Klein, G. A., & Hoffman, R. R. (1992). Seeing the invisible: Perceptual-cognitive aspects of expertise. In M. Rabinowitz (Ed.), *Cognitive science foundations of instruction* (pp. 203-226). Mahwah, NJ: Erlbaum.

9 Seeing the Invisible: Perceptual–Cognitive Aspects of Expertise

Gary A. Klein
Klein Associates Inc.

Robert R. Hoffman
Adelphi University

"SEEING" AND "SEEING BEYOND"

Novices see only what is there; experts can see what is not there. With experience, a person gains the ability to visualize how a situation developed and to imagine how it is going to turn out. In this chapter we examine how experts can use their knowledge to visualize tasks. Our emphasis is not on rules, or strategies, or size of knowledge base per se, but on the perceptual and cognitive qualities of experience—experts do not seem to perceive the same world that other people do.

Novices see only what is there. As long as people have a general sense of what is going on, they are alert to what is happening in front of them. Chi, Feltovich, and Glaser (1981) found that both expert and novice physics students could identify the critical cues in a physics problem; the advantage that experts had was in perceiving the interactions among the cues. Brezovic, Klein, and Thordsen (1987) studied the decision making of tank platoon leaders and found that during field exercises novices who had just been introduced into the tanks could list as many different cues as the tank instructors. The novices were not overwhelmed with information, and they were reasonably aware of important items of information. Cue acquisition did not appear to be a critical component of expertise. Elstein, Shulman, and Sprafka (1978) found that accuracy of medical diagnoses was minimally related to thoroughness of cue acquisition. For most tasks, higher levels of performance do not necessarily depend on more powerful strategies for acquiring information that is directly perceivable.

Experts can see what is not there. Their experience lets them notice when something is missing. Consider the following example. A new employee in an organization made the observation that projects were usually left for the last moment. She drew this conclusion from watching how several technical reports were completed with considerable strain just before their deadlines. When she shared her observation, however, it was hotly disputed. In fact, the organization had just completed 10 proposals within the past 6 weeks, with such little effort that she had not even known about them. In previous years, proposal writing had been a major burden and so extra care had been taken to prepare in advance. People who had been in the organization for several years could see the difference. The new employee had no way to detect the *absence* of frantic activities. She could notice the times when the system broke down, but not the times when it worked.

What nonobservable events can an expert detect? Only with experience can you visualize how a course of events is likely to unfold, so that you can see the expected outcomes even in the beginning. Only with experience can you form expectancies. Only with experience can you notice when the expectancies are violated, when something that was supposed to happen did not. And only with experience can you acquire the perceptual skills to make fine discriminations.

In the sections of this chapter, we first describe what we mean by "expert performance." Next, we consider traditional theories of expertise that focus on higher order strategies and on the development of a larger knowledge base. Following that is a discussion of perceptual–cognitive aspects of expertise, which leads to a concluding discussion of training issues.

THE DEVELOPMENT OF EXPERTISE

In this chapter, we are interested in perceptual–cognitive phenomena. For example, if a fireground commander is able to judge that a fire is hotter than would be expected, implying the presence of potentially toxic chemicals, we cannot clearly articulate what is being noticed, but we can assert that the color of the fire and the apparent pressure driving the smoke are the perceptual dimensions being used to make the judgment. We want to examine the ability of experts to perceive things that novices and journeymen cannot detect, in order to determine how these "perceptual" abilities can be used to assess proficiency level and to drive training requirements.

We can distinguish among experts, journeymen, and novices. In this chapter, our focus is on the achievement of high levels of proficiency. Swets and Bjork (1990) have asserted that the training of expertise represents a national need. Because of factors such as turnover and technology change, it is essential that we learn how to "bring people up to speed" as rapidly as possible—not just to the level of adequate performance, but to high levels of proficiency. Accordingly, our prime focus is on the difference between competent performers (journeymen) and experts. We are less interested in the transition from novice to journeyman, because the majority of existing training programs are designed to help people achieve mediocre levels of performance. There has been little work in moving people from competent levels of performance to expert performance.

Before going on, we need to be a bit clearer about what we mean by expertise. The achievement of expertise requires a large amount of experience, but simple accumulation of practice is not sufficient. If you endlessly repeat the same exercises, you will not develop very far. In research on firefighters (Klein, Calderwood, & Clinton-Cirocco, 1986), we observed that 10 years with a rural volunteer fire department were not as valuable for skill development as 1 year in a decaying inner city. Although some minimum amount of time is necessary it must be accompanied by a chance to accumulate a varied set of experiences.

When a person attains a high level of proficiency, we expect to see certain characteristics of performance. We expect the person to be able to make judgments and discriminations that are difficult for most other people—the expert's judgments are significantly more accurate and reliable than those of a journeyman. The expert must be able to apply the experience base to a wide range of tasks encountered in the domain, including nonroutine cases that would stymie people who were merely competent. The best experts will set the standards of ideal performance for a domain.

Glaser (1976) has suggested that when people develop expertise the following changes occur:

- Variable awkward performance becomes consistent, accurate, complete, and relatively fast.
- Individual acts and judgments are integrated into overall strategies.
- Perceptual learning occurs so that a focus on isolated variables shifts to perception of complex patterns.
- There is increased self-reliance and ability to form new strategies when required.

It is intriguing to speculate about the way a person develops expertise. Although there are several accounts of this phenomenon (e.g., Glaser, 1976), currently the most widely cited account was provided by Dreyfus and Dreyfus (Dreyfus, 1972; 1986).

Dreyfus and Dreyfus have argued that the achievement of expertise follows a reliable progression of five levels of expertise, ranging from the beginner to the master. Perceptual skills, rather than analytical skills, are assigned a central role in the progression. Table 9.1 presents a description of these five levels of expertise, adapted from Benner (1984).

The five levels of expertise described in Table 9.1 are sometimes hard to apply, for a number of reasons. People rarely perform at the same level on all tasks in a domain. Presumably, someone who is proficient may behave at the level of expert for a few tasks, would perform at the level of proficiency for many

TABLE 9.1
Levels of Expertise

Novice
Beginners have had little experience of the situation in which they are expected to perform. Their initial learning about the situation is in terms of objective attributes—those that are measurable. These are features of the task world that can be recognized without situational experiences. Novices are limited in their understanding to context-free rules that guide action—this means their behavior is limited and inflexible.
Advanced Beginner
They have coped with enough real situations to note (or have pointed out to them) recurring, meaningful situational components. At this level, understanding of aspects of the situation is limited to global characteristics that reflect prior experience in actual situations. Advanced beginners need help setting priorities, because they operate on general guidelines and are only beginning to perceive recurrent, meaningful patterns.
Competent
Performers at a journeyman's level can see their actions in terms of long-range goals or plans. They are consciously aware of formulating, evaluating, and modifying goals–plans. The competent performer is able to generate plans in terms of current and contemplated future aspects that are most important, and those that are not. The competent performer lacks the speed and flexibility that emerges at higher levels of expertise but has a sense of mastery and the ability to cope with and manage a variety of types of situations.
Proficient
Proficient performers perceive situations as wholes, rather than in terms of situational components. Their performance is guided by "maxims." Perception is key. The perspective is NOT thought out but "presents itself" based upon experience. The proficient performer has learned what typical events to expect in a given situation and how plans need to be modified in accord with these events. This also means that she or he can recognize when the expected typical picture does not materialize and can modify plans and goals accordingly. Situational aspects stand out as more or less important in *this* situation.
Expert
Expert performers no longer rely on analytic principles (rules, guidelines, maxims) to connect their understanding of the situation to an appropriate action. The expert, with an enormous background of experience, has an intuitive grasp of each situation and zeros in on the accurate region of the problem without wasteful consideration of a large range of unfruitful, alternative diagnoses and solutions. The performer is no longer aware of features and rules, and his or her performance becomes fluid and flexible and highly proficient.

tasks, and would perform at the level of competence for some tasks that were relatively unfamiliar. So it might be a mistake to label a person as *expert* and expect consistent performance at that level. Rather, an expert simply is capable of handling a wider range of tasks nonanalytically, compared to people with less experience. Usually an expert will be proficient at all or most tasks in a domain, but we would not expect equivalent levels of mastery for all subtasks.

A different approach is to model skill development on the cognitive maturation concepts of Piaget, because these are designed to accept variability and fluctuations in performance. The research of Campbell, Brown, and DiBello (1991) on the development of expertise in computer programmers relied heavily on structured interviews with experts and trainees. The subjects were asked "meta-level" questions such as: How do you plan projects? How do you recognize problems? Can you compare how you do things now with how you did things when you were a beginner? Can you tell from a program how expert the programmer was? How did your knowledge of languages help when you learned a new one? The researchers also conducted a longitudinal study, using audiotaped diaries of programmers who were learning Smalltalk, an object-oriented language. The research showed the limitations of trying to distinguish experts from novices without considering intermediate stages of skill development.

Campbell et al. identified a number of developmental milestones in the learning of different program languages. With regard to Smalltalk, the researchers were able to specify seven distinct developmental levels. At each level, performance was distinguishable in operational terms (e.g., specific tasks the person could or could not perform well). This sets the developmental approach of Campbell et al. apart from others (e.g., Dreyfus & Dreyfus, 1986), which distinguish stages solely at a conceptual level, unanchored to empirical markers. The Campbell et al. approach may be a useful strategy for identifying milestones in skill development, and for allowing more useful assessment procedures.

These characteristics help us to tell that a person has acquired expertise, but how are changes brought about? That is the topic of the next three sections.

EXPERTISE AS THE DEVELOPMENT OF HIGHER ORDER STRATEGIES

Many researchers have hoped to find process differences between experts and novices; that is, we would like to show that people jump to new levels of ability because they use different processes, different strategies in solving problems. It would be nice to demonstrate that novices use simple rules, journeymen use more complex rules, and experts use even higher order rules. If we could show that experts use better strategies, then it would be possible to develop tests for such strategies. It might also enable us to teach these strategies. Even better, it might be possible to show that some strategies generalize across domains, in which case we could design generic evaluation and training methods. Because of the high potential payoff for finding strategy differences, a great deal of research has been devoted to this attempt. Much of this work has been done with children, to explain why children are not as competent as adults. Researchers attempted to show that adults were able to call on effective problem-solving strategies that children lacked (e.g., Flavell, 1971; Kail & Hagen, 1977; Ornstein, 1978).

It seemed reasonable to search for strategy differences because the information-processing tradition in cognitive psychology is a framework for understanding thinking in terms of acquiring, and performing operations on, information. It is obvious that experts acquire a more extensive knowledge base, but it also seemed possible that experts have learned different and more powerful operators, better ways of manipulating information.

Chi et al. (1981) studied expert and novice physicists and found that the knowledge bases, rather than the reasoning strategy, accounted for performance differences. Both experts (graduate students—we would call them journeymen) and novices (postintroductory level physics undergraduates) utilized top-down and bottom-up processing. They differed primarily in their "knowledge schemata." Experts possessed a large number of schemata that enabled them to categorize problems according to underlying concepts and laws and then apply well-known basic approaches for solving problems of a given type. In other words, the knowledge base determined the reasoning strategy. There is no evidence that the experts had learned higher order strategies.

We agree with this conclusion that experts do not necessarily use different strategies than novices. We cannot envision training people to become experts by showing them the importance of top-down processing and analysis of deep structure. Hoffman, Burton, Shanteau, and Shadbolt (1991) have also concluded that both experts and novices rely to some extent on top-down and bottom-up reasoning. Both utilize a divide and conquer strategy, and a cycle of forming and testing mental models. Both experts and novices rely on analogies and metaphors. All the various general strategies appear to some extent in almost all forms of reasoning.

EXPERTISE AS A FUNCTION OF KNOWLEDGE BASE

This traditional approach asserts that expertise is a function of the knowledge base itself, and that as people develop richer knowledge bases they are able to represent problems in more powerful ways, and to take more advantage of stronger reasoning strategies. For example, Chi et al. (1981) and Larkin, McDermott, Simon, and Simon (1980) have shown that expert physicists represent problems in a different way from novices. The experts' representations are conceptually richer and more organized than those of the novices. Novices use hastily formed, concrete, superficial problem representations, whereas experts use abstract representations that rely on "deep" knowledge—imaginal and conceptual understanding of functional relations and physical principles that relate concepts (e.g., conservation of energy, laws of mechanics, etc.). Furthermore, experts are better able to gauge the importance of different kinds of information, and the difficulty of problems. They are also more likely to know the conditions that require the use of particular knowledge and procedures (e.g., if there is acceleration, use Newton's second law; Chi, Glaser, & Rees, 1982). Gentner (1988) has also found that experts encode problems using deep structure, whereas novices use surface features.

It is not terribly clear how to operationalize top-down and bottom-up reasoning. What, specifically, counts as an instance of each? Beyond that, we cannot look solely at the nature of the processing (top-down or bottom-up) to directly infer level of expertise. We must consider the knowledge base that is available to the novice, journeyman, and expert. It is often the knowledge base that drives (or limits) the type of processing that is available for use, rather than an individual's choice of processing strategy. Novices may rely more on bottom-up strategies because they lack the knowledge to make a top-down approach work, not because they have not learned the more effective top-down strategy!

The accumulation of information focuses on the knowledge base itself. We have no quarrel with this account—our intent is to try to supplement it by examining the effect of an expanded knowledge base on the perceptual–cognitive aspects of experience. We suggest that the knowledge base and accumulated experiences may change the way people view their worlds.

EXPERTISE AS PERCEPTUAL–COGNITIVE DIFFERENCES

Experts can see things other people cannot. In this section we examine three aspects of expert perception: the ability to see typicality, the ability to see distinctions, and the ability to see antecedents and consequents. We are not asserting that experts use different strategies than journeymen. Clearly, journeymen attempt to judge typicality, to make fine distinctions, and to be sensitive to antecedents and consequents. The problem is that without an adequate knowledge base it is difficult to rely on these sources of power. Gaps in the knowledge base can result in a misleading idea of what is typical, or what is distinct, or what is likely to happen. Sometimes, journeymen will find it safe to use more analytical strategies rather than trust their experience base. It is not that they lack a strategy for using perceptual–cognitive skills. Rather, they have the metacognitive ability to sense when their experience base is not reliable. In most cases, we feel novices, journeymen, and experts each attempt to use perceptual–cognitive skills. They differ in their success, as a function of their experience base.

The Ability to See Typicality

A common example of the ability to judge typicality is the Secretary problem. A new office manager, faced with the need to hire a secretary, may have to conduct many interviews in order to learn the shape of the skill distribution curve, and to tell which job applicants are superior and which are inferior. In contrast, an experienced office manager will already have a sense of the typicality and variability of secretarial skills so that, if an outstanding applicant shows up as the first one interviewed, a job offer can be made immediately without having to conduct additional interviews.

There is no way for a novice to judge what is normal and what is an exception. Consider a study by Chi, Hutchinson, and Robin (1988), in which descriptions of novel dinosaurs were presented to children who were "dinosaur expert" or novices. These novel dinosaurs were designed to be either typical or atypical. The experts, of course, realized immediately that a dinosaur was typical of a class of dinosaurs and were then able to attribute all the relevant features from the family to which the novel "typical" dinosaur belonged. The experts were equally proficient at determining that a novel dinosaur that was not typical did not belong to any of the familiar families. Novices lacked this ability to judge typicality and to use it to infer other characteristics. A journeyman is likely to have a general sense of typicality, but we feel that there is a clear difference between the journeyman and the expert in ability to rapidly "size up" a situation.

Klein et al. (1986) found that expert firefighters could rapidly size up a situation by identifying it as typical, and that this judgment of typicality evoked several important types of knowledge: recognizing plausible goals, recognizing relevant cues, recognizing events that are expected, and recognizing feasible courses of action. These observations were supported by subsequent studies with tank platoon instructors and design engineers (see Klein, 1989, for a summary description of these projects).

Here is another example of how a proficient firefighter used his ability to size up a situation rapidly. He arrived at the scene of a fire in an apartment building and found that the nearby hydrants were not functional, and that the tanker trucks contained the only water available to him. He did not know the exact source of the fire, but he judged that if he did not act immediately the apartment building would be beyond saving. So he ordered his crews to direct all the water on the basement site that appeared to be the most likely source of the fire. He was successful in extinguishing the fire; if he had been wrong, he would have had to watch helplessly while the building burned down. His rapid situation assessment provided him with a goal (hit the fire immediately), a sense of critical information (limited resources and a likely candidate as the site of the fire), expectancies (waiting would just make the fire uncontrollable), and a reasonable course of action (direct the water at the basement where the smoke was thickest and darkest).

The ability to judge typicality enables experts to perform more effectively and efficiently than journeymen. By quickly seeing which goals are feasible, experts can direct their actions and not waste any effort. By recognizing which cues are relevant, experts can avoid information overload. By anticipating what events to expect, experts can rapidly notice if they have misperceived the situation. And by recognizing a typical course of action, experts can respond rapidly. This type of recognitional decision making enables experts to handle complex cases under time-compressed conditions where analytical methods would not be possible.

The ability of experts to size up situations is also seen in the classic research program conducted by de Groot (1946/1965), who showed chess players difficult chess problems. Grandmasters were able to recognize the best move, often as the first one they considered. Moderately skilled players rarely even considered the best move. De Groot carefully reviewed this think-aloud protocol and concluded that grandmasters were able to recognize threats as deviations from a norm. Such recognition was usually part of the initial evaluation, the first perception of a problem. In contrast, players at lower skill levels had to carefully look for the threats. De Groot identified "anticipation of urgency," "anticipation of solvability," and "intuitive faith in a possibility" as aspects of how a grandmaster could size up situations quickly and accurately, whereas players at lower skill levels had to search to analyze the dynamics of a position.

Moreover, according to de Groot, the sense of typicality carries with it a basis for forming expectations—for judging the level of outcome that is possible in a situation. In this way, a strong chess player can consider a move and judge it as favorable or disappointing in terms of what it yields. In contract, a journeyman cannot sense the level of outcome that should be possible and is forced to consider many moves in order to form a yardstick for evaluating any individual move.

The ability to judge typicality should enable people to detect when expectancies are violated (e.g., when expected events did not occur). Recently, Fraser, Smith, and Smith (1990) reviewed the literature on heuristics and decision errors. One of the topics they addressed was the negativity bias—the difficulty people have in using missing information, or events that were supposed to occur but did not. They found some evidence for such a tendency, in studies of inexperienced subjects (although there were still no compelling reasons interpreting them as errors). But a study by Christensen-Szalanski and Bushyhead (1981) found that experienced decision makers—physicians—used the absence of symptoms as efficiently as their presence. This supports our thesis that expertise permits the effective use of invisible cues.

If experts are able to size up situations quickly, we would expect to see evidence for such ability under time pressure. Calderwood, Klein, and Crandall (1988) compared chess masters with Class B players (journeymen). Games were played under regulation time (approximately 150 seconds per move) and blitz conditions (approximately 6 seconds per move). For complex moves, the masters showed little performance degradation with time pressure, whereas the move quality decreased sharply for the Class B players. These data suggest that the journeymen could not rapidly size up a situation the way experts could, and that this ability of experts made them less vulnerable to time compression. Is the ability to play blitz chess well a unique skill, unrelated to general chess ability? No, it seems that blitz chess calls on many of the same skills as regulation chess. Charness (1981) has reported data showing that performance under regulation time is highly correlated ($r = .85$) with performance under blitz conditions, further suggesting that expertise depends on rapid perception–action links.

The use of typicality judgments is important for problem solving as well as decision making. Elstein et al. (1978) studied the way physicians made diagnoses and found that the physicians rarely used a purely inductive method of letting the data drive the inferences. Even though the physicians were trained to reserve judgment rather than contaminate the process, they still could not resist forming early impressions. Elstein et al. referred to this strategy as a hypothetico-deductive method, because the early judgments were hypotheses that helped to direct the subsequent information gathering. Without early hypotheses, the information gathering would have been inefficient and interminable. Weitzenfeld, Klein, Riedl, Freeman, and Musa (1991) observed the same phenomenon in a study of expertise in software troubleshooting. Experts were able to formulate initial hypotheses, or stories of how the problem might have arisen, and could use these hypotheses to direct the search for more evidence. Where do these hypotheses come from? Presumably, they involve the same mechanism of situation assessment as was discussed earlier—using experience to judge typicality.

How does this ability to judge typicality develop? In time, the knowledge base develops, and a person may be able to recognize situations that are similar to previous experiences. It is tempting to infer that experts retrieve specific instances and use these to guide their behavior. Currently, the field of analogical problem solving is receiving a great deal of attention, and it is likely that this research will generate useful findings. However, Klein et al. (1986) tried to identify specific analogues that helped firefighters make difficult decisions. We obtained little evidence for the use of analogical reasoning in this domain. Doubtless there are task conditions in which analogical retrieval plays an important role—low time pressure, and relatively low experience levels—which make individual incidents more informative (see Klein & Calderwood, 1988). And when our experts (average amount of experience was 23 years) did use analogues, the analogues were judged to be very helpful. Nevertheless, we suspect that analogues are more prevalent for journeymen than for novices. When people are able to gain a great deal of experience, particularly direct experience, we hypothesize that individual incidents become less vivid and blend together in memory, enabling the expert to make additional types of judgments—detecting typicality and variability.

Another possible explanation for the ability to judge typicality, besides analogical reasoning, is pattern matching. The hypothesis is that experts are better at pattern matching. Chase and Simon (1973) demonstrated that expert chess players used content knowledge to recognize patterns, rather than their possessing superior factual memory. And Simon and Gilmartin (1973) estimated that chess masters store a large repertoire of between 10,000 to 100,000 patterns.

But the story is not so simple. If expertise were just a matter of pattern matching, why would experts also be better at handling nonroutine events and complex variations? De Groot (1990, personal communication), Holding (1985), and Elstein et al. (1978) have all been very critical of the idea that experts simply possess a large repertoire of specific patterns, enabling them to rapidly match events to prior experiences. We entirely agree with their criticisms. Experts must be recognizing complex invariants nested within complex and often irrelevant data. Pattern matching is perhaps a component of judging typicality (depending on the task), but it is not sufficient—something must provide a basis for inferences. In simple situations pattern matching might be enough, but for complex cases it is difficult to see how pattern matching would allow the expert to derive a sense of the situation ("I should be able to get out of this with only the loss of a pawn") and of promising courses of action ("How can I take advantage of my well-placed knight?").

At present, there is no satisfactory explanation of how experts can quickly size up situations and judge typicality. This ability sets experts apart from novices and from journeymen and represents one way that experts can use their experience to see subtle and critical aspects of a situation.

The Ability to See Distinctions

Experts are particularly better than novices and journeymen in making fine perceptual discriminations. The effects of perceptual learning can be seen in common experience. Consider televised broadcasts of Olympics events such as gymnastics and diving, where expert analysts notice aspects of performance that we novices can detect only when shown the slow-motion replay. Lesgold, Rubinson, Feltovich, Klopfer, and Glaser (1985) have shown the importance of perceptual skills for radiology diagnosis. And Shanteau (1985) has demonstrated perceptual skill differences for experts versus novices and journeymen in a wide variety of fields. Most of Shanteau's early work focused on an expert's "hit rate" within the framework of research on judgment and decision making. Shanteau found that expert judgment can appear unreliable in domains such as psychiatric diagnosis (where there are no correct answers, so feedback is problematic), whereas in domains such as soil and livestock judging, experts can be very reliable (Phelps & Shanteau, 1978). The reader who is interested in additional discussions of perceptual learning and expertise should consult Chi, Glaser, and Farr (1988) and Hoffman, Burton, Shanteau, and Shadbolt (1991).

The Ability to See Antecedents and Consequents

Experts can visualize how a situation developed into its current state, and they can visualize how it will continue to develop. Kahneman and Tversky (1982) have described a simulation heuristic that people use to arrive at judgments, and Klein (1989) has discussed the importance of this heuristic for decision making and problem solving.

Mental simulation is an important source of power. It enables people to judge how a situation may have developed. Pennington and Hastie (in press) have shown that jurors rely on mental simulations to determine whether a given argument appears plausible. Mental simulation also enables decision makers to evaluate a course of action without comparing it to others. By imagining how the course of action will be carried forward, the decision maker can respond to the quality of the end states that are anticipated. De Groot (1946/1965) has used the term *progressive deepening* to describe how chess players follow a sequence of moves. Holding (1985) has reviewed a wide range of studies on chess expertise and notes that experts are able to "look deeper" into a line of play as well as recognize a better set of moves to consider in the first place. In other words, experts start out with an advantage. They can consider a better set of options, but they also use their understanding of the domain to trace these options to a greater depth, searching for pitfalls and for opportunities. Nonexperts perform the same functions, but without the benefit of a large experience base.

We feel that skill at mental simulation sets experts apart from journeymen and novices. In the study of tank platoon leaders discussed earlier (Brezovic et al., 1987), whereas the novices attended to the same range of cues as the experts, the experts were also visualizing where the adversary's tanks were moving, on the other side of hills and behind ridges. The experts were also visualizing the location and progress of sister platoons. The novices showed minimal awareness of any event that was not physically present.

We assert that experts have an important advantage over journeymen in being able to visualize how situations have developed, and how they are going to evolve. This is especially salient in the skills of expert aerial photo interpreters who, from a single photo, can perceive the long processes that led to the formation of the given terrain (Hoffman, 1984). Novices and journeymen have difficulty in seeing anything other than the current state of a situation, and for this reason they are often unclear about the dynamics of a situation. Novices and journeymen also have difficulty in keeping up with situations, because they lack a basis for anticipating changes and generating expectancies. To use an aviation term, they are usually *flying behind the airplane*. It is no wonder that their responses are variable, awkward, and unintegrated.

IMPLICATIONS AND APPLICATIONS OF A PERCEPTUAL–COGNITIVE VIEW OF EXPERTISE

Theory and application both benefit when they are linked. Without an applied focus, it is easy for theory to lose focus, and to address easy problems rather than hard problems that demand greater complexity. Application without theory can also lose shape and become a disorganized collection of practices. If our perspective on expertise is to develop, it must have value in shaping practice.

Some Implications for Training

How long does it take to become an expert? Simon and Gilmartin (1973) estimated that chess expertise depends on acquiring the knowledge base that allows one to distinguish among approximately 10,000 to 100,000 patterns, and they estimated that this might take up to 10 years. Medical expertise also seems to take about 10 years; and the interpretation of aerial photographs has been estimated to require about 5 years minimum for expertise. In our own work with expert nurses, pilots, and fireground commanders, moderate proficiency (the achievement of journeyman status) seems to take at least 2 years to develop, given continual practice. Many more years are necessary before the person is considered an expert.

There is little reason to believe that we can help people achieve expertise by training them to use the same reasoning strategies as experts. Means, Salas, Crandall, and Jacobs (in press) have reviewed the literature on attempts to teach people generic strategies for decision making and for cognitive development and found no evidence for the effectiveness of such attempts. Means et al. did not find any unambiguous success stories but were able to find many instances of minimal success or outright failure, with one exception: It does seem possible to train people to improve metacognition. Otherwise, it seems fairly clear that there is usually little value in attempting to teach generic reasoning skills.

So how should we approach the job of training people to become experts? If expertise is so dependent on learning to perceive the world differently, then we should look at ways to sharpen perceptual skills, rather than ways to simply add to the experience base. This would enable people to make more rapid and accurate judgments about the nature of the situations they are in.

We can identify at least four strategies for increasing the development of perceptual–cognitive skills: personal experiences, directed experiences, manufactured experiences, and vicarious experiences.

Personal Experiences

The acquisition of personal experiences is straightforward, but inefficient. It is obvious how to gain 10 years of experience—devote 10 years to it. But time-on-the-job is not the key. Rather, the opportunity to be continually challenged seems central.

There is a difference between performing a job automatically and being sensitive to changes in context. There is a difference between performing a job by using standard operating procedures and developing a feel for the situation. Dreyfus and Dreyfus (1986) have argued that an expert learns to give up analytical strategies, including the standard operating procedures that are taught to beginners. In some domains, such as terrain analysis, all cues are always systematically analyzed (Hoffman, 1984). In other domains, the process of explicit analysis can interfere with necessary perceptual learning. Gallwey (1984) has suggested for such domains that someone who insists on analyzing each incident will make less progress and may even be interfering with progress, compared to someone who is learning to rely on perceptual capabilities.

So, the straightforward way to achieve expertise is through direct experience, but this involves more than time-on-the-job. It involves the number, range, and difficulty of challenges faced, and it involves the way a person is able to learn from each incident, along with factors such as degree of engagement with the task.

Directed Experiences

Surveys indicate that over 70% of organizations in the United States use some form of one-on-one instruction to train employees (Lee, 1991). On-the-Job Training (OJT) frequently depends on tutoring from co-workers and supervisors; yet there has been little research on the skills necessary for effective tutoring. Recently, Crandall, Kyne, Militello, and Klein (1991) performed a knowledge elicitation study of master tutors in two domains—critical care nursing and musical instruction. Whereas these domains had different requirements, a list of common skills was compiled: observation of performance, assessment, modelling, guiding motivation and attitudes, relieving anxiety, developing a professional identity.

There appears to be a clear value for organizations to carefully examine their reliance on tutoring, and to adopt measures for more carefully selecting and training tutors. Typically, tutors in the workplace are given no training, and the result may be to reduce the value of OJT. It is the perceptual–cognitive skills of the tutor that the trainee needs to acquire, and these are precisely the skills that are most difficult to communicate. A perceptual–cognitive approach to defining tutor skills and helping tutors learn ways to pass these skills along to trainees could have great value.

It is worth noting that many skills of master tutors are difficult or impossible to incorporate into Intelligent Tutoring Systems (ITSs), thereby setting bounds on the effectiveness of an ITS within a workplace setting.

Here is an example of tutoring expertise, using an incident studied by Crandall et al. (1991):

> A music tutor described how she assessed a young woman who was becoming proficient on the viola. The tutor was surprised to find that the student's tone was off. Upon probing, we found that it was not simply that the student was not hitting the notes, since most viola players are off a bit. Rather, what struck the tutor was that the student was not adjusting her fingering. Even for long notes, the student was not sliding her fingers to make corrections. So the critical cue wasn't what the tutor heard, but what she didn't hear. Her expertise enabled her to have a clear expectancy of how a student at this level should be adjusting. The expectancy was so clear that the tutor could immediately notice that the playing was flawed.

This incident demonstrates the use of expertise to form and use expectancies, and to focus on what was not happening as a cue to improving performance.

Manufactured Experiences

It may be inefficient or impractical to rely on personal experiences, and in such cases it may be useful to present simulations. Computer technology is increasingly able to provide low-cost and high-fidelity simulations that, if used properly, can speed the acquisition of expertise.

One great advantage of training simulations is that they allow one to sharpen discriminations. By facing stimulus configurations that are similar, the trainee has to learn how to make finer distinctions. Also, a person can practice with a larger and more varied set of configurations. So, simulations can allow a person to develop situation assessment skills, and to quickly size up a situation (Lintern, 1991).

Such training need not be extremely expensive. To increase the ability to rapidly size up a situation, simple visual displays could be used for presenting task and cue configurations. This would give practice in determining the different aspects of situation assessment—plausible goals, critical cues, expectancies, and courses of action that are likely to succeed.

For instance, if expertise depends on experience with tough cases, we could increase the rate at which tough cases are encountered. Rothkopf (1982) provided an illustrative anecdote: Mechanical looms break down in a number of ways, some of which are rare. It used to take 5 to 7 years to become an expert loom operator because some loom breakdowns occurred only once every 5 to 7 years. To expedite matters, a training loom was constructed that could be deliberately made to show each breakdown type. The time needed to achieve expertise was dramatically reduced.

Vicarious Experiences

A third approach is to use vicarious experiences. For example, stories are accounts of the experiences of others and are often sufficiently vivid to serve as additions to the experience base. Connelly and Clandinin (1990), Howard (1991), and Schank (1990) are only a few researchers who have recently examined storytelling, and Wilensky (1983) has provided a carefully worked-out presentation of issues concerning story grammars.

Klein (1989) has described a strategy that can be used for training—the Critical Decision method. Experts are asked to recount unusual and particularly challenging episodes, or stories, and the interviewer then uses probe questions to investigate what the expert was trying to do, what mistakes would have been made by someone with less ability, and so on. (The initial story account usually glosses over many important details about the role of expertise in the incident.) These accounts are documented in writing and made available as training materials; sometimes a videotape of the interview itself serves as training material.

Essentially, the critical incidents function as stories, and every attempt is made to maintain the account in the respondent's own words. But the respondent may be confused about sequence, so a second sweep is used to pin events to a timeline. Because people often take aspects of expertise for granted, it is necessary to take a third sweep through the incident account, using probes to find out how the situation was perceived, what goals were judged as feasible, and so on. Usually, there will also be a fourth sweep through the account, attempting to identify possible mistakes. These data are appended to the story account. Sometimes it is possible to incorporate these details into the narrative account itself. And sometimes the incident account is diagrammed, and the probe responses are represented as branches from the main theme. Here is an example cited by Crandall and Gamblian (1991), who collected and formatted stories to be used in training nurses working in a neonatal intensive care unit (NICU). The incident was presented by an experienced nurse working in the NICU:

> Recently I had a primary baby, a 28-weeker, who had been born at home. The type of birth meant the baby was open to infection, but he had had an antibiotic series and had stabilized nicely. This was about two weeks after we admitted him: I came on duty and the first thing I thought was—his color's different. It wasn't any dramatic change or anything, not the gross gray-green they get later. More like a shadow across him. It was really subtle. I checked his chart—blood count looking okay, and no change otherwise. Monitors looked okay. He was sleeping so I went ahead with care for the other baby I had that night and got back to him around midnight.
>
> I had found in caring for him that if I spoke quietly as I opened the isolette and then touched him gently that he would rouse, open his eyes and turn his head a bit. I often held his foot, rubbing the sole as he was waking, and he would curl his toes around my finger. This night he had much more trouble rousing. He didn't respond to my touching his feet, and he couldn't seem to get his eyes fully open. He seemed floppy, with less muscle tone than I was used to.
>
> When I went to weigh him, I really began to worry. Many kids hate being weighed, but he liked it. Typically, he would really come to an alert state, opening his eyes wide, moving his legs and feet, and turning his head toward me. I would talk to him, it was a time for interacting with this baby that I looked forward to. This night as I put him on the bedside scale he sneezed a couple of times, turned his head away, and then his whole body went flaccid. It struck me that he was really taxed, and was shutting out. I put him in his isolette and he immediately went into a

sound sleep. I charted all this, and made a mental note to really keep an eye on him. I didn't yet think he was septic, just that something was going on. His O_2 was stable, not increasing, and I had not seen any dramatic color changes. I thought, maybe he's worn out after the evening visit with his folks—I'll see if he recovers.

We drew blood at 4:00. Although he was usually very reactive to the heel prick, this time he roused, fussed briefly, and then went right back to sleep. His blood gasses were not as good, but not remarkably different. He still seemed very lethargic. Towards early morning, he began having some apnea and bradycardia. Again, nothing severe, but a change from earlier in the shift. By the time the day nurse came on, he was looking blue around the mouth. I was relieved that the nurse who took over for me had taken care of him before. She knew him, and I also knew she would relay my concerns to the doctors during rounds. I told her I was really worried about him, that I thought he was getting sick. She took a look at him and said, "Gee, he really does look different." It turned out that he had a massive infection. He was really sick. I think if we had waited even one or two days longer he would have been in a badly deteriorated state, and might not have made it. As it was, he pulled through.

What Can Go Wrong?

One of the difficulties of spotting sepsis very early is that the signs may be extremely subtle. A nurse, even a highly experienced one, may misread the cues or be misled when some cues are present but others are lacking. Sometimes a nurse may think a baby is becoming septic when that's not the case. And sometimes, she may miss those early, subtle cues. What follows are incident accounts that illustrate these situations, along with some ideas about what may have lead the nurse to an erroneous conclusion. We offer them in the spirit in which they were shared with us—that we learn from mistakes, and that even the most skilled, vigilant, and caring nurse misses occasionally. (pp. 10–11)

This is a story about perceptual learning and experience—being able to size up a situation. Crandall and Gamblian (1991) found that stories such as these were evaluated as very helpful for expanding the expertise of experienced nurses who were new on the neonatal intensive care unit, and who needed to come up to speed quickly. They also found that cues embedded within such stories showed high rates of retention weeks after being presented to nursing students.

Chi, Bassok, Lewis, Reimann, and Glaser (1989) studied the use of examples by good and poor college students who had not taken a course in college physics; the task was to solve problems in mechanics. The good students were able to benefit from examples because they were monitoring their understanding of the examples and were actively trying to expand their comprehension of general principles. In contrast, the poor students did not try to generate more global comprehension of examples and did not monitor their own understanding; they relied heavily on examples but were apt to misapply these examples. In other words, poor students may prefer to learn from analogues, but analogical learning is risky because students often do not know what lessons to extract from the analogues.

The use of vicarious experiences treats expertise as a resource. Klein (1992) has suggested that we think about a general discipline of knowledge engineering, akin to petroleum engineering. If knowledge is a resource, we can develop techniques for locating critical knowledge in an organization—identifying the experts. We can develop techniques for eliciting the knowledge, and for processing or codifying it. Finally, we can develop strategies for applying the knowledge, as in training. Knowledge engineering has been discussed in the context of expert systems, and indeed expert systems are a prime example of how to engineer knowledge. But there is no reason why it is necessary to use only expert systems for knowledge engineering. The Critical Decision method is a low-technology strategy that does not depend on computers at all.

Thus far, we have used the Critical Decision method as a knowledge engineering strategy in several different domains. Crandall and Gamblian (1991) used the method to capture and communicate subtle perceptual skills needed by nurses working in a neonatal intensive care unit. Weitzenfeld et al. (1991) used the method to capture important aspects of the expertise of highly skilled computer programmers working for AT&T and developed a training course for improving the troubleshooting and debugging skills of journeyman programmers. Crandall and Klein (1989) used the method to explicate the skills of expert scientists and program managers working at Wright-Patterson Air Force Base, in order to train entry-level scientists and engineers. Projects such as these illustrate how the study of expertise can lead directly to applications for training, particularly to bridge the gap between journeyman and expert levels of performance.

Before turning to a discussion of assessment, we should note that we have not considered those ideas about the training of expertise that can be derived from traditional psychological research on learning (e.g., learning is enhanced by motivation, or by informational feedback; learning is enhanced by an appropriate mix of "part" and "whole" training or massed and distributed practice; temporary memory should not be overloaded, and so on). Our reason for not exploring these ideas is that we find them to be underspecific and therefore of little use when actually applied to a "real-world" domain. For example, to encourage perceptual learning and pattern recognition, one might want to deliberately overload temporary memory, to force the trainee out of an analytical mode. For discussions of training issues from a traditional learning–human factors perspective, see Duncan (1974), Hagman and Rose (1983), and Holding (1987). Also, see Schneider (1985) for a skeptical appraisal of conventional wisdom in this field.

THE ASSESSMENT OF EXPERTISE

One dimension for assessing expertise involves reasoning strategies. If expertise depended on advanced forms of reasoning, it would be clear how to perform assessments. We would ask people to describe the strategy they used in solving a problem, either an actual problem they remembered, or a problem with which they were presented. If there was evidence for the use of more sophisticated strategies, and evidence that these strategies were used skillfully, we could infer that the person showed signs of expertise.

However, as we have argued earlier, there is little evidence that experts use different, let alone more powerful, reasoning strategies than novices, or than journeymen. Experts can draw upon a larger experience base, and this may enable them to use particular strategies, but the strategies themselves do not confer expertise. Nevertheless, it may be possible to examine the strategies people use to take advantage of their declarative knowledge, perceptual–cognitive knowledge, and other aspects of their experience base.

We generally know who the experts are. They notice the subtle but critical cues that others miss. They can reliably make discriminations that are opaque to others. They have clear judgments of the appropriate ways to act in a situation. They can anticipate what is supposed to happen next, and their expectancies are so clear that they quickly notice when they are wrong, so they can rethink their interpretation of what is going on.

We can suggest several ways to distinguish experts from novices and journeymen: performance, content knowledge, and developmental milestones.

Performance

There are clear ways to distinguish expert and journeyman performance. Turning to the suggestions of Glaser (1976), we should be able to distinguish expert performance in terms of variability–consistency, accuracy, completeness, and speed. We would expect expert performance to be smoother, whereas journeymen should exhibit discontinuities in shifting and lurch from one subtask to the next. Even these differences will not always emerge—experts may take longer than journeymen in sizing up a task, and impulsive novices may be very fast. But holding accuracy constant, experts will generally be able to perform more quickly.

A prime example of a performance assessment method is the use of chess ratings, which reliably differentiate between players of different strengths and can be calibrated to predict the proportion of games that will be won by players of unequal strengths. Although this is close to an ideal strategy, in most domains it is not feasible to achieve such precision.

We have speculated that experts are better able to anticipate future events, compared to journeymen. Klein and Peio (1989) employed a prediction paradigm to study this. The application domain was chess. It was argued that strong chess players (i.e., those rated as Class A or as experts) should be able to anticipate the moves in a game played between two experts, whereas mediocre players (rating of 1300 or below) should have much less success in predicting the moves in the same games. The study demonstrated that it was possible to significantly discriminate between experts and novices in terms of their prediction accuracy. This prediction paradigm should be applicable to a variety of domains.

Content Knowledge

There are also ways to distinguish experts and journeymen in terms of their content knowledge. The focus here is on what experts know, rather than on their performance.

Expertise can be assessed in terms of content knowledge, as in tests of declarative knowledge. Expertise can also be studied in terms of the way people see relationships between concepts, especially causal relationships. The work of Schvaneveldt (1991) shows how a simple approach called a conceptual graph can be useful to illustrate how someone understands a domain. Multidimensional scaling techniques were used to generate graphic representations of knowledge structures. This representational format may support methods for distinguishing expert from journeyman conceptual relationships. Multidimensional scaling techniques usually require a great deal of effort to apply, and instructional designers have learned that much of the same benefit can be obtained from informal use of semantic nets (e.g., the concept maps described by McFarren, 1987, and McNeese et al., 1990).

In performing and interpreting research studies, it is important to be sensitive to the dependent variable being used. Even if the design is rigorous, an abstract and tangential measure can mislead us about the phenomenon being studied. For example, Bateson, Alexander, and Murphy (1987) studied expert and novice computer programmers and found that program recall was not as good a test of knowledge as more naturalistic tasks like writing programs or explaining programs. Feltovich, Spiro, and Coulson (1989) said: "Memory for material that has been taught is not the same as learning from instruction" (p. 118).

In this context, we should also mention a project recently completed that used the Critical Decision method for defining the cues used by commercial aviation checkpilots, as they evaluated pilots for certification to fly different aircraft (Kaempf & Klinger, 1992); the goal was to define the cues and to help to standardize the assessment procedures in a way that was sensitive to the perceptual cues relied on by expert evaluators of pilot performance.

Developmental Milestones

Another way to assess expertise is to use progress along skill development milestones. The model presented by Dreyfus and Dreyfus (1986) provides a potential basis for assessing skill. The Piagetian model of Campbell, Brown, and DiBello (1991) may be an even more useful framework for determining a person's progress in mastering a domain.

Our conclusion is that the assessment of expertise can focus on performance, content acquisition, or developmental milestones. Each of these has its own limitations. We have identified a number of dimensions that could be used for assessment and provided some examples of assessments that have been carried out.

CONCLUSIONS

This chapter has examined a view of expertise centered around the idea that experts can see aspects of a situation that are not accessible to people at lower skill levels. Traditional accounts of expertise have tended to emphasize the strategies that experts use, and the knowledge base that experts have compiled. One attractive hypothesis is that expertise is a function of learning more powerful generic strategies and higher order rules (e.g., Flavell, 1971; Kail & Hagen, 1977). Such a view is attractive, because we should be able to identify these strategies and train people to use them. It is usually a good idea to search for generic strategies, because these can be leveraged across many different situations. Unfortunately, this hypothesis has not been a very successful explanation. It does not appear that experience enables people to learn more powerful problem-solving strategies.

A second view of expertise is that the larger knowledge base of experts permits the use of powerful strategies that are also known to journeymen and to novices (Chi et al., 1981). Clearly, experts have acquired more knowledge, so it makes good sense to posit a difference in knowledge base between experts and nonexperts. We have no great quarrel with this approach, but we find it to be somewhat disappointing in terms of what it tells us about the functions of the knowledge base, and in terms of applications to training and to assessment. It emphasizes knowledge rather than the way knowledge is put into action. For these reasons, we feel it may be useful to examine a somewhat different perspective, one that centers around the way that experts perceive tasks and situations:

1. There are several ways that experts can see things that others cannot: Experts can use their knowledge base to recognize typicality, experts can use perceptual learning to make fine discriminations, and experts can use mental simulation to represent antecedents and consequents.

2. Experts can also use their knowledge base to apply higher level rules, such as top-down processing. Such rules are also in the repertoire of novices and journeymen. People at lower skill levels infrequently use higher level strategies because they lack the experience base to make such strategies work.

3. Therefore, trying to train general strategies for thinking like experts may not be worthwhile. Attempts to teach such strategies are not useful in developing expertise, and research does not demonstrate the effectiveness of such attempts.

4. Training can address methods for sharpening perceptual skills and changing the way situations are experienced. There are ways of providing personal experiences, manufactured experiences, and vicarious experiences to accelerate the growth of expertise. We have shown how models of expertise can form the basis of knowledge engineering programs for training specialists in areas as diverse as nursing, computer programming, and research management.

5. It may be useful to develop a discipline of knowledge engineering, to direct us in eliciting and applying expert knowledge.

ACKNOWLEDGMENTS

Support was received for Dr. Klein's work from the U.S. Army Research Institute, Contract MDA903-89-C-0032. The support of Judith Orasanu, Michael Drillings, and Michael Kaplan is gratefully acknowledged. Also appreciated were the comments and reviews from Beth Crandall, Caroline Zsambok, and Michelene Chi.

REFERENCES

Bateson, A. G., Alexander, R. A., & Murphy, M. (1987). Cognitive processing differences between novice and expert computer programmers. *International Journal of Man–Machine Studies, 26*, 649–660.

Benner, P. (1984). *From novice to expert: Excellence and power in clinical nursing practice*. Menlo Park, CA: Addison-Wesley.

Brezovic, C. P., Klein, G. A., & Thordsen, M. (1987). *Decision making in armored platoon command* (KATR-858(b)-87-05F). Yellow Springs, OH: Klein Associates. Prepared under contract MDA903-85-C-0327 for U.S. Army Research Institute, Alexandria, VA.

Calderwood, R., Klein, G. A., & Crandall, B. W. (1988). Time pressure, skill, and move quality in chess. *American Journal of Psychology, 101*, 481–493.

Campbell, R. L., Brown, N. R., & DiBello, L. A. (1991). The programmer's burden: Developing expertise in programming. In R. R. Hoffman (Ed.), *The cognition of experts: Psychological research and empirical AI* (pp. 269–294). New York: Springer-Verlag.

Charness, N. (1981). Search in chess: Age and skill differences. *Journal of Experimental Psychology: Human Learning and Memory, 7*, 467–476.

Chase, W. B., & Simon, H. A. (1973). The mind's eye in chess. In W. G. Chase (Ed.), *Visual information processing* (pp. 215–281). New York: Academic Press.

Chi, M. T. H., Bassok, M., Lewis, M. W., Reimann, P., & Glaser, R. (1989). Self-explanations: How students study and use examples in learning to solve problems. *Cognitive Science, 13*, 145–182.

Chi, M. T. H., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science, 5*, 121–152.

Chi, M. T. H., Glaser, R., & Farr, M. J. (1988). *The nature of expertise*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Chi, M. T. H., Glaser, R., & Rees, E. (1982). Expertise in problem solving. In R. J. Sternberg (Ed.), *Advances in the psychology of human intelligence* (pp. 7–25). Hillsdale, NJ: Lawrence Erlbaum Associates.

Chi, M. T. H., Hutchinson, J. E., & Robin, A. F. (1988). *How inferences about domain-related concepts can be constrained by structured knowledge*. Pittsburgh, PA: Learning Research and Development Center, University of Pittsburgh.

Christensen-Szalanski, J. J. J., & Bushyhead, J. B. (1981). Physicians' use of probabilistic information in a real clinical setting. *Journal of Experimental Psychology, 7*, 928–935.

Connelly, F. M., & Clandinin, D. J. (1990). Stories of experience and narrative inquiry. *Educational Researcher, 19*(5), 2–14.

Crandall, B. W., & Gamblian, V. (1991). *Guide to early sepsis assessment in the NICU*. Fairborn, OH: Klein Associates.

Crandall, B., & Klein, G. A. (1989). *Organizational expertise at WRDC: A roadmap to survival and success* (Final Rep. prepared under Contract F33615-88-C-1707 from MacAulay Brown, Inc.). Yellow Springs, OH: Klein Associates.

Crandall, B., Kyne, M., Militello, L., & Klein, G. A. (1991). Describing expertise in one-on-one instruction (Final Tech. Rep.). Contract MDA903-91-C-0058, U.S. Army Research Institute, Alexandria, VA.

de Groot, A. D. (1946/1965). *Thought and choice in chess* (1st ed.). New York: Mouton.

Dreyfus, H. L. (1972). *What computers can't do: A critique of artificial reason.* New York: Harper & Row.

Dreyfus, H. L., & Dreyfus, S. E. (1986). *Mind over machine: The power of human intuitive expertise in the era of the computer.* New York: Free Press.

Duncan, K. D. (1974). Analytical techniques in training design. In E. Edwards & F. P. Lees (Eds.), *The human operator in process control* (pp. 283–319). London: Taylor & Francis.

Elstein, A. S., Shulman, L. S., & Sprafka, S. A. (1978). *Medical problem solving: An analysis of clinical reasoning.* Cambridge, MA: Harvard University Press.

Feltovich, P. J., Spiro, R. J., & Coulson, R. L. (1989). The nature of conceptual understanding in biomedicine: The deep structure of complex ideas and the development of misconceptions. In D. A. Evans and V. L. Patel (Eds.), *Cognitive science in medicine: Biomedical modeling* (pp. 113–172). Cambridge, MA: MIT Press.

Flavell, J. H. (1971). What is memory development the development of? *Human Development, 14*, 225–286.

Fraser, J. M., Smith, P. J., & Smith, J. W. (1990). *A catalog of errors.* Columbus, OH: The Ohio State University, Department of Industrial and Systems Engineering.

Gallwey, W. (1984). *Inner game of tennis.* New York: Bantam Books.

Gentner, D. R. (1988). Expertise in typewriting. In M. T. H. Chi, R. Glaser, & M. J. Farr (Eds.), *The nature of expertise* (pp. 1–21). Hillsdale, NJ: Lawrence Erlbaum Associates.

Glaser, R. (1976). Cognition and instructional design. In D. Klahr (Ed.), *Cognition and instruction* (pp. 303–315). Hillsdale, NJ: Lawrence Erlbaum Associates.

Hagman, J. D., & Rose, A. M. (1983). Retention of military tasks: A review. *Human Factors, 25*(2), 199–213.

Hoffman, R. R. (1984). *Methodological preliminaries to the development of an expert system for aerial photo interpretation* (Report ETL-0342). Ft. Belvoir, VA: Engineer Topographic Laboratories, U.S. Army Corps of Engineers.

Hoffman, R. R., Burton, A. M., Shanteau, J., & Shadbolt, N. R. (1991). *Eliciting knowledge from experts: A methodological analysis.* Unpublished manuscript. Garden City, NY: Department of Psychology, Adelphi University.

Holding, D. H. (1985). *The psychology of chess skill.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Holding, D. H. (1987). Concepts of training. In L. G. Salvendy (Ed.), *Handbook of human factors* (pp. 939–962). New York: Wiley.

Howard, G. S. (1991). Culture tales: A narrative approach to thinking, cross-cultural psychology, and psychotherapy. *American Psychologist, 46*(3), 187–197.

Kaempf, G. L., & Klinger, D. W. (1992). *Integrated measurement of crew resource management and technical flying skills.* Final Report under contract DTRS-57-89-D-0086 for the Federal Aviation Administration. Fairborn, OH: Klein Associates Inc.

Kahneman, D., & Tversky, A. (1982). The simulation heuristic. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 201–210). Cambridge, MA: Cambridge University Press.

Kail, R. V., & Hagen, J. W. (Eds.). (1977). *Perspectives on the development of memory and cognition.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Klein, G. A. (1989). Recognition-primed decisions. In W. Rouse (Ed.), *Advances in man–machine systems research* (Vol. 5, pp. 47–92). Greenwich, CT: JAI Press.

Klein, G. A. (1992). Using knowledge engineering to preserve corporate memory. In R. R. Hoffman (Ed.), *Psychology of expertise: Cognitive research and empirical AI*. New York: Springer–Verlag.

Klein, G. A., & Calderwood, R. (1988). How do people use analogues to make decisions? *Proceedings of a workshop on case-based reasoning*, 209–223. San Mateo, CA: Morgan Kauffman Publishers.

Klein, G. A., Calderwood, R., & Clinton-Cirocco, A. (1986). Rapid decision making on the fire ground. *Proceedings of the Human Factors Society 30th Annual Meeting, 1*, 576–580. Dayton, OH: Human Factors Society.

Klein, G. A., & Peio, K. J. (1989). The use of a prediction paradigm to evaluate proficient decision making. *American Journal of Psychology, 102*(3), 321–331.

Larkin, J., McDermott, J., Simon, D. P., & Simon, H. A. (1980). Expert and novice performance in solving physics problems. *Science, 208*, 1335–1342.

Lee, C. (1991, October). Who gets trained in what. *Training*. Minneapolis: Lakewood.

Lesgold, A., Rubinson, H., Feltovich, P., Klopfer, D., & Glaser, R. (1985). *Impasses in complex perceptual learning*. University of Pittsburgh, PA: Learning Research and Development Center.

Lintern, G. (1991). An informational perspective on skill transfer for human–machine systems. *Human Factors, 33*, 251–266.

McFarren, M. R. (1987). *Using concept maps to define problems and identify key kernels during the development of a decision support system*. Unpublished master's thesis, Air Force Institute of Technology, Dayton, OH.

McNeese, M. D., Zapf, B. S., Peio, K. J., Snyder, D. E., Duncan, J. C., & McFarren, M. R. (1990). *An advanced knowledge and design acquisition methodology: Application for the pilot's associate* (Final Report AAMRL-TR-90-060). WPAFB, OH: Armstrong Aerospace Medical Research Laboratory.

Means, B., Salas, E., Crandall, B. W., & Jacobs, T. O. (in press). Training decision makers for the real world. In G. A. Klein, J. Orasanu, R. Calderwood, & C. E. Zsambok (Eds.), *Decision making in action: Models and methods*. Norwood, NJ: Ablex.

Orstein, P. A. (Ed.). (1978). *Memory development in children*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Pennington, N., & Hastie, R. (in press). A theory of explanation-based decision making. In G. Klein, J. Orasanu, R. Calderwood, & C. E. Zsambok (Eds.), *Decision-making in action: Models and methods*. Norwood, NJ: Ablex.

Phelps, R. H., & Shanteau, J. (1978). Livestock judges: How much information can an expert use? *Organizational Behavior and Human Performance, 21*, 109–219.

Rothkopf, E. (1982, October). *Basic research and applied psychology: The quest for symbiosis*. Paper presented at the First Annual Adelphia University Conference on Applied Experimental Psychology, Adelphi University, Garden City, NY.

Schank, R. C. (1990). *Tell me a story*. New York: Scribner.

Schneider, W. (1985). Training high-performance skills: Fallacies and guidelines. *Human Factors, 273*, 285–300.

Schvaneveldt, R. W. (Ed.). (1991). *Pathfinder associative networks*. Norwood, NJ: Ablex.

Shanteau, J. (1985). Psychological characteristics of expert decision makers. *Applied Experimental Psychology Series, 85*(2). Kansas State University, Manhattan, KS.

Simon, H. A., & Gilmartin, K. (1973). A simulation of memory for chess positions. *Cognitive Psychology, 5*, 29–46.

Swets, J. A., & Bjork, R. A. (1990). Enhancing human performance: An evaluation of "New Age" techniques considered by the U.S. Army. *Psychological Science, 1*, 85–96.

Weitzenfeld, J. S., Klein, G. A., Riedl, T. R., Freeman, J. T., & Musa, J. (1991, October). Knowledge elicitation for software engineering expertise. *Proceedings of the Fifth SEI Conference on Software Engineering Education*, Pittsburgh, PA.

Wilensky, R. (1983). *Planning and understanding: A computational approach to human reasoning*. Reading, MA: Addison-Wesley.