perceived structure of the environment. Thus, in projects students learn first to find a problem and then, ideally, to use the constraints of the embedding context to help solve it. This is the process of "problem finding" identified by Getzels and Csikszentmihalyi (1976), while studying artists and the notion of "emergent goals" identified by Scardamalia and Bereiter (1985) in the writing process.

Reading and writing instruction might be situated in the context of an electronic message system, where students are sending each other questions and advice, as in the computer club described by Levin (1982). Dewey created a situated learning environment in his experimental school by having the students design and build a clubhouse (Cuban, 1984), a task which emphasizes arithmetic and planning skills.

2. *Culture of expert practice* refers to the creation of a learning environment in which the participants actively communicate about, and engage in, the skills involved in expertise, where expertise is understood as the practice of solving problems and carrying out tasks in a domain. A culture of expert practice helps situate and support learning in several ways. First, a culture focussed on expert practice provides learners with readily available models of expertise-in-use; as we have discussed, the availability of such models helps learners build and refine a conceptual model of the task they are trying to carry out. However, a learning environment in which experts simply solve problems and carry out tasks, and learners simply watch, is inadequate to provide effective models for learning, particularly in cognitive domains, where many of the relevant processes and inferences are tacit and hidden. Thus, if expert modelling is to be effective in helping students internalize useful conceptual models, experts must be able to identify and represent to students the cognitive processes that they engage in as they solve problems. Drawing students into a culture of expert practice in cognitive domains involves teaching them how to "think like experts." The focus of much current cognitive research is to understand better what is really meant by such a goal and to find ways to communicate more effectively about the processes involved. However, even without a thorough theoretical understanding and formulation of expert processes, such mechanisms as group problem solving are helpful in externalizing relevant processes and reasoning, so that students can observe and enact them. Thus, the creation of a culture of expert practice for learning should be understood to include focussed interactions among learners and experts for the purpose of solving problems and carrying out tasks.

Activities designed to engender a culture of expert practice for reading might engage students and teacher in reading and discussing how they interpret and use what they've read for a wide variety of purposes, including the variety of learning needs that arise in other classes or domains.

3. *Intrinsic motivation.* Related to the issue of situated learning and the creation of cultures of expert practice is the need to promote intrinsic motivation for learning. Lepper and Greene (1978) and Malone (1981) discuss the importance of creating learning environments in which students perform tasks because they are intrinsically related to an interesting or at least coherent goal, rather than for some extrinsic reason like getting a good grade or pleasing the teacher. There is some evidence that when an extrinsic reward is provided for performing a task like reading, students are less likely to perform the task on their own. In general, the methods of modelling-coaching-fading, insofar as they promote acquisition of integrated skills in the service of a coherent overall activity, are supportive of intrinsic motivation. But equally important is that students attempt to carry out realistic tasks in the spirit and for the purposes that characterize adult expert practice. In reading, for example, intrinsic motivation might be achieved by having students communicate with students in another part of the world by electronic mail (Collins, 1986; Levin, 1982) or by playing a game that requires a lot of reading (e.g., *Dungeons and Dragons*).
4. *Exploiting cooperation* refers to having students work together in a way that fosters cooperative problem solving. Learning through cooperative problem solving is both a powerful motivator and a powerful mechanism for extending learning resources. As we discussed earlier, cooperative learning and problem solving provides students with an additional source of scaffolding, in the form of knowledge and processes distributed throughout the group. One crucial aspect of distributed knowledge concerns the multiple roles that a problem solver must play in order to successfully carry out a complex task and which students may have difficulty integrating. For example, in order to write effectively, students must alternate between the roles of producer and critic. By taking turns writing and reading each other's writing, students can get practice in both roles. Moreover, as students learn complex processes, they will grasp different aspects of a problem and of the methods needed to solve it. Cooperative problem solving enables them to share their knowledge and skills, giving students additional opportunities to grasp the relevant conceptual and other aspects of an overall process. In addition, students are often able to help each other grasp the rationale for, or distinguishing characteristics of, some new concept or skill because they are closer to the problem of learning about it. Said differently, a student may have a better internal model of another student's difficulties and how to address them because they have recently had the same or a similar difficulty themselves. Finally, cooperative learning helps foster the situated articulation of processes and concepts, thus helping students to gain conscious access to, and control of, cognitive and metacognitive processes and the ways these employ conceptual and factual knowledge.

In reading, activities to exploit cooperation might involve having students break up into pairs, where one student articulates his thinking process while reading, and the other student questions the first student about why he made different inferences.

5. *Exploiting competition* refers to the strategy of giving students the same task to carry out, and then comparing what each produces. One of the important effects of comparison is that it provides a focus for students' attention and efforts for improvement by revealing the sources of strengths and weaknesses. However, for competition to be effective for this purpose, comparisons must be made not between the products of student problem solving, but between the processes, and this is rarely the case. Moreover, while competition is a powerful motivator and organizer of learning for some students, it presents a number of thorny issues for educators. For example, there is evidence that many students are inhibited, rather than motivated, by competitive situations. Competition raises difficult emotional issues for some students, thus introducing potentially confusing or confounding factors into classroom interactions. And some people feel that competition encourages behavior and attitudes that are socially undesirable and even unethical.

We suspect that at least some of the ill effects of competition have to do with attitudes toward, and beliefs about, errors (Brown & Burton, 1978). If students believe that making errors or being wrong about some process makes them "dumb," then comparative, competitive situations will be profoundly discouraging to weaker students. Another factor that makes competition seem problematic is that under many forms of teaching, students lack the means, in the form of an understanding of the underlying processes, strategies, and heuristics involved in solving problems, for improving their performance. In these cases, the motivation to improve that might be engendered by competition is blocked, leaving students inevitably frustrated and discouraged.

It may be that at least some of these ill effects can be reduced by blending cooperation and competition; for example, individuals might work together in groups in order to compete with other groups. In such cases, students can take advantage of the scaffolding provided by the group to learn and strengthen their performance. For example, in reading, different groups might compete in trying to find some obscure information by searching through the library.

This summarizes our framework for the design of learning environments. The framework was evolved partly through a close consideration of the three success models discussed in the first sections of the paper, as well as other models of apprenticeship learning, e.g., tennis (Braden & Bruns, 1977; Gallwey, 1974), skiing (Burton, Brown, & Fischer, 1984), computational skills (Lave, Murtaugh, & de la Rocha, 1984) and Dewey's experimental school (Cuban, 1984). In turn, the framework provides a critical lens for evaluating the strengths and weaknesses of different learning environments and teaching methods.

#### 4. Conclusion

Apprenticeship learning is the way we learn most naturally. It characterized learning before there were schools, from learning one's language to learning how to run an empire. We now have three very successful models of how apprenticeship methods, in all their dimensions, can be applied to teaching the school curriculum of reading, writing, and mathematics.

These models, and the framework we have developed, help point the way toward the redesign of schooling so as to help students acquire true expertise and robust problem-solving skills, as well as an improved ability to learn throughout life. Perhaps less obviously, we believe that the core techniques of modelling, coaching and fading can be formalized and embedded in tomorrow's powerful personal computers, thereby fostering a renewal of apprenticeship-style learning in our schools. Obviously a number of advances in research are required before this dream can become a widespread reality. Current work on developing explicit, cognitive theories of domain skills, metacognitive skills, and tutoring skills is making the crucial first steps in the right direction.

We believe the thrust toward computer-aided learning is an important development in education for several reasons. First, computers make it possible to give more personal attention to individual students, without which the coaching and scaffolding of an apprenticeship-style learning are impossible. It is precisely in human-resource-intensive settings, such as tennis coaching, learning foreign languages at Berlitz, or receiving training in medical diagnosis, that apprenticeship methods are still used. Appropriately designed computer-based modelling, coaching, and fading systems can make cost-effective and widely available a style of learning that was previously severely limited. Of course, apprenticeship-based computer systems need not take on the total responsibility. Instead, they only need to *augment* the master teacher in a way that amplifies and makes her efforts more cost-effective.

Second, and perhaps more importantly, research aimed at building computer-based apprenticeship learning environments can act as a forcing function to encourage the more precise formulation, not only of the processes and knowledge that students require for expertise, but also of the knowledge that we, as teachers, require in order to effectively diagnose student difficulties, give useful hints, sequence learning activities, and so on. This sort of knowledge can obviously have fruitful implications not only for the design of electronic learning environments, but also for teacher training, curriculum design, and educational policy generally.

## References

- Baker, L., & Brown, A. L. (1980). Metacognitive skills of reading. In P. D. Pearson (Ed.), *Handbook of reading research*. New York: Longman.
- Bereiter, C., & Scardamalia, M. (in press). *The psychology of written composition*. Hillsdale, NJ: Erlbaum.
- Braden, V., & Bruns, W. (1977). *Vic Braden's Tennis for the Future*. Boston: Little Brown and Company.
- Brown, A. L., & Palincsar, A. S. (in press). Reciprocal teaching of comprehension strategies: A natural history of one program for enhancing learning. In J. B. Borkowski & J. D. Day. (Eds.), *Intelligence and cognition in special children: Comparative studies of giftedness, mental retardation, and learning disabilities*. Norwood, NJ: Ablex.
- Brown, J. S. (1985a). Process versus product: A perspective on tools for communal and informal electronic learning. *Journal of Educational Computing Research*, 1, 179-201.
- Brown, J. S. (1985b). Idea-amplifiers: New kinds of electronic learning. *Educational Horizons*, 63, 108-112.
- Brown, J. S., & Burton, R. (1978). Diagnostic models for procedural bugs in basic mathematical skills. *Cognitive Science*, 2, 155-192.
- Brown, J. S., & VanLehn, K. (1980). Repair Theory: A generative theory of bugs in procedural skills. *Cognitive Science*, 4, 379-426.
- Burton, R., Brown, J. S., & Fischer, G. (1984). Skiing as a model of instruction. In B. Rogoff & J. Lave (Eds.), *Everyday cognition: Its development in social context*. Cambridge, MA: Harvard University Press.
- Collins, A. (1986). Teaching reading and writing with personal computers. In J. Orasanu (Ed.), *A decade of reading research: Implications for practice*. Hillsdale, NJ: Erlbaum.
- Collins, A., & Brown, J. S. (in press). The computer as a tool for learning through reflection. In H. Mandl and A. Lesgold (Eds.), *Learning issues for intelligent tutoring systems*. New York: Springer.
- Collins, A., Brown, J. S., & Larkin, K. M. (1980). Inference in text understanding. In R. J. Spiro, B. C. Bruce, & W. F. Brewer (Eds.), *Theoretical issues in reading comprehension*. Hillsdale, NJ: Erlbaum.
- Collins, A., & Smith, E. E. (1982). Teaching the process of reading comprehension. In D. K. Detterman & R. J. Sternberg (Eds.), *How much and how can intelligence be increased?* Norwood, NJ: Ablex.
- Collins, A., & Stevens, A. L. (1982). Goals and strategies of inquiry teachers. In R. Glaser (Ed.), *Advances in instructional psychology* (Vol. 2). Hillsdale, NJ: Erlbaum.
- Collins, A., & Stevens, A. L. (1983). A cognitive theory of interactive teaching. In C. M. Reigeluth (Ed.), *Instructional design theories and models: An overview*. Hillsdale, NJ: Erlbaum.
- Cuban, L. (1984). *How teachers taught*. New York: Longman.

- Gallwey, W. T. (1974). *The inner game of tennis*. New York: Random House.
- Getzels, J., & Csikszentmihalyi, M. (1976). *The creative vision: A longitudinal study of problem finding in art*. New York: Wiley.
- Hayes, J. R., & Flower, L. (1980). Identifying the organization of writing processes. In L. W. Gregg & E. R. Steinberg (Eds.), *Cognitive processes in writing*. Hillsdale, NJ: Erlbaum.
- Lave, J. (in preparation). *Tailored learning: Education and everyday practice among craftsmen in West Africa*.
- Lave, J., Murtaugh, M., & de la Rocha, O. (1984). The dialectic of arithmetic in grocery shopping. In B. Rogoff & J. Lave (Eds.), *Everyday cognition: Its development in social context*. Cambridge, MA: Harvard University Press.
- Lepper, M. R., & Greene, D. (1978). *The hidden costs of reward*. Hillsdale, NJ: Erlbaum.
- Levin, J. A. (1982). Microcomputer communication networks for education. *The Quarterly Newsletter of the Laboratory of Comparative Human Cognition*, 4, No. 2.
- Malone, T. (1981). Toward a theory of intrinsically motivating instruction. *Cognitive Science*, 4, 333-369.
- Palincsar, A. S., & Brown, A. L. (1984). Reciprocal teaching of comprehension-fostering and monitoring activities. *Cognition and Instruction*, 1, 117-175.
- Papert, S. (1980). *Mindstorms: Children, computers, and powerful ideas*. New York: Basic Books.
- Polya, G. (1945). *How to solve it*. Princeton, NJ: Princeton University Press.
- Scardamalia, M., & Bereiter, C. (1983a). Child as co-investigator: Helping children gain insight into their own mental processes. In S. Paris, G. Olson, & H. Stevenson (Eds.), *Learning and motivation in the classroom*. Hillsdale, NJ: Erlbaum.
- Scardamalia, M., & Bereiter, C. (1983b). The development of evaluative, diagnostic, and remedial capabilities in children's composing. In M. Martlew (Ed.), *The psychology of written language: A developmental approach* (pp. 67-95). London: Wiley.
- Scardamalia, M., & Bereiter, C. (1985). Fostering the development of self-regulation in children's knowledge processing. In S. F. Chipman, J. W. Segal, & R. Glaser (Eds.), *Thinking and learning skills: Research and open questions*. Hillsdale, NJ: Erlbaum.
- Scardamalia, M., Bereiter, C., & Steinbach, R. (1984). Teachability of reflective processes in written composition. *Cognitive Science*, 8, 173-190.
- Schoenfeld, A. H. (1983). Problem solving in the mathematics curriculum: A report, recommendations and an annotated bibliography. *The Mathematical Association of America*, MAA Notes, No. 1.
- Schoenfeld, A. H. (1985). *Mathematical problem solving*. New York: Academic Press.

VanLehn, K., & Brown, J. S. (1980). Planning nets: A representation for formalizing analogies and semantic models for procedural skills. In R. E. Snow, P. A. Federico, & W. E. Montague, (Eds.), *Aptitude learning and instruction, (Vol 2): Cognitive process analyses of learning and problem-solving*. Hillsdale, NJ: Erlbaum.

White, B. Y. (1984). Designing computer games to help physics students understand Newton's laws of motion. *Cognition and Instruction, 1*, 69-108.

White, B. Y., & Frederiksen, J. (in press). Progressions of qualitative models as a foundation for intelligent learning environments. *Artificial Intelligence*.



TABLE 1

The Acquisition of Question-Asking by a Weak Seventh-Grade Student  
Using Reciprocal Teaching (from Palincsar & Brown, 1984)

Day 1

*TEXT:* The water moccasin, somewhat longer than the copperhead, is found in southeastern states. It lives in swampy regions. It belongs, as do also the copperhead and the rattlesnakes, to a group of poisonous snakes called pit vipers. They have pits between their eyes and their nostrils which, because they are sensitive to heat, help the snakes tell them when they are near a warm-blooded animal. Another name for the water moccasin is "cottonmouth." This name comes from the white lining of the snake's mouth.

1. C: What is found in the southeastern snakes, also the copperhead, rattlesnakes, vipers--they have. I'm not doing this right.
2. T: All right. Do you want to know about pit vipers?
3. C: Yeah.
4. T: What would be a good question about the pit vipers that starts with the word "why?"
5. C: (No response)
6. T: How about, "Why are the snakes called pit vipers?"
7. C: Why do they want to know that they are called pit vipers?
8. T: Try it again.
9. C: Why do they, pit vipers in a pit?
10. T: How about, "Why do they call the snakes pit vipers?"
11. C: Why do they call the snakes pit vipers?
12. T: There you go! Good for you.

Day 4:

*TEXT:* Spinner's mate is much smaller than she, and his body is dull brown. He spends most of his time sitting at one side of her web.

13. C: (No question)
14. T: What's this paragraph about?
15. C: Spinner's mate. How do spinner's mate...
16. T: That's good. Keep going.
17. C: How do spinner's mate is smaller than...How am I going to say that?

18. T: Take your time with it. You want to ask a question about spinner's mate and what he does, beginning with the word "how."
19. C: How do they spend most of his time sitting?
20. T: You're very close. The question would be, "How does spinner's mate spend most of his time?" Now, you ask it.
21. C: How does spinner's mate spend most of his time?

*Day 7:*

*TEXT:* Perhaps you are wondering where the lava and other volcanic products come from. Deep within our earth there are pockets of molten rock called magma. Forced upward in part by gas pressure, this molten rock continually tries to reach the surface. Eventually--by means of cracks in the crustal rocks or some similar zone of weakness--the magma may break out of the ground. It then flows from the vent as lava, or spews skyward as dense clouds of lava particles.

22. C: How does the pressure from below pushes the mass of hot rock against the opening? Is that it?
23. T: Not quite. Start your question with, "What happens when?"
24. C: What happens when the pressure from below pushes the mass of hot rock against the opening?
25. T: Good for you! Good job.

*Day 11:*

*TEXT:* One of the most interesting of the insect-eating plants is the Venus's flytrap. This plant lives in only one small area of the world--the coastal marshes of North and South Carolina. The Venus's flytrap doesn't look unusual. Its habits, however, make it truly a plant wonder.

26. C: What is the most interesting of the insect eating plants, and where do the plants live at?
27. T: Two excellent questions! They are both clear and important questions. Ask us one at a time now.

*Day 15:*

*TEXT:* Scientists also come to the South Pole to study the strange lights that glow overhead during the Antarctic night. (It's a cold and lonely world for the few hardy people who "winter over" the polar night.) These "southern lights" are caused by the Earth acting like a magnet on electrical particles in the air. They are clues that may help us understand the Earth's core and the upper edges of its blanket of air.

28. C: Why do scientists come to the south pole to study?
29. T: Excellent question! That is what this paragraph is all about.

## TABLE 2

Planning Cues for Opinion Essays  
(From Scardamalia et. al., 1984)

### *NEW IDEA*

An even better idea is...  
An important point I haven't considered yet is...  
A better argument would be...  
A different aspect would be...  
A whole new way to think of this topic is...  
No one will have thought of...

### *IMPROVE*

I'm not being very clear about what I just said so...  
I could make my main point clearer...  
A criticism I should deal with in my paper is...  
I really think this isn't necessary because...  
I'm getting off the topic so...  
This isn't very convincing because...  
But many readers won't agree that...  
To liven this up I'll...

### *ELABORATE*

An example of this...  
This is true, but it's not sufficient so...  
My own feelings about this are...  
I'll change this a little by...  
The reason I think so...  
Another reason that's good...  
I could develop this idea by adding...  
Another way to put it would be...  
A good point on the other side of the argument is...

### *GOALS*

A goal I think I could write to...  
My purpose...

### *PUTTING IT TOGETHER*

If I want to start off with my strongest idea I'll...  
I can tie this together by...  
My main point is...

### TABLE 3

#### Example of Teacher Modelling in Response to a Student-Suggested Writing Assignment

#### ASSIGNMENT

Write an essay on the topic, "Today's Rock Stars are More Talented Than Musicians of Long Ago."

#### THINKING-ALOUD EXCERPT

I don't know a thing about modern rock stars. I can't think of the name of even one rock star. How about, David Bowie or Mick Jagger . . . But many readers won't agree that they are modern rock stars. I think they're both as old as I am. Let's see, my own feelings about this are . . . that I doubt if today's rock stars are more talented than ever. Anyhow, how would I know? I can't argue this . . . I need a new idea . . . An important point I haven't considered yet is . . . ah . . . well . . . what do we mean by talent? Am I talking about musical talent or ability to entertain--to do acrobatics? Hey, I may have a way into this topic. I could develop this idea by . . .

Note: Underlined phrases represent selection from planning cues similar to those shown in Table 2.

TABLE 4

An Example of Expert Modelling in Mathematics (from Schoenfeld, 1983)

**Problem**

Let  $P(x)$  and  $Q(x)$  be two polynomials with "reversed" coefficients:

$$\begin{aligned}
 P(x) &= a_n x^n + a_{n-1} x^{n-1} + \\
 &\dots + a_2 x^2 + a_1 x + a_0, \\
 Q(x) &= a_0 x^n + a_1 x^{n-1} + \\
 &\dots + a_{n-2} x^2 + a_{n-1} x + a_n,
 \end{aligned}$$

where  $a_n \neq 0 \neq a_0$ . What is the relationship between the roots of  $P(x)$  and those of  $Q(x)$ ? Prove your answer.

**Expert Model**

What do you do when you face a problem like this? I have no general procedure for finding the roots of a polynomial, much less for comparing the roots of two of them. Probably the best thing to do for the time being is to look at some simple examples, and hope I can develop some intuition from them. Instead of looking at a pair of arbitrary polynomials, maybe I should look at a pair of quadratics: at least I can solve those. So, what happens if

$$P(x) = ax^2 + bx + c, \text{ and}$$

$$Q(x) = cx^2 + bx + a?$$

The roots are

$$\frac{-b \pm \sqrt{b^2 - 4ac}}{2a} \text{ and } \frac{-b \pm \sqrt{b^2 - 4ac}}{2c}, \text{ respectively.}$$

That's certainly suggestive, since they have the same numerator, but I don't really see anything that I can push or that'll generalize. I'll give this a minute or two, but I may have to try something else...

Well, just for the record, let me look at the linear case. If  $P(x) = ax + b$  and  $Q(x) = bx + a$ , the roots are  $-b/a$  and  $-a/b$  respectively.

They're reciprocals, but that's not too interesting in itself. Let me go back to quadratics. I still don't have much of a feel for what's going on. I'll do a couple of easy examples, and look for some sort of a pattern. The clever thing to do may be to pick polynomials I can factor; that way it'll be easy to keep track of the roots. All right, how about something easy like  $(x + 2)(x + 3)$ ?

Then  $P(x) = x^2 + 5x + 6$ , with roots  $-2$  and  $-3$ . So,

$$Q(x) = 6x^2 + 5x + 1 = (2x + 1)(3x + 1), \text{ with roots } -1/2 \text{ and } -1/3.$$

Those are reciprocals too. Now that's interesting. How about

$$P(x) = (3x + 5)(2x - 7) = 6x^2 - 11x - 35? \text{ Its roots are } -5/3 \text{ and } 7/2:$$

$$Q(x) = -35x^2 - 11x + 6 = -(35x^2 + 11x - 6) = -(7x - 2)(5x + 3).$$

All right, the roots are  $2/7$  and  $-3/5$ . They're reciprocals again, and this time it can't be an accident. Better yet, look at the factors: they're reversed! What about

$$P(x) = (ax + b)(cx + d) = acx^2 + (bc + ad)x + bd? \text{ Then}$$

$$Q(x) = bdx^2 + (ad + bc)x + ac = (bx + a)(dx + c).$$

Aha! It works again, and I think this will generalize...

At this point there are two ways to go. I hypothesize that the roots of  $P(x)$  are the reciprocals of the roots of  $Q(x)$ , in general. (If I'm not yet sure, I should try a factorable cubic or two.) Now I can try to generalize the argument above, but it's not all that straightforward: not every polynomial can be factored, and keeping track of the coefficients may not be that easy. It may be worth stopping, re-phrasing my conjecture, and trying it from scratch:

Let  $P(x)$  and  $Q(x)$  be two polynomials with "reversed" coefficients. Prove that the roots of  $P(x)$  and  $Q(x)$  are reciprocals.

All right, let's take a look at what the problem asks for. What does it mean for some number, say  $r$ , to be a root of  $P(x)$ ? It means that  $P(r) = 0$ . Now the conjecture says that the reciprocal of  $r$  is supposed to be a root of  $Q(x)$ . That says that  $Q(1/r) = 0$ . Strange. Let me go back to the quadratic case, and see what happens.

Let  $P(x) = ax^2 + bx + c$ , and  $Q(x) = cx^2 + bx + a$ . If  $r$  is a root of  $P(x)$ , then  $P(r) = ar^2 + br + c = 0$ . Now what does  $Q(1/r)$  look like?

$$Q(1/r) = c(1/r)^2 + b(1/r) + a = \frac{c + br + ar^2}{r^2} = \frac{P(r)}{r^2} = 0$$

So it works, and this argument will generalize. Now I can write up a proof.

Proof:

Let  $r$  be a root of  $P(x)$ , so that  $P(r) = 0$ . Observe that  $r \neq 0$ , since  $a_0 \neq 0$ . Further,  $Q(1/r) = a_0(1/r)^n + a_1(1/r)^{n-1} + \dots + a_{n-2}(1/r) + a_n = (1/r^n)(a_0 + a_1r + a_2r^2 + \dots + a_{n-2}r^{n-2} + a_{n-1}r^{n-1} + a_nr^n) = (1/r^n)P(r) = 0$ , so that  $(1/r)$  is a root of  $Q(x)$ .

Conversely, if  $S$  is a root of  $Q(x)$ , we see that  $P(1/S) = 0$ . Q.E.D.

All right, now it's time for a post-mortem. Observe that the proof, like a classical mathematical argument, is quite terse and presents the results of a thought process. But where did the inspiration for the proof come from? If you go back over the way that the argument evolved, you'll see there were two major breakthroughs.

The first had to do with understanding the problem, with getting a feel for it. The problem statement, in its full generality, offered little in the way of assistance. What we did was to examine special cases in order to look for a pattern. More specifically, our first attempt at special cases -- looking at the quadratic formula -- didn't provide much insight. We had to get even more specific, as follows: Look at a series of straightforward examples that are easy to calculate, in order to see if some sort of pattern emerges. With luck, you might be able to generalize the pattern. In this case we were looking for roots of polynomials, so we chose easily factorable ones. Obviously, different circumstances will lead to different choices. But that strategy allowed us to make a conjecture.

The second breakthrough came after we made the conjecture. Although we had some idea of why it ought to be true, the argument looked messy and we stopped to reconsider for a while. What we did at that point is important, and often overlooked: we went back to the conditions of the problem, explored them, and looked for tangible connections between them and the results we wanted. Questions like 'what does it mean for  $r$  to be a root of  $P(x)$ ?', 'what does the reciprocal of  $r$  look like?' and 'what does it mean for  $(1/r)$  to be a root of  $Q(x)$ ?' may seem almost trivial in isolation, but they focused our attention on the very things that gave us a solution.

**TABLE 5**

**Characteristics of Ideal Learning Environments**

**Content**

- Domain knowledge
- Heuristic strategies
- Control strategies
- Learning strategies

**Methods**

- Modelling
- Coaching
- Scaffolding and fading
- Articulation
- Reflection
- Exploration

**Sequence**

- Increasing complexity
- Increasing diversity
- Global before local skills

**Sociology**

- Situated learning
- Culture of expert practice
- Intrinsic motivation
- Exploiting cooperation
- Exploiting competition





