4

Modeling Expertise

Modeling an expert literally means watching and imitating what an expert does. For instance, in learning a motor task such as serving in tennis, it would mean a novice's observing how the tennis teacher throws her elbow back and how far she reaches behind her back with her racket. The observation itself may in fact be more accurate than the teacher's own analysis of how she actually serves. Conversely, in correcting the student's serve, the teacher can observe the outcome of the serve—where the ball landed—and conclude what might have been the postural cause and provide feedback accordingly. For "transparent" tasks such as motor tasks (playing tennis) or other physical skills (such as weaving), the connection between the purpose and process of an action taken and its outcome is direct, and correction is straightforward, although often difficult to achieve.

There are two obvious difficulties in using direct modeling for complex cognitive tasks. First, the rationale for the performance of the tasks is not only opaque to observers, but may also be implicit for the experts: they may not be able to describe their own thought processes or the rationale for them, even though they can perform the tasks. Second, in order to properly coach a novice, an expert may have to formulate an accurate mental model of the novice's understanding of the task (sometimes called the student model). But a novice's understanding of a task is not always obvious to the expert. These two key problems—which are woven throughout this chapter—suggest that there might be a limitation on the extent to which direct modeling of complex cognitive skills can be done.

For example, suppose that a trainee is asked to learn to diagnose a fault when the instrument panel shows three problems: warning light A is flashing (condition A), dial B is off (condition B), and the sound system is generating a warning tone (condition C). The expert arrives and on seeing the three conditions, takes the following three actions: checks the power supply (action X), swaps the board (action Y), and resets a gauge (action Z). By observing this sequence of actions, what can the trainee learn? The trainee could learn a set of individual local rules, such as: if condition A holds, then take action X; if condition B holds, take action Y, and so forth. In order to learn this simple sequence of conditional action rules, the trainee has to accomplish two processes. First, he or she must learn to identify what the precise condition is for each action. Was it the flashing light that mattered (condition A), or was it the yellow flashing light rather than a blue flashing light that mattered? Second, the trainee has to link the condition with the action. For example, dial B, saying "off," initiates resetting a gauge. Inducing these simple local rules are by no means trivial, but once learned, they can be quite powerful in troubleshooting a number of problems.

It may be, however, that the expert's actions were not meant to be interpreted as a sequence of local conditional rules, but that a conjunctive rule should have been learned: if the pattern of A, B, and C conditions occurs, then take actions X, Y, and Z. This means that if conditions A, B, and F occur, one would not take action X, action Y, and some alternative third action that is an appropriate action to condition F in isolation; instead, perhaps when the pattern of conditions A, B, and F, occurs, some totally different action should be taken. Alternatively, perhaps conditions A, B, and C lead the expert to realize that this is a special kind of fault Q, which normally can be solved by taking actions X, Y, and Z. If this is the case, then the trainee has to learn to identify what kind of conditions typically can be considered to be a type Q kind of fault so that the sequence of actions X, Y, and Z are appropriate. Direct modeling, in which a simple linking of actions to conditions is induced, typically results in nontransferable skills: the trainee has not understood and learned the reasons behind the actions, nor has the trainee accumulated the knowledge necessary to recognize the conditions which then generate the actions. (This simplistic example assumes that direct modeling is not accompanied by verbal explanations. The nature of effective explanations is itself an important and complex topic.)

In the first of this chapter's four sections, we consider direct modeling of complex cognitive skills. We focus on a body of contemporary research on cognitive apprenticeship. In the second section we review the ways in which experts excel, considering their possession of a greater

knowledge base, the way that knowledge is organized, and processing strategies to solve problems. Since expertise is based on knowledge, which in turn generates the actions that experts take, we focus in the third section on the difficult process of extracting an expert's knowledge. The last section focuses on how knowledge extracted from experts can be imparted to novices.

COGNITIVE APPRENTICESHIP

There is one active body of research that promotes a technique for direct modeling of complex cognitive tasks, called cognitive apprenticeship (Collins et al., 1989). Cognitive apprenticeship borrows heavily from traditional apprenticeship, which is quite successful in teaching physical skills. Traditional apprenticeship involves three key components: observation, coaching, and practice. Observation means the apprentice participates as a spectator, observing a master or expert executing the target skill. Coaching refers to the guidance that the expert provides while the apprentice attempts to perform the task. Coaching physical skills involves two key features. First, the coaching or feedback is given in a continuous, on-line fashion. For example, as an apprentice in weaving is weaving the threads, the master might guide the apprentice's hands (Rogoff, 1986); it should be noted that the guidance provided in traditional apprenticeship is often physical demonstration, not verbal instruction. Second, the master provides conceptual "scaffolding," that is the support, in the form of reminders and help, necessary for the apprentice to perform an approximation of the composite task. The degree of scaffolding provided depends on the extent of help the apprentice needs. As the apprentice improves in his or her skill, the scaffolding can be "faded." The expert, therefore, must monitor the apprentice's "zone of proximal development" (Vygotsky, 1978), or "region of sensitivity to instruction" (Wood and Middleton, 1975). The zone of proximal development is the distance between the developmental levels at which children can perform a task alone and the level at which they can perform it with some assistance. Third, the apprentice practices with the master present. In this fashion, the apprentice begins by executing piecemeal aspects of the skill and yet enjoys the reward of the entire skill. The entire learning situation is embedded in practice or guided practice.

Cognitive apprenticeship consists of six key components: modeling, coaching, scaffolding (and fading), articulation, reflection, and exploration (Collins et al., 1989). The first three are provided mostly by the teacher (expert), and the last three are exercised by the learner (novice).

Modeling

Modeling cognitive tasks requires that the expert explain ("externalize") as much of the process that underlies his or her actions as possible. For example, in solving a mathematics problem, the expert should talk out loud for the novice while coming up with the equation. This externalization is comparable to the think-aloud protocols that cognitive psychologists have been collecting as their data for analyzing the processes of problem solving (see Newell and Simon, 1972). Ideally, the train of thought should reveal all the knowledge that the expert is using, not only the factual and conceptual knowledge that is relevant to the substantive domain, but also such strategic knowledge as special heuristics for solving problems (such as "decomposing" the problem into subproblems or considering the problem as a "simplified or special case"); control strategies in making decisions, such as knowing which path to search; and learning strategies, such as scanning the table of contents before reading a book to get a general idea of what the book is about. The basic idea is to expose the complete thought processes of the expert, including searching the wrong paths, arriving at incorrect solutions, and so forth. Externalizing thought processes in this way should allow a novice to learn about heuristics that may be useful for solving problems.

There are two aspects to this procedure of externalizing one's thought processes. One aspect, which can be made known readily, is to make overt the solution trace that the expert is undertaking. That is, the expert simply articulates the operations he or she is carrying out as they are executed. The other aspect is more tacit and may not be exposed as readily: the expert explains the rationale for his or her selection of operations. The first aspect, exposing the "solution trace," tends to reveal a strategy that the expert is following; the second aspect, exposing the rationale, may reveal the expert's knowledge structure. This procedure implies that modeling would be more effective in domains in which the use of a few strategies can promote learning and would be less effective in domains that require deep domain knowledge.

Schoenfeld's (1983) teaching of mathematics illustrates this distinction. An excerpt of his protocol on finding the relationship between the roots of two polynomials with "reversed" coefficients shows the application of a special heuristic that he attempts to teach students to use. The heuristic is called "special cases," and the idea is to solve simpler cases, such as to find the roots of simpler quadratic equations rather than complex polynomials and to look for relations between the roots. If one fails to "see" any relationship among the roots, then the next step is to further reduce the polynomials to linear cases and solve for their roots. The expert thus reduces the complex polynomials to simpler equations:

first to quadratic, then to linear equations. The expert then observes a pattern in the solution of linear cases—that the roots are reciprocals of each other.

The key strategy that Schoenfeld wants to get across is to generate and test for a series of straightforward examples and then to see if some sort of pattern emerges. The processes of doing so are exposed to the student by having the expert talk out loud as he or she solves problems. Thus, basically, the expert is modeling the strategies of generate-and-test and is reducing the problem to simpler ones. What the expert is not modeling is the rationale for rejecting certain roots after testing or how to select a set of roots to consider in the first place. For example, the first pair of roots that Schoenfeld obtained from the quadratic equation were

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

He then dismisses these by saying, "I don't really see anything that I can push or that'll generalize." How did Schoenfeld know that these roots are uninteresting, without the potential for generalization? It is clear that by carefully grading the exercise problems for the students, they will have learned the strategy of reducing the problem to simple ones and of trying to look for patterns. However, no pattern can be detected unless a student has recognizable patterns already stored in memory. The point is that one can easily model the strategy that an expert is undertaking, but it is more difficult to model and expose the tacit knowledge that the expert might be using to carry out the strategy.

In sum, this kind of modeling (exposing the reasoning processes) is far superior to traditional instruction in which a student is simply given a solution (such as a worked-out example) and expected to induce the steps that were undertaken to arrive at the solution; however, some aspects of the rationale of an expert's reasoning processes may still remain tacit and not be easy for students to follow.

Coaching and Scaffolding

Coaching consists of observing students who are performing a task and offering hints, feedback, reminders, new tasks, or redirecting a student's attention to a salient feature—all with the goal of making the students' performance approximate the expert's performance as closely as possible. In traditional apprenticeship, coaching is fairly constant and continuous. From the protocol excerpt cited in Collins et al. (1989), the coaching provided by the teacher was predominantly in the form of prompts for students either to summarize, ask questions, make predictions, or clarify

difficulties. Interspersed with prompts were feedback (such as remarks about the quality of the summary or definitions of terms) and modeling (used here more in the sense of imitation rather than exposing the reasoning process). As mentioned above, however, it is not clear how a teacher knows which prompts and guidelines to give—or what feedback to provide—unless the teacher has an accurate idea of the student's mental model.

Scaffolding refers to the support a teacher provides so that the student can succeed in performing the task. This support can be suggestions, hints, actual execution of parts of the task, or physical props such as providing cue cards (Scardamalia et al., 1984). Scaffolding thus involves cooperative execution by an expert and a student, in a way that allows the novice to take an increasingly larger burden for performing the task. In both scaffolding and coaching, the expert is required to monitor the student's progress and understanding. Such monitoring processes are demanding and complex, and they are not well understood by cognitive psychologists. It is not now clear which aspects of cognitive apprenticeship are more effective in promoting learning—modeling or scaffolding and coaching. It is likely that the advantage gained results from an interactive effect. An excellent example of a successful apprenticeship learning program is provided by Brown and Palincsar (1989), in which they used modeling, coaching, and scaffolding to teach students four strategic skills for reading comprehension: formulating questions about the text, summarizing, making predictions about what will come next in the text, and clarifying difficulties with the text. The students' acquisition and use of these skills improved their comprehension tremendously.

Articulation, Reflection, and Exploration

The purpose of modeling in cognitive apprenticeship is to encourage students to undertake three activities: articulation, reflection, and exploration. Articulation, defined broadly, involves any kind of overt constructive activity carried out by a student, such as explaining the reasoning in an example problem (Chi et al., 1989b), summarizing, predicting, clarifying and asking questions (Palincsar and Brown, 1984), posing problems (Sharples, 1980), editing or revising the text and formulating hypotheses (Collins and Stevens, 1983) or goals (Schuell, 1991). The cognitive outcome of these activities is the integration, synthesis, and linking of knowledge in memory.

Reflection refers to the process of evaluating one's own problem-solving processes and comparing them with another student's or the teacher's processes. Such comparison presumably leads to the perception

of ambiguity, conflicts, and so forth, which will allow the student to modify his or her own problem-solving or decision-making processes. Reflection seems to involve complex processes: since it is known that confrontation by a teacher is not necessarily an effective means of instruction, it is not clear why self-perceived conflicts as obtained through reflection would be effective at promoting learning.

Exploration refers to the pursuing of new goals, formulating and testing new hypotheses, conducting experiments, and so forth. Although it may not be clear why each of these processes might work, the common characteristic of articulation, reflection, and exploration is the pursuit of activities that can foster an integration and synthesis of a student's newly acquired knowledge, so that additional knowledge can be inferred and constructed.

It is important to emphasize that in cognitive apprenticeship one cannot separate the role of the expert from the role of the learner. The expert's role entails modeling the target skill, preferably accompanying it with explanation of the rationales; modeling the use of learning and monitoring strategies; monitoring the student's progress and understanding; and providing feedback. How all this is accomplished in a social context is not yet understood. In fact, studies of tutoring show that experts in a tutoring situation often do not formulate ideas of students' mental models nor tailor feedback to the student's current level of understanding (Putnam, 1987). Hence, an important aspect of cognitive apprenticeship lies in the constructive activities in which the student or apprentice engages. We return to this point below.

HOW EXPERTS EXCEL

The cognitive science literature on expertise in the last two decades has concentrated primarily in demonstrating how experts excel in the domain for which they have expertise. The focus has been on three issues: how experts perform on laboratory tasks such as recall, recognition, problem solving, and decision making; how the knowledge of experts and novices differ; and how the strategies that experts and novices use to perform the tasks that they are given differ.

Measures in Laboratory Tasks

One way to try to assess the abilities of an expert is by means of standard laboratory tasks used by experimental psychologists, such as recall, recognition, comprehension, speed of performance, monitoring accuracies, automaticity, flexibility, and so forth. We illustrate the first four of these below.

One of the most widely used laboratory measures of expertise is recall. Experts such as chess masters can remember a far greater number of chess pieces from a chess board than can a beginning player. For example, if a chess board with pieces that define a midgame position is shown to both expert and novice chess players for 5 seconds, the expert player can reproduce the locations of about 25 chess pieces, while novices can reproduce only 5 or 6 pieces (Chase and Simon, 1973). This measure holds true for any game domain in which recall can be easily measured, such as the game of GO or bridge (Charness, 1979; Reitman, 1976)

Experts are also better than others at recognizing the important features of a stimulus, such as an X-ray, a diagram of an electronic circuit (Egan and Schwartz, 1979), or a topological map. Basically, they can detect whatever salient features or patterns exist in the stimulus (Lesgold et al., 1988) and can recognize during later trials which sets of stimuli have been presented earlier. Chase and Ericsson (1981) showed, for example, that a person who had managed to develop an outstanding memory span—the ability to give back over 100 digits presented at the rate of 1 digit per second—could also recognize which subsequences of digits were the ones that had been presented in a previous memory trial.

Not surprisingly, experts can also better comprehend passages in a domain for which they have expertise. Chiesi et al. (1979) have shown that experts can also understand event sequences much better than can nonexperts. In baseball, for example, such greater comprehension is demonstrated by the ability to generate action outcomes that could happen on the next play, given a specific game state. In addition, the actions that they generate are important to the game's goal structure.

Monitoring accuracy refers to how accurately experts can monitor or assess their own cognitive states, such as how well they have understood a passage, how accurately they can predict how much effort they need to spend studying, or what they think they can remember. Again, experts are far more accurate than novices in predicting and assessing their own state of comprehension. In a study looking at how good and poor problem solvers study physics examples taken from a textbook, Chi et al. (1989b) found that the good solvers were much more accurate at assessing whether they had understood a specific line in an example solution. Moreover, when experts do not understand something, they know more specifically what aspect of the example they do not understand, so that the questions they ask about an example are more specific and targeted. Poor solvers are much more inaccurate in their monitoring of their own comprehension; they almost always think they have understood the example. Likewise, skilled chess players are more accurate than less skilled players at predicting how many times they need to look

at the chessboard before they will remember the locations of all the chess pieces (Chi, 1978).

An example from physics problem solving shows another dimension of students' knowledge. Because experts can "see" beyond the surface description of a problem, and "know" the solution procedure, they tend to solve a problem by the appropriate method, such as using the force law. A student, however, might solve a pulley problem the same way he or she solved a previous pulley problem, simply because they both involve pulleys, without realizing that these two problems may require two different solution methods.

In sum, experts excel on laboratory tasks in their domain of expertise (see also Chi et al., 1988; Hoffman, 1991). There are two ways to explain why they excel: in terms of their knowledge structures or in terms of the strategies and heuristics they use. Each of these is discussed below.

Organized Knowledge Structures

Experts are able to perform more efficiently than novices on the types of tasks summarized because they have more knowledge. However, it is not merely the presence of knowledge that is responsible for the demonstration of skilled performance; rather, it is how that knowledge is organized. More knowledge, if improperly organized, actually leads to a deterioration of performance (Anderson, 1974).

There are various ways to portray the organization of experts' knowledge, corresponding more or less to the kind of knowledge that is being used during task performance. This knowledge organization can be characterized by schemas (mostly for declarative knowledge), production rules (for procedural knowledge), and mental models (for both). Schemas and production rules allow an expert to use knowledge flexibly and efficiently, to represent problems in a deep way so as to access automatic processes, to recognize patterns and chunks, and so forth.

Schemas involve both conceptual and factual knowledge. Experts have many more concepts, as well as many more features associated with each concept, than do novices. More important, their concepts are interrelated in meaningful ways, as reflected in their category structures. Chi and Koeske (1983) have shown, for example, that for a given partial set of concepts that is available to both experts and novices, the experts' concepts are interrelated in meaningful ways that correspond to family types (in this research the subject domain is dinosaur knowledge); novices' concepts are not interrelated in meaningful categories. The existence and importance of meaningful patterns of associations have been known for quite some time in visual-perceptual domains, such as chess

playing. Chase and Simon (1973) showed that chess masters "see" patterns of interrelated chess pieces on the chess board. It is these familiar patterns of pieces that enable them to select the best move because these familiar patterns are associated with optimal moves.

Experts know many more facts than do nonexperts. In a domain such as physics, it is obvious that students must learn many facts—such as "A point mass may be considered to be a body" in Newton's force law, or that g is the gravitational constant. In another example of knowing more facts, of taxi drivers' knowledge of routes, Chase (1983) found that expert drivers can generate a far greater number of secondary routes (lesser known streets) than novice drivers; furthermore, they can use this knowledge by generating short-cuts to reach a target destination when there is an impasse on the main routes.

Again, the organization of facts is what facilitates their efficient use and what distinguishes expertise. A commonly discussed representation for factual knowledge is the schema or mental model, which is a knowledge structure that captures the essential features of concepts, categories, situations, or events. Thus, one way that factual knowledge about a category or an event is made coherent is through the structure imposed by a schema. For example, a schema about the concept of dog would contain prototypical instances, as well as key features that are true of a generic dog. Each feature can be thought of as a dimension with a range of acceptable values. Similarly, a complete coherent mental model of a city permits expert cab drivers to take efficient secondary routes when they encounter road blocks on the primary routes.

Mental models are analogous to a "structured analog of the world" (Johnson-Laird, 1983:165). They represent not only the objects and properties comprising a system or event, but, more importantly, they represent the structural, functional, and causal relations among the components. Moreover, a mental model can be "run" so that it can capture the dynamic aspects of a system. In many ways, one can conceive of experts' superior skill as having more accurate mental models.

Procedural knowledge generally refers to how one performs a task. There are two kinds of procedural knowledge. Domain-specific procedural rules are specific how-to rules that have actions attached to the conditions of application in the domain. For instance, in the domain of electronic troubleshooting, seeing a flashing light might trigger a rule of checking the power switch or some other switch. Domain-specific procedural knowledge, as the name implies, is related to a specific domain and cannot be used in a different domain; expert chemists, for instance, cannot solve agriculture problems (Voss et al., 1983).

It is clear that experts possess a large quantity of procedural rules that enable them to perform a task automatically and efficiently, as soon as

they "see" the conditions. It should be stressed that the conditions are visible for all to see, but only the experts realize the import or meaningfulness of what is seen. Experts can perform some tasks automatically because many of these procedural rules, over time and with repeated use, become combined into one larger rule. Thus, several rules would have all their conditions concatenated, and upon seeing the conditions, a sequence of actions is automatically taken. Procedural knowledge is generally represented by condition-action rules, and the entire set of rules is called a production system.

Metaknowledge is knowledge about what one knows, as well as one's ability to monitor one's own comprehension state. This kind of knowledge has been tapped by tasks such as asking people to judge the difficulty of a problem and to allocate time and resources efficiently, which necessitates being sensitive to their capabilities and limitations. For example, experts are much more accurate than novices at judging how difficult it is to solve a given physics problem (Chi et al., 1982); good students are better at allocating their time for studying; and good students also are better at knowing when they understand the material and when they do not (Chi et al., 1989b).

Strategies of Problem Solving and Reasoning

Besides having superior knowledge structures, experts are believed to excel in their use of general strategies. For instance, solving a problem by going backward from the unknown to the givens is usually called a means-ends strategy. With a means-end strategy, a person's entire trace of problem-solving actions is guided by a strategy for solution in which the problem solver starts by using the unknowns as the goals and then searches for equations that can satisfy the unknowns. This strategy can be used in many domains, whether physics, algebra, or taxi routes. Two other additional general strategies are described below.

Experts are often said to have complex general reasoning skills, such as analogical comparison or reasoning. Scientific discoveries are often made by using analogies, but it is easy to demonstrate that reasoning by analogy is a heuristic that is available to everyone, even young children. The difference between experts and novices lies in how one elicits a particular case as an analogical one. Experts tend to be superior at picking the right case for making the comparison. This ability, again, can be attributed to the richness of their knowledge base. For example, Chi et al. (1989a) have shown that for children aged 4-7, those who have expertise in the domain of dinosaurs reason in much the same way as those who are novices. The difference is that when expert children elicit analogical cases to help them understand the features of a novel dinosaur,

the cases are drawn from other dinosaurs, in fact, dinosaurs from the same family; in contrast, novices use other animals as comparative cases from which to reason analogically. Similarly, Schauble et al. (1989) looked at how good and poor solvers predict the outcome of a simple circuit and found that the solvers' understanding of the circuit exhibited the use of analogies. Both the good and the poor solvers showed the same frequency of analogy usage. The success of the good solvers, however, seemed to depend on their ability to pick the right analogy, one that shows deep understanding of both the analogy and the target domain. Thus, it appears that the availability of general strategies per se is not the deciding factor in determining an expert's superior performance; rather, it is the application of a general strategy to a well-organized knowledge domain that is critical.

One type of a general "divide-and-conquer" strategy is simplification, or decomposition. The strategy involves breaking down a problem into its subparts, solving each subproblem, and, finally, putting all the subproblem solutions together. Jeffries et al. (1981) found that the way expert software designers solve a complex problem is by decomposing it into simpler problems for which they already know the solutions. Another variant of the divide-and-conquer strategy is "subgoaling," which is simply picking out an intermediate state on the solution path as the temporary goal. In effect, subgoaling also divides a problem into two or more subproblems, thus reducing the entire search space into a smaller one after the subgoal is reached.

One common interpretation of these results is that experts' performance is far superior to that of novices because they use these strategies. One could equally argue, however, that the use of these strategies is a manifestation of the content knowledge that the experts have, and that these general strategies are available to everyone, novices and experts alike. The strategies are not used efficiently by the novices because their inadequate knowledge base does not permit such use. Thus, adopting strategies alone is not adequate for guiding a novice's execution of a solution, even though the trace of an expert's execution of the solution reflects the use of such a strategy.

A domain-specific reasoning strategy could either be a specific procedural rule that an expert knows in a given domain (e.g., if light A flashes, push switch X) or the application of a general strategy to a specific piece of knowledge. For instance, when an expert radiologist reads an X-ray efficiently and correctly, it need not imply that he or she used a more sophisticated scanning strategy, but, rather, that the radiologist recognized one shaded area as malignant, which suggested that another area might also be malignant (Lesgold et al., 1988). That is, the scanning path was dictated by the knowledge of what related areas may

be malignant. This declarative knowledge would lead an expert to scan the two targeted area consecutively. Similarly, chess masters scan the chessboard very efficiently, in that they focus on the key important pieces on the board, such as the locations of the queen, king, rook, etc. (Simon and Barenfeld, 1969). Although much research has been devoted to showing how efficiently experts use specific strategies, the committee concludes that such use of efficient strategies is often a manifestation of an underlying coherent and rich knowledge base. This conclusion suggests that direct instruction of general strategies for novices may not always be helpful.

When Do Experts Not Excel?

In general, experts excel only when they can use their rich domain knowledge. In the classic chess experiment referred to above, expert and beginner players were asked to recall chess pieces after having looked at a chessboard position for 5 seconds. The experts exhibited spectacular recall of chess pieces when these pieces were placed on the board according to meaningful game positions. When the position of the same pieces on a board were determined randomly, however, expert's recall levels dropped to the level of novices. Thus, the experts excelled only when the chess pieces formed meaningful configurations that corresponded to patterns that they had developed in memory over the course of years (Chase and Simon, 1973).

Psychological measures of expertise—such as amount of recall, time to complete a task, or efficiency in sorting and recognizing a variety of problems and patterns—all show that experts excel largely because the stimuli used for the tasks are familiar and prototypical so that the processing is fairly automatic. That is, the processing involves a recognition or match between stored memory representations and the external stimuli, as in the case of X-ray or circuit diagrams. However, when experts are presented with nonroutine cases (e.g., there are no stored representation of the problem or its solution), they use a deliberate reasoning heuristic, such as a means-ends strategy. Similarly, their speed and efficiency also deteriorate. Norman et al. (1989) found that expert dermatologists' recognition time for typical disease cases is faster and more accurate than that of novices. However, for atypical cases, experts' reaction time was actually slower than that of novices. Obviously, for nonroutine cases, experts cannot invoke automated perceptual skills that are derived from a rich knowledge base. In these atypical cases, they must resort to the step-by-step deliberation that novices generally have to undertake, relying on their general strategies rather than domain-specific procedural rules.

ELICITING KNOWLEDGE FROM EXPERTS

Because experts' knowledge is complex, multifaceted, and often tacit, the problem of eliciting what experts know is a nontrivial problem. Different techniques yield different forms of knowledge, some declarative and some procedural. In addition, because some knowledge is more easily verbalized than others, different techniques of knowledge elicitation may bias what is being elicited.

A variety of methods can be used to elicit experts' declarative knowledge. To elicit experts' knowledge of concepts, the interrelations among the concepts, and perhaps even the coherence of the concepts, a common technique is a concept-listing task (Cooke and McDonald, 1987). In such a task, the subject is simply asked to list all the concepts related to a topic. Some assessment of the organization can be gleaned from this simple task by examining the order of the listing, as well as the content of the concepts listed. For example, prototypical items of a category tend to be listed most frequently (Rosch et al., 1976). In order to elicit subjects' organization about a specific domain of knowledge (such as the form and content of schemata in memory), ordered recall and grouping are often used. The groups that subjects either indicate explicitly (such as the perceptually related pieces on the board of a GO game; Reitman, 1976), or indicate implicitly (by order of recall), reveal the frequency of associations and interrelations among concepts. These kinds of methods can reveal the tacit conceptual structures among concepts (Chi and Koeske, 1983). Yet another common method of eliciting experts' declarative knowledge is through interviews, question and answers, and explanations (Chi et al., 1989b) or asking for the basic approach (Chi et al., 1982). In the last method, experts and novices are simply asked to state their basic approach to solving a problem; their answers indicate the kinds and content of their schemata.

All these methods mentioned do not typically reveal the kind of procedural knowledge that is attached to conceptual knowledge. One method for eliciting procedural knowledge from experts or novices is simply to ask them to perform the task, such as troubleshooting, solving problems, or making decisions, and note the sequence of actions they take. The problem with this approach is that the solver's trace (or sequence of actions) does not tell us what knowledge generated the particular trace. Some computer simulation models that are sensitive to this issue try to design a model in which the trace can be derived from the knowledge. However, the knowledge that is modeled to generate the derivation of such a trace is never unique, so that many models can be constructed that generate the same trace. This outcome suggests that observing overt

actions is too coarse a method to capture the underlying procedural knowledge. Nevertheless, overt behavior in the course of performing a task is more likely to provide knowledge of the procedures than experimental laboratory-type tasks. Hanisch et al. (1988) showed, for example, that when expert telephone system operators are asked to identify features of a telephone system, they tend to generate declarative knowledge. However, when they were asked to act in specified situations, their actions produced procedural knowledge.

Whatever the method of knowledge elicitation a great deal of effort can be expanded on the analyses of the data, especially those that are collected by interviews (structured or unstructured) and think-aloud protocols. Although it is generally true that different methods of elicitation tend to uncover different kinds of knowledge, with ingenuity one might be able to capture both declarative and procedural knowledge. Alternatively, one might be able to devise more comprehensive methods of eliciting knowledge from experts. Klein (1990) has a taxonomy of 10 types of knowledge that one could elicit from an expert: procedures; specific details; declarative knowledge; physical relations; interpersonal knowledge; perceptual-cognitive skills; perceptual-motor skills; goals; precedents (special cases or incidents); cultural knowledge. In order to elicit all these types of knowledge, he has adapted a method introduced by Flanagan (1954) that he calls the critical decision method. One of the two key features of this method is that it focuses on nonroutine cases, which is assumed to expose experts' knowledge more thoroughly than routine cases. The second key feature is that the critical decision interviews are accompanied by specific probes that should elicit knowledge corresponding more or less to each of the types listed above. So, for instance, in order to elicit the knowledge about goals, the subject may be asked directly, "What were your specific goals at this time?"

But even if all the knowledge an expert has can be thoroughly elicited, a major problem remains: how to represent that knowledge in a way that can be transmitted to a novice. For example, expert systems, which are computer programs that embody expert knowledge, typically represent the expert knowledge by rules. Such rule-based codings may not be adequate for representing cultural practice, interpersonal knowledge, or even perceptual knowledge. This is a substantial problem and the focus of much current research.

IMPARTING EXPERTS' KNOWLEDGE TO TRAINEES

Once we know how to elicit expert's knowledge, then we need to know how to teach this knowledge to novices. The goal is to instruct

skills in a way to ensure transfer and learning with understanding. This section discusses three approaches: direct instruction, computer-aided support systems, and cognitive apprenticeship.

Direct Instruction

The most traditional way of instructing skills is through a formal classroom, using textbooks and teachers' lectures. We need not recount here the failures that have been encountered by this approach. The most serious problem is that knowledge learned in a training setting or classroom does not transfer to the criterion setting—the real world. The appropriate conditions under which such a skill needs to be accessed and executed are not attached to the actions during training. For instance, the first thing fire fighters are typically taught to do is search and rescue. However, they are not usually taught the conditions under which such a routine should be disregarded, such as when a fire is getting started, and there could be a good chance of extinguishing it (Klein, 1990). Skill instruction typically focuses on a set of executable actions, and the relation between the actions and the conditions are usually not specified in the classroom. If instruction takes place in the targeted context, however, actions can be associated with the proper conditions.

One approach to textbook learning in the classroom that may meet with more success is to focus attention on the worked-out solutions or examples in the textbook. Zhu and Simon (1987) have even shown a clear advantage if students are given only examples and problems to solve, as opposed to standard instruction with a textbook and instructor's presentations: a 3-year mathematics course can be reduced to 2 years. In laboratory studies, there is evidence showing that students prefer to rely on examples. Pirolli and Anderson (1985) found that novices rely heavily on analogies to examples in the early stages of learning: 18 of their 19 students used the example presented in the text as a template for solving their first programming problem. Futhermore, Chi et al. (1989b) further showed that not only are there quantitative differences in how frequently examples are used as a function of the students' skill at solving problems, but also that there is a qualitative difference in how examples are used. Successful solvers of physics problems used the examples about 2½ times per problem solution; less successful solvers used the examples 6½ times per problem solution. Successful solvers used the example only as a reference; less successful solvers reread virtually all of the example each time it was consulted, which suggests that they were searching aimlessly for a solution. Thus, beginning students rely on examples for learning to solve problems, even though there are differences among students in the extent of their reliance.

LeFevre and Dixon (1986) found that subjects actually prefer to use the example information and ignore the written instruction when learning a procedural task. One of the reasons that examples are so popular and useful, even though they are often inadequate in explaining the rationale behind each action taken in the solution, is that the examples contain many of the procedural actions that are not explicated in texts (Chi and VanLehn, 1991): the text usually presents concepts and principles; the examples contain the actions that are executed to solve the problem.

Although, in general, classroom instruction has not been very successful, there is room for improvement. As suggested above, two promising directions are to link the actions with the appropriate conditions and to focus more attention on examples in a text. A third approach is group discussion, which has been shown recently to improve learning in the classroom (Brown and Palincsar, 1989; Lampert, 1986; Minstrell, 1989).

Computer-Aided Support Systems

Numerous efforts have been applied to the development of computer-aided tools or systems for instruction. This approach has its obvious promise in its ability to record a student's responses, to replay the student's solution trace, and to display motion and invisible mechanisms on the screen, to name a few. It is also believed that evolving interactive learning environments and intelligent tutoring systems can be adapted to instruction of individual students. Three types of computer-based instruction are practice environments, embedded expert systems, and intelligent tutoring systems.

Computers can be used to create a variety of practice environments for simulation-based training. The most common kind are direct simulations of a real-world environment, such as a flight simulator. These simulations possess physical fidelity: they can look and feel like the real situation. Such simulations provide risk free and inexpensive practice. Simulations that forgo physical fidelity but retain cognitive fidelity have also been implemented. "Steamer," for example, simulates the operation of a steam plant (Williams et al., 1981). Although it does not preserve physical fidelity, its cognitive fidelity exposes the processes of a steam plant, such as the flow of the fluids. Finally, one can implement practice environments that present abstracted alternatives, such as navigating a rocket in a different gravitational field (Abelson and diSessa, 1980).

Practice environments can be made more sophisticated by incorporating an expert system. Such a system not only provides simulation of an environment, but can also solve problems within that environment. This kind of system can have different types of intelligence. For example, a

system such as SHERLOCK (Lajoie and Lesgold, 1989) can solve a specified set of problems, and a system such as the Intelligent Maintenance Training System (Towne and Monroe, 1988) can troubleshoot any problem that the student poses to it. Thus, the effect of having a built-in expert system is that it can provide feedback (the solution) to the student. A drawback, however, is that the feedback is not sensitive to the student's understanding.

A full-fledged intelligent tutoring system has four components: a practice environment for solving problems, an expert system that knows how to solve all the problems, a model of the student (in terms of understanding how the student solves a problem, rather than whether the student's solution approximates the expert's solution), and a system for pedagogy. At the present time, there is no system that completely incorporates all these components. The best approximations are Burton and Brown's (1982) "How the West Was Won" system, and Clancey's (1983) Guidon system. Burton and Brown's system is designed to teach children basic arithmetic operations in the context of a game that is a variant of Chutes and Ladders. The system observes students as they play the game and gives them hints or advice at critical moments. Since it is a very simple game, the essence of the system is based on an analysis of the task, which divides it into subskills. The subskills are arranged in an hierarchy and each player's move is analyzed to determine the subskill to which it corresponds. Erroneous moves are compared with the subskills to see which one has been overlooked. Thus, the computer coach delivers feedback as a function of the match or mismatch between the student's moves and the subskill hierarchy. The coach also has heuristics for how to deliver instruction. Such pedagogical heuristics include rules of thumb, such as only give unsolicited advice when the player has made three or four moves without asking for advice. Hence, the system contains a practice environment; an expert system, in that the system knows how to play the game; a simple student model; and a pedagogy system, in its ability to monitor students' "understanding" of the game and deliver feedback selectively on the basis of the student's moves.

Cognitive Apprenticeship Revisited

In the first section of this chapter we noted that apprenticeship training is a popular method of instruction that embodies cognitive modeling. It relies on the teacher or the expert exposing the cognitive processes he or she is undertaking while solving a problem, making a decision, or diagnosing a fault. As we noted, certain strategies of learning can be easily modeled, and thereby perhaps learned by students with relative ease. But the rationale for an expert's choices of strategies or solution

routes may not be obvious, especially if a great deal of content knowledge is needed. Furthermore, monitoring a student's own understanding demands complex processing, and it is not clear that this is undertaken normally by tutors. Indeed, there is some evidence that teachers in a tutoring situation do not necessarily provide "intelligent" coaching and scaffolding, presumably because they have not monitored the student's understanding accurately. In fact, Putnam (1987) observed six teachers tutoring students in a one-on-one situation and found that they did not systematically diagnose students' misunderstanding and tailor feedback accordingly. Instead, they seemed to follow a structured and prescribed sequence of presenting subject-matter content to their students.

Despite the lack of complete understanding of the processes undertaken by the tutor or expert, it appears that modeling in cognitive apprenticeship has had some success in teaching reading and writing, perhaps because it encourages students to engage in guided participation, actively articulating, reflecting, and exploring the skill that they are learning. There is evidence to suggest that these activities may be the critical source of learning in the apprenticeship situation. These activities can be viewed as self-construction, including self-explaining, asking questions, summarizing, providing critiques, posing problems, formulating hypotheses, and revising and editing. We noted above that there is evidence that these activities may facilitate learning and problem solving. There are also research findings that can be interpreted as providing indirect evidence that these activities are beneficial. For example, it is often noted that students improve their problem-solving performance if they participate in pairs. The most profitable dyadic situations are ones in which one member is more expert than the other (Radziszewska and Rogoff, 1988). It appears that the less expert partner offers self-explanations in these dyadic situations more than in the control situations in which each partner solves problems alone (Azmitia, 1988). It is tempting to speculate that what promotes the improvement of the less expert partner in the dyadic situations is the opportunity to construct new knowledge through articulation. Thus, to the extent that modeling promotes self-construction, it can be an effective instructional tool for cognitive skills.

SUMMARY AND CONCLUSIONS

An expert's knowledge is not only extensive, but much of it may be tacit, so that it takes a considerable amount of effort to elicit it. Experts may not always be aware of what their knowledge entails nor why they proceed in a specific way. The tacitness of expert knowledge may pose a problem for the utility of cognitive modeling for complex tasks requiring

a great deal of domain knowledge. However, it has been shown to be quite effective in enhancing the learning of basic skills, such as reading (Palincsar and Brown, 1984) and writing (Scardamalia et al., 1984). In the context of a complex task, modeling in the sense of simply exposing an expert's own cognitive processes may present two additional complexities from the standpoint of instruction: to what extent does the student understand the expert's processes and to what extent should the expert attempt to understand the student's cognition in order to resolve the discrepancies between the two sets of processes. For example, it is often not clear to a student why his or her own approach to solving a problem is faulty or inappropriate. It appears that in order for cognitive modeling to work in a knowledge-rich complex domain, the expert has to be sensitive not only to his or her own cognitive processes, but also to how the student thinks and approaches a problem and to be able to explain in those terms why the student's approach may not work. To do so requires that an expert simultaneously is aware of and can articulate his or her own problem-solving processes and actively monitors his or her model of the student. Without further research on the cognitive processes of a tutor, it is not clear how successfully a tutor can achieve this dual role in a complex domain.

Although the idea of cognitive apprenticeship is promising and enticing, we cannot determine at this time whether the expert-modeling component is the critical one, since cognitive apprenticeship involves many other components, such as the requirement that instruction be situated in context and that students participate actively. Cognitive modeling also involves extensive explanations, and it seems safe to conclude that the nature of explanations is critical in assuring how well students can learn from an expert. Moreover, the self-constructions undertaken by the students while engaged in cognitive apprenticeship may also be an effective factor. Nevertheless, the success of cognitive apprenticeship for reading comprehension and writing suggests that cognitive modeling may be a feasible approach to explore for training basic skills.

REFERENCES

Abelson, H., and A.A. diSessa 1980 *Turtle Geometry: The Computer as a Medium for Exploring Mathematics.* Cambridge, Mass.: MIT Press.

Anderson, J.R. 1974 Retrieval of propositional information from long-term memory. *Cognitive Psychology* 6:451-474.

Azmitia, M. 1988 Peer interaction and problem solving: when are two heads better than one? *Child Development* 59:87-96.

Brown, A.L., and A.S. Palincsar 1989 Guided, cooperative learning and individual knowledge acquisition. Pp. 393-452 in L.B. Resnick, ed., *Knowing, Learning, and Instruction: Essays in Honor of Robert Glaser.* Hillsdale, N.J.: Erlbaum.

Burton, R.R., and J.S. Brown 1982 An investigation of computer coaching for informal learning activities. Pp. 79-98 in D. Sleeman and J.S. Brown, eds., *Intelligent Tutoring Systems.* New York: Academic Press.

Charness, N. 1979 Components of skill in bridge. *Canadian Journal of Psychology* 33:1-16.

Chase, W.G. 1983 Spatial representation of taxi drivers. Pp. 391-411 in D.R. Rogers and J.A. Sloboda, eds., *The Acquisition of Symbolic Skills.* New York: Plenum Press.

Chase, W.G., and K.A. Ericsson 1981 Skilled memory. In J.R. Anderson, ed. *Cognitive Skills and Their Acquisition.* Hillsdale, N.J.: Erlbaum.

Chase, W.G., and H.A. Simon 1973 Perception in chess. *Cognitive Psychology* 5:55-81.

Chi, M.T.H. 1978 Knowledge structures and memory development. In R.S. Siegler, ed., *Children's Thinking: What Develops?* Hillsdale, N.J.: Erlbaum.

Chi, M.T.H., and R. Koeske 1983 Network representation of a child's dinosaur knowledge. *Developmental Psychology* 19:29-39.

Chi, M.T.H., and K.A. VanLehn 1991 The content of self-explanations. *Journal of the Learning Sciences.* In press.

Chi, M.T.H., R. Glaser, and E. Rees 1982 Expertise in problem solving. Pp. 7-76 in R. Sternberg, ed., *Advances in the Psychology of Human Intelligence.* Vol. 1. Hillsdale, N.J.: Erlbaum.

Chi, M.T.H., R. Glaser, and M.J. Farr, eds. 1988 *The Nature of Expertise.* Hillsdale, N.J.: Erlbaum.

Chi, M.T.H., J. Hutchinson, and A.F. Robin 1989a How inferences about novel domain-related concepts can be constrained by structured knowledge. *Merrill-Palmer Quarterly* 25:27-62.

Chi, M.T.H., M. Bassok, M. Lewis, Reimann, and R. Glaser 1989b Self-explanations: how students study and use examples in learning to solve problems. *Cognitive Science* 13:145-182.

Chiesi, H.L., G.J. Spilich, and J.F. Voss 1979 Acquisition of domain-related information in relation to high and low domain knowledge. *Journal of Verbal Learning and Verbal Behavior* 18:257-273.

Clancey, W. 1983 Guidon. *Journal of Computer-Based Instruction* 10:8-14.

Collins, A., and A.L. Stevens 1983 A cognitive theory of interactive teaching. In C.M. Reigeluth, ed., *Instructional Design Theories and Models: An Overview.* Hillsdale, N.J.: Erlbaum.

Collins, A., Brown, J.S.; and S.E Newman 1989 Cognitive apprenticeship: Teaching the crafts of reading, writing, and mathematics. L. Resnick, ed., *Knowing, Learning, and Instruction: Essays in Honor of Robert Glaser.* Hillsdale, N.J.: Erlbaum.

Cooke, N.M., and J.E. McDonald 1987 The application of psychological scaling techniques to knowledge elicitation for

knowledge-based systems. *International Journal of Man–Machine Studies* 26:533-550.

Egan, D.E., and B.J. Schwartz 1979 Chunking in the recall of symbolic drawings. *Memory and Cognition* 7:149-158.

Flanagan, J.C. 1954 The critical incident technique. *Psychological Bulletin* 51:327-358.

Hanisch, K.A., A.F. Kramer, C.L. Hulin, and R. Schumacher 1988 Novice–expert differences in the cognitive representation of computing systems: mental models and verbalizable knowledge. Pp. 219-223 in *Proceedings of the Human Factors Society 3rd Annual Meeting.* Santa Monica, Calif.: Human Factors Society.

Hoffman, R.R., ed. 1991 *The Psychology of Expertise: Cognitive Research and Empirical Al.* New York: Springer-Verlag.

Jeffries, R., A. Turner, P. Polson, and M. Atwood 1981 The processes involved in designing software. Pp. 255-283 in R.J. Anderson, ed., *Cognitive Skills and Their Acquisition.* Hillsdale, N.J.: Erlbaum.

Johnson-Laird, P.N. 1983 *Mental Models.* Cambridge, England: Cambridge University Press.

Klein, G.A. 1990 Knowledge engineering: beyond expert systems. *Information and Decision Technologies* 16:27-41.

Lajoie, S.P., and A. Lesgold 1989 Apprenticeship training in the workplace: computer-coached practice environment as a new form of apprenticeship. *Machine-Mediated Learning* 3 (1):7-28.

Lampert, M. 1986 Knowing, doing, and teaching multiplication. *Cognition and Instruction* 3:305-342 .

LeFevre, J., and P. Dixon 1986 Do written instructions need examples? *Cognition and Instruction* 3:1-30.

Lesgold, A.M., H. Robinson, P. Feltovich, R. Glaser, D. Klopfer, and Y. Wang 1988 Expertise in a complex skill: diagnosing X-ray pictures. Pp. 3121-3142 in M.T.H. Chi, R. Glaser, and M.J. Farr, eds., *The Nature of Expertise.* Hillsdale, N.J.: Erlbaum.

Minstrell, J. 1989 Teaching science for understanding. In L.B. Resnick and L.E. Klopfer, eds., *Toward the Thinking Curriculum: ASCD Yearbook.* Alexandria, Va./Hillsdale, N.J.: Association for Supervision and Curriculum Development/ Erlbaum.

Newell, A., and H.A. Simon 1972 *Human Problem Solving.* Englewood Cliffs, N.J.: Prentice-Hall.

Norman, G.R., D. Rosenthal, L.R. Brooks, S.W. Allen, and L.J. Muzzin 1989 The development of expertise in dermatology. *Archives of Dermatology* 125:1063-1068.

Palincsar, A.S., and A.L. Brown 1984 Reciprocal teaching of comprehension-fostering and monitoring activities. *Cognition and Instruction* 1:117-175.

Pirolli, P.L., and J.R. Anderson 1985 The role of learning from examples in the acquisition of recursive programming skills. *Canadian Journal of Psychology* 39:240-272.

Putnam, R. 1987 Structuring and adjusting content for students: a study of live and simulated tutoring of addition. *American Educational Research Journal* 24:13-48.

Radziszewska, B., and B. Rogoff 1988 Influence of adult and peer collaborators on children's planning skills. *Developmental Psychology* 24:840-848.

Reitman, J. 1976 Skilled perception in GO: deducing memory structures from inter-response times. *Cognitive Psychology* 8:336-356.

Rogoff, B. 1986 Adult assistance of children's learning. In T.E. Raphael, ed. *The Contexts of School-Based Literacy.* New York: Random House.

Rosch, E., C.B. Mervis, W.D. Gray, D.M. Johnsen, and P. Boyes-Graem 1976 Basic objects in natural categories. *Cognitive Psychology* 8:382-440.

Scardamalia, M., C. Bereiter, and R. Steinbach 1984 Teachability of reflective processes in written composition. *Cognitive Science* 8:173-190.

Schauble, L., K. Raghavan, R. Glaser, and M. Reiner 1989 Causal Models and Processes of Discovery. Technical Report. Learning Research and Development Center, University of Pittsburgh.

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.

Schuell, T.J. 1991 Designing instructional computing systems for meaningful learning. In P.H. Winne and M. Jones, eds., *Foundations and Frontiers in Instructional Computing Systems.* New York: Springer-Verlag.

Sharpies, M. 1980 A Computer Written Language Lab. DAI Working Paper No. 134. Artificial Intelligence Department, University of Edinburgh, Scotland.

Simon, H.A., and M. Barenfeld 1969 Information-processing analysis of perceptual processes in problem solving. *Psychological Review* 76:473-483.

Towne, D.M., and A. Monroe 1988 The intelligent maintenance training system. Pp. 479-530 in J. Psotka, L.D. Massey, and S.A. Mutter, eds.,*Intelligent Tutoring Systems: Lessons Learned.* Hillsdale, N.J.: Erlbaum.

Voss, J.F., J.R. Greene, T.A. Post, and B.C. Penner 1983 Problem solving skill in the social sciences. Pp. 165-213 in G.H. Bower, ed., *The Psychology of Learning and Motivation: Advances in Research and Theory,* Vol. 17. New York: Academic Press.

Vygotsky, L.S. 1978 *Mind in Society: The Development of Higher Psychological Processes.* Collected writings by Vygotsky. Edited and transcribed by M. Cole, V. John-Steiner, S. Scribner, and E. Souberman. Cambridge, Mass.: Harvard University Press.

Williams, M., J. Hollan, and A. Stevens 1981 An overview of STEAMER: an advanced computer assisted instructional system for propulsion engineering. *Behavior Research Methods and Instrumentation* 13:85-90.

Wood, D., and D. Middleton 1975 A study of assisted problem-solving. *British Journal of Psychology* 66:181-191.

Zhu, X., and H.A. Simon 1987 Learning mathematics from examples and by doing. *Cognition and Instruction* 4:137-166.