CHAPTER 4

How Do Students Develop Mastery?

A Sum of Their Parts

I worked in industry for over twenty years before coming to academia, and I know how critical teamwork is, so in my Industrial Management course I assign a number of group projects in addition to individual projects. Students generally do well on their individual projects, and since the group assignments and individual assignments require more or less the same content knowledge, you would think that students would do even better on the group projects: after all, there are more people to share the work and generate ideas. Instead, it is just the reverse. Not only do my student groups fail to meet deadlines, but their analyses are also superficial and their projects lack internal coherence. I am not sure what the problem is, but at this point I am tempted to scrap the group projects and go only with individual projects. I just wish someone could explain to me why these groups are *less*, not *more*, than a sum of their parts.

Professor Fritz Solomon

Shouldn't They Know This by Now?

I just came from the second meeting of my acting class, and I have never felt so frustrated. This is an upper-level course, so by the time students get to my course they have already taken a number of courses in speech, voice, and movement. In other words, they *should* have a solid grounding in the fundamentals. Yet they make the most elementary mistakes! To give an example, I assigned students an easy scene from a Tennessee Williams play, something they should be able to handle with ease. And yet, a good proportion of the class mangled the Southern accents, dropped props, or mumbled their lines. Not only that, but they completely disregarded two things I know their instructors have emphasized over and over again in the introductory classes: the importance of doing vocal warm-ups and phonetically transcribing all their lines. How can they not know this stuff by now? I know they have learned it, because I have sat in on some of the first- and second-year classes and have been impressed by their skills. So why do they seem to have forgotten everything when they get to my course?

Professor Pamela Kozol

WHAT IS GOING ON IN THESE STORIES?

The instructors in these two stories believe that their students have the skills and knowledge necessary to perform well on the assigned tasks, yet their students' performance is disappointing, and neither instructor knows why. What is happening in each case that can help explain why these students fail to meet their instructor's expectations?

In fact, the tasks these instructors have assigned may require more from students than the instructors realize, and their students may be less prepared than their instructors assume. In the first story, for example, Professor Solomon expects the quality of group projects to be higher than the quality of individual projects because there are more people "to share the work and generate ideas." This seems like a reasonable assumption and is one that many instructors make. However, it is predicated on the expectation that students will know how to work effectively in groups. In fact, successful teamwork requires not only content skills and knowledge, but also an additional and qualitatively different set of process skills, such as the ability to delegate tasks, coordinate efforts, resolve conflicts, and synthesize the contributions of group members. When students possess the process skills necessary to manage the particular challenges of teamwork, the quality of work they produce in teams may indeed surpass the quality of the work they do individually. But when students lack these key component skills, it can seriously impede their performance.

Professor Kozol's students, in contrast, appear to have the necessary component skills. They have taken classes in and apparently mastered fundamental movement, voice, and speech skills. Yet when assigned a task that requires these skills, their performance is characterized by mistakes and omissions. Why? There are several possible explanations. First, although students have come to Professor Kozol's class with a solid grounding in movement, voice, and speech, they practiced these skills in classes targeting each skill area separately. Consequently, they may not have had sufficient practice using the complete set of skills in combination—especially while acting out an entire scene. If so, it is not the component skills they lack, but rather the ability to integrate them effectively.

Another possible explanation is that Professor Kozol's students did not recognize the relevance of phonetic transcriptions and vocal warm-ups—practices they had learned in previous courses—to the task they were assigned in her class. They may have failed to make this connection if their understanding of the underlying function of these practices was superficial or if they associated them entirely with the contexts (voice and speech classes) in which they had originally learned them. If so, the problem was not that students lacked component skills or that they were unable to integrate them successfully, but that they could not transfer them successfully to a new context and apply them appropriately.

WHAT PRINCIPLE OF LEARNING IS AT WORK HERE?

As the stories above suggest, tasks that seem simple and straightforward to instructors often involve a complex combination of skills. Think back to when you learned to drive. You had to keep in mind a sequence of steps (for example, adjust the mirrors, apply the brakes, turn the key in the ignition, put the car in reverse, check the rear view mirror, release the brake, press the accelerator), a set of facts (for example, traffic rules and laws, the meaning of street signs, the functions of the car's controls and gauges), and a set of skills (for example, accelerating smoothly, parallel parking, performing a three-point turn). You also had to learn how to integrate all of these component skills and knowledge, such as checking your mirror and moving into another lane. Finally, you had to recognize the appropriate context for certain knowledge and skills, such as adapting speed and braking behavior when driving on icy or clear roads.

To an experienced driver, driving is effortless and automatic, requiring little conscious awareness to do well. But for the novice driver it is complex and effortful, involving the conscious and gradual development of many distinct skills and abilities. A similar process exists in the development of mastery in academic contexts, as described in the following principle.

Principle: *To develop mastery, students must acquire component skills, practice integrating them, and know when to apply what they have learned.*

Mastery refers to the attainment of a high degree of competence within a particular area. That area can be narrowly or broadly defined, ranging from discrete skills (for example, using a Bunsen burner) or content knowledge (for example, knowing the names of all U.S. presidents) to extensive knowledge and skills within a complex disciplinary domain (for example, French theater, thermodynamics, or game theory). For students to achieve mastery within a domain, whether narrowly or broadly conceived, they need to develop a set of key component skills, practice them to the point where they can be combined fluently and used with a fair degree of automaticity, and know when and where to apply them appropriately (see Figure 4.1).

WHAT DOES THE RESEARCH TELL US ABOUT MASTERY?

Common sense suggests that having achieved mastery within a domain should position an instructor well to help novices develop mastery. But this is not necessarily the case. In the following sections we examine why expertise can potentially be a problem for teachers; we then explore research relevant to each element of mastery and discuss implications for teaching.

Expertise

Ironically, expertise can be a liability as well as an advantage when it comes to teaching. To understand why, consider the model of

Figure 4.1. Elements of Mastery

mastery offered by Sprague and Stuart (2000) and illustrated in Figure 4.2. It describes a four-stage developmental trajectory from novice to expert focused on two dimensions: competence and consciousness.

As illustrated in the diagram below, novice students are in a state of *unconscious incompetence*, in that they have not yet developed skill in a particular domain, nor do they have sufficient knowledge to recognize what they need to learn. Put simply, they do not know what they do not know. As they gain knowledge and experience, they advance to a state of *conscious incompetence*, where

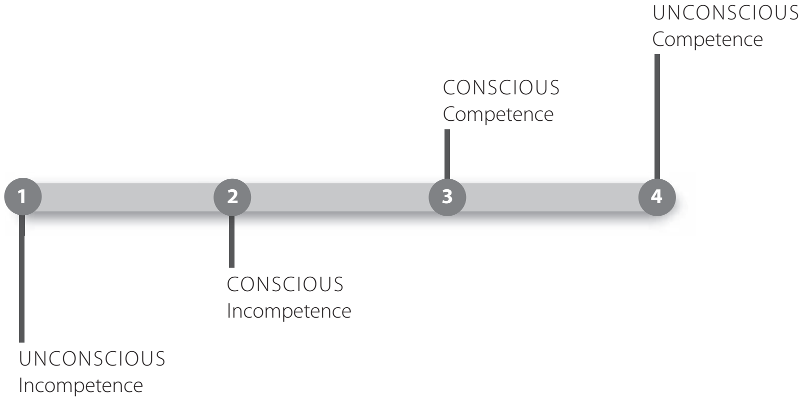

Figure 4.2. Stages in the Development of Mastery

they are increasingly aware of what they do not know and, consequently, of what they need to learn. As their mastery develops, students advance to a state of *conscious competence* wherein they have considerable competence in their domain, yet still must think and act deliberately and consciously. Finally, as students reach the highest level of mastery, they move into a state of *unconscious competence* in which they exercise the skills and knowledge in their domain so automatically and instinctively that they are no longer consciously aware of what they know or do. As this model suggests, while competence develops in a more-or-less linear way, consciousness first waxes and then wanes, so that novices (in stage one) and experts (in stage four) operate in states of relative unconsciousness, though for very different reasons.

It is easy to see why novices lack conscious awareness of what they do not know, but less obvious why experts lack conscious awareness of what they *do* know. Research on expert-novice differences helps to illuminate the issue, however. Experts, by definition, possess vastly more knowledge than novices, but they also organize, access, and apply their knowledge very differently (see

Chapter Two on organization of knowledge; Ericsson & Smith, 1991; Ericsson & Lehmann, 1996). For instance, experts organize knowledge into large, conceptual "chunks" that allow them to access and apply that knowledge with facility (Chase & Simon, 1973b; Chase & Ericsson, 1982; Koedinger & Anderson, 1990). Moreover, because experts immediately recognize meaningful patterns and configurations based on their previous experiences, they are able to employ shortcuts and skip steps that novices cannot (DeGroot, 1965; Anderson, 1992; Chase & Simon, 1973a; Koedinger & Anderson, 1990; Blessing & Anderson, 1996). Also, because experts have extensive practice in a narrowly defined area (for example, planning a problem-solving strategy or critiquing a theoretical perspective), they can perform with ease and automaticity tasks that are much more effortful for novices (Smith & Chamberlin, 1992; Lansdown, 2002; Beilock, Wierenga, & Carr, 2002). Finally, experts link specific information to deeper principles and schemas and are consequently better able than novices to transfer their knowledge across contexts in which those principles apply (see Chapter Two; Chi, Feltovich, & Glaser, 1981; Larkin et al., 1980; Boster & Johnson, 1989).

These attributes of expertise are an obvious advantage when instructors are working within their disciplinary domains, but they can be an obstacle to effective teaching. For example, the way instructors chunk knowledge can make it difficult for them to break a skill down so that it is clear to students. Moreover, the fact that instructors take shortcuts and skip steps with no conscious awareness of doing so means they will sometimes make leaps that students cannot follow. In addition, the efficiency with which instructors perform complex tasks can lead them to underestimate the time it will take students to learn and perform these tasks. Finally, the fact that instructors can quickly recognize the relevance of skills across diverse contexts can cause them to overestimate students' ability to do the same.

When expert instructors are blind to the learning needs of novice students, it is known as *expert blind spot* (Nickerson, 1999; Hinds, 1999; Nathan & Koedinger, 2000; Nathan & Petrosino, 2003). To get a sense of the effect of expert blind spot on students, consider how master chefs might instruct novice cooks to "sauté the vegetables until they are done," "cook until the sauce is a good consistency," or "add spices to taste." Whereas such instructions are clear to chefs, they do not illuminate matters to students, who do not know what "done" entails, what a "good consistency" is, or what spices would create a desired taste. Here we see the unconscious competence of the expert meeting the unconscious incompetence of the novice. The likely result is that students miss vital information, make unnecessary mistakes, and function inefficiently. They may also become confused and discouraged. Although they might muddle through on their own, it is unlikely that they will learn with optimal efficiency or thoroughness.

As instructors, we are all susceptible to expert blind spot. However, we can reduce the problems it poses for student learning by becoming more consciously aware of three particular elements of mastery that students must develop: (1) the acquisition of key component skills, (2) practice in integrating them effectively, and (3) knowledge of when to apply what they have learned.

Component Skills

As the driving and cooking examples above suggest, tasks that seem fairly simple to experts can hide a complex combination of component skills. For example, the ability to analyze a case study requires component skills such as the capacity to identify the central question or dilemma of the case, articulate the perspectives of key actors, enumerate constraints, delineate possible courses of action, and recommend and justify a solution. Similarly, problem solving might involve a number of component skills including (but not limited to) representing the problem, determining an appropriate solution strategy, doing the calculations necessary to execute that strategy, and evaluating the result. These component skills are particularly difficult to identify when they involve purely cognitive processes (for example, recognizing, planning, and formulating) that are not directly visible.

If students lack critical component skills—or if their command of those skills is weak—their performance on the overall task suffers (Resnick, 1976). This is demonstrated in a number of studies in which researchers decompose complex tasks, identify weak or missing component skills, and track the effect of those gaps on student performance. Lovett's (2001) research with introductory statistics students, for instance, identified two key skills involved in statistical data analysis: the ability to recognize the relevant variables and the ability to categorize them according to types. Lovett found that when students lacked these component skills, they were less able to choose appropriate forms of analysis and their performance on the overall problem-solving task was compromised (Lovett, 2001). We see a similar phenomenon in the first story at the beginning of the chapter: while Professor Solomon's students possess many of the component skills necessary for their group projects—as evidenced by their performance on individual assignments—their lack of teamwork skills erodes their overall performance.

In order to teach complex skills systematically—without missing pieces—instructors must be able to "unpack" or decompose complex tasks. This can be challenging because of expert blind spot, but there are tangible payoffs for student learning. Indeed, research indicates that when instructors identify and reinforce weak component skills through targeted practice, students' performance on the overall task often improves significantly. For example, Koedinger and Anderson (1990) found that, relative to experts, novice geometry students lacked the ability to plan problem-solving strategies. After assigning students exercises to specifically reinforce this skill within the context of the larger task, the researchers found that students became much more adept problem-solvers (Koedinger & Anderson, 1993). Lovett (2001) found that if beginning students were given a mere 45 minutes of practice identifying statistical problem types, and were given feedback on this particular skill, they were able to select appropriate analyses as adeptly as students who had had a semester-long course. In other words, even a small amount of focused practice on key component skills had a profound effect on overall performance. This same effect is demonstrated in research on cognitive tutors (computer-based tutoring programs), which are designed to detect the component skills that students lack and direct them to exercises that strengthen their abilities in those areas (Anderson et al., 1995; Singley, 1995; Ritter et al., 2007; Anderson, Conrad, & Corbett, 1989).

While we know that students need to practice component skills in order to improve their performance on the complex tasks involving those skills, the question of whether students should practice component skills in isolation or in the context of the whole task is more complicated. The advantage to practicing a component skill in isolation is that it allows students to focus their attention solely on the skill that needs work. Think, for example, of the benefits to a basketball player of drills that emphasize dribbling or shooting. Drilling these component skills in isolation gives players more repeated practice with each skill than they could ever get in the context of a game or scrimmage, and allows them to devote their energy and concentration exclusively to the skill in question. The advantage to practicing the whole task, on the other hand, is that students see how the parts fit into the whole in a context that is authentically complex. Think, for example, how much more difficult it is to shoot under defensive pressure in a game situation than when taking practice shots during a drill!

Whether or not students benefit more from practicing component skills in isolation or in the context of the overall task depends to a large extent on the nature of the task. Although the research results are mixed, it seems generally true that whole-task practice is preferable if the overall task is fairly simple or if components cannot be realistically extracted from the whole (Wightman & Lintern, 1985; Naylor & Briggs, 1963; Teague, Gittelman, & Park, 1994). However, if the task is highly complex and can be easily divided into component parts, students often learn more effectively if the components are practiced temporarily in isolation, and then progressively combined (White & Frederickson, 1990; Wightman & Lintern, 1985; Salden, Paas, & van Merrienboer, 2006). The extent to which isolated practice facilitates learning also depends in part on the skill level of the student. Studies have shown that explicit instruction and isolated practice of component skills, while helpful for novice learners (Clarke, Ayres, & Sweller, 2005), might be counterproductive for advanced learners if they have already integrated these components into a coherent whole (Kalyuga, Ayres, Chandler, & Sweller, 2003). Finally, the extent to which isolated practice is beneficial depends on the learning objectives of the class. For example, if a central objective of a course like Professor Solomon's is to help students build teamwork skills, then it might make sense to focus on specific skills in isolation. One example might be to reinforce students' abilities to reconcile intra-group differences of opinion by having them role-play responses to hypothetical conflicts.

Implications of This Research In order to build new skills systematically and to diagnose weak or missing skills, instructors must be able to break complex tasks down into their component parts. Decomposing complex tasks helps instructors pinpoint skills that students need to develop through targeted practice. However, in designing practice opportunities to reinforce component skills, instructors should consider whether their learning goals are best accomplished through isolated practice, whole-task practice, or some combination of the two.

Integration

Acquiring component skills does not by itself prepare students to perform complex tasks. This is because mastering complex tasks requires not only the temporary decomposition of subskills and the opportunity to practice them separately, but also their eventual recomposition and the opportunity to practice them in combination. Integrating component skills can be difficult and demanding, as is evidenced in the second story at the beginning of this chapter in which Professor Kozol's students struggle to integrate and use in combination skills they have learned separately.

The performance deficits that Professor Kozol's students exhibit when attempting to combine skills are not unusual. Many studies have shown that people's performance tends to degrade when they are asked to do more than one task at a time (Kahnemann, 1973; Navon & Gopher, 1979; Wickens, 1991). This degradation occurs because performing multiple tasks simultaneously tends to require attention to and processing of a great deal of information, and yet people have a limit to how much they can attend to and process at once. In other words, the total information-processing demands imposed by a given task or set of tasks—also known as *cognitive load*—can easily exceed what people can manage. When people's limit is exceeded, they are left with insufficient attention and other cognitive resources to complete the task effectively. For example, Strayer and Johnston (2001) found that when they asked adults to perform a simulated driving task, various measures of performance (for example, the number of traffic signals obeyed and reaction time for braking at red lights) declined when a cell-phone conversation task was added to the driving task. Furthermore, as the complexity of the cell-phone task increased, driving performance worsened. In other words, although the participants in this study likely had sufficient cognitive resources to perform well on the driving task in isolation, the more resources that were demanded by the secondary (cell phone) task, the fewer resources there were left for driving—leading to worse driving performance.

The same phenomenon often occurs when people perform a single complex task, because complex tasks require people to perform multiple skills in concert, which can similarly overload people's limited cognitive resources. Thinking back to Professor Kozol's acting class, it appears that her students could manage the cognitive load of voice, speech, or movement individually in classes devoted to each of those skill areas. However, the cognitive load of executing and coordinating these skills all at once—while incorporating new acting skills—may have been too much for them to manage, as revealed in their errors and mistakes.

Interestingly, experts do not suffer as much as novices when performing complex tasks or combining multiple tasks. Because experts have extensive practice within a circumscribed domain, the key component skills in their domain tend to be highly practiced and more automated. Each of these highly practiced skills then demands relatively few cognitive resources, effectively lowering the total cognitive load that experts experience. Thus, experts can perform complex tasks and combine multiple tasks relatively easily (Smith & Chamberlin, 1992; Lansdown, 2002; Beilock, Wierenga, & Carr, 2002). This is not because they necessarily have *more* cognitive resources than novices; rather, because of the high level of fluency they have achieved in performing key skills, they can *do* more with what they have. Novices, on the other hand, have not achieved the same degree of fluency and automaticity in each of the component skills, and thus they struggle to combine skills that experts combine with relative ease and efficiency.

Because instructors, as experts, do not experience the same cognitive load as novices, they may have performance expectations for students that are unrealistically high. This can lead to the kind of astonishment and frustration Professor Kozol experiences as her students struggle with an assignment she perceives as easy. For her, combining speech, voice, movement, and other acting skills is not terribly cognitively demanding, so her students' mistakes seem inexplicable. Fortunately, as students gain mastery over time, the knowledge and procedures required for complex tasks become automatized and thus require fewer cognitive resources. Thus, with practice, students gain greater fluency in executing individual subskills and will be better prepared to tackle the complexity of multiple tasks.

How then can we help students manage cognitive load as they learn to perform complex tasks? One method that has proved effective in research studies is to allow students to focus on one skill at a time, thus temporarily reducing their cognitive load and giving them the opportunity to develop fluency before they are required to integrate multiple skills. For example, Clarke, Ayres, and Sweller (2005) found that math students who knew little about spreadsheets learned less and performed less well when they were taught new mathematical concepts in the context of spreadsheets. This is because they had to learn both the spreadsheet skills and the math concepts concurrently, and they became overwhelmed. However, when these students first learned spreadsheet skills and *then* used those skills to learn the mathematics, learning and performance improved. Another method to emerge in the research is to support some aspects of a complex task while students perform the entire task (Sweller & Cooper, 1985; Cooper &

Sweller, 1987; Paas & van Merrienboer, 1994). For example, Sweller and Cooper (1985) demonstrated this with students learning to solve problems in a variety of quantitative fields from statistics to physics. They found that when students were given typical word problems, it was possible for them to solve the problems without actually learning much. This is because the problems themselves were sufficiently demanding that students had no cognitive resources available to learn from what they did. But when students were given "worked-examples" (such as presolved problems) interspersed with problems to solve, studying the worked-examples freed up cognitive resources that allowed students to see the key features of the problem and to analyze the steps and reasons behind problem-solving moves. The researchers found this improved students' performance on subsequent problem solving. This result, called the *worked-example effect,* is one example of a process called *scaffolding,* whereby instructors temporarily relieve some of the cognitive load so that students can focus on particular dimensions of learning. (For more discussion on scaffolding, see Chapter Seven.)

A subtle but important point to mention here is that some reductions in cognitive load promote learning while others do not (Paas, Renkl, & Sweller, 2003, 2004). The key to reducing cognitive load effectively lies in identifying which of the demanding aspects of a task are related to the skills students need to learn and which may be disruptive to (or distracting from) those learning goals. Research has shown that removing extraneous load—that is, aspects of a task that make it difficult to complete but that are unrelated to what students need to learn—is helpful. In contrast, reducing load that is germane to what students need to learn will naturally be counterproductive in that students will not have a chance to practice what they need to learn. To illustrate this distinction between extraneous and germane load, consider engineering students who are having difficulty solving practice problems. They have been introduced to a number of different formulas over the course of the semester and are having trouble keeping them straight. Now, if the instructor's goal is for students to learn to select and apply the appropriate formula for each of the problems, then giving students a sheet listing all the relevant formulas might be a legitimate choice: it would reduce extraneous load because students would no longer have to spend their time and cognitive resources *remembering* the relevant formulas and could focus instead on the skills of selection and application. However, if the instructor's goal is for students to be able to remember the formulas and then apply each one when told to do so, a sheet listing all the formulas would obviously thwart the learning goal.

Implications of This Research Performing complex tasks can be cognitively demanding for students, particularly when they have not yet developed fluency or automaticity in all the component skills. Thus, instructors should have reasonable expectations about the time and practice students will need, not only to develop fluency in component skills but also to learn to integrate those skills successfully. It can be helpful under some circumstances for instructors to strategically lighten aspects of the task that introduce extraneous cognitive load so that students can focus their cognitive resources on the aspects of a task most germane to the learning objectives. Several specific ways to do this are discussed in the Strategies section.

Application

As we have seen, mastery requires component skills *and* the ability to integrate them successfully. However, it also requires that students know when and where to use what they have learned. When students acquire skills but do not learn the conditions of their appropriate application, they may fail to apply skills that are relevant to a task or problem, or, alternatively, apply the wrong skill for the context.

The application of skills (or knowledge, strategies, approaches, or habits) learned in one context to a novel context is referred to as *transfer*. Transfer is said to be *near* if the learning context and transfer context are similar, and *far* when the contexts are dissimilar. For example, various dimensions of farness come into play when a student is given a task in his Public Policy course that requires him to apply a statistics formula he learned two semesters previously in Statistics 101. Not only has the knowledge domain changed from statistics to public policy, but so too have the physical and temporal contexts (a new class, two semesters later). If the transfer task were in a different functional context altogether, say outside academia, additional transfer distance would be introduced (for a discussion of different dimensions of transfer, see Barnett & Ceci, 2002).

Far transfer is, arguably, the central goal of education: we want our students to be able to apply what they learn beyond the classroom. Yet most research has found that (a) transfer occurs neither often nor automatically, and (b) the more dissimilar the learning and transfer contexts, the less likely successful transfer will occur. In other words, much as we would like them to, students often do not successfully apply relevant skills or knowledge in novel contexts (Singley & Anderson, 1989; McKeough, Lupart, & Marini, 1995; Thorndike & Woodworth, 1901; Reed, Ernst, & Banerji, 1974; Singley, 1995; Cognition and Technology Group at Vanderbilt, 1994; Singley & Anderson, 1989; Holyoak & Koh, 1987). In this section, we examine why this is the case by exploring issues that can affect transfer negatively and positively.

There are a number of reasons students may fail to transfer relevant knowledge and skills. First, they may associate that knowledge too closely with the context in which they originally learned it and thus not think to apply it—or know how to apply it—outside that context. This is called *overspecificity* or *context dependence* (Mason Spencer & Weisberg, 1986; Perfetto, Bransford, & Franks, 1983). To illustrate: students in a statistics course might perform well on their chapter quizzes but perform poorly on a final exam involving questions of precisely the same type and difficulty, but from a number of different chapters. If students relied on superficial cues to figure out which formula to apply on chapter quizzes (for example, if it is chapter 12, it must be a T-test), then in the absences of these cues, they may have been unable to identify the salient features of each problem and select an appropriate statistical test. Their knowledge, in other words, was overly context dependent and thus not flexible. Context dependence may also account for why students in Professor Kozol's class failed to phonetically transcribe their lines. If they associated phonetic transcription narrowly with the physical context in which they learned it (speech class), it may not have occurred to them to carry this practice over to their acting class.

Second, students may fail to transfer relevant skills, knowledge, or practices if they do not have a robust understanding of underlying principles and deep structure—in other words, if they understand what to do but not why. This might explain some of the problems Professor Kozol encountered in the story at the beginning of this chapter. If Professor Kozol's students understood some of the functions of vocal warm-ups (for example, to prevent vocal strain when singing) but not others (such as to relax the voice for greater emotional expressivity), they might not have recognized the applicability of this practice to the assigned task. In other words, an incomplete understanding of the functions of this practice might have affected their ability to apply it appropriately in new contexts.

Fortunately, much of the same research that documents transfer failure also suggests instructional approaches that can bolster transfer. For example, studies have shown that students are better able to transfer learning to new contexts when they can combine concrete experience within particular contexts and abstract knowledge that crosscuts contexts (Schwartz et al., 1999). A classic study by Schoklow and Judd (in Judd, 1908) illustrates this point. The researchers asked two groups of students to throw darts at a target twelve inches under water. Predictably, the performance of both groups improved with practice. Then one group was taught the abstract principle of refraction, while the other was not. When asked to hit a target four inches under water, the group that knew the abstract principle adjusted their strategies and significantly outperformed the other group. Knowing the abstract principle helped students transfer their experiential knowledge beyond the immediate context in which it was learned and to adjust their strategies for new conditions. Similarly, when students have the opportunity to apply what they learn in multiple contexts, it fosters less context-dependent, more "flexible" knowledge (Gick & Holyoak, 1983).

Structured comparisons—in which students are asked to compare and contrast different problems, cases, or scenarios—have also been shown to facilitate transfer. For example, Loewenstein, Thompson, and Gentner (2003) asked two groups of management students to analyze negotiation training cases. One group analyzed each case individually; the other group was asked to compare cases. The researchers found that the group that compared cases demonstrated dramatically more learning than the group that considered them individually. Why? Because when students were asked to compare cases, they had to recognize and identify the deep features of each case that would make it analogous or non-analogous to other cases. Having identified those deep features, students could link the cases to abstract negotiation principles, which then allowed them to learn more deeply and apply what they learned more effectively. Other methods that have been found to facilitate transfer include analogical reasoning (Gentner, Holyoak, & Kokinov, 2001; Catrambone & Holyoak, 1989; Holyoak & Koh, 1987; Klahr & Carver, 1988), using visual representations to help students see significant features and patterns (Biederman & Shiffrar, 1987), and asking students to articulate causal relationships (Brown & Kane, 1988).

Finally, research indicates that minor prompts on the part of the instructor can aid transfer. In Gick and Holyoak's (1980) study, college students were presented with a passage describing a military conundrum in which an army is trying to capture a fortress and must ultimately divide into small groups, approach from different roads, and converge simultaneously on the fortress. After memorizing this information, students were presented with a medical problem that required a similar solution (the use of multiple laser beams coming from different angles and converging on a tumor). Despite having just encountered the military solution, the large majority of students did not apply what they had learned to the medical problem. Even though the physical, social, and temporal contexts were the same, the knowledge domains (military strategy versus medicine) and functional contexts (storming a fortress versus treating a tumor) were sufficiently different that students did not recognize their analogous structures or think to apply knowledge from one problem to the other. However, when students were asked to think about the medical problem in relation to the military one, they could solve it successfully (Gick & Holyoak, 1980). Similar results have been shown in other studies as well (Perfetto et al., 1983; Klahr & Carver, 1988; Bassok, 1990). A little prompting, in other words, can go a long way in helping students apply what they know.

Implications of This Research Transfer does not happen easily or automatically. Thus, it is particularly important that we "teach for transfer"—that is, that we employ instructional strategies that reinforce a robust understanding of deep structures and underlying principles, provide sufficiently diverse contexts in which to apply these principles, and help students make appropriate connections between the knowledge and skills they possess and new contexts in which those skills apply. We consider some specific strategies under the heading, "Strategies to Facilitate Transfer," later in this chapter.

WHAT STRATEGIES DOES THE RESEARCH SUGGEST?

The following strategies include those faculty can use to (1) decompose complex tasks so as to build students' skills more systematically and to diagnose areas of weakness, (2) help students combine and integrate skills to develop greater automaticity and fluency, and (3) help students learn when to apply what they have learned.

Strategies to Expose and Reinforce Component Skills

Push Past Your Own Expert Blind Spot Because of the phenomenon of expert blind spot, instructors may have little conscious awareness of all the component skills and knowledge required for complex tasks. Consequently, when teaching students, instructors may inadvertently omit skills, steps, and information that students need in order to learn and perform effectively. To determine whether you have identified all the component skills relevant for a particular task, ask yourself: "What would students have to know—or know how to do—in order to achieve what I am asking of them?" Keep asking this question as you decompose the task until you have identified all the key compo-