



Are Cognitive Skills Context-Bound?

Author(s): D. N. Perkins and Gavriel Salomon

Source: *Educational Researcher*, Vol. 18, No. 1 (Jan. - Feb., 1989), pp. 16-25

Published by: [American Educational Research Association](#)

Stable URL: <http://www.jstor.org/stable/1176006>

Accessed: 17/03/2011 17:13

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=aera>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



American Educational Research Association is collaborating with JSTOR to digitize, preserve and extend access to *Educational Researcher*.

<http://www.jstor.org>

Are Cognitive Skills Context-Bound?

D. N. PERKINS

GAVRIEL SALOMON

Once upon a time, an astute and beneficent leader in a remote country anticipated increasing aggression from a territory-hungry neighbor nation. Recognizing that the neighbor had more military might, the leader concluded that his people would have to out-think, rather than overpower, the enemy. Undistinguished in its military armament and leadership, the country did have one remarkable resource: the reigning world chess master, undefeated for over twenty years. "Aha," the leader said to himself, "we will recruit this keen intellect, honed so long on the whetstone of chess, teach him some politics and military theory and then outmaneuver the enemy with the help of his genius."

A fanciful tale, to be sure, but consider the leader's plan for a moment: Is it disastrously naive, possibly helpful, or a pretty good bet? In fact, the tale has no definite conclusion. Rather, it is the beginning of another tale, a tale about psychology and the human intellect that research is gradually spinning. Questions like that of the chess master's political and military potential stand in the center of one of the most puzzling and important issues that cognitive psychology has addressed: the roles of general and of context-specific knowledge in thinking.

Within the discipline of psychology and across three decades, very different voices might be heard sizing up the chances of the chess master. One says, "Basically, the chess master plays chess well because he knows the moves of the game well. There's no reason at all why that knowledge should carry over powerfully to political or military matters." Another voice counters: "Well, there are analogies to be mined between chess and matters of political and military strategy. Control of the center, for example—that's a principle important in chess, but also in politics and war." Still another voice emphasizes not the transferrable aspects of chess skill but general problem solving abilities: "Above all, a chess player is a problem solver, needing to plan ahead, explore alternatives, size up strategic options, just as a politician or military tactician does. So we might expect a lot from the chess master."

Which of these voices speaks with most authority? Or do we need to listen for another voice altogether, stating some more complicated opinion? On this question has hung a good deal of psychological research, as psychologists have

Effective problem solving, sound decision making, insightful invention—do such aspects of good thinking depend more on deep expertise in a specialty than on reflective awareness and general strategies? Over the past thirty years, considerable research and controversy have surrounded this issue. An historical sketch of the arguments for the strong specialist position and the strong generalist position suggests that each camp, in its own way, has oversimplified the interaction between general strategic knowledge and specialized domain knowledge. We suggest a synthesis: General and specialized knowledge function in close partnership. We explore the nature of this partnership and consider its implications for educational practice.

sought to understand the factors that underlie cognitive skills in domains like chess play, problem solving in mathematics and physics, medical diagnosis, musical composition, and more. Let us see how that story has unfolded and, at the end, appraise the chess master's chances.

The Heart of the Issue

Some sharpening of the problem is needed at the outset. At issue is the generality of cognitive skill. Is skillful

thought-demanding performance relatively context-bound, or does it principally reflect use of general abilities of some sort?

There can be little doubt that some aspects of cognitive skill are quite general: IQ and *g* for general intelligence measure a side of human intellectual functioning that correlates with effective performance over a wide range of academic and nonacademic tasks. For this aspect of cognitive skill, the answer is in, and favors generality. By way of qualification, however, arguments can be made that giftedness in particular domains, such as music, reflects neurologically based, relatively inborn aspects of intelligence (Gardner, 1983). These arguments have been somewhat controversial but certainly have a considerable following.

At any rate, neither *g* nor any more specialized aspect of giftedness speaks directly to the role of *knowledge* in intellectual functioning. And it's obvious that knowledge counts for a lot. Without considerable experience, the most gifted individual cannot play chess, repair a car, play the violin, or prove theorems. Indeed, recent research on *g* argues that it wields its influence on performance *by way of* knowledge: People with high *g* tend to perform well because they have a rich knowledge base, the direct determinant of performance (Hunter, 1986). And people with a lower *g* but more knowledge than those with high *g* will usually perform better—it's the knowledge that counts, rather than *g*.

The question is, which kind of knowledge counts most—

DAVID PERKINS is at the Harvard University Graduate School of Education, Cambridge, Massachusetts 02138. GAVRIEL SALOMON is at the Department of Communication, University of Arizona, Tucson, Arizona 85721.

general knowledge of how to think well, or specific knowledge about the detailed ins and outs of a field? General knowledge includes widely applicable strategies for problem solving, inventive thinking, decision making, learning, and good mental management, sometimes called *autocontrol*, *autoregulation*, or *metacognition*. In chess, for example, very specific knowledge (often called local knowledge) includes the rules of the game as well as lore about how to handle innumerable specific situations, such as different openings and ways of achieving checkmate. Of intermediate generality are strategic concepts, like *control of the center*, that are somewhat specific to chess but that also invite far-reaching application by analogy.

There is an obvious partial answer to the question, "what counts most?" It's plain that some local knowledge is necessary; one can't play chess without knowing the rules of the game, after all. But that partial resolution misses the real issue—*where is the bottleneck in attaining mastery?* Does it lie in acquiring a deep and detailed knowledge of chess, whereas anyone can learn whatever general thinking strategies are needed? Or does it lie in becoming reflective and cultivating the general thinking strategies, whereas anyone can learn the relevant particulars of the game?

These different theories write different endings to the chess master's story. If he is masterful in virtue of his general savvy about the use of his mind, the chess master might carry it over to the political and military realms. At the opposite extreme, if his mastery depends on richly developed local knowledge of chess, the chances seem slender.

Such enigmas arise in every domain and bring with them fundamental questions about educational design. Should we teach entirely for richly developed local knowledge, subject matter by subject matter? Or should we invest a significant portion of educational resources in developing general skills of problem solving, self-management, and so on? Or, indeed, does this dichotomy obscure some important factors? To work toward an answer, let us examine the controversy, adopting a broad-stroke historical perspective without pretense of reconstructing events in detail.

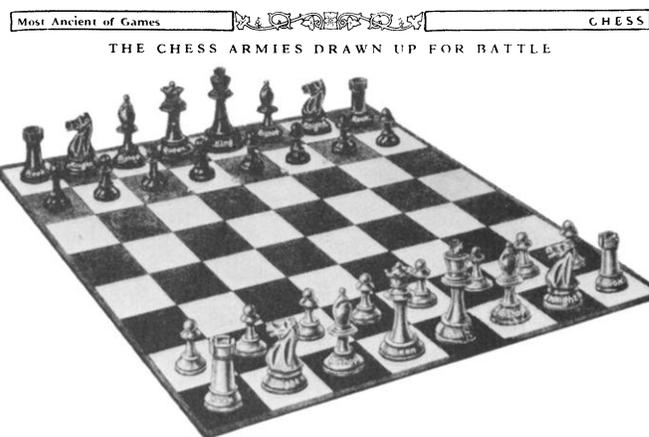
Before the Fall: The Golden Age of General Heuristics

Thirty years ago, it was widely thought that good problem solving and other intellectual performances reflected general strategies (supported by *g*) operating on whatever database of knowledge happened to be needed. True ability resided in the general strategies, with the database an incidental necessity.

One source of this perspective was the mathematician Gyorgy Polya's analysis of mathematical problem solving (Polya, 1954, 1957). Polya argued that the formalities of mathematical proof and derivation had little to do with the real work of problem solving in mathematics. Although such formalities were the evening dress of journal publication, success in finding solutions depended on a repertoire of *heuristics*, general strategies for attacking a problem that did not guarantee a solution, but often helped. Polya discussed such heuristics as breaking a problem into subproblems, solving simpler problems that reflected some aspect of the main problem, using diagrams to represent a problem in different ways, and examining special cases to get a feel for a problem. Polya spoke to mathematical problem solving specifically, but many of the heuristics he emphasized were plainly applicable to problems of all sorts, which encour-

aged the notion that problem solving could be viewed as a general ability and mathematical problem solving simply a special case.

Another source of encouragement was early work on Artificial Intelligence (AI), the design of computer programs to carry out processes such as chess playing or theorem proving that, in a human being, would be considered intelligent. The "General Problem Solver" was one outstanding example (Ernst & Newell, 1969; Newell & Simon, 1972). Developed around 1957 by Alan Newell, J. P. Shaw, and Herbert Simon, this program relied on a flexible heuristic called *means-end analysis*. Input to the program included information about a beginning state, an end state (the goal)



Source *Compton's Pictured Encyclopedia* © 1924

and allowable operations on states, all in a compact notation. Many simple puzzles and problems in logic could be cast into this form. The program pursued a chain of operations for transforming the beginning state into the end state. It did so by comparing and contrasting the beginning with the end state and seeking an operation that would reduce the contrast—a means that would bring the beginning state closer to the end state. After executing that operation, the program would seek another operation to reduce the contrast yet further, and so on. If it encountered a cul-de-sac that forbade further progress, the program would back up and try another path. There were other sophisticated features as well. Here again, as in the perspective of Polya, it appeared that problem solving power lay in some rather general principles, systematically applied to whatever the relevant database of knowledge happened to be.

A host of factors—its generating interesting data; its accord with intuitions about the value of analytic ability; its economy, elegance, and availability to testing in computer models—reinforced the position that good thinking depended in considerable part on a repertoire of rather general heuristic knowledge. Many such heuristics were identified, heuristics for problem solving, memorizing, inventive thinking, decision making, general mental management, and so on (cf. Nickerson, Perkins, & Smith, 1985). As to local knowledge, the part of knowledge specific to a domain like chess or mathematics, it was thought not very important. Of course, one had to have it. But there really wasn't much to it beyond a few rules in the case of chess,

a few axioms in the case of a mathematical system, and so on. There didn't seem to be *enough to know* about such databases to make them central to thinking ability.

The Power of the Particular

The golden age could not last. Even then, certain results in the literature gave warning that all was not well with this picture of general heuristics driving intellectual performance. In the years to come, a wave of compelling findings would cast profound doubt on the centrality of general ability in human thinking, particularly ability based on heuristics. The gathering force of contrary findings falls neatly into three parts: the argument from expertise, the argument from weak methods, and the argument from transfer.

The argument from expertise. Investigators even during the golden age were discovering that the seeming smallness of the database demanded by chess, symbolic logic, and other favorite areas of research was deceptive. To be sure, chess, for example, looked like a game of general reasoning applied to a few specific rules. It seemed that all a player needed to do was to know the rules and reason well about options and consequences: "If I move there, my opponent might move there, but then I could. . . , but then my opponent could. . .", and so on. But close observation showed that there was much more to it.

Research on the games of grand master chess players showed that their tactics depended on an enormous knowledge base of important patterns of chess pieces—not only the standard patterns such as pins, forks, and rooks on open files, but far more, with a diversity and complexity not recognized by the chess masters themselves. Expert chess players reasoned about the game using these chunk-like configurations, rather than thinking about one piece at a time, and so had much more power to think ahead and devise strategies than a simple command of the rules would afford.

The classic experiments demonstrating this began with examinations of the reputed ability of grand master players to memorize the layout of pieces on a chess board at a glance (Chase & Simon, 1973; de Groot, 1965). The experiments showed that experts could indeed do this—but only if the chess pieces' positions had emerged in the natural course of play, not if the same pieces were arranged randomly on the board. Beginning players did just as well as the grand masters in recalling random layouts and, significantly, their recall did *not* improve on the layouts that emerged in the course of a game. These results showed that the grand masters knew something very powerful, but *very specific to chess*, else they would have done well on the random layouts too. Chase and Simon (1973) used certain approximations to estimate the grand master chess player's repertoire of something like fifty thousand chess-specific configurations, or *schemata*, as they are usually called, that provide the "chunks" that the grand master thinks with.

The experiments in chess inspired similar studies in a number of areas, with parallel findings. A general profile of expertise began to emerge (cf. Glaser, 1984; Rabinowitz & Glaser, 1985): Expert performance entailed (a) a large knowledge base of domain-specific patterns (for example, typical configurations of pieces in chess, typical uses of conservation laws in physics); (b) rapid recognition of situations where these patterns apply; and (c) reasoning that moves

from such recognition directly toward a solution by working with the patterns, often called *forward reasoning*.

In contrast, novices tended not to see the relevant patterns, because they did not know them or lacked rapid recognition-like access to them. Novices often based their reasoning on superficial problem content, for instance treating inclined plane problems similarly when different physics principles applied. Novices often solved problems by focussing first on the unknown and seeking equations or rules that bridged back from the unknown toward the givens. If they found equations or rules, then they plugged in the givens to determine the unknown. This *backward reasoning* ran opposite to experts' forward reasoning from givens toward the unknown. These contrasts between experts' and novices' performances emerged in such domains as physics problem solving (e.g. Chi, Feltovich, & Glaser, 1981; Larkin, 1982; Larkin, McDermott, Simon, & Simon, 1980a; Larkin, McDermott, Simon, & Simon, 1980b), mathematical problem solving (Schoenfeld & Herrmann, 1982), computer programming (Ehrlich & Soloway, 1984), and medicine (Elstein, Shulman, & Sprafka, 1978; Patel & Groen, 1986). The investigations came to be known as research on expertise, because the account of proficient performances was so compelling.

These studies of expertise revealed the naivete in a key premise of the golden age. To be sure, chess, symbolic logic, Newtonian physics, and so on, each involved a fairly parsimonious foundation of basic rules or axioms. Nonetheless, experts depended on a much richer database, an elaborate superstructure of ramifications erected on top of the parsimonious foundation. General heuristics appeared to be no substitute for the rich database of ramifications, stored in memory, accessed by recognition processes, and ready to go. Indeed, the broad heuristic structure of expert as contrasted to novice problem solving—the reasoning forward rather than reasoning backward—seemed attributable not to any heuristic sophistication on the part of the experts, but to the driving influence of the experts' rich database. General heuristics no longer looked as central or as powerful.

The argument from weak methods. Work in AI, although it initially supported the idea that general heuristics drive skillful problem solving, also began to take a different turn. To be sure, programs like the General Problem Solver could solve some rather simple formal problems, such as those in elementary symbolic logic. But these generic programs seemed quite helpless in complex problem solving domains such as chess play, integrating mathematical expressions, or medical diagnosis. In contrast, programs designed specifically for those knowledge domains scored significant successes (Boden, 1977; Rich, 1983). In the late 1960s and early 1970s, the AI community became increasingly aware of these successes, and many investigators began to lay their bets differently as they tried to construct powerful artificial intelligence systems (Gardner, 1985, pp. 160–161).

Investigators in the AI community came to refer to general heuristics such as means-end analysis as *weak methods* (see Rich, 1983, section 3.6). When new to a domain, all a computer or a human could do was deploy weak methods that turned out weak results. Real power in problem solving emerged over time, as application of weak methods created the opportunity to learn and store up the ramifications of particular moves in the domain and build the rich database.

This database would become the real power behind good problem solving, leaving the weak methods behind. Investigators spoke of the "power-generality tradeoff," the more general the method, the weaker the method. Seeking to make the best of the situation and taking a cue from the work of psychologists on expertise, many AI researchers turned to developing *expert systems*, which sought to simulate the intelligence of an expert in a domain through manipulating a massive domain-specific knowledge base in areas such as medical diagnosis (cf. Rich, 1983; Wenger, 1987).

Although the argument from weak methods derived principally from AI, little happened in the psychological community to make a counterexample. In particular, a number of investigators sought to teach Polya's heuristics for mathematical problem solving with little success. Students exhibited just exactly the difficulties expected, given the results of the research on expertise: They didn't know *what to do* with the heuristics. They understood the heuristics in broad terms but didn't seem to understand the mathematics well enough to apply them in the rather complex and context sensitive ways required. Local knowledge, more than general problem-solving heuristics, appeared to be the bottleneck (Schoenfeld, 1985, pp. 71-74).

The argument from transfer. A third line of argument seemed to drive the last nail in the coffin of general cognitive skills. According to the premise of the "golden age," much of the knowledge acquired in a particular domain is inherently general, at least implicitly, and should lead to transfer to other areas. Thus, learning the logic imbedded in mathematics or in Latin should, for example, yield improved scores on standard IQ tests or better learning in other seemingly unrelated fields. Similarly, learning to program computers in a powerful language such as LOGO should improve students' reasoning and planning abilities.

A variety of studies, initiated as far back as the turn of the century, generally failed to uphold these predictions. E. L. Thorndike (e.g., 1923) and Thorndike and Woodworth (1901) reported experiments, some on a large scale, showing that training in such fields as Latin and math has no measurable influence on other cognitive functions, thus dispelling a then prevalent belief in the training of the mind's "faculties."

More recent studies have yielded similar findings. Such studies suggest, for example, that training on one version of a logical problem has little if any effect on solving an isomorphic version, differently represented (Hayes & Simon, 1977); that becoming literate with no schooling does not improve mastery of general cognitive skills (Scribner & Cole, 1981); or that teaching children to use general, context-independent cognitive strategies has no clear benefits outside the specific domains in which they are taught (for a summary, see Pressley, Snyder, and Cariglia-Bull, 1987). Findings from research on the cognitive effects of programming have generally been negative (Pea & Kurland, 1984; Salomon & Perkins, 1987).

Overall, research on transfer suggests the same conclusion as the arguments from expertise and weak methods: Thinking at its most effective depends on specific, context-bound skills and units of knowledge that have little application to other domains. To the extent that transfer does take place, it is highly specific and must be cued, primed, and

guided; it seldom occurs spontaneously. The case for generalizable, context-independent skills and strategies that can be trained in one context and transferred to other domains has proven to be more a matter of wishful thinking than hard empirical evidence (Pressley et al., 1987).

The Skeleton of a Synthesis

We said that the argument from transfer might be the last nail in the coffin of general cognitive skills. But the skeleton is restless. Some people seem generally smart—not just knowledgeable, but insightful no matter the subject.

For instance, if you have mixed some with academic philosophers, you may have noticed that they have an unsettling habit: You mention some casual claim, and they often smack you with a counterexample. Moreover, the discussion does not have to deal with a topic in academic philosophy. You may be discussing politics, family life, the dangers of nuclear power plants, or the latest best-seller. It almost seems as though the philosophers have a general cognitive skill: the strategy of looking for counterexamples to test claims.

Is this a general cognitive skill? Recalling the arguments from expertise, weak methods, and transfer, you might object this way: "What has the appearance of a general reasoning strategy in these philosophers' remarks is really a highly contextualized strategy. The philosophers can only construct counterexamples in domains where they have a good knowledge base."

"More than that," your objection might continue, "certain domains bring with them special criteria for what counts as a counterexample. A counterexample to a mathematical claim would have to be constructed appropriately from the premises of the mathematical system; a counterexample to a legal claim would have to be the result of prior due process. This is a special case of a point that Toulmin (1958), among others, has emphasized: Different domains share many structures of argument, but bring with them somewhat different criteria for evidence."

These points ought to be granted at once. Yet there is something disturbing about casting them as an objection: They treat *general* and *contextualized* as though they were exclusive of one another. The heart of the synthesis we would like to suggest challenges this dichotomy. There *are* general cognitive skills; but they always function in contextualized ways, along the lines articulated in considering the philosophers' habit of mind (cf. Perkins & Simmons, 1987; Perkins, Schwartz, & Simmons, in press).

Granting the need for contextualization through a knowledge base, what argues that the philosophers' move of seeking counterexamples is nonetheless a general, learnable, and worthwhile cognitive skill? First of all, seeking counterexamples is a strategy for which philosophers show *seeming use*: That is, it certainly looks as though they are applying a general strategy, although perhaps their thinking is entirely contextualized and only appears to be general. Second, the seeking of counterexamples itself appears to play an *important role* in the philosophers' reasoning: It allows them to detect the flaws in claims that otherwise might be missed. Third, the seeking of counterexamples seems to be *transferable*: Apparently, philosophers pick it up from their philosophical studies and apply it widely to other domains. Fourth, the move of seeking counterexamples is *commonly absent*: Everyday experience suggests that most people do

not reflexively seek counterexamples. Moreover, research on everyday reasoning shows that seeking any sort of evidence on the other side of the case is a relatively rare move, even in educated populations (Perkins, 1985; Perkins, in press).

Of course, this is only one case, informally argued through everyday observation, and subject to several objections. Its real purpose is not to mount a compelling argument for general cognitive skills but to illustrate what a general cognitive skill might look like—and rattle the skeleton that met an early and unceremonious end. To flesh out that skeleton, we would have to find patterns of information processing that (a) show seeming use, (b) play an important role, (c) are demonstrably transferrable, and (d) are commonly absent. It would be reasonable to call a pattern of information processing that satisfies those conditions a general cognitive skill.

Generality on the Rebound

Throughout the period of “the fall,” considerable interest, in some quarters, continued to focus on the nature of general cognitive skills and the potentials of teaching such skills. In recent years, results have begun to emerge that challenge the picture of expert performance as driven primarily by a rich knowledge base of highly context-specific schemata. One by one, the arguments from expertise, weak methods, and transfer have begun to show cracks.

We take those arguments up again, reexamining each in light of new findings. In doing so, we enter the region of recent history and contemporary work, where results are scattered, and replications are few. Nonetheless, from our perspective, a new outline is emerging.

When experts face unfamiliar problems. Most of the research on expertise has examined experts addressing standard problems in a domain—typical chess positions, physics problems, programming problems, and so on. In these circumstances, the experts’ behavior appears to be strongly driven by local knowledge. But this picture could be misleading. What happens when experts tackle atypical problems—not problems outside the domain, but problems less “textbookish?” Might more general kinds of knowledge play a more prominent role?

One response to this question would be to dismiss it from the outset. What does it matter how experts respond to atypical problems? Expertise certainly should be assessed and examined with problems typical of the domain. But such a response takes a narrow view of expertise. Presumably, in many domains, people become experts not to function as technicians solving new variants of the classic problems but to open the field further. From this standpoint, atypical problems are just the right test of truly flexible expertise.

John Clement, working at the University of Massachusetts, Amherst, has examined experts’ responses to atypical problems. The results are provocative (Clement, 1982, in press; see also Johnson, Ahlgren, Blount, & Petit, 1980). As in other work on expertise, the experts addressing such problems certainly use their rich physics knowledge base, trying to see the deep structure of the problem and deploying principles like conservation of energy. But, because these unusual problems do not yield to the most straightforward approaches, the experts also apply many general strategies.

For example, the experts faced with an unfamiliar problem will often: (a) resort to analogies with systems they understand better; (b) search for potential misanalogies in the analogy; (c) refer to intuitive mental models based on visual and kinesthetic intuition to try to understand how the target system would behave; (d) investigate the target system with “extreme case” arguments, probing how it would work if various parameters were pushed to zero or infinity; (e) construct a simpler problem of the same sort, in hopes of solving that and importing the solution to the original problem.

There are just a few of the Polya-like strategies that seem to appear in Clement’s protocols; no doubt others could be identified as well. Such results suggest that a number of general heuristics not apparent when experts face typical problems play a prominent role when experts face atypical problems.

How does Clement’s evidence speak to the four conditions needed for the synthesis: seeming use, important role, transferrable, and common absence? These studies give distinct evidence of seeming use of heuristics. They also give clear evidence of an important role: In the protocols, the heuristics often constitute crucial steps along a subject’s path to a solution.

However, Clement’s studies offer no evidence of transfer. Although it may seem plausible that a problem solver acquainted with, let us say, extreme case arguments from physics would sometimes carry them over to chemistry or mathematics problems, that remains to be shown. Also, Clement’s studies give no direct evidence of common absence. It’s plausible that many weaker students of physics fail to pick up the “extreme case” pattern of argument simply from normal learning in the domain, but, again, that remains to be shown. (These points are not criticisms of Clement’s studies, which were designed to address other issues.)

It’s notable that the general heuristics seemingly used by Clement’s subjects certainly do not substitute for domain knowledge. On the contrary, the general heuristics operate in a highly contextualized way, accessing, and wielding sophisticated domain knowledge. In particular, conservation of energy, conservation of momentum, and other deep structure principles of physics are brought to bear, held in the pincers, so to speak, of these general heuristics.

When weak methods work. Recall that the argument from weak methods complained that general heuristics appeared not to work very well, either in instructional experiments or in AI. The years have brought changes in that appraisal. Mathematician and educator Alan Schoenfeld, in extensive work on teaching mathematical problem solving, has demonstrated that heuristic instruction can yield dramatic gains in college students’ mathematical problem solving (Schoenfeld, 1982; Schoenfeld & Herrmann, 1982). Schoenfeld emphasizes that this success requires teaching many of the heuristics in a very contextualized way, so the heuristics make good contact with students’ knowledge base in the domain (Schoenfeld, 1985, Chapter 3). At the same time, an important thrust of Schoenfeld’s approach is fostering a seemingly quite general level of control or problem management. Students learn to monitor and direct their own progress, asking questions such as, “What am I doing now,” “Is it getting me anywhere,” “What else could I be doing instead?” This general metacognitive level helps

students to avoid perseverating in unproductive approaches, to remember to check candidate answers, and so on.

Again, it's worth asking how Schoenfeld's work speaks to the four elements of the synthesis position. His research gives evidence of students' use of Polya-like heuristics (Schoenfeld, 1985), and Schoenfeld's experiments demonstrated that students indeed acquired the use of these heuristics (Schoenfeld, 1982; Schoenfeld & Herrmann, 1982). Regarding the heuristics' important role, Schoenfeld's studies also demonstrated that the better performance of the students on posttesting directly depended on their active use of the heuristics. Protocol analysis disclosed that those students who used the heuristics performed better, whereas those who did not failed to show gains. Regarding common absence, the impressive gains students exhibited showed that they did not already possess the heuristics they acquired.

However, Schoenfeld's studies offer no evidence of transfer. To be sure, many of the Polya-like heuristics seem straightforwardly applicable to chemistry and physics, and the general problem management strategy seems even more widely relevant. However, such observations do not make the empirical case that these skills can be decontextualized and applied more broadly.

Another recent effort has focussed on teaching general reading skills to poor readers. Palincsar and Brown (1984) developed and evaluated a method called *Reciprocal Teaching* that through a process of modeling, guiding, and group participation, has helped young poor readers learn to monitor and direct their reading. The intervention encourages and refines four key cognitive activities: *questioning* about the main points of a paragraph, *clarifying* to try to resolve difficulties of understanding, *summarizing* to capture the essence of a text, and *predicting*, to forecast what might happen next in the text. Palincsar and Brown's approach yielded dramatic gains in the students' reading comprehension, transfer to in-class reading in science and social studies, and long-term retention.

The direct cultivation of these reading strategies and the resultant gains give evidence of seeming use, important role, and common absence (in the poor reader population). The matter of transfer can be looked at in two ways: One might say that transfer across domains was demonstrated, because students showed reading gains *in situ* in school subject-matters. Or, one might say that transfer across domains was not addressed, because students were taught and tested on reading performance specifically. The two perspectives seem equally defensible because reading is what might be called a "tool domain," like writing or arithmetic: We learn reading, writing, and arithmetic in order to apply them to various content domains, such as literature, history, or biology.

It's worth noting that the very existence of tool domains that enhance thinking and learning in content domains, in itself, constitutes evidence for general cognitive skills of a sort. Reading is a general cognitive skill, which people routinely transfer to new subject matters, beginning to read in a domain with their general vocabulary and reading tactics and, as they go along, acquiring new domain-specific words, concepts, and reading tactics. However, reading as a general cognitive skill does not much resemble those skills that have been at issue over the past thirty years—such

strategy-like skills as Polya's heuristics, for instance.

There have been several other seemingly successful efforts to teach cognitive skills of some generality in recent years, for example, the development and testing of Project Intelligence, a general course to teach skills of problem solving, decision making, inventive thinking, and other sorts (Herrnstein, Nickerson, Sanchez, & Swets, 1986) and the *guided design* perspective developed by Wales and his colleagues (Wales & Nardi, 1984; Wales & Stager, 1978). A general resource reviewing many such programs is Nickerson et al. (1985). The collection edited by Segal, Chipman, and Glaser (1985) offers somewhat earlier assessments of several programs. Resnick (1987) has authored a monograph appraising the promise of work in this area, with cautiously optimistic conclusions. Likewise, there is considerable evidence from more basic investigations that learning in human beings depends on the deployment of general learning strategies (e.g., recent findings by Bereiter & Tinker, 1988; Chan & Burtis, 1988; Ng, 1988; Ogilvie & Steinbach, 1988).

In artificial intelligence, although work has continued on expert systems, investigators have also returned to the challenge of producing more general models of mind. Two systems in particular, ACT* (Anderson, 1983) and SOAR (Laird, Rosenbloom, & Newell, 1984) have been developed as general models of cognitive processing. Both are learning systems that learn by trying to solve problems. Given a new class of problems, they commence by applying week methods. As they work, they search for and store shortcuts in the solution process and so gradually build up a repertoire of domain-specific chunks, much as human beings do, through extended experience in a domain.

Moreover, the functioning of SOAR, for example, is in some ways not unlike the functioning of Clement's physicists facing an unfamiliar kind of physics problem. SOAR tries the specific moves compiled into its library through experience. But, if SOAR encounters an "impasse," as it is called, in which the specialized techniques it has compiled do not work, it resorts to more general methods.

Important instructional applications are being built that continue this AI tradition. An example is GUIDON 2 (Clancey, 1986, 1987), a medical diagnostic expert system that combines specific medical expertise with more general reasoning strategies. The latter teach tactics for the management of diagnostic hypotheses whereby cases can be grouped and then differentiated into finer categories. The heuristics used are quite general and applicable to other domains of problem solving requiring heuristic classification. It is expected that through interaction with the program, students might become better problem solvers in medicine as well as in other domains.

ACT*, SOAR, and GUIDON 2 represent a provocative re-engagement with issues concerning the interaction between general and local knowledge.

When transfer happens. During the fall, negative findings on transfer generally were interpreted as showing that skill depends mostly on local knowledge and that we have little ability to decontextualize knowledge and apply it in different domains. However, a more careful examination of the research discloses that the findings that support these conclusions allow other explanations altogether. These other explanations accord general knowledge more potency, with-

out challenging the idea that local knowledge has great importance.

A casual look at the research on transfer might suggest that our cognitive apparatus simply does not incline very much to transfer. But this would be a misapprehension. On the contrary, when faced with novel situations, people routinely try to apply knowledge, skills, and specific strategies from other, more familiar domains. In fact, people commonly ignore the novelty in a situation, assimilating it into well-rehearsed schemata and mindlessly bringing to bear inappropriate knowledge and skill, yielding negative transfer (Langer, in press). In other cases, although people fail to apply purely logical, abstract, or syntactical rules to formally presented problems (e.g., Wason, 1966), they clearly do employ analogous inferential rules to more everyday versions of such problems (e.g. Cheng & Holyoak, 1985).

Moreover, recent research shows that, when general principles of reasoning are taught together with self-monitoring practices and potential applications in varied contexts, transfer often is obtained (e.g. Nickerson, et al., 1985; Palincsar & Brown, 1984; Schoenfeld, 1978, 1982; Schoenfeld & Herrmann, 1982). Relatedly, Lehman, Lempert, and Nisbett (1988) have recently demonstrated that graduate students in such fields as psychology and medicine show clear transfer of probabilistic and methodological reasoning to everyday problems.

Brown and her associates (e.g., Brown & Kane, 1988; Brown, Kane, & Long, in press; Brown & Palincsar, in press) have recently shown in a series of laboratory and classroom studies that transfer to new problems does take place, even among three- and four-year-olds, when (a) learners are shown how problems resemble each other; (b) when learners' attention is directed to the underlying goal structure of comparable problems; (c) when the learners are familiar with the problem domains; (d) when examples are accompanied with rules, particularly when the latter are formulated by the learners themselves; and perhaps most importantly, (e) when learning takes place in a social context (e.g., reciprocal teaching), whereby justifications, principles, and explanations are socially fostered, generated, and contrasted. It becomes evident from this research, as it does from that of others (e.g., Gick & Holyoak, 1987), that transfer is possible, that it is very much a matter of how the knowledge and skill are acquired and how the individual, now facing a new situation, goes about trying to handle it. Given appropriate conditions, such as cueing, practicing, generating abstract rules, socially developing explanations and principles, conjuring up analogies (e.g., Strauss, 1987), and the like, transfer from one problem domain to another can be obtained. General skills and bits of knowledge taught within a specific context can become transferable.

Specifically, we have proposed two different mechanisms by which transfer of specific skill and knowledge takes place (Perkins & Salomon, 1987; Salomon & Perkins, in press). One mechanism, called the "low road" to transfer, depends on extensive and varied practice of a skill to near automaticity (see also Anderson, 1983, on automaticity). A skill so practiced in a large variety of instances becomes applied to perceptually similar situations by way of response or stimulus generalization. For example, having driven different cars under a variety of conditions allows us to shift to driving a truck fairly easily. Unfortunately, learning in many natural

settings and in many laboratory experiments does not meet the conditions for low road transfer: *much practice, in a large variety of situations, leading to a high level of mastery and near-automaticity.* For example, these conditions were not met by the Vai literates studied by Scribner and Cole (1981), or the young programmers studied by Pea and Kurland (1984).

The second mechanism, called the "high road," depends on learners' deliberate mindful abstraction of a principle. People sometimes abstract principles in advance, keeping them in mind in anticipation of appropriate opportunities for application, or, in a new situation, reach back to prior experiences and abstract from them principles that might be relevant. For an example of the latter, in a recent partial replication of Gick and Holyoak's (1987) analogy studies, Salomon and Globerson (1987) showed that college students who were urged to formulate an abstract principle from two problems did not show more transfer to a new, analogous problem than students who were given the principle ready-made. However, the former (but not the latter) showed impressive transfer when urged to search their memories for an appropriate principle that they may have encountered before.

Likewise, the expert chess player mobilized to save his country in our opening story would be expected to mine the context of chess for chess-bound principles such as "get hold of the board's center," decontextualize them, and apply them in forms like "let's capture or destroy the enemy's command centers." Unfortunately, in many real-world situations and many laboratory experiments on transfer, there is nothing to provoke the active decontextualization of knowledge, so the high-road mechanism does not operate.

But it can be activated. The transfer findings of Lehman and his colleagues (1988) are a case in point. In the treatment employed, graduate students did not just absorb statistical principles and practice them to near automaticity; they were urged to comprehend the logic behind them and mindfully generate abstractions, applying them in a variety of learning situations. Similarly, Salomon, Globerson, and Guterman (Salomon, 1988) have found that children can acquire reading strategies involving self-monitoring from a computerized Reading Aid and apply them a month later to essay writing, a clear case of high-road transfer of a generalized ability (see also Brown & Palincsar, 1988).

In summary, recent research and theorizing concerning transfer put the negative findings cited earlier in a different light. These findings do not imply either that people have little ability to accomplish transfer or that skill is almost entirely context bound. Rather, the negative results reflect the fact that transfer occurs only under specific conditions, which often are not met in everyday life or laboratory experiments (Brown, Kane, & Long, in press). When the conditions *are* met, useful transfer from one context to another often occurs.

So Are Cognitive Skills Context-Bound?

As the psychological tale has unfolded, the answer to the question looks to be, "Yes and no." The tale is one of neglected complexities. Early advocacy of general cognitive skills overlooked the importance of a rich knowledge base, took it for granted that general heuristics would make ready contact with a person's knowledge base, and had few worries about transfer, which was supposed to happen more

or less spontaneously. Mistakes all three, these oversights led to considerable skepticism about general cognitive skills, the view that cognitive skills in the main were context bound, and interesting developments in the psychology of expertise as well as artificial intelligence work on expert systems.

But more recent results suggest that this trend had its blind spots too, in neglecting how general heuristics help when experts face atypical problems in a domain, how general heuristics function in contextualized ways to access and deploy domain specific knowledge, and how lack of conditions needed for transfer, rather than domain specificity, is to blame for many cases of failure of transfer. These more recent results point toward the synthesis that we now think might be fleshed out.

What general cognitive skills are like. In the synthesis, general cognitive skills do not function by somehow taking the place of domain-specific knowledge, nor by operating exactly the same way from domain to domain. Rather, cognitive skills are general tools in much the way the human hand is. Your hands alone are not enough; you need objects to grasp. Moreover, as you reach for an object, whether a pen or a ball, you shape your hand to assure a good grip. And you need to learn to handle different objects appropriately—you don't pick up a baby in the same way you pick up a basket of laundry.

Likewise, general cognitive skills can be thought of as general gripping devices for retrieving and wielding domain-specific knowledge, as hands that need pieces of knowledge to grip and wield and that need to configure to the kind of knowledge in question. Remember, for instance, the case of thinking of counterexamples. As you learn a new subject matter, trying to think of counterexamples to claims surely is a good critical posture to maintain. But you have to accumulate knowledge in the domain with which to find or build counterexamples. And you have to develop a sense of what *counts* as a counterexample in the domain. Similarly, in applying to this new domain a reading strategy that asks you to summarize, you have to develop a sense of what counts as relevant. Or, in applying an extreme case heuristic to the new domain, you have to find out what dimensions are significant, so that you will know how to push a proposition to an extreme meaningful in that domain.

Of course, none of this need to learn and to adjust implies that the cognitive gripper you are using lacks generality. All specific applications of anything general need to configure to the context. This approach acknowledges the importance of domain-specific adjustments, which indeed often are challenging, while maintaining the reality and power of general cognitive skills.

Completing the case. It should be acknowledged that the findings supporting this synthesis paint a partial and scattered picture. Indeed, the four conditions for generality mentioned earlier—seeming use, important role, transferable, and common absence—offer a map of the kinds of empirical work needed to test the matter further.

Regarding *seeming use*, more protocol studies are needed that examine experts addressing atypical problems within their domain of expertise, to check for seeming use of general strategies. Also, more experiments in teaching heuristics are needed that test whether gains in problem solving can

be attributed directly to the use of the heuristics. Both sorts of studies would also address the question of *important role*, because they can show general strategies figuring crucially in finding solutions. Teaching experiments can also address *common absence*, by documenting that students lack certain strategies before intervention and gain from their use after intervention.

Supposing that positive results accrue on seeming use, important role, and common absence, then the issue of generality hangs on the question of transfer, where considerably more work is needed. Transfer can prove more or less robust, even the least robust case providing some evidence of generality. In the strongest case, a person mastering a general method by contextualizing it to a domain would spontaneously transfer it to other domains. Lacking spontaneous transfer, the person might show transfer if a teacher or other source alerted the learner to its relevance in the new domain. If even that failed, it might still be so that instruction systematically helping the learner to contextualize the method in the new domain would go more quickly because of prior experience in the original domain. As mentioned in the section on transfer, some contemporary results show one or another of these patterns. But much careful systematic work on the question has yet to be done.

Educating Memories Versus Educating Minds

We are fairly confident in the synthesis position outlined here, not only because it makes sense of both the negative and positive findings so far, but also because it makes sense of everyday observations—such as the philosophers' wont to pick apart claims with counterexamples. But what is its import for education?

Despite many efforts to refashion educational practices to cultivate more thoughtful learning within and across domains, the fact of the matter is that most educational practice remains doggedly committed to imparting facts and algorithms. Regrettably, E. D. Hirsch (1987) and other educators have even taken the negative arguments from expertise, weak methods, and transfer as reasons to eschew attention to higher order skills so that more time is given to building students' factual knowledge base in a domain.

This seems particularly unfortunate. To be sure, general heuristics that fail to make contact with a rich domain-specific knowledge base are *weak*. But when a domain-specific knowledge base operates without general heuristics, it is brittle—it serves mostly in handling formulaic problems. Although we don't want the weak results of the kind of attention to general heuristics that neglects knowledge base, we also don't want the brittle competency forged by exclusive attention to particularized knowledge! We would hope for more from education. And, according to the synthesis theory, we can get more.

As noted earlier, several contemporary experiments in the direct teaching of cognitive skills have yielded very positive results. Moreover, guidelines are available for classroom practices that can foster the transfer of knowledge and skills (Perkins & Salomon, 1987, 1988). The fact remains, however, that most efforts to cultivate general cognitive skills have not focussed on bringing together context-specific knowledge with general strategic knowledge. Rather, they have taken the form of courses or minicourses segregated from the conventional subject matters and make little effort to link up to subject matter or to nonacademic applications (cf.

Nickerson et al., 1985; Segal et al., 1985).

In contrast, the approach that now seems warranted calls for the intimate intermingling of generality and context-specificity in instruction. A few methodologies and educational experiments have addressed exactly that agenda (e.g. Mirman & Tishman, 1988; Palincsar & Brown, 1984; Perkins, 1986; Schoenfeld, 1985; Wales & Stager, 1978). We believe that this direction in education is promising and provocative: It gets beyond educating memories to educating minds, which is what education should be about.

The Chess Master's Chances

It's high time to return to the chess master's chances of becoming an insightful political and military counsel. In the golden age of general cognitive skills, many might have said, "Give it a try!" albeit with some caveats. After the fall, most would have said, "No way!" Now, what should we say?

The right response seems to be that we should first gather more information about this chess master. Does he already have some general principles ("control the center—any center") rather than entirely contextualized principles ("control the middle squares of the chess board")? How meta-cognitive is his thinking about his chess play and other life activities? Does he tend to do what high-road transfer calls for: mindfully decontextualize principles? Or, in contrast, is he a gifted intuitive player of chess, with an enormous fund of experience but little predilection to reflect and generalize? Depending on the answers to such questions, we might forecast his chances as ranging from "No way!" to "There's some hope." Although recognizing the great importance of years of experience in a domain, the latter would be far from a sure bet.

If there is a sure bet to be had, it is that, with the polarized debate about general as opposed to local knowledge quieting down, we are open to learning much more about how general and local knowledge interact in human cognition. And, of course, we can put that understanding to use in educational contexts. We forecast that wider scale efforts to join subject-matter instruction and the teaching of thinking will be one of the exciting stories of the next decade of research and educational innovation.

Acknowledgements: The writing of this paper was partially supported by a grant given jointly to Gavriel Salomon and to the late Tamar Globerson by the Spencer Foundation.

Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.

Bereiter, C., & Tinker, G. (1988, April). *Consistency of constructive learning effort across domains of high and low cultural familiarity*. Paper presented at the annual meeting of the American Educational Research Association, New Orleans, LA.

Boden, M. (1977). *Artificial intelligence and natural man*. New York: Basic Books.

Brown, A. L., & Kane, M. J. (1988, April). *Cognitive flexibility in young children: The case for transfer*. Symposium paper presented at the Annual Meeting of the American Educational Research Association, New Orleans, LA.

Brown, A. L., Kane, M. J., & Long, C. (in press). Analogical transfer in young children: Analogies as tools for communication and exposition. *Applied Cognitive Psychology*.

Brown, A. L., & Palincsar, A. S. (in press). Guided, cooperative learning and individual knowledge acquisition. In L. Resnick (Ed.),

Knowing and learning: Essays in honor of Robert Glaser. Hillsdale, NJ: Erlbaum.

Chan, C., & Burtis, J. (1988, April). *Level of constructive effort, prior knowledge, and learning*. Paper presented at the annual meeting of the American Educational Research Association, New Orleans, LA.

Chase, W. C., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, 4, 55-81.

Cheng, P. W., & Holyoak, K. J. (1985). Pragmatic reasoning schemas. *Cognitive Psychology*, 17, 391-416.

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

Clancey, W. J. (1986). From GUIDON to NEOMYCIN and HERACLES in twenty short lessons. ONR Final Report, 1979-1985. *AI Magazine*, 7, 40-60.

Clancey, W. J. (1987). *The knowledge engineer as student: Metacognitive bases for asking good questions* (Technical Report STAN-CS-87-1183). Stanford, CA: Department of Computer Science, Stanford University.

Clement, J. (1982). Analogical reasoning patterns in expert problem solving. *Proceedings of the Fourth Annual Conference of the Cognitive Science Society*. Ann Arbor, MI: University of Michigan.

Clement, J. (1982). Students' preconceptions in introductory mechanics. *American Journal of Physics*, 50, 66-71.

Clement, J. (in press). Nonformal reasoning in physics: The use of analogies and extreme cases. In J. Voss, D. N. Perkins, & J. Segal (Eds.), *Informal reasoning*. Hillsdale, NJ: Lawrence Erlbaum Associates.

de Groot, A. D. (1965). *Thought and choice in chess*. The Hague: Mouton.

Ehrlich, K., & Soloway, E. (1984). An empirical investigation of the tacit plan knowledge in programming. In J. Thomas & M. L. Schneider (Eds.), *Human Factors in Computer Systems*. Norwood, NJ: Ablex.

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

Ernst, G. W., & Newell, A. (1969). *GPS: A case study in generality and problem solving*. New York: Academic Press.

Gardner, H. (1983). *Frames of mind*. New York: Basic Books.

Gardner, H. (1985). *The mind's new science*. New York: Basic Books.

Gick, M. L., & Holyoak, K. J. (1987). The cognitive basis for knowledge transfer. In S. M. Cormier & J. D. Hagman (Eds.), *Transfer of learning* (pp. 81-120). New York: Academic.

Glaser, R. (1984). Education and thinking: The role of knowledge. *American Psychologist*, 39, 93-104.

Hayes, J. R., & Simon, H. A. (1977). Psychological differences among problem isomorphs. In N. J. Castellan, Jr., D. B. Pisoni, & G. R. Potts (Eds.), *Cognitive theory* (pp. 21-41). Hillsdale, NJ: Erlbaum.

Herrnstein, R. J., Nickerson, R. S., Sanchez, M., & Swets, J. A. (1986). Teaching thinking skills. *American Psychologist*, 41, 1279-1289.

Hirsch, E. D. (1987). *Cultural literacy: What every American needs to know*. Boston, MA: Houghton-Mifflin.

Hunter, J. E. (1986). Cognitive ability, cognitive aptitudes, job knowledge, and job performance. *Journal of Vocational Behavior*, 29, 340-362.

Johnson, P. E., Ahlgren, A., Blount, J. P., & Petit, N. J. (1980). Scientific reasoning: Garden paths and blind alleys. In J. Robinson (Ed.), *Research in science education: New questions, new directions*. Colorado Springs, CO: Biological Sciences Curriculum Study.

Laird, J. E., Rosenbloom, P. S., & Newell, A. (1984). Towards chunking as a general learning mechanism. In *Proceedings of the National Conference on Artificial Intelligence*, (pp. 188-192). Los Altos, CA: W. Kaufman, Inc.

Langer, E. (in press). *Mindfulness*. Reading, MA: Addison-Wesley.

Larkin, J. H. (1982). The cognition of learning physics. *American Journal of Physics*, 49, 534-541.

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

Larkin, J. H., McDermott, J., Simon, D. P., & Simon, H. A. (1980b). Modes of competence in solving physics problems. *Cognitive Science*, 4, 317-345.

Lehman, D. R., Lempert, R. O., & Nisbett, R. E. (1988). The effects of graduate training on reasoning: Formal discipline and thinking about everyday-life problems. *American Psychologist*, 43, 431-442.

Mirman, J., & Tishman, S. (1988). Infusing thinking through "Connections." *Educational Leadership*, 45(7), 64-65.

- Newell, A., & Simon, H. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Ng, E. (1988, April). *Three levels of goal-directedness in learning*. Paper presented at the annual meeting of the American Educational Research Association, New Orleans, LA.
- Nickerson, R., Perkins, D. N., & Smith, E. (1985). *The teaching of thinking*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Ogilvie, M., & Steinbach, R. (1988, April). *The development of skill across domains: The role of learning strategies*. Paper presented at the annual meeting of the American Educational Research Association, New Orleans, LA.
- Palincsar, A. S., & Brown, A. L. (1984). Reciprocal teaching of comprehension-fostering and comprehension-monitoring activities. *Cognition and Instruction*, 1, 117-175.
- Patel, V. L., & Groen, G. J. (1986). Knowledge based solution strategies in medical reasoning. *Cognitive Science*, 10, 91-116.
- Pea, R. D., & Kurland, D. M. (1984). On the cognitive effects of learning computer programming. *New Ideas In Psychology*, 2, 137-168.
- Perkins, D. N. (1985). Postprimary education has little impact on informal reasoning. *Journal of Educational Psychology*, 77(5), 562-571.
- Perkins, D. N. (1986). *Knowledge as design*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Perkins, D. N. (in press). Reasoning as it is and could be. In D. Topping, D. Crowell, & V. Kobayashi (Eds.), *Thinking: The third international conference*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Perkins, D., & Salomon, G. (1987). Transfer and teaching thinking. In D. N. Perkins, J. Lochhead, & J. Bishop (Eds.), *Thinking: The second international conference* (pp. 285-303). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Perkins, D. N., & Salomon, G. (1988). Teaching for transfer. *Educational Leadership*, 46(1), 22-32.
- Perkins, D. N., Schwartz, S., & Simmons, R. (in press). Toward a unified theory of problem solving: A view from computer programming. In M. Smith (Ed.), *Toward a unified theory of problem solving*. Hillsdale, NJ: Erlbaum.
- Perkins, D. N., & Simmons, R. (1987). Patterns of misunderstanding: An integrative model of misconceptions in science, mathematics, and programming. In J. D. Novak (Ed.), *Proceedings of the second international seminar on misconceptions and educational strategies in science and mathematics* (Vol. 1, pp. 381-395). Ithaca, NY: Cornell University.
- Polya, G. (1954). *Mathematics and plausible reasoning* (2 vols.). Princeton, NJ: Princeton University Press.
- Polya, G. (1957). *How to solve it: A new aspect of mathematical method* (2nd ed.). Garden City, NY: Doubleday.
- Pressley, M., Snyder, B. L., & Cariglia-Bull, T. (1987). How can good strategy use be taught to children? Evaluation of six alternative approaches. In S. M. Cormier & J. D. Hagman (Eds.), *Transfer of learning* (pp. 81-120). New York: Academic.
- Rabinowitz, M., & Glaser, R. (1985). Cognitive structure and process in highly competent performance. In F. D. Horowitz & M. O'Brien (Eds.), *The gifted and talented: Developmental perspectives* (pp. 75-98). Washington, DC: American Psychological Association.
- Resnick, L. B. (1987). *Education and learning to think*. Washington, DC: National Academy Press.
- Rich, E. (1983). *Artificial intelligence*. New York: McGraw-Hill.
- Salomon, G. (1988) AI in reverse: Computer tools that turn cognitive. *Journal of Educational Computing Research*, 4, 123-139.
- Salomon, G., & Globerson, T. (1987, October). Rocky roads to transfer. *The second Annual Report to the Spencer Foundation*. Israel: Tel-Aviv University.
- Salomon, G., & Perkins, D. N. (1987). Transfer of cognitive skills from programming: When and how? *Journal of Educational Computing Research*, 3, 149-169.
- Salomon, G., & Perkins, D. N. (in press). Rocky roads to transfer: Rethinking mechanisms of a neglected phenomenon. *Educational Psychologist*.
- Schoenfeld, A. H. (1978). Presenting a strategy for indefinite integration. *American Mathematical Monthly*, 85, 673-678.
- Schoenfeld, A. H. (1982). Measures of problem-solving performance and of problem-solving instruction. *Journal for Research in Mathematics Education*, 13(1), 31-49.
- Schoenfeld, A. H. (1985). *Mathematical problem solving*. New York: Academic Press.
- Schoenfeld, A. H., & Herrmann, D. J. (1982). Problem perception and knowledge structure in expert and novice mathematical problem solvers. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 8, 484-494.
- Scribner, S., & Cole, M. (1981). *The psychology of literacy*. Cambridge, MA: Harvard University Press.
- Segal, J. W., Chipman, S. F., & Glaser, R. (Eds.). (1985). *Thinking and learning skills, Volume 1: Relating instruction to research*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Strauss, S. (1987). Educational-development psychology and school learning. In L. S. Liben (Ed.), *Development and learning: Conflict or congruence?* (pp. 133-158). Hillsdale, NJ: Erlbaum.
- Thorndike, E. L. (1923). The influence of first year Latin upon the ability to read English. *School Sociology*, 17, 165-168.
- Thorndike, E. L., & Woodworth, R. S. (1901). The influence of improvement in one mental function upon the efficiency of other functions. *Psychological Review*, 8, 247-261.
- Toulmin, S. E. (1958). *The uses of argument*. Cambridge, England: Cambridge University Press.
- Wales, C. E., & Nardi, A. (1984). *Successful decision-making*. Morgantown, WV: West Virginia University, Center for Guided Design.
- Wales, C. E., & Stager, R. A. (1978). *The guided design approach*. Englewood Cliffs, NJ: Educational Technology Publications.
- Wason, P. C. (1966). Reasoning. In B. M. Foss (Ed.), *New horizons in psychology*. Harmondsworth, England: Penguin.
- Wenger, E. (1987). *Artificial intelligence and tutoring systems: Computational and cognitive approaches to the communication of knowledge*. Los Altos, CA: Morgan Kaufmann Publishers.

Economic Perspectives on Education

A Special Issue of the *Educational Researcher*, May 1989

Gary Becker, Consulting Editor
Henry Levin, Introductory Essayist

Articles by

Eric Hanushek, University of Rochester
on the impact of differential
expenditure on school performance

Dale Jorgenson, Harvard University
on comparisons of investments in
education with other types of capital
formation

Jacob Mincer, Columbia University
on job training's effect on the labor
market and economic growth's effect
on job training

Finis Welch, University of California-
Los Angeles
on time series changes in income
differentials between college and
high school graduates