March 17, 1992

Professor Keith Holyoak
Department of Psychology
University of California at Los Angeles
Los Angeles, California 90024

Dear Keith:

In penance for my inability to get to your lecture, I thought I might send you some rough notes I have put down on reading your chapter in the Ericsson-Smith collection. You will observe that I don't at all agree with your three-generation story, much less with your third-generation proposal.

I would be interested in your comments, particularly explaining what it is that I am missing.

As I may have mentioned to you, we (Richman, Gobet and I) have just received an NSF grant to extend the EPAM system, especially to accommodate trained expertise like that studied by Chase-Ericsson-Polson-Staszewski. It is pretty far along, and when it is completed, it will explain some of the things that didn't fit too well with the Chase-Simon chunking hypothesis in its original form. As soon as we have something on paper (it may be six months), I will be glad to send it on to you. We will be writing the program in LISP, so that EPAM will become, for the first time, an exportable program that can be run on a PC or Macintosh.

It was a pleasure to see you and have dinner with you last week.

Cordially yours,

Herbert A. Simon

HAS:eda
enclosure

Comments on
"Symbolic Connectionism"

The "three-generation" account of cognitive science described by the author, though widely accepted, is dubious history. The early systems dealt with knowledge-poor tasks (e.g., puzzles, logic, chess), not because knowledge was thought unimportant, but because computer memories were too small to allow simulation of knowledge-rich systems of any size. During the first couple of decades I dissuaded several graduate students from undertaking theses involving large knowledge bases that simply could not be handled with the computers then available. The NSS chess program (1958) provides good evidence that the relevance of knowledge to expertise was clearly understood from the beginning. Feigenbaum was quite correct, however, in claiming that DENDRAL was the first system actually to attempt AI for a knowledge-rich task environment.

Even more dubious is the characterization of the prospective third generation, and especially the ways in which it is supposed to deviate from its predecessors. On page 303, Holyoak lists some 12 "predicted characteristics of expertise" that have received some empirical support, then argues that "none provides a universal characterization of expert performance." He then recites some "exceptions that have been reported" under each of these twelve headings.

Experts sometimes achieve mediocrity. "Expert decision-makers appear to do remarkably well, 'generating hypotheses and inducing complex decision rules . . . Unfortunately, their knowledge and rules have little impact on experts' performances. Sometimes experts are more accurate than novices (though not always), but they are rarely better than simple statistical models.'"

This objection is based on studies of expert clinical decision making. I am not aware of any other domain where it is true (although I accept that it is true in this domain). It is best handled simply by denying that there can be expertise in this domain until such time as there are verified regularities in the phenomena that aspiring experts can learn. In the same way, I would argue that the term "expert" is inappropriately applied to people who claim to have acquired proficiency in ESP, dousing, or perhaps even economic or weather prediction. No knowledge; no potential for expertise.

Experts sometimes feel more pain. Experts do not always accomplish with ease what novices do only with difficulty. The evidence comes from studies of writing, where experts are observed to strive harder than novices. This is hardly a refutation of the theory, since the quality of the product is not held constant. Experts generally also think harder about chess moves than rank novices do, but they also make far better moves. And of course, people who accomplish difficult tasks (e.g., world-class scientists or mountaineers who ascend Everest) strive very hard to do it. Given the same task, novices work harder than experts or don't complete (or even start) the task.

Means-ends analysis can impair learning. The experiments cited in support of this proposition did not use expert subjects. Nor is there any claim in the usual theory that learning by doing (which ordinarily requires means-ends analysis) is the sole or best way of acquiring expert levels of skills. Hence the evidence is simply irrelevant.

Conditions and actions cometimes can be flexibly recoupled. It is simply untrue that "In second-generation theories, expertise is viewed as the result of automatic evocation of specialized actions in response to specialized conditions." The correct claim is that experts have available large numbers of productions permitting such "automatic" responses, and use these as *part* of their repertory of responses, thereby allowing them to focus attention on the more difficult parts of their tasks.

For example, it has been pointed out that chess grandmasters can play simultaneous games against large numbers of opponents, almost wholly on the basis of the moves evoked by their productions, but that they play more poorly than in games in which they move much more slowly and deliberately and have time to look ahead and engage in means-ends analysis. (Holding also has this misconception, that the received theory of expertise denies the role of look-ahead in chess play.) The automation of various subparts of the search processes with the help of stored knowledge (productions) enhances the power of the more deliberate processes in whose service these automated processes operate.

Expertise sometimes can be decoupled from memory performance. The correct claim is that experts have superior memory for those things that are *relevant* to superior performance in their domains. Chunking relates to the structure of the indexed knowledge base they acquire, and irrelevant things (e.g., details of codes that can readily be reconstructed) don't fit this structure. Several people have failed to replicate the Holding and Reynolds findings. None of this refutes the frequently observed fact that experts remember most situations in their domains far better (e.g., factors of five or more) than novices, but does call for more exact specification of what we mean by "better memory."

Expertise sometimes can be decoupled from pattern perception. The Allard and Starkes study simply shows, again, that the authors misjudged what are the relevant cues in a volleyball game. Allard and Starkes essentially admit this in their explanation of their results.

Expert search strategies are extremely varied. I am not aware that this is denied by the accepted theory. The whole idea that experts often work forward, instead of backward, is that -- for problems that are easy for them -- they proceed by applying automated, or nearly automated, procedures without a lot of thinking ahead. It has long been shown that for harder problems they revert to procedures that look much more "novice-like," e.g., make use of means-ends analysis. Moreover, they classify problems -- that is, diagnose the nature of the problems they are confronted with -- and then apply methods appropriate to that class of problem (the well-known paper of Chi et al.). Hence their strategies vary with the characteristics of the problem. All this is part of the "second-generation" theory.

Performance may not show continuous improvement with practice. I do not know that the received theory, or the mechanisms on which it is based, make the contrary claim.

Learning need not require goals or feedback. The evidence cited here does not seem to prove the point. "simply by exposure to music" says nothing of what is going on in the child's head. "Feedback" is not at all limited to explicit feedback from a teacher. Patterned material (like music) provides its own feedback from success or failure to recognize the patterns. Kotovsky and I have explored this point (as have many others) in our study of the Thurstone letter series completion task.

Knowledge can be transferred across domains. I know of no evidence that knowledge can be transferred across domains unless there are "common elements" shared by the domains. Has any second-generation theorist doubted that someone skilled in the calculus may be able to use that skill in both physics and economics -- might, that is to say, be an "expert" in the calculus? Analogy, of course, is an example of common elements, and perhaps the mystery is not why it is so often effective as why it is so often ineffective (not evoked when relevant).

Teaching expert rules may not yield expertise. The received theory does not claim that expertise is achieved by memorizing or learning to repeat verbal rules, but by learning to connect sensory cues with knowledge stored in memory, so that the cues will be noticed when present and the knowledge accessed. These are two nearly independent matters. That is one of many reasons why excellent coaches are often not expert performers.

Rules elicited from experts may not predict their performance. This point relates to the previous one. Possessing a production that enables one to notice appropriate cues and respond to them is not equivalent (either way) to being able to verbalize either the cue or the response.

Summary. For the reasons given above, the "disconcerting lack of constancy in the correlates of expertise" seems to be largely an illusion produced by talking about these matters in vague and inaccurate terms, and without reference to the rigorous simulations of expert skills that nail the relevant terms down rigorously and demonstrate the phenomena.

It is claimed that "there appears to be no single 'expert way' to perform all tasks." This is undoubtedly true if it means that experts in different domains do not have the same knowledge -- the same indexed encyclopedia. It is almost surely false if it means that experts do not exhibit their expertise by employing a large body of knowledge organized in production-system form (recognize and access), combined with modest inferential capabilities (means-ends analysis and other weak methods). The productions will be most evident when an expert is performing an easy task rapidly and "intuitively," the inferential processes will be in greater evidence when difficulties are encountered that require the expert to supplement his/her "automatic" knowledge-based responses with heuristic search.

If we took as our second-generation theory of expertise, a combination of EPAM, GPS, and UNDERSTAND (and this is what I have been using since at least 1972, and except for the latter, since 1960), I think we would see that the theory is not refuted by any of the evidence cited in this article. (My current version of the theory adds the mechanism of adaptive production systems, as an additional learning mechanism. We will shortly have a new version of EPAM that handles experiments on expert memory (i.e., the Chase-Ericsson-Staszewski experiments) more gracefully.)