

Overview

Robert Glaser

Micheline T. H. Chi
University of Pittsburgh

Expertise became an intriguing subject for investigation as a result of work in the mid- to late sixties, largely due to developments in artificial intelligence (AI) and cognitive psychology. Research in AI and attempts to simulate human capabilities had failed to construct programs that could outperform humans, even though computers were by then equipped with powerful search heuristics and essentially limitless search capabilities. Even in programs using selective search, such as Greenblatt's (Greenblatt, Eastlake, & Crocker, 1967) chess program, the best "plausible" move was still selected on the basis of an extensive evaluation, whereas human experts do not engage in particularly extensive searches or elaborate analyses, as shown by findings in cognitive psychology. Investigations of chess playing, for example, the early work of deGroot (1966) and the later extended work of Chase and Simon (1973), demonstrated that what distinguishes strong from weak players are their abilities to correctly reproduce large patterns of chess positions after a few seconds of viewing, rather than their searching more deeply or broadly than weaker players. Clearly, specialized structures of knowledge were strongly implicated, but the nature of this knowledge and of its interactions with general heuristic processes required further analysis.

Newell and Simon (1972) described the chess master's "perceptual" ability as follows:

Clusters of related pieces in a position are recognized as familiar constellations; hence, each cluster is stored as a single symbol. Less skilled players have to describe the board as a larger number of simpler chunks — hence cannot

hold all of the information required to reproduce the board in short-term memory. When the same number of pieces is arranged on the board at random, few of the resulting configurations are familiar even to grandmasters. They then need more symbols to describe the position that can be held simultaneously in short-term memory, hence, they perform as poorly as weaker players. (p. 781)

By using the concept of chunks to explain the chess master's pattern recognition superiority, it became necessary to identify experimentally the structure and size of chunks in the knowledge base, because a chunk appeared to be a defining unit of knowledge structure. Hence, early in its history, the study of expertise provided evidence of a knowledge-competence dimension as a primary focus.

In AI research, it became widely acknowledged that the creation of intelligent programs did not simply require the identification of domain-independent heuristics to guide search through a problem space; rather, that the search processes must engage a highly organized structure of specific knowledge for problem solving in complex knowledge domains. This shift in AI was characterized by Minsky and Papert (1974) as a change from a power-based strategy for achieving intelligence to a knowledge-based one. They wrote:

The *Power* [italics added] strategy seeks a generalized increase in computational power. It may look toward new kinds of computers ("parallel" or "fuzzy" or "associative" or whatever) or it may look toward extensions of deductive generality, or information retrieval, or search algorithms. . . . In each case, the improvement sought is intended to be "uniform" — independent of the particular data base. The *Knowledge* strategy sees progress as coming from better ways to express, recognize, and use diverse and particular forms of knowledge. (p. 59)

This point of view has since been reiterated in the textbooks and handbooks on building expert systems (e.g., Hayes-Roth, Waterman, & Lenat, 1983). These texts point out that the principal developments in AI fostered the current emphasis on knowledge-based expert systems and the related field of knowledge engineering. Machines that lack knowledge can perform only intellectually trivial tasks. Those that embody knowledge and apply it can approximate the performance of human experts. As a consequence, expert-system building has concentrated on the knowledge that underlies human expertise and given less emphasis to the significance of domain-independent problem-solving heuristics.

Thus, the seeds of the study of the characteristics of highly competent expert performance were sown in the fertile ground of Newell and Simon's 1972 book, *Human Problem Solving*, although the topic *expertise* was not listed. In the ensuing years, the need for research in expertise has been

recognized, and much research in cognitive psychology has been devoted to this topic.¹ In the following pages, we briefly summarize some key characteristics of experts' performances that this research has uncovered. These findings are robust and generalizable across the various domains that have been studied (Glaser, 1988). We also highlight other relevant findings, and speculate briefly on the nature of the organization of the knowledge base that generates each characteristic.

1. Experts Excel Mainly in Their Own Domains. There is little evidence that a person highly skilled in one domain can transfer the skill to another. As Minsky and Papert (1974) noted: "A very intelligent person might be that way because of specific local features of his knowledge-organizing knowledge rather than because of global qualities of his 'thinking'" (p. 59). Evidence for such a conclusion can be drawn from the work of Voss and Post (this volume) on problem solving in political science. In that work, nondomain experts (chemists) solved political science problems much like novices, describing the causes for the problem at a very concrete and specific level, whereas domain experts described more abstract causal categories.

The obvious reason for the excellence of experts is that they have a good deal of domain knowledge. This is easily demonstrated; for example, in medical diagnosis, expert physicians have more differentiations of common diseases into disease variants (Johnson et al., 1981). Likewise, in examining taxi drivers' knowledge of routes, Chase (1983) found that expert drivers can generate a far greater number of secondary routes (i.e., lesser known streets) than novice drivers.

2. Experts Perceive Large Meaningful Patterns in Their Domain. As mentioned, this is apparent in chess, where it is well known that chess masters excel in their recall of the clusters of pieces that they see. This perceptual superiority has been replicated in several other domains, such as in the game of GO (Reitman, 1976), in reading circuit diagrams (Egan & Schwartz, 1979), in reading architectural plans (Akin, 1980), and in interpreting x-ray plates (Lesgold et al., this volume). It should be pointed out, however, that this ability to see meaningful patterns does not reflect a generally superior perceptual ability; rather, it reflects an organization of the knowledge base. Programmers, for example, can recall key programming language words in meaningful clusters (McKeithen, Reitman, Reuter, & Hirtle, 1981), and expert programmers can also recognize and recall familiar subroutines (see Soloway, Adelson, & Ehrlich, this volume).

¹The topic of expertise first appears in major textbooks in cognitive psychology in 1985, in John Anderson's second edition of *Cognitive Psychology and Its Implications*.

3. Experts are Fast; They Are Faster than Novices at Performing the Skills of Their Domain, and They Quickly Solve Problems with Little Error. An easy way to observe the skill of master chess players is to watch them play "lightning chess," where they have only a few seconds to decide on a move. Although studies in the literature actually find experts slower than novices in the initial phases of problem solving, experts solve problems faster overall.

There are at least two ways to explain experts' speed. For simple tasks, such as typing, the speed that experts have acquired comes with many hours of practice, which makes the skill more automatic and frees up memory capacity for processing other aspects of the task (see Gentner, this volume). Thus, they can be fast because they are actually faster at the skill itself or because they have more capacity to perform the total task. The expert typists in Gentner's study were fast because their fingers moved quickly (there were more overlapping movements), as well as because they could free up resources to perform related tasks such as typing degraded pseudowords, whereas novices had few resources available for attending to pseudowords.

A further possible explanation for experts' speed in solving problems rests on the idea emphasized earlier that experts can often arrive at a solution without conducting extensive search. The patterns that chess experts see on the board suggest reasonable moves directly, presumably because, through many hours of playing, they have stored straightforward condition-action rules in which a specific pattern (the condition) will trigger a stereotypic sequence of moves. Cab drivers, for instance, will recognize a shorter route while traveling to their destination, even though they may not have generated this shorter route in the laboratory (Chase, 1983).

4. Experts Have Superior Short-Term and Long-Term Memory. With recently presented materials, experts' recall seems to exceed the limits of short-term memory. This is not because their short-term memory is larger than other humans', but because the automaticity of many portions of their skills frees up resources for greater storage. Experts seem to excel in long-term recall as well. For example, in chess, it is not uncommon for chess masters to recognize plays from certain well-known games.

Chase and Ericsson's (1982) study demonstrated experts' superiority in both short-term and long-term recall. They found that their trained memory expert could remember more than 80 digits in a short-term memory serial recall task. They also found, however, that he could recognize over 80–90% of the digit groups that had been presented to him for recall a week earlier.

5. Experts See and Represent a Problem in Their Domain at a Deeper (More Principled) Level than Novices; Novices Tend to Represent a Problem at a Superficial Level. An easy and robust way to demonstrate this is to ask experts and novices to sort

problems and analyze the nature of their groupings. Using physics problems, Chi, Feltovich, and Glaser (1981) found that experts used principles of mechanics to organize categories, whereas novices built their problem categories around literal objects stated in the problem description. Similar results have been found in the domain of programming (Weiser & Shertz, 1983); when expert and novice programmers were asked to sort programming problems, the experts sorted them according to solution algorithms, whereas the novices sorted them according to areas of application (e.g., whether the program was supposed to create a list of employees' salaries or whether it was supposed to keep a file of current user identifications). These results indicate that both experts and novices have conceptual categories, but that the experts' categories are semantically or principle-based, whereas the categories of the novices are syntactically or surface-feature oriented.

6. Experts Spend a Great Deal of Time Analyzing a Problem Qualitatively. Protocols show that, at the beginning of a problem-solving episode, experts typically try to "understand" a problem, whereas novices plunge immediately into attempting to apply equations and to solve for an unknown. What do the experts do when they qualitatively analyze a problem? Basically they build a mental representation from which they can infer relations that can define the situation, and they add constraints to the problem. Paige and Simon's (1966) well-known example illustrates this by asking students to solve simple algebra word problems, such as: A board was sawed into two pieces. One piece was two thirds as long as the whole board and was exceeded in length by the second piece by four feet. How long was the board before it was cut? Paige and Simon found that some students immediately applied equations, which then resulted in their coming up with a negative length; others, however, remarked that the problem was meaningless because one cannot have a board with a negative length. One can conclude that those students who paused had formed a mental model of the situation and made some inferences about the relation between the boards.

The utility of qualitative analysis for adding constraints to a problem can be seen most clearly in ill-defined problems. Voss and Post (this volume) presented economic problems, such as: Imagine you are the Minister of Agriculture for the Soviet Union. Crop productivity has been too low for the past several years. What would you do to increase crop production? About 24% of the experts' solution protocols (those of political scientists specializing in the Soviet Union) were elaborations on the initial state of the problem, as opposed to 1% of the novices' protocols. By elaborating the initial state, the experts identified possible constraints, such as Soviet ideology and the amount of arable land. (Adding constraints, in effect, reduced the search space. For example, introducing the constraint of the amount

of arable land eliminated the solution of increasing planting, and considering the constraint of the Soviet ideology precluded the solution of fostering private competition — a capitalistic solution.) Other examples of adding constraints can be seen in the work of Lawrence (this volume) on magistrates' decision-making processes.

7. Experts Have Strong Self-Monitoring Skills. Experts seem to be more aware than novices of when they make errors, why they fail to comprehend, and when they need to check their solutions. For example, the expert physics-problem solver in Simon and Simon's study (1978) would often check his answer, and the expert physics-problem solver in Larkin's study (1983) would often abandon solution attempts before carrying out the mathematical details. Experts' self-knowledge is also manifested in their being more accurate than novices in judging the difficulty of a physics problem (Chi, Glaser, & Rees, 1982). Expert chess players are more accurate than novice players at predicting how many times they will need to see a given board position before they can reproduce it correctly (Chi, 1978). Experts ask more questions, particularly when the texts from which they have to learn are difficult (Miyake & Norman, 1979). Novice learners, on the other hand, ask more questions on the easier materials.

We can argue that, in each of the above examples, the superior monitoring skills and self-knowledge of experts reflect their greater domain knowledge as well as a different representation of that knowledge. We illustrate this dependence on domain knowledge with an example from our own work on physics. As stated, we found that expert physicists were more accurate than novices in predicting which physics problems will prove more difficult to solve. If we probe further and look at the bases on which they made such judgments, we see that they relied on the same knowledge of principles in this task as they used to sort problems into categories. Although about a third of both experts' and novices' decisions about problem difficulties were based on the problems' characteristics (such as "the problem is simplified because it is frictionless"), another third of the experts' judgments were based on the underlying physics principle governing the solution (such as "it's a straightforward application of Newton's second Law"). Only 9% of the novices' judgments were based on the underlying principle. In addition, novices used nonproblem related characteristics (such as, "I've never done problems like this before") about 18% of the time as compared to 7% for the experts (Chi, 1987). The ability of experts to predict accurately which problems were difficult and which were easy enabled them to monitor accurately how they should allocate their time for solving problems. Thus, the monitoring skills of experts appear to reflect their greater underlying knowledge of the domain, which allowed them to predict problem difficulty on the basis of the physics principles rather than less relevant surface features.

Summary

The short history of research on expertise might be written as follows: Information-processing studies of problem solving in the 1960s and 1970s and early work in AI and expert systems accepted a tradition of concentrating primarily on basic information-processing capabilities that humans employ when they behave more and less intelligently in situations in which they lack any specialized knowledge and skill. The pioneering work of Newell and Simon and others richly described these general heuristic processes, but they also offered crucial beginning insight on the learning and thinking of experts, processes that require a rich structure of domain-specific knowledge. In recent years, research has examined knowledge-rich tasks — tasks that require hundreds and thousands of hours of learning and experience. These studies of expertise, together with theories of competent performance and attempts at the design of expert systems, have sharpened this focus by contrasting novice and expert performances. These investigations into knowledge-rich domains show strong interactions between structures of knowledge and processes of reasoning and problem solving. The results force us to think about high levels of competence in terms of the interplay between knowledge structure and processing abilities. They illuminate the set of critical differences highlighted in this overview between individuals who display more and less ability in particular domains of knowledge and skill. We interpret these differences as primarily reflecting the expert's possession of an organized body of conceptual and procedural knowledge that can be readily accessed and used with superior monitoring and self-regulation skills.

Now research needs to go beyond this stage of analysis. We must better understand the properties of domain structure and integrated knowledge; the mechanisms of problem-space definition with minimal search through rapid pattern recognition; and the processes involved in redefining the space of ill-structured and difficult problems. To do so, we should investigate the forms of reasoning and problem-solving strategies that structured knowledge facilitates. We also need to understand how expertise is acquired, how it can be taught, and how beginning learners can be presented with appropriate experience. The papers in this volume consider these themes and represent the type of research that is presently being carried out that investigates both human and artificial expertise.

The Contents of This Volume

The majority of the chapters in this volume were presented at a conference held at the Learning Research and Development Center at the University of Pittsburgh, sponsored by the Personnel and Training Research Program,

