

Problem Solving and Human Expertise

T J Nokes and C D Schunn, University of Pittsburgh, Pittsburgh, PA, USA

M T H Chi, Arizona State University, Tempe, AZ, USA

© 2010 Elsevier Ltd. All rights reserved.

Glossary

Absolute methods – A research methodology that consists of an in-depth examination of high-level experts, usually in a task specific to their domain of expertise (e.g., playing a game of chess for chess experts).

Chunk – A unit of knowledge that is composed of several smaller units of information.

Declarative knowledge – This consists of descriptions about the world including facts, strategies, and principles, and is commonly referred to as knowing that.

Knowledge compilation – A cognitive mechanism hypothesized to interpret declarative knowledge into a set of specific procedural rules given a particular goal.

Procedural knowledge – This consists of information for how to perform particular actions to accomplish task goals, and is commonly referred to as knowing how.

Relative methods – A research methodology that consists of comparing more- to less-experienced participants, often in a neutral task not typically practiced in their domain of expertise (e.g., recalling chess positions for chess experts).

Schema – A hierarchical knowledge structure that includes prototypical information about the type of problem, including declarative knowledge of objects, facts, strategies, and constraints and may also include procedural operators for solving the problem.

Problem solving is a critical cognitive activity that permeates many aspects of our day-to-day lives. We solve problems at home, school, and work ranging from the simple – such as figuring out the tip on a bill – to the complex – such as planning the logistics of a family trip. Problems sometimes have clear goals and steps you can take (as in algebra problems), but sometimes have vague goals or ambiguity about what solution methods are possible. The latter are called ill-structured problems and are considered much more difficult than well-structured problems. Developing expertise in problem solving is critical to the success of a wide range of human activities, including pursuits in science, art, business, and politics. As our society becomes ever-more technologically diverse and

sophisticated, experts are sought in more and more specialized fields. Having a scientific explanation of expert performance is needed to understand its development and to facilitate its acquisition. Knowing what to teach influences the methods of teaching. Expertise research is an area that provides a basis for determining what needs to be taught. Our purpose in writing this article is to provide an integrative review of the psychological research on expert problem solving by taking a close look at what it is, how it is acquired, and the implications for education and instruction.

We structure the article around two interrelated themes. First, that expertise can be understood from an information-processing perspective by focusing on the role of knowledge, its content, and the cognitive processes that bring that knowledge to bear during problem solving. Second, that expert performance is acquired through deliberate practice (Ericsson *et al.*, 1993). This view that expertise can be decomposed into a set of knowledge structures that are learned has implications for how to structure learning environments in order to facilitate its acquisition. In the rest of this article, we explore these themes beginning with a brief review of the methods used to examine expertise, followed by a detailed analysis of how expertise impacts each stage of problem solving. We then review the research on its acquisition with a focus on the underlying cognitive processes. In the final section, we discuss current directions and implications for instruction.

Methods

Researchers have typically used one of two approaches to study expertise, what Chi (2006) has called absolute and relative methods. Absolute methods consist of an in-depth examination of high-level experts, usually in a task specific to their domain of expertise, such as playing a game of chess for chess experts. Defining the level of expertise occurs through established criteria for a particular domain. For some domains, there are written criteria (a rating or scoring system) to determine rank, such as in chess. In other domains, expertise is determined by a certain level of professional achievement, such as becoming a professional ballet dancer, physicist, or a commercial airline pilot. The absolute approach is aimed at providing an in-depth description of the knowledge and cognitive processes underlying

expert performance. This approach includes both observational studies as well as historical analyses of famous cases (e.g., James Clerk Maxwell; Nersessian, 1992). Relative methods involve comparing more- to less-experienced participants, often in a neutral task outside of their domain expertise, such as recalling chess positions for chess experts. The advantage of this approach is that it can uncover the structures and processes of performing the task, and not merely the ways that experts can excel. Both approaches have made extensive use of verbal protocols to obtain detailed data as to the thinking processes that accompany expert (and novice) performance (Ericsson, 2006).

These approaches have produced a wealth of findings on the nature of expertise (for general reviews see the section titled 'Further reading'). In the next section, we draw upon this literature to examine the impact of expertise on problem solving. We begin by describing a general theory of problem solving and then at each stage of the process describe the differences between expert versus novice performance and the explanations to account for those differences.

Expert Problem Solving: Major Findings

Most theories of human problem solving consist of some formulation of the following seven stages:

1. problem categorization,
2. construction of a mental representation of the problem,
3. search for the appropriate problem-solving operators (e.g., strategies or procedures),
4. retrieval and application of those operators to the problem,
5. evaluation of problem-solving progress and solution,
6. iterating stages 1–4 if not satisfied with progress/solution, and finally
7. storage of the solution (e.g., Newell and Simon, 1972).

These stages may not be strictly sequential, but may be iterative. In the following subsections, we describe the expertise findings relevant to each stage and discuss the theories proposed to account for them (see Figure 1 for an illustration of the problem-solving stages and the impact of expertise on each one).

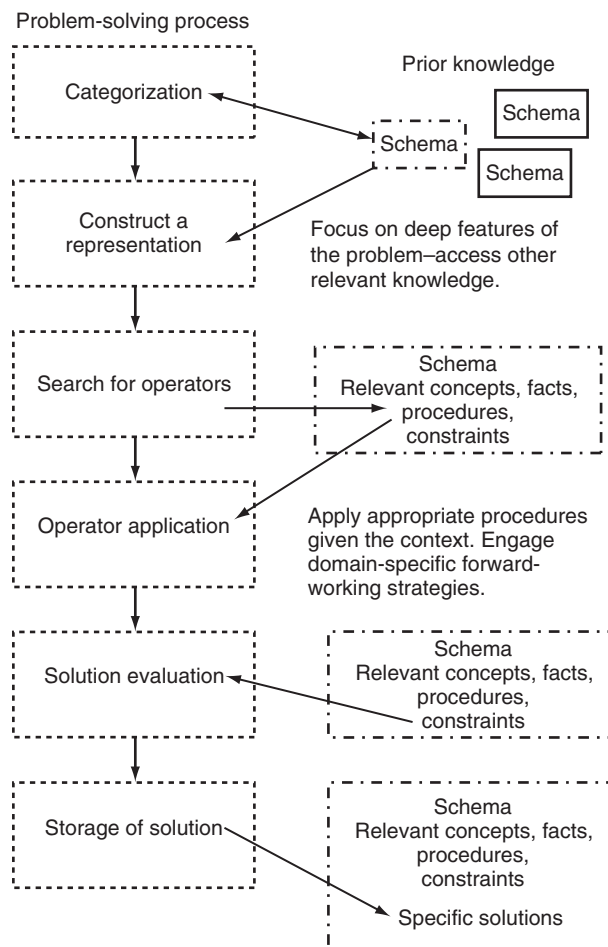


Figure 1 The impact of expertise on each stage of the problem-solving process.

Problem Categorization

The first stage involves the categorization of the problem. This stage is critical as it impacts all subsequent problem-solving processes, such as determining what knowledge to use and what strategies are relevant. For example, after a statistician categorizes a statistics problem as a permutations problem, she or he can proceed by retrieving and applying the appropriate formula to solve it. Much research has shown that experts' domain knowledge actually influences problem perception. When experts are presented a problem or task relevant to their domain of expertise, they see the problem in terms of prior meaningful patterns of information. For example, Chase and Simon (1973a, 1973b) found that expert chess players recalled more than novices on a memory task in which they were briefly presented a game scenario that they had to reconstruct. The experts recalled approximately four times as many pieces as the novices but only for scenarios that were from real games; when the scenario consisted of randomly placed pieces, experts performed at the same level as the novices. It was hypothesized that the experts' prior knowledge facilitated the recognition and recall of domain-relevant patterns, or chunks, of information from the scenarios (see Figure 2 for an example of chunks in chess). These chunks provide experts useful ways to perceive and reason about large amounts of domain-relevant information.

Similar effects have been shown in research on medical expertise. For example, Lesgold *et al.* (1988) compared expert to novice physicians as they diagnosed X-ray films of the lungs. The physicians were asked to draw on the X-rays to identify the important features of their diagnosis. Both groups noticed abnormalities associated with a collapsed lung. However, experts were much more likely to identify the correct shape and size of the abnormality, whereas novices identified abnormalities that were approximately half the size of those identified by the experts. This work shows that experts and novices can perceive a problem very differently even when looking at the exact same stimulus. This finding that expert knowledge impacts problem perception has been found in a variety of tasks and domains including: architecture (Akin, 1980), mathematics (Silver, 1979), and naturalistic decision-making (NDM) tasks such as a fireman determining the safety of a room in a burning building (Klein, 1998).

A related effect is the finding that experts are more likely than novices to categorize problems at a deep level of abstraction (or function), whereas novices are more likely to categorize problems based on the surface features. For example, in the seminal work by Chi *et al.* (1981) experts and novices were asked to sort physics word problems based on their similarity. Experts sorted them according to their underlying physics principles, such as Newton's second

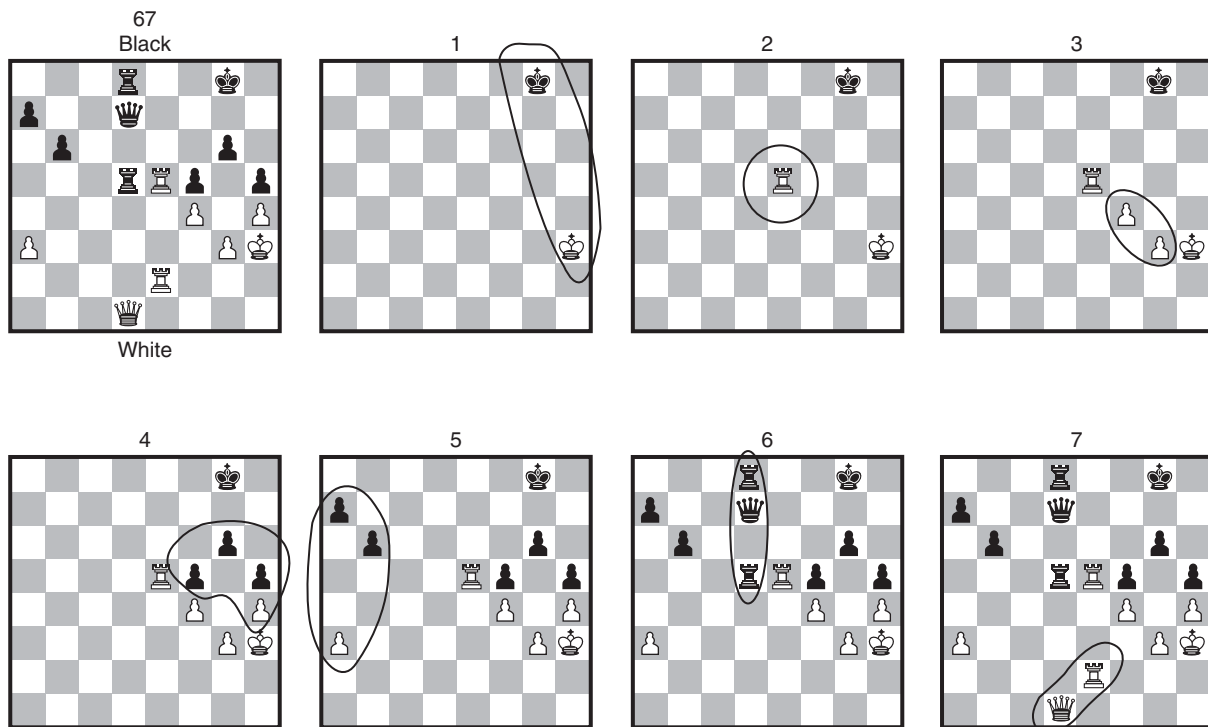


Figure 2 An example of a chess master's chunk-by-chunk recall of position 67 from Reinfeld (1945) with each new chunk circled. From Chase, W. G. and Simon, H. A. (1973). The minds eye in chess. In Chass, W. G. (ed.) *Visual Information Processing*, pp 215–281. New York: Academic Press. With permission from Elsevier.

law, whereas novices sorted them based on their surface features such as inclined planes or pulleys. Similar results have been shown in mathematics, where novices categorized algebra problems on the basis of the problem content (e.g., river problems), whereas more experienced students categorized them based on the underlying equation or principle (Silver, 1979). This effect has been found in a number of domains, including computer programming (Adelson, 1981), medicine (Groen and Patel, 1988), and engineering design (Moss, *et al.*, 2006) among others.

These results have typically been explained by the hypothesis that experts' problem schemas are organized differently than novices. Schemas are hierarchical knowledge structures that include prototypical information about the type of problem, including declarative knowledge of objects, facts, strategies, and constraints, and may also include the procedural operators for solving the problem (Marshall, 1995). Expert schemas are hypothesized to include many principle or structural features of the problem type, whereas novice schemas include few structural features and shallower, surface features. Schemas play a critical role in categorizing a problem. See Figure 1 for the interactive role schemas play in both problem perception and construction of the problem representation. Next, we discuss how experts and novices construct a mental representation of the whole problem.

Construction of a Representation

After a problem has been categorized, the problem solver can begin to elaborate their mental representation that goes beyond the given information of the task environment. For some simple or well-practiced tasks (e.g., puzzle-type tasks, NDM tasks, and procedural skills), this step happens as rapidly as categorization but for other complex, multistep problems, such as those in physics or ill-structured tasks such as design tasks, constructing a mental representation is an iterative process that takes time to develop. Constructing a representation involves specifying the important features of the task such as the relevant objects, operators, and constraints. When experts are solving complex or ill-structured problems, they approach them qualitatively, first examining and elaborating the givens of the problem and then refining that representation. For example, Voss and Post (1988) examined how experts and novices in political science solved an open-ended problem on how to increase crop productivity in the Soviet Union. They showed that experts spent more time than novices in developing their representation of the problem by elaborating the history and causal factors underlying the problem.

These results are consistent with the schema account of expertise in which experts' schematic knowledge provides access to additional knowledge and strategies to

help elaborate and develop the initial problem representation. This process has been hypothesized to be highly interactive (Chi *et al.*, 1981). Based on the initial categorization, the activated schema can provide additional information, strategies, constraints, and expectations to further characterize and elaborate the problem representation that may in turn activate other relevant schemas. For very complex problems, this process may take several iterations. Consistent with the categorization results described earlier, McDermott and Larkin (1978) (see also Reimann and Chi, 1989) have proposed that physicists construct problem representations at different levels of abstraction including: literal, naive, scientific (qualitative), and algebraic (see Table 1 for a description of each level). Differences in levels of representation have also been shown in medicine where experts represent text descriptions of patient cases with an abstract situation model, whereas novices represent them more at the text-based (or surface) level (Groen and Patel, 1988).

Not only does expert knowledge facilitate the development and elaboration of the problem representation, but research also shows these representations are very durable. Experts have been shown to quickly encode problems and are able to easily access that representation even after disruption, whereas novices often take much longer to encode and re-represent a problem (Ericsson and Kintsch, 1995). The durability of expert memory and encoding has been shown in a variety of domains, including bridge (Charness, 1979), medicine (Norman *et al.*, 1989), and computer programming (McKeithen *et al.*, 1981). To account for these findings, Ericsson and Kintsch postulate that experts have developed effective long-term working memories that use very specific cues in the task environment to reliably retrieve prior knowledge structures (chunks and schemas).

Table 1 Four different levels of abstraction in representing physics problems

<i>Representation level</i>	<i>Description</i>
Literal	Representations containing keywords from the text.
Naive	Representations containing literal objects and their spatial relationships, often accompanied by a sketch of the situation.
Scientific	Representations containing idealized objects (points, bodies) and physical concepts (forces, momenta).
Algebraic	Equations containing physical concepts and their relationships.

From Reimann, P. and Chi, M. T. H. (1989). Human expertise. In Gilhooly, K. J. (ed.) *Human and Machine Problem Solving*, pp 161–191. New York: Plenum, with permission from Springer.

Application of Problem-Solving Procedures

After a problem representation has been constructed, the problem solver can then access and apply the appropriate problem-solving strategies and procedures to solve it. Experts have been shown to have more reliable access than novices to domain-specific solution procedures for well-practiced problem types. For simple problems, they make decisions faster and more accurately than novices. Research has also shown that experts and novices use different types of strategies when solving simple problems. Experts are more likely to use forward-working strategies for well-practiced problems, whereas novices use backward-working strategies. Forward-working strategies consist of working toward the solution from the domain principles. For example, physics experts first identify the principles for the task and then apply the domain-specific strategies and procedures, working step-by-step toward the solution (Simon and Simon, 1978). In contrast, novices have been shown to use general problem-solving heuristics, such as means–ends analysis to work backward from the problem goal (e.g., a sought value in physics or math). However, strategy use for both experts and novices critically depends on the relationship between prior knowledge and the task. Experts may also use general problem-solving methods and backward-working strategies when solving very novel tasks in the domain (e.g., physicists in their own research).

Solution Evaluation and Storage

Solution evaluation is the process of assessing a problem solution. Research has shown that experts spend more time than novices evaluating their solutions to make sure they satisfy task constraints (Groen and Patel, 1988; Voss and Post, 1988). Experts are also more likely than novices to identify and correct errors. For example, historians given a problem outside their subdomain are more likely than novices to seek additional resources and information to revise their initial framing of the problem, whereas novices are more likely to proceed with their initial incorrect assumptions (Wineburg, 1998). This research suggests that experts have developed better meta-cognitive skills (i.e., reflective monitoring) than novices for domain-relevant tasks. These skills may be particularly useful when adapting their knowledge to novel tasks in the domain.

After a solution has been generated it can be stored for later use. Much research shows that prior knowledge has a large impact on what is learned. For example, it is easier for experts to acquire new knowledge in the domain than for novices. Baseball experts have better recall than novices after listening to the broadcast of a novel baseball game (Spilich *et al.*, 1979) and expert pilots recall more than novices after listening to new air traffic control messages (Morrow *et al.*, 2001). Experts' rich, well-organized knowledge structures

enable them to easily incorporate (assimilate) new information into their prior knowledge.

Summary

Theoretical accounts of expert–novice differences are primarily articulated in the representation and organization of expert knowledge. Not only do experts have more conceptual and procedural knowledge than novices, but their knowledge is also organized in ways that facilitate effective problem solving. They are able to quickly recognize large chunks of domain-relevant information, see the deep features of the problem, and effectively elaborate their initial problem representations. They can apply domain-specific strategies, efficiently monitor their problem-solving progress by refining and correcting solutions, and can learn new domain-relevant information easier than novices. In the next section, we briefly review the theoretical accounts of how this knowledge is acquired.

Acquisition of Expertise

Much research shows that a minimum of 10 years of daily deliberate practice is necessary to develop expertise in most domains (Ericsson *et al.*, 1993). Ericsson and colleagues refer to deliberate practice as

repeated experience in which the individual can attend to the critical aspects of the situation and incrementally improve her or his performance in response to knowledge of results, feedback, or both from a teacher (Ericsson *et al.*, 1993: 368).

This perspective emphasizes how the type and structure of practice is critical to the acquisition of expert performance. In contrast to this perspective is the view that expertise is due to some talent or innate ability. The talent perspective, originally proposed by Galton (1869), is the notion that psychological traits, like physical traits, are inherited and family lineage (i.e., genes) strongly influences the person who achieves expert performance. Most modern formulations of this perspective hypothesize that expertise is the result of a complex interaction between genetic dispositions and experience (e.g., Simonton, 1999). Given that the talent perspective has received limited empirical support (see Howe *et al.*, 1998 for a discussion and commentary) and much research shows that expert advantages are due to their domain knowledge (and not general reasoning or memory abilities), we focus on the cognitive learning processes that give rise to this knowledge.

Expert knowledge is composed of both declarative and procedural components. Declarative knowledge consists of descriptions about the world, including facts, strategies, and principles, and is commonly referred to as knowing that. Procedural knowledge consists of information for how to

perform particular actions to accomplish task goals, and is commonly referred to as knowing how. Different learning processes have been hypothesized to account for the acquisition of these two types of knowledge. Learning declarative knowledge has been hypothesized to occur through observation, comprehension processes for oral and written discourse, induction, analogy, inference, and self-explanation (see Chi and Ohlsson, 2005 for a recent review of the learning mechanisms that lead to the acquisition of complex declarative knowledge). The key point is that declarative knowledge can be acquired through a number of reflective cognitive processes. Learning environments (e.g., classroom instruction) can be structured to facilitate its acquisition by including and improving these processes.

The acquisition of procedural knowledge or skill is hypothesized to occur through the repeated practice of a particular task or problem (Anderson, 1982, 1987). Fitts (1964) has characterized skill acquisition into three stages of performance, including the cognitive, associative, and automatic stages. During the cognitive stage, a person applies declarative knowledge to solve a problem and performance is characterized as being slow, effortful, and error prone. In domains such as mathematics and physics, novices rely heavily on declarative knowledge from prior examples to solve new problems (e.g., VanLehn, 1998). Students often apply this knowledge by making an analogy between the current problem they are solving and a previous problem that was solved similarly or had similar content. In the associative stage, the skill is practiced and performance becomes faster, more accurate, and less susceptible to interference. In this stage, students rely less on examples and more on applying learned rules to solve the problem. In the automatic stage, the skill has become proceduralized and is characterized by the fast application of the knowledge (or rules) with little or no errors and requires minimal cognitive resources.

Research on skill acquisition has revealed a power-law relationship between the amount of practice and performance. Generally, it shows that performance improves most when first learning a task, followed by decreasing learning gains as practice continues until performance asymptotes. However, the pattern of learning is more specific than the fast-then-slow pattern: when plotted on a logarithmic scale, the power-law relationship is revealed as an exact straight line. This exact relationship has been shown to be a very general phenomenon and has been observed in a variety of activities from learning to roll cigars to learning to solve math problems (see Proctor and Dutta, 1995 for a review). See Figure 3 for a real-world example of the power-law relationship.

One mechanism hypothesized to account for procedural learning is knowledge compilation (Anderson, 1987). Knowledge compilation acts as a translation device that interprets, or compiles, declarative knowledge into a set of specific procedural rules given a particular goal. As those

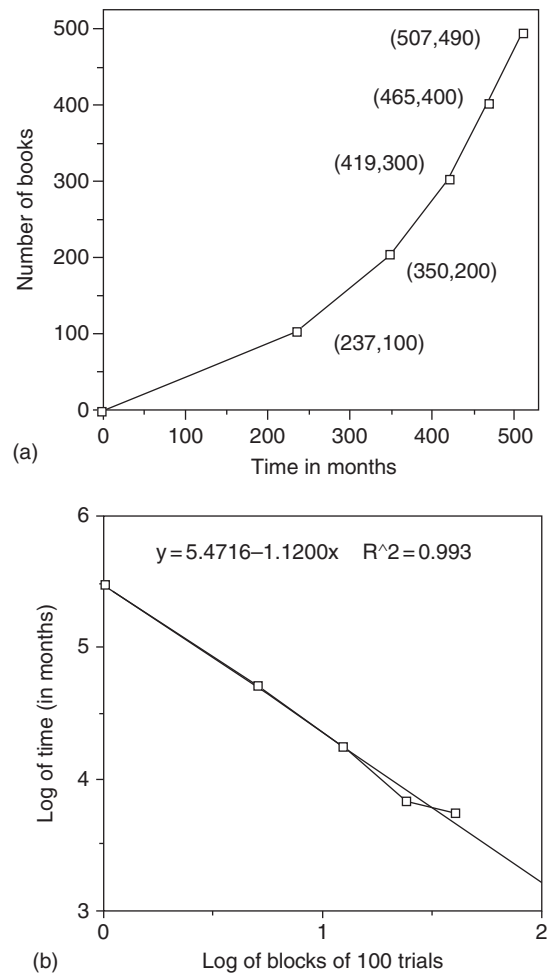


Figure 3 An illustration of the power-law relationship for the development of Professor Asimov's professional writing skills. (a) The number of books Professor Isaac Asimov wrote as a function of time in months. (b) The time to complete 100 books as a function of practice, plotted with logarithmic coordinates on both axes. From Ohlsson, S. (1992). The learning curve for writing books: Evidence of Professor Asimov. *Psychological Science* 3, 380–382. With permission from Wiley-Blackwell.

procedures (rules) get repeatedly applied they become concatenated or chunked together into more compact rules. This mechanism shows how cognitive processing changes from relying on the interpretation and retrieval of declarative knowledge to embedding that knowledge into a set of procedural rules that become more compact with use. The result is a context-specific representation of the skill that can be quickly and efficiently executed.

In sum, research has shown that the acquisition of expert performance requires extended deliberate practice in the domain. Expert knowledge is composed of both declarative and procedural knowledge and research on learning has shown that declarative knowledge can be acquired through multiple cognitive pathways, whereas procedural knowledge comes from the repeated practice of a task. This view suggests that the type and structure of

the learning environment are critical to the acquisition of expert performance. In the final section, we discuss two extensions to the traditional paradigm for research on expertise.

Current Directions

Current research extends the traditional paradigm in a number of ways. In this section, we focus on two: collaborative expertise and using expert–novice differences to determine targets of learning. In recent work, Schunn and colleagues (Tollinger *et al.*, 2006) had the unique opportunity to examine how over 50 NASA scientists worked together to plan the day-to-day operations of the two Mars rovers (Mars Explorer Rover Mission). The scientists' daily task was to analyze the data from the previous day and then come up with a plan for what experiments the rovers would conduct on the next day. They found that the amount of planning decreased across days and followed a learning curve similar to those typically observed for the acquisition of individual expertise, suggesting that expertise can also be acquired at the group level. Initial analyses suggest that the speedup in planning was due to both cognitive factors, such as individual knowledge chunking, plan reuse, and reducing task uncertainty, as well as social factors, such as coordinating information with others and the effect of leadership on the group.

In other recent work, Nokes *et al.* (2006) conducted a laboratory experiment on the effect of expertise on collaborative problem solving. They examined both expert and novice pilots' problem-solving performance when either working alone or with another participant of the same level of expertise. They found that experts working in pairs showed much larger collaborative benefits than novices working together, particularly for complex problem-solving tasks. Analysis of verbal protocols revealed that expert collaborative performance was supported by both domain knowledge (e.g., elaborating each other's contributions) and collaborative skill (e.g., acknowledging and restating the partner's contributions). The pilot and NASA scientist work extends the traditional paradigm and asks how expertise impacts cognitive and social processes at both the individual and the group level.

A second direction focuses on using expertise research to help identify targets of learning for novices. For example, Mestre and colleagues have used some of the classic findings in physics expertise (e.g., Chi *et al.*, 1981) to develop an instructional intervention to help students adopt similar strategies to that of the experts (Dufresne *et al.*, 1992; Mestre *et al.*, 1993). In one study, students were instructed to perform conceptual analyses vis-à-vis a computer interface that was based on the way experts strategize and solve problems, by first identifying the appropriate principles, justifying the use of those principles, and then

articulating the solution procedures. They found that this type of strategizing improved student's conceptual understanding and subsequent problem solving compared to control conditions where students used more traditional approaches to solve problems (e.g., textbook instruction). This research provides one example for how findings from the expertise literature can be used to help improve instructional techniques.

Conclusions

In this article, we reviewed the psychological research on expertise in human problem solving. We saw that expert knowledge impacts each stage of the problem-solving process from problem perception to solution storage. Expert knowledge is composed of both declarative and procedural knowledge and is organized into knowledge structures (e.g., chunks and schemas) that facilitate the categorization and construction of a mental representation of the problem, support the selection of appropriate strategies and procedures, provide constraints to evaluate problem-solving progress, and provide a framework to effectively store new information about the domain. These knowledge structures are acquired through deliberate practice, and learning environments can be designed to facilitate their acquisition. Future work should build upon this rich knowledge base to further advance theories of learning and instruction.

Acknowledgments

Preparation of this paper was supported by Grant SBE0354420 from the Pittsburgh Science of Learning Center and Grant R305B070085 from the Institute of Education Sciences.

See also: Concept Learning; Memory; Metacognition.

Bibliography

- Adelson, B. (1981). Problem solving and the development of abstract categories in programming languages. *Memory and Cognition* **9**, 422–433.
- Akin, O. (1980). *Models of Architectural Knowledge*. London: Pion.
- Anderson, J. R. (1982). Acquisition of cognitive skill. *Psychological Review* **89**, 369–403.
- Anderson, J. R. (1987). Skill acquisition: Compilation of weak-method problem solutions. *Psychological Review* **94**, 192–210.
- Charness, N. (1979). Components of skill in bridge. *Canadian Journal of Psychology* **33**, 1–50.
- Chase, W. G. and Simon, H. A. (1973a). The minds eye in chess. In Chase, W. G. (ed.) *Visual Information Processing*, pp 215–281. New York: Academic Press.
- Chase, W. G. and Simon, H. A. (1973b). Perception in chess. *Cognitive Psychology* **4**, 55–81.

- Chi, M. T. H. (2006). Two approaches to the study of experts' characteristics. In Ericsson, K. A., Charness, N., Feltovich, P. J., and Hoffman, R. R. (eds.) *The Cambridge Handbook of Expertise and Expert Performance*, pp 21–30. Cambridge, UK: Cambridge University Press.
- Chi, M. T. H., Feltovich, P. J., and Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science* **5**, 121–152.
- Chi, M. T. H. and Ohlsson, S. (2005). Complex declarative learning. In Holyoak, K. J. and Morrison, R. G. (eds.) *The Cambridge Handbook of Thinking and Reasoning*, pp 371–399. Cambridge: Cambridge University Press.
- Dufresne, R. J., Gerace, W. J., Hardiman, P. T., and Mestre, J. P. (1992). Constraining novices to perform expertlike analyses: Effects on schema acquisition. *Journal of the Learning Sciences* **2**, 307–331.
- Ericsson, K. A. (2006). Protocol analysis and expert thought: Concurrent verbalizations of thinking during experts' performance on representative tasks. In Ericsson, K. A., Charness, N., Feltovich, P. J., and Hoffman, R. R. (eds.) *The Cambridge Handbook of Expertise and Expert Performance*, pp 223–241. Cambridge, UK: Cambridge University Press.
- Ericsson, K. A. and Kintsch, W. (1995). Long term working memory. *Psychological Review* **102**, 211–245.
- Ericsson, K. A., Krampe, R. T., and Tesch-Romer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review* **100**, 363–406.
- Fitts, P. M. (1964). Perceptual-motor skill learning. In Melton, A. W. (ed.) *Categories of Human Learning*, pp 243–285. New York: Academic Press.
- Galton, F. (1869). *Hereditary Genius: An Inquiry into Its Laws and Consequences*. London: Julian Friedman Publishers.
- Groen, G. J. and Patel, V. L. (1988). The relationship between comprehension and reasoning in medical expertise. In Chi, M. T. H., Glaser, R., and Farr, M. J. (eds.) *The Nature of Expertise*, pp 287–310. Hillsdale, NJ: Erlbaum.
- Howe, M. J. A., Davidson, J. W., and Sloboda, J. A. (1998). Innate talents: Reality or myth? *Behavioral and Brain Sciences* **21**, 399–442.
- Klein, G. (1998). *Sources of Power: How People Make Decisions*. Cambridge, MA: MIT Press.
- Lesgold, A., Rubinson, J., Feltovich, P., et al. (1988). Expertise in a complex skill: Diagnosing X-ray pictures. In Chi, M. T. H., Glaser, R., and Farr, M. J. (eds.) *The Nature of Expertise*, pp 311–342. Hillsdale, NJ: Erlbaum.
- Marshall, S. (1995). *Schemas in Problem Solving*. New York: Cambridge University Press.
- McDermott, J. and Larkin, J. H. (1978). Re-representing textbook physics problems. In *Proceedings of the 2nd National Conference of the Canadian Society for Computational Studies of Intelligence*. Toronto: University of Toronto Press.
- McKeithen, K. B., Reitman, J. S., Reuter, H. H., and Hirtle, S. C. (1981). Knowledge organization and skill differences in computer programmers. *Cognitive Psychology* **13**, 307–325.
- Mestre, J. P., Dufresne, R., Gerace, W. J., Hardiman, P. T., and Touger, J. S. (1993). Promoting skilled problem solving behavior among beginning physics students. *Journal of Research in Science Teaching* **30**, 303–317.
- Morrow, D. G., Menard, W. E., Stine-Morrow, E. A. L., Teller, T., and Bryant, D. (2001). The influence of expertise and task factors on age differences in pilot communication. *Psychology and Aging* **16**, 31–46.
- Moss, J., Kotovsky, K., and Cagan, J. (2006). The role of functionality in the mental representations of engineering students: Some differences in the early stages of expertise. *Cognitive Science* **30**, 65–93.
- Nersessian, N. J. (1992). How do scientists think? Capturing the dynamics of conceptual change in science. In Giere, R. N. (ed.) *Cognitive Models of Science: Minnesota Studies in the Philosophy of Science*, vol. XV, pp 3–44. Minneapolis, MN: University of Minnesota Press.
- Newell, A. and Simon, H. A. (1972). *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Nokes, T. J., Meade, M. L., Morrow, D. G., and Stine-Morrow, E. A. L. (2006). Investigating the effect of domain knowledge on collaborative problem solving. In Sun, R. and Miyake, N. (eds.) *Proceedings of the 28th Annual Cognitive Science Society*, pp 2572. Vancouver, CA: Cognitive Science Society.
- Norman, G. R., Brooks, L. R., and Allen, S. W. (1989). Recall by expert medical practitioners and novices as a record of processing attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **15**, 1166–1174.
- Proctor, R. W. and Dutta, A. (1995). *Skill Acquisition and Human Performance*. Thousand Oaks, CA: Sage.
- Reimann, P. and Chi, M. T. H. (1989). Human expertise. In Gilhooly, K. J. (ed.) *Human and Machine Problem Solving*, pp 161–191. New York: Plenum.
- Silver, E. A. (1979). Student perceptions of relatedness among verbal problems. *Journal of Research in Mathematics Education* **10**, 195–210.
- Simon, D. P. and Simon, H. A. (1978). Individual differences in solving physics problems. In Siegler, R. (ed.) *Thinking: What Develops?* Hillsdale, NJ: Erlbaum.
- Simonton, D. K. (1999). Talent and its development: An emergent and epigenetic model. *Psychological Review* **106**, 435–457.
- Spilich, G. J., Vesonder, G. T., Chiesi, H. L., and Voss, J. F. (1979). Text processing of domain-related information for individuals with high and low domain knowledge. *Journal of Verbal Learning and Verbal Behavior* **14**, 506–522.
- Tollinger, I., Schunn, C. D., and Vera, A. H. (2006). What changes when a large team becomes more expert? Analyses of speedup in the Mars exploration rovers science planning process. *Paper Presented at the Human Computer Interaction Consortium*. Frasier, CO.
- VanLehn, K. (1998). Analogy events: How examples are used during problem solving. *Cognitive Science* **22**, 347–388.
- Voss, J. F. and Post, T. A. (1988). On the solving of ill-structured problems. In Chi, M. T. H., Glaser, R., and Farr, M. J. (eds.) *The Nature of Expertise*, pp 261–285. Hillsdale, NJ: Erlbaum.
- Wineburg, S. (1998). Reading Abraham Lincoln: An expert/expert study in the interpretation of historical texts. *Cognitive Science* **22**, 319–346.

Further Reading

- Chi, M. T. H., Glaser, R., and Farr, M. J. (eds.) (1988). *The Nature of Expertise*. Hillsdale, NJ: Erlbaum.
- Ericsson, K. A. (ed.) (1996). *The Road to Excellence: The Acquisition of Expert Performance in the Arts and Sciences, Sports, and Games*. Mahwah, NJ: Erlbaum.
- Ericsson, K. A., Charness, N., Feltovich, P. J., and Hoffman, R. R. (eds.) (2006). *The Cambridge Handbook of Expertise and Expert Performance*. Cambridge, UK: Cambridge University Press.
- Ericsson, K. A. and Smith, J. (1991). Prospects and limits of the empirical study of expertise: An introduction. In Ericsson, K. A. (ed.) *Toward a General Theory of Expertise*, pp 1–38. Cambridge, UK: Cambridge University Press.