Psychological Review
1998, Vol. 105, No. 1, 33–57

0033-295X/98/$3.00

An Ecological Theory of Expertise Effects in Memory Recall

Kim J. Vicente and JoAnne H. Wang
University of Toronto

Previous research has shown a significant correlation between domain expertise and memory recall performance after a very brief exposure time. Despite the large number of such studies, several findings in the literature have no satisfactory theoretical explanation. A novel theory based on an ecological approach is proposed to explain these results. This *constraint attunement hypothesis* provides a framework for identifying and representing the various levels of goal-relevant constraint in a domain. The theory predicts that there will be a memory expertise advantage in cases in which experts are attuned to the goal-relevant constraints in the material to be recalled and that the more constraint available, the greater the expertise advantage can be. The theory explains a number of diverse empirical findings in the literature in a coherent, unique, and parsimonious fashion and suggests a number of promising issues for future research.

In this article we describe a novel product theory of expertise effects in memory recall tasks. Research in this area began at least 52 years ago with the seminal work of de Groot (1946/1965) on problem solving in chess. In that recall experiment, four chess players of various levels of expertise were asked to reconstruct meaningful board positions after having been exposed to them for only a few seconds. De Groot found that the Master and Grandmaster level players performed this task with near-perfect accuracy, whereas the performance of the lesser players was not nearly as impressive. In a subsequent experiment, Chase and Simon (1973a, 1973b) found that when the board positions consisted of randomly placed pieces, the recall performance of a Master plummeted so that it was not significantly different from that of a beginner, thereby indicating that the Master's superior performance on meaningful positions was not a result of better overall memory.[1]

Further research has replicated this basic pattern of results many times. There are at least 51 studies of expertise effects in at least 19 different domains. Heavily studied domains include the following: chess (Charness, 1976, 1981; Chase & Simon, 1973b; Chi, 1978; Cooke, Atlas, Lane, & Berger, 1993; de Groot, 1946/1965; Frey & Adesman, 1976; Goldin, 1978; Holding & Reynolds, 1982; Jongman, 1968; Lane & Robertson, 1979; Pfau & Murphy, 1988; Reynolds, 1982; Saariluoma, 1984, 1985), computer programming (Adelson, 1981; Barfield, 1986; Bateson, Alexander, & Murphy, 1987; Di Persio, Isbister, & Shneiderman, 1980; McKeithen, Reitman, Reuter, & Hirtle, 1981; Shneiderman, 1976; Vessey, 1987), and medical diagnosis (Claesen & Boshuizen, 1985; Coughlin & Patel, 1987; Hassebrock, Johnson, Bullemer, Fox, & Moller, 1993; Norman, Brooks, & Allen, 1989; Patel & Frederiksen, 1984; Patel, Groen, & Arocha, 1990; Patel, Groen, & Frederiksen, 1986). Games other than chess include the following: bridge (Charness, 1979; Engle & Bukstel, 1978), go (Eisenstadt & Kareev, 1975; Reitman, 1976), gomoku (Eisenstadt & Kareev, 1975), and othello (Billman & Shaman, 1990; Wolff, Mitchell, & Frey, 1984). Expert recall in sports has been examined as well and includes the following: baseball (Chiesi, Spilich, & Voss, 1979; Spilich, Vesonder, Chiesi, & Voss, 1979), basketball (Allard, Graham, & Paarsalu, 1980), field hockey (Starkes, 1987), figure skating (Deakin & Allard, 1991), football (Garland & Barry, 1991), and soccer (Williams, Davids, Burwitz, & Williams, 1993). Miscellaneous other domains include the following: algebra (Sweller & Cooper, 1985), ballet (Starkes, Deakin, Lindley, & Crisp, 1987), circuit diagrams (Egan & Schwartz, 1979), maps (Chang, Lenzen, & Antes, 1985; Gilhooly, Wood, Kinnear, & Green, 1988), musical notation (Sloboda, 1976), and process control (Moray et al., 1993; Vicente, 1992). Despite the wide range of domains and the variety of methods used in these studies, memory recall performance on meaningful stimuli has almost always been found to be correlated with domain expertise (see Vicente, 1988, for a review). How can this finding be explained? As the review in the Existing Theories section shows, a satisfactory answer to this question has yet to appear in the literature.

Kim J. Vicente and JoAnne H. Wang, Department of Mechanical and Industrial Engineering, University of Toronto, Ontario, Canada.

This article was prepared, in part, while Kim J. Vicente was a visiting scientist in the Cognitive Systems Group of Risø National Laboratory, Roskilde, Denmark. Subsequent funding was provided by a research grant from the Natural Sciences and Engineering Research Council of Canada.

We are indebted to Adriaan de Groot, John Flach, Alex Kirlik, and Robert Shaw for fruitful discussions and advice and Bill Brewer, Anders Ericsson, Fernand Gobet, Ken Hammond, Kurt Joseph, Randy Mumaw, Ed Reed, David Rubin, Herbert Simon, Jim Staszewski, and Daniel Willingham for their constructive comments on versions of this article.

Correspondence concerning this article should be addressed to Kim J. Vicente, Department of Mechanical and Industrial Engineering, 5 King's College Road, University of Toronto, Ontario, Canada M5S 3G8. Electronic mail may be sent via Internet to benfica@mie.utoronto.ca. Kim J. Vicente's World Wide Web address is http://www.mie.utoronto.ca/labs/cel/.

[1] This random control condition has often incorrectly been attributed to de Groot (Vicente & Brewer, 1993; Vicente & de Groot, 1990).

Existing Theories

The best known theory of expertise effects in memory recall is the perceptual chunking account proposed by Chase and Simon (1973a, 1973b). In their view, "Chess skill depends in large part upon a vast, organized long-term memory of specific information about chessboard patterns" (Chase & Simon, 1973b, p. 279). These chunks of chess pieces, that Chase and Simon referred to as *perceptual structures,* are said to be stored in memory as relational structures in propositional form. Although beginners must abstract from what they see on the board in order to encode a position, Masters can rely on the contents of their memory to recognize and to encode rapidly the position.

Simon and Gilmartin (1973) developed an information-processing model that instantiates the perceptual chunking theory of memory recall in the form of a computer simulation. The simulation assumes that chess Masters hold a very large number of perceptual structures in long-term memory (LTM), each corresponding to a familiar pattern of pieces. During exposure to a given chess position, the pieces on the board serve as cues that allow the Master to recognize the labels corresponding to chunks in LTM. These labels are retrieved into short-term memory (STM) and are used during recall to derive the information about the location of pieces stored in the chunk in LTM.

Subsequent research has demonstrated that Chase and Simon's (1973a, 1973b) original theory is inadequate. For example, Charness (1976) and Frey and Adesman (1976) both found that recall of chess positions was not affected by interference tasks that disrupt the contents of STM. These findings indicate that, contrary to what Chase and Simon supposed, the chunks are not encoded in STM. Knowledge structures from LTM seem to be involved in expert chess memory, indicating that participants are attending to the meaning of the position they are recalling. As a result of these and other findings, there is a consensus in the literature that the original theory is seriously flawed (e.g., Gobet & Simon, 1996c; Holding, 1985).

However, as we discuss later, Chase and Simon's (1973a, 1973b) theory has been modified recently on the basis of findings from the *skilled-memory theory* developed by Chase and Ericsson (1981, 1982), which postulates that exceptional memory performance can be accounted for by efficient coding and retrieval of information in LTM. Skilled-memory theory has received a fair amount of empirical scrutiny and has done a good job of accounting for the findings that have been generated (see Ericsson & Staszewski, 1989). Moreover, it has recently been extended into the theory of long-term working memory (LTWM), a framework that provides a process account of how people are able to deal with large demands on working memory in the face of the well-known limitations on STM (Ericsson & Kintsch, 1995). According to Ericsson and Kintsch, exceptional memory performance is possible because experts acquire the ability to quickly and effectively store the products of intermediate stages of processing as *retrieval structures* in LTM. This information is maintained in a directly accessible state by the use of retrieval cues in STM. In this way, experts learn to expand the functional capacity of working memory by relying on LTM as a temporary storage buffer.

Ericsson and Kintsch's (1995) LTWM theory extends the basic ideas of skilled-memory theory to a broader set of phenomena in memory in both text comprehension and expert performance. However, the examples of expert memory that they account for in detail come almost exclusively from domains in which memorizing stimuli is an *intrinsic* task (i.e., a task that is a definitive feature of that domain of expertise). Examples of such domains include blindfolded chess (Ericsson & Staszewski, 1989), memorizing dinner orders (Ericsson & Polson, 1988), and memorizing digits (Chase & Ericsson, 1981, 1982). It is important to distinguish such domains from those in which memorizing stimuli is a *contrived* task (i.e., a task that is not part of that domain of expertise). For example, "Chess players' main occupation is not to perform recall experiments, but to play chess" (Gobet, 1994, p. 45). The distinction between intrinsic tasks that experts usually perform and contrived tasks that experts do not usually perform, but that are used to study experts' reasoning processes, is an important one in both the expertise and knowledge elicitation literatures (Ericsson & Lehmann, 1996; Hoffman, Shadbolt, Burton, & Klein, 1995, respectively). The distinction is important here as well because, according to skilled-memory theory, retrieval structures are deliberately acquired and fully available to consciousness (Gobet & Simon, 1996c). Therefore, retrieval structures are only of obvious relevance in domains in which there is a need for deliberate memory enhancement, not in domains in which the de Groot (1946/1965) memory recall task is a contrived task. For example, why would a basketball player deliberately acquire a retrieval structure to remember a basketball situation or event? While the question may seem facetious, it is not when one considers that most, if not all, of the cases of expertise effects in memory recall tasks documented earlier are in domains in which memory recall is contrived, not intrinsic. It is not clear how LTWM can explain these results. Thus, LTWM does not provide a comprehensive explanation of expertise effects in memory recall.

Recently, Richman, Staszewski, and Simon (1995) have rehabilitated the chunking theory of Chase and Simon (1973a, 1973b). They developed a computer model, EPAM IV, that builds on the earlier family of EPAM models and retains many of the insights of the original Chase and Simon theory, but it also incorporates the retrieval structures posited by skilled-memory theory. Richman et al.'s model does an impressive job of quantitatively accounting for the data obtained from a longitudinal study of an expert increasing his ability to recall long sequences of digits. However, again, this is a domain in which memory recall is an intrinsic task.

Is there any empirical evidence to indicate that EPAM IV succesfully accounts for all of the expertise effects in domains in which memory recall is a contrived task? The answer to this question seems to be no because "the range of tasks studied in detail [with EPAM IV] is very limited to date" (Richman et al., 1995, p. 327). The only example in which the EPAM IV architecture has been applied to a contrived memory recall task is the work of Gobet (1993; de Groot & Gobet, 1996) in chess, and even in this one domain "the [computer] program does a poor job at simulating the recall performance of human Masters" (de Groot & Gobet, 1996, p. 243).

In summary, skilled-memory theory, LTWM theory, and EPAM IV cannot provide an adequate theoretical explanation for expertise effects in memory recall for the considerable number of domains in which memory recall is a contrived task.

An Alternative Tack

As a next step in developing a more comprehensive theory, it may be useful to reflect on the research strategy that has been adopted in trying to understand expertise effects in memory recall. An important characteristic of the theories just reviewed is that they are *process* theories. That is, they are directed at the question of *how* (i.e., by what psychological mechanisms and knowledge structures) expertise effects in memory recall come about. It is possible to formulate a theory of the same phenomenon at a different level of analysis; namely, a *product* (i.e., input–output) theory that provides an accurate functional account of the phenomenon but that does not commit to a particular psychological mechanism. The goal of this latter type of theory is to predict and understand *what* expertise effects in memory recall are observed under various conditions. These two accounts, what and how, are complementary descriptions of the same phenomenon. However, as J. J. Gibson (1966) observed three decades ago, the question of what logically precedes that of how. That is, to develop an adequate theory of how expertise effects in memory recall are generated, it may be very useful, and perhaps even necessary, to first have an accurate account of exactly what those effects are because the latter can put strong constraints on the former. The value of Gibson's research strategy has since been recognized by some cognitive scientists (e.g., Anderson, 1990; Marr, 1982; Neisser, 1987), but it has not widely been adopted in psychology.

Perhaps one of the reasons then that researchers have yet to develop an adequate process theory of expertise effects in memory recall is that the nature of those effects has yet to be clearly ascertained. The level of understanding made possible by current theories can briefly be summarized as follows: When presented with "natural" material, experts remember more than novices, but when the material is random, the recall performance of experts plummets to the level of novices. Such a statement fails to address several basic issues. First, the adjective *natural* was deliberately used as a placeholder because the literature is unclear as to exactly when one should expect to see an expertise advantage. For example, Gobet and Simon (1996c) stated that such an advantage occurs in chess when participants are presented with an "unfamiliar" position (p. 2). Literally speaking, this is correct because the participants had never seen the position before, but the same can be said for random positions as well. Elsewhere, Gobet and Simon (1996b) have described random chess positions as "situations containing some infrequently observed features" (p. 159), presumably suggesting that it is the fact that game positions contain features that are frequently observed that makes them easier for skilled players to remember. Ericsson and Staszewski (1989) stated that expertise advantages occur because the stimulus is "familiar" (p. 254) or "meaningful" (p. 256), suggesting that these two adjectives are synonymous. Thus, there is a lack of consensus on this point in the literature. Different definitions are given for the boundary conditions under which an expertise effect is to be expected. Moreover, familiarity and meaningfulness have been used interchangeably but they are not identical concepts (Vicente, 1988). For example, a stimulus can be meaningful in that it can be interpreted by participants with their existing knowledge base, yet it can be unfamiliar in the sense that participants have never been exposed to it before. Conversely, a stimulus can be familiar because it has been seen before but may not be meaningful because participants may not understand its significance (cf. Ericsson & Harris, 1990). Because these terms have sometimes been used interchangeably, previous theories are not clear as to when an expertise effect is to be expected.

Second, several experiments have shown that the magnitude of expertise effects in memory recall can vary significantly as a function of task manipulations, even when nonrandom stimuli are presented (e.g., Chase & Simon, 1973b; Frey & Adesman, 1976). However, existing theories cannot account for all of these observed differences, as the remainder of our article shows.

Finally, there is uncertainty as to what the functional equivalent of a random chessboard is, particularly in medical diagnosis (Coughlin & Patel, 1987). Even in chess, Gobet (1994) has pointed to "the necessity of defining precisely what constitutes a random position" (p. 37). This uncertainty makes it difficult to predict the conditions under which expertise advantages will occur in various domains.

That such fundamental issues have not been resolved may well explain why an adequate process theory of expertise effects in memory recall has remained elusive. This lack of understanding points to the value of developing a product theory of expert recall that can then serve to constrain the development of more viable process theories. Only mechanisms that are capable of reproducing the input–output behaviors described by the product theory then need be considered. It is important to emphasize that the tack of developing a product theory is a methodological choice, not an ontological claim. Internal mechanisms and structures specified by a process theory add valuable knowledge that is complementary to that offered by a product theory. The claim made here is that stepping back to develop a product theory may provide important insights that can lead to the revision, and thus improvement, of existing process theories.

In this article we address these issues by presenting a product theory of expertise effects in memory recall. The remainder of the article is organized as follows. First, the details of the theory are presented. Second, the ability of the theory to account for a number of anomalous findings in the literature is evaluated. To anticipate, this analysis demonstrates that the theory proposed here provides a coherent, unique, and parsimonious account of the vast majority of the expertise effects observed in memory recall experiments. Finally, the article ends with a discussion of the implications of the theory for a number of important issues.

The Constraint Attunement Hypothesis

There are at least three related issues that need to be addressed by a viable product theory of expertise effects in memory recall.

1. How should one represent the constraints that the environment (i.e., the problem domain) places on expertise?
2. Under what conditions will there be an expertise advantage?
3. What factors determine how large that advantage will be?

This section describes a theory adapted from ecological theories of skill acquisition[2] that addresses these questions, thereby building upon and extending a large body of existing research, albeit from a different area of psychology.

An Ecological Approach to Skill Acquisition

The ecological approach to skill acquisition, led by the seminal work of E. J. Gibson (1969, 1991), has generated a *specificity theory* that views perceptual learning as the "education of attention." Cognitive economy and goal-relevance play an important part in this theory; with experience, people learn the strategy that is most economical for the task at hand and thereby focus on the minimal number of distinctive features (invariants) that will successfully discriminate among the events of interest. Thus, skill acquisition consists of changing what one attends to, the goal being to identify diagnostic high-order information that can be used to satisfy task goals. Training of attention is accomplished by abstraction, filtering, and optimization of perceptual search (see E. J. Gibson, 1969, 1991, for more details).

Thus, from an ecological perspective, the process of skill acquisition consists of adaptation to the constraints imposed by the environment. Both the structure of the environment (i.e., the problem domain) and the person's behavior become important objects of inquiry. Although the theories reviewed earlier acknowledge the role of the stimulus, none of them has an explicit theory of the environment.

Constraint Attunement in Memory Recall

An ecological approach to skill leads to a concise explanation of expertise effects in memory recall. There can be expertise effects when there are goal-relevant constraints (i.e., relationships pertinent to the domain) that experts can exploit to structure the stimuli. The more constraint available, the greater the expertise advantage can be. Fully random stimuli have no constraints, so no expertise advantage would be expected. To realize these potential advantages, experts must be attuned (i.e., they must attend) to the goal-relevant constraints in question. If they do not pick up on this information, then no expertise advantage is expected. This *constraint attunement hypothesis* is a novel theoretical explanation for expertise effects in memory recall, but it has existed for years as the cornerstone of ecological theories of skill acquisition and has received considerable empirical support (E. J. Gibson, 1969, 1991).

Describing the Constraints in the Environment

For the constraint attunement hypothesis to have any potency, there must be a way to identify the goal-relevant constraints in a domain of expertise. Some ecological psychologists have used the concept of an *affordance* (J. J. Gibson, 1979) to describe the environment in a goal-relevant manner. An object's affordances are the action possibilities it offers an actor with a given set of capabilities. Affordances are environmental constraints on goal-directed action. Although this concept continues to generate a great deal of productive research, it is too limited for the purposes of this article. Affordances describe properties of individual objects, events, and places. What is needed here is a framework for identifying, and describing the relationships between, the many affordances making up an entire problem domain of expertise.

Rasmussen (1985) has developed the concept of an *abstraction hierarchy* for this purpose. The abstraction hierarchy provides a framework for developing a hierarchical description of the goal-relevant constraints for a problem domain. Although it was developed independently of ecological psychology, the abstraction hierarchy can be viewed as a nested hierarchy of affordances (Vicente & Rasmussen, 1990). Each level in the hierarchy represents a normative model of the goal-relevant constraints in the world. Note that the abstraction hierarchy represents the problem domain and is therefore not a task analysis. This distinction is key: A task describes or prescribes human problem-solving activity, whereas a problem domain is the object of that activity. In the remainder of this section we describe the abstraction hierarchy in more detail, as well as providing a simplified example of the role it plays in the constraint attunement hypothesis.

What kind of hierarchy? Different types of hierarchical structures have frequently been used to represent complex systems in a variety of disciplines (e.g., Korf, 1987; Mesarovic, Macko, & Takahara, 1970; Pattee, 1972; Simon, 1981). However, there are certain properties that distinguish Rasmussen's abstraction hierarchy from other types of hierarchies. More specifically, the abstraction hierarchy belongs to the class of *stratified hierarchies* described by Mesarovic et al. (1970), the properties of which are listed below.

1. Each stratum, or level, deals with the very same system, the only difference being that different strata provide different descriptions, or different models, for observing the system.

2. Each stratum has its own unique set of terms, concepts, and principles.

3. The selection of strata for describing a particular system depends on the observers and their knowledge and interest in the control of the system.[3] For many systems, however, there may be some strata that appear to be natural or inherent.

---

[2] The ecological approach to psychology (Brunswik, 1956; J. J. Gibson, 1979) must be distinguished from the specific theory of direct perception (J. J. Gibson, 1979) because the former is much broader than the latter (Reed, 1988). Our theory is ecological because it explicitly represents the influence of environmental constraints on human behavior. However, it is not based on the theory of direct perception and therefore should not be identified with it.

[3] This is not to suggest that the construction of an abstraction hierarchy is an arbitrary process. The hierarchy is intended to be a normative representation, so it must faithfully describe a particular problem domain—it is not the case that anything goes. Each level in the abstraction hierarchy below the top level describes the set of means that is available to achieve the ends at the level above. Thus, one cannot just construct an abstraction hierarchy in an arbitrary manner because this means–ends relation must be preserved throughout the hierarchy. This, in essence, provides a guide for constructing abstraction hierarchies and a test for evaluating whether they are in fact correct or not. Having said this, it is in principle possible for two observers to develop different abstraction hierarchies for the same system, if they adopt different system purposes. For example, a saboteur would develop a different abstraction hierarchy for a nuclear power plant than a systems engineer because the properties of the plant that are relevant to the intended purposes of the plant are different (although probably overlapping) for each observer. For the purposes of the present article, such idiosyncratic differences

4. The requirements for proper system functioning at any level appear as constraints on the meaningful operation of lower levels, whereas the evolution of the state of the system is specified by the effect of the lower levels on the higher levels.

5. Understanding of the system increases by crossing levels: By moving up the hierarchy, one obtains a deeper understanding of system significance with regard to the goals that are to be achieved; whereas in moving down the hierarchy, one obtains a more detailed explanation of the system's functioning in terms of how those goals can be carried out.

In addition, the structure of the abstraction hierarchy is further specified by a means–end relationship between levels (Rasmussen, 1985). This defining feature is in contrast to other types of hierarchies that are often defined by attributes that are not explicitly related to goals (e.g., spatial scale and temporal scale). As we show later, this explicitly goal-oriented nature has important psychological implications.

The properties just described define a family of representations for a modeler. Thus, the abstraction hierarchy is not a specific representation but rather a framework that a modeler can use to develop representations for various problem domains. The number of levels and their content will vary as a function of the types of constraints in each domain. Regardless of the domain, however, an abstraction hierarchy representation will have the properties described above.

Psychological advantages. One important psychological property of the abstraction hierarchy is that higher levels are less detailed than lower levels. Shifting one's representation from a low, detailed level to a higher level of abstraction with less resolution makes complex systems look simpler. In effect, this provides a mechanism for coping with complexity. Metaphorically, moving up one or more levels allows one to "see the forest for the trees." Thus, part of the psychological relevance of the abstraction hierarchy is that it allows resource-bounded agents, as people are, to deal with systems that would be unmanageable if they had to observe a whole system in full detail all at once.

This advantage is not unique to the abstraction hierarchy, however (cf. Simon, 1981). Virtually all hierarchies allow one to observe systems at a less detailed level. From a psychological point of view, the unique and important characteristic of the abstraction hierarchy is that it is explicitly goal oriented. The various levels in the hierarchy are linked by a means–end relation. This relationship provides a very important source of constraint that can be exploited in problem solving. Thus, search can be constrained by initiating the problem-solving process at a high level of abstraction, deciding which system function is relevant to the current situation, and then concentrating on the subtree of the hierarchy that is connected to the function of interest. This type of search is efficient (cf. Korf, 1987) because all system parts not pertinent to the function of interest can be ignored.

In summary, an abstraction hierarchy representation supports goal-directed problem solving in a computationally economic manner. It should allow people to (a) structure their overall problem-solving process, (b) frequently start their problem-solving activities at a higher level of abstraction to avoid detail, and (c) iteratively "zoom in" on lower levels of abstraction to examine selectively only those parts of the domain that are relevant to the goal or function of current interest (cf. Egan & Schwartz, 1979). Each of these claims is empirically falsifiable.[4] In fact, these predictions have empirically been confirmed by verbal protocol data obtained from experts engaged in problem-solving activities in a number of complex problem domains under representative conditions, including electronic troubleshooting (Rasmussen, 1986), nuclear power plants (Itoh, Yoshimura, Ohtsuka, & Masuda, 1990), and process control simulations (Vicente, Christoffersen, & Pereklita, 1995). These predictions are also consistent with many studies of expert problem solving that are not based on the abstraction hierarchy. Such studies have repeatedly shown that experts spend a great deal of their time analyzing the functional structure of a problem at a high level of abstraction before narrowing in on more concrete details (Glaser & Chi, 1988). Some notable examples include Selz (1922; see Frijda & de Groot, 1981, for an English account), Duncker (1945), and de Groot (1946/1965).

The psychological advantages listed earlier are not enjoyed by other types of hierarchies (e.g., part–whole or classification hierarchies). The links between levels in these other representation formats are not explicitly related to system purposes or functions. Although it is possible to examine the system at a high level of the hierarchy to choose a subsystem of interest, the subtree of the hierarchy that is connected to that subsystem may not necessarily contain system components that are relevant to the purposes or functions of the selected subsystem. In other words, other hierarchies constrain search but not in a way that is explicitly related to system purposes. It is this latter type of constraint that is needed to understand goal-directed behavior.

An example. A simplified example illustrating the role of the abstraction hierarchy in the constraint attunement hypothesis is presented next (two more detailed examples are presented in the Appendix). This example was inspired by the work of Chiesi et al. (1979) and Spilich et al. (1979) on memory recall of text describing events in baseball. Like other domains of expertise, baseball is subject to various types of constraint that can be exploited by experts in memory recall. The left side of Figure 1 identifies the levels of an abstraction hierarchy for baseball from the perspective of a team manager, suggesting plausible layers of goal-relevant constraint connected by means–end links between levels (not shown in the left side of Figure 1). This representation can be viewed as a normative problem space for managers because it identifies the constraints that they should consider in making decisions during a game.

At the lowest level, *players*, there are constraints on the number and positions of the players that constitute a fielded team

are not encountered in practice because agreement on purpose is straightforward for various domains of expertise (e.g., in chess, the purpose is to win the game; in medical diagnosis, the purpose is to identify the patient's ailment).

[4] The abstraction hierarchy itself is not an empirically falsifiable hypothesis. Rather, it is a framework for representing domain constraints. It is a demonstrable fact that one can develop abstraction hierarchy representations for problem domains that have a known structure. However, the predictions derived from the rationale for the abstraction hierarchy can be, and have been, directly tested empirically.

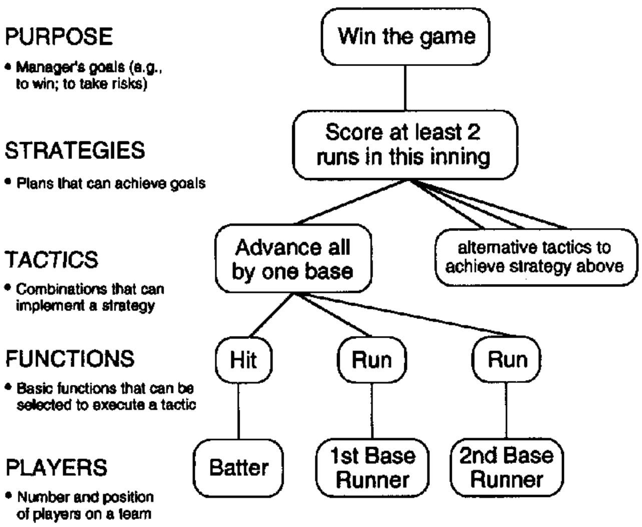

Figure 1. The left side of the figure identifies the levels of goal-relevant constraint in an abstraction hierarchy for baseball. The right side illustrates in more detail the subset of this hierarchy that is activated by Vignette 1 (see text for description).

(e.g., pitcher and batter). There are certain configurations that are not allowed by the rules of the game (e.g., having more than nine players on defense). Thus, the players level puts a minimum amount of constraint on what constitutes a meaningful baseball event. At the second level, *functions,* the constraints arising from the legal functions that can be performed are represented. Because of the rules of the game, there are only a small number of primitive functions (i.e., strike, ball, hit, catch, throw, and run). The *tactics* level adds another layer of constraint. Examples of offensive tactics include squeeze play, sacrifice out, stolen base, and hit and run. Examples of defensive tactics include intentional walk, double play, forced play, and pick off a base runner. The next level of constraint is *strategies.* Baseball managers do not constrain their moves by the rules of the game or tactics alone. Their actions are also constrained by the strategies that they have adopted to achieve a win (e.g., if a team is down by one run in the bottom of the ninth inning, the strategy is to score at least one run). Finally, at the level of *purpose,* there is only a small set of meaningful purposes. Managers can choose to play a certain style of game (e.g., risky or conservative), but the end will always be to win.

The preceding paragraph has focused solely on the constraints within each level of abstraction. However, there are also constraints between levels that are defined by the means–ends linkages spanning adjacent levels of abstraction. For example, the means–ends links between the players and functional levels also impose meaningful constraints on events (e.g., a batter cannot throw a strike and a pitcher cannot steal a base while on defense). Also, the links between functions and tactics also impose constraints because primitive functions are not arbitrarily combined in all possible ways but instead are selectively chosen to implement the activated tactics. Similar relations connect other adjacent levels in the hierarchy.

This abstraction hierarchy for baseball is not intended to be definitive (extensive analysis would be required to develop a comprehensive and accurate representation). Its goals are much more modest; namely, to provide a plausible example of the role that the abstraction hierarchy plays in the constraint attunement hypothesis. This goal can be realized by showing how the presence, or absence, of goal-relevant constraints is expected to affect memory recall performance.

With this goal in mind, a number of vignettes describing events in baseball are presented. Vignette 1, adapted from Spilich et al. (1979), describes a set of events that is consistent with all of the constraints identified in Figure 1.

> The Springfield Senators lead the Bloomington Bucks 4–2 in the bottom of the ninth. The Bucks are at bat with one out and runners Smith and Jones on first and second, respectively. Whitcomb is at bat. He knows that the next batter is their best power hitter, so Whitcomb's goal is to hit a single to load the bases. The pitcher looks towards first and second for wide lead-offs by the runners, then throws a high fastball for a strike. The catcher returns the ball and signals the next pitch.

The right side of Figure 1 illustrates the subsection of the abstraction hierarchy that is activated by Vignette 1. The purpose is to win the game, of course. Given the current state, the only viable strategy is to score at least two runs in this inning. A few tactics could have been chosen to implement this strategy, but the one that has been selected is to advance the two base runners and the batter by one base. Given this tactic, the functions required and the mapping between functions and players are both completely constrained: the batter must hit a single and the two base runners must safely run to their next base. Figure 1 shows just how structured this vignette is. According to the constraint attunement hypothesis, it is this very structure that allows experts to exhibit exceptional recall after only a brief exposure to the material.

Vignette 1 represents a baseline, which was used to generate other vignettes. For instance, Vignette 2 provides an example violating a means–ends link between tactics and strategies (only the modified portion of the vignette is quoted, with the underlined text indicating the change from Vignette 1): "He knows that the next batter is their weakest hitter, so Whitcomb's goal is to hit a single to load the bases." Given that the next batter is the weakest on the team, the chosen tactic is not an effective means of achieving the desired end.[5]

Vignette 3 provides an example violating a means–ends link between tactics and functions: "The pitcher . . . throws a high fastball for a strike. Smith tries to steal second, and is called out when he runs into Jones who is still on second." The primitive functions that have been chosen (base runner on first stealing

[5] It is possible, although extremely unlikely, that the next batter (the weakest on the team) could hit a home run, for instance, and win the game for the Bucks. In domains like baseball in which analytical models are not available, the means–ends links are probabilistic, not deterministic. Thus, the domains can be modeled by using information theory, which allows one to quantify the degree of uncertainty in the means–ends relations (see the Appendix, Example 2). For baseball, the statistical databases that professional teams develop would be very useful in developing an abstraction hierarchy representation by using information theoretic measures.

second, with a base runner still on second) do not serve as means to implement the identified tactic.

Vignette 4 shows an example violating a means–ends link between functions and players: "The shortstop looks towards first and second for wide lead-offs by the runners, then throws a high fastball for a strike." According to the rules of baseball, a shortstop is not allowed to throw a pitch; only the pitcher can do so.

Finally, Vignette 5 represents multiple violations of goal-relevant constraints across levels of abstraction.

> The home-team Springfield Senators lead the Bloomington Bucks 4–2 in the top of the first. The Senators are at bat with three outs and two pitchers Smith and Jones on first. Whitcomb is on deck. He knows that the next batter is their weakest hitter, so Whitcomb's goal is to throw a strike to load the bases. The shortstop looks towards the outfield and home for wide lead-offs by the fielders, then throws a pitch out for a strike. The base runners return the ball and signal the next hit to the pitcher.

To anyone who is familiar with baseball, this vignette is obviously less structured than any of the others, perhaps disturbingly so.

According to the constraint attunement hypothesis, expertise advantages should be larger when experts are attuned to the goal-relevant constraints in highly structured stimuli (such as Vignette 1), smaller in slightly less structured stimuli (such as Vignettes 2–4), and very much smaller in significantly less structured stimuli (such as Vignette 5). Only when there is no goal-relevant structure at all is an expertise advantage not expected. Thus, the abstraction hierarchy also provides a defensible way of defining a random stimulus.

In addition to making strong predictions of this type, the abstraction hierarchy also represents a useful methodological tool because it allows experimenters to identify, and therefore selectively eliminate, specific constraints or classes of constraints. Without such a tool, experimenters have to resort to breaking structure by less precise methods, such as randomizing the order of sentences in a meaningful passage (as Chiesi et al., 1979, did in their study of baseball). Thus, in addition to identifying the constraints that allow one to evaluate empirically the claims of the constraint attunement hypothesis, the abstraction hierarchy is also a useful experimental tool.

Empirical Evidence

In this section we review some of the evidence pertinent to the constraint attunement hypothesis. In evaluating the worth of a new theory, it is customary to compare its predictions with those of alternative theories. In this case, however, the alternatives are process theories, whereas the constraint attunement hypothesis is a product theory emphasizing a different set of questions (e.g., under which conditions will an expertise effect be observed and what determines the magnitude of the effect? vs. what psychological mechanisms and knowledge structures are responsible for these expertise advantages?). Therefore, the type of evidence that is relevant to evaluating each can differ. Process theories can be evaluated by using either process measures (e.g., eye movement data and duration of pauses in serial recall) or product measures (e.g., recall accuracy), but product theories do not make predictions about process and thus can only be evaluated by the latter type of measure. Consequently, these two types of theories can only be compared by using product measures of performance. Even so, such a comparison is difficult because process theories only make predictions about product criteria in an indirect, rather than a direct, fashion (i.e., by *generating* an output based on the psychological mechanisms postulated by the theory and the task conditions of interest). To take an arbitrary example, the only way to determine what predictions EPAM IV would make regarding the magnitude of expertise effects in the recall experiments reviewed later would be to give the model the large number of inputs it requires and then to simulate it under the appropriate set of task conditions. For these reasons, it is sometimes difficult to contrast our predictions with those of other theories.

Richman et al. (1995) also found it difficult to compare their theory against others, albeit for very different reasons. As a result, they decided to "comment on other theories only cursorily" (p. 306). Our solution to this dilemma is to state explicitly whether the results accounted for by our product theory had been predicted by the researchers that generated those results. Note that even if an experiment is designed to evaluate a process theory, it can still make predictions regarding product criteria of performance. Typically, this occurs when two competing psychological mechanisms lead to different predictions about when an expertise effect will occur, or how large that effect will be. Because researchers can use process theories to make predictions about product criteria for the conditions in their studies, we can determine if the results accounted for by the constraint attunement hypothesis were anticipated by those researchers.

In five cases out of the nine that we reviewed, the results had not been predicted. In three of these, the experimenters offered post-hoc explanations (Barfield, 1986; Coughlin & Patel, 1987; Myles-Worsley, Johnston, & Simons, 1988), but none of these explanations was integrated into a general theory. In another of the five cases (Chase & Simon, 1973b), no explanation was offered. The fifth case (Vicente, 1992) led to a post hoc explanation—the constraint attunement hypothesis. In the remaining four cases in which the results had been predicted (Frey & Adesman, 1976; Gobet & Simon, 1994; Lane & Robertson, 1979; Reynolds, 1982), we determined whether the prediction was based on a theory that accounted for all of the other evidence. None of the cases satisfied this criterion because the predictions were primarily based on literature reviews, not on a comprehensive theory of expertise effects in memory recall. Thus, most of the results we account for had not been anticipated, and no other theory accounts for all of the evidence accounted for by the constraint attunement hypothesis.

For the sake of brevity, a complete abstraction hierarchy for each domain is not presented for each study reviewed below. Instead, only a brief description of the constraints that are required to explain the results of each study is presented.

Vicente (1992)

The first experiment to be reviewed was conducted by Vicente (1992) in the domain of process control. The layers of goal-relevant constraint in this domain are described in the Appendix. Participants were presented with a dynamic scenario showing the behavior of a computer-simulated thermal–hydraulic process plant over a period of 25–30 s during each trial. The participants' task was to recall the final state of the process variables and to diagnose the type of trial. An important aspect of this recall study is that there were three types of trials rather than the usual two types used in chess: normal, in which the process was operating the way it should; fault, in which one fault was introduced into the process simulation; and pseudorandom, in which the process variables were driven virtually randomly and independently of each other (these scenarios were not physically meaningful).

Extrapolating from the results observed in most other recall studies, one would predict an expertise advantage for normal trials but not for pseudorandom trials. However, what about the fault trials? On the one hand, using the criteria defined by Gobet and Simon (1996b), one could argue that fault trials have "infrequently observed features" and thus should be more similar in nature to random trials than to normal trials. One would reach a similar conclusion by using one of the criteria identified by Ericsson and Staszewski (1989) because the fault trials were also unfamiliar (participants had never seen them before). In either of these cases, one would expect that there would not be an expertise advantage for fault trials because they are analogous to random stimuli. Conversely, using the other criterion identified by Ericsson and Staszewski, one would predict that fault trials would be more similar in nature to normal trials than random trials because they are meaningful in that they are physically realizable. If this were true, then one would expect an expertise advantage for fault trials. There is no single criterion that one can apply from the memory recall literature to determine which of these hypotheses is correct. Thus, Vicente's (1992) data are of particular interest because they provide unique information concerning the boundary conditions of expertise effects.

The constraint attunement hypothesis makes specific predictions for this experiment. In fault scenarios, only one or two constraints that govern the system under normal operating conditions are violated. The many intact constraints should allow experts to structure the scenario and thereby enable them to remember more than novices, despite the fact that the event is abnormal and has never been observed before. Thus, there should be an expertise advantage for fault trials that is slightly less in magnitude than that observed for normal trials. The pseudorandom trials were not fully random, as a few of the constraints governing the system under normal circumstances were still operating. The constraint attunement hypothesis predicts an expertise advantage for these trials, albeit slight, because there are still some constraints for experts to pick up on. However, the magnitude of that expertise advantage should be much less than that observed on fault and normal trials because there is much less constraint in the former than in the latter two.

The results of this study, shown in Figure 2, reveal a significant Expertise × Trial Type interaction. First, there was a strong expertise advantage for normal trials, as predicted. Second, there was a weak but nevertheless significant expertise advantage on the pseudorandom trials. Third, a significant expertise effect on the fault trials was also observed. The magnitude of this expertise effect was slightly less than that obtained on normal trials,

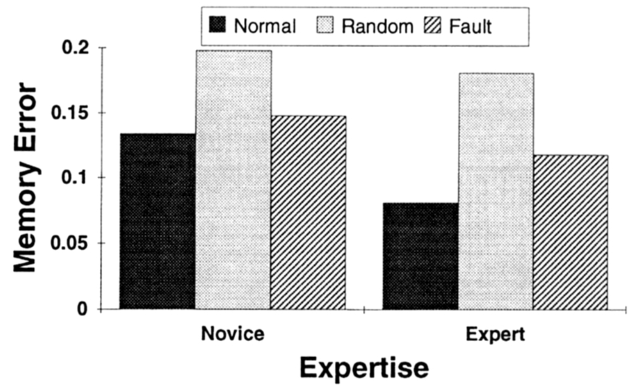

Figure 2. Mean recall error for three different trial types in a process control simulation as a function of level of expertise in Vicente (1992).

as predicted. Thus, all of the results are consistent with the predictions made by the constraint attunement hypothesis.

Frey and Adesman (1976)

Frey and Adesman (1976) conducted an experiment in the domain of chess, in which the constraints include purpose, strategies, tactics, paths, and board (see the Appendix). Frey and Adesman compared memory for static random chess positions, static meaningful positions, and move-by-move *sequences* in which the moves leading up to the final meaningful position were also shown to participants. The first two conditions are those typically used in memory recall experiments in chess. The more interesting condition is the third one. What should the relative magnitude of the expertise effect be for this sequence condition?

As usual, the static, meaningful condition should result in an expertise advantage, whereas the random condition should not. The prediction for the sequence condition is perhaps less obvious. Presenting a move-by-move sequence adds more constraint than is available in the static, meaningful condition because the preceding moves provide a source of redundant information about the tactics and strategies that led to the position to be recalled. This extra constraint, unavailable in the other two conditions, can be used to facilitate the recall of the final position. As a result, the constraint attunement hypothesis predicts expertise advantages for the meaningful and move-by-move conditions, with a greater advantage for the latter.

Participants were assigned to one of three groups, according to their chess rating (see Figure 3). For those not familiar with chess ratings, it may useful to point out that higher ratings reflect greater expertise. Master level players, who have a rating of at least 2200, were not included in this particular study.

As shown in Figure 3, there was an expertise advantage for the meaningful positions but not for the random condition. More important, however, there was also an expertise effect for the move-by-move condition that was greater than the advantage generated by the static, meaningful condition. These results confirm the predictions made by the constraint attunement hypothesis.

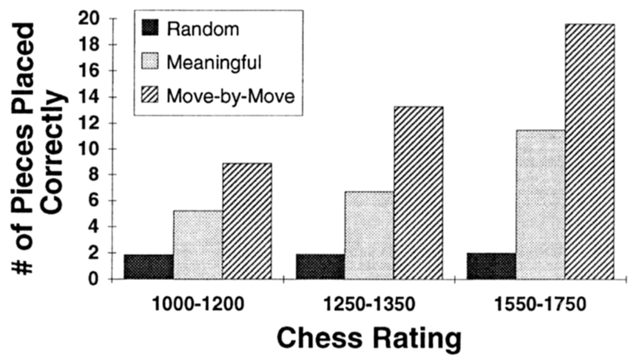

Figure 3. Mean number of chess pieces recalled and placed correctly in three different presentation conditions as a function of level of expertise in Frey and Adesman (1976).

Chase and Simon (1973b)

Chase and Simon (1973b, pp. 261–264) asked participants from three different levels of chess expertise to recall a sequence of moves rather than just a static position. There were two stimuli types: moves taken from actual games, and moves that were randomly generated from the set of legal moves that could be made from the existing board configuration. On the basis of their previous results with static board positions, Chase and Simon expected that the meaningful-moves condition would produce an expertise effect, whereas the random-moves condition would not.

The predictions made by the constraint attunement hypothesis must be derived by analyzing the amount of goal-relevant constraint in each experimental condition. The sequence of moves from real games is constrained by the board and path constraints resulting from the rules of chess (see the Appendix) and tactical, strategic, and purpose considerations; whereas the randomly selected moves are constrained by the rules of chess alone. Thus, the constraint attunement hypothesis predicts that the expertise advantage for sequences of moves from real games would be greater than for randomly selected moves. Moreover, the theory also predicts an expertise advantage for the random moves because there is still some constraint in the stimulus (i.e., the rules defining a legal move) that can be exploited by experts. Note that these predictions directly conflict with those made by Chase and Simon (1973b).

The results from this experiment, illustrated in Figure 4, support both of the predictions generated from the constraint attunement hypothesis. The memory for sequences of moves from real games was indeed better than for randomly selected moves. There was also an expertise advantage, although smaller in magnitude, for the random-moves condition. Chase and Simon (1973b) did not expect this result, and they could not provide a post hoc explanation for it either. They just stated, "Apparently, the skilled players were able to find some meaning in the randomly generated moves" (Chase & Simon, 1973b, p. 26). Chase and Simon did not state what that source of "meaning" was, how one could describe it systematically, what other sources of meaning there could be in the stimulus, or what the relationship is between these different types of meaning and the magnitude of expertise advantages in memory recall. These theoretically important issues are, however, all addressed by the constraint attunement hypothesis.

Reynolds (1982)

One of the most direct tests of the constraint attunement hypothesis can be found in a study conducted by Reynolds (1982). By reanalyzing the verbal protocols presented in de Groot (1946/1965), Reynolds found that Grandmasters and Masters directed their attention to squares on the board that are affected by many pieces. In contrast, lesser players tended to direct their attention to the squares on which the pieces themselves are located. Assuming that better players direct their attention to the functionally relevant features of the board, these findings can be interpreted as indicating that the degree to which squares are affected by other pieces is an important source of goal-relevant information. Reynolds then went on to design an ingenious memory recall experiment that varied the amount of goal-relevant information presented to participants of three levels of expertise. Three classes of chess positions were randomly generated so as to differ only in the degree to which the pieces converged on the same squares. On the basis of de Groot's protocols, the amount of goal-relevant constraint in the stimulus is operationally defined as the degree to which squares are affected by multiple pieces; the greater the number of pieces affecting each square, the greater the degree of goal-relevant constraint in the stimulus. Consequently, the constraint attunement hypothesis predicts that the expertise advantage between the three participant groups should increase as a function of the degree to which various pieces affect common squares.

The experimental results, shown in Figure 5, support this prediction. The only statistically significant effect of expertise occurs in the condition with the highest order grouping (the third-order grouping in Figure 5). Master chess players show superior recall for positions with distributions concentrated about the squares that are affected by the greatest number of pieces. In the two lesser constrained positions, there is no significant effect of expertise, although the expertise effect for the second-order grouping is greater in magnitude than that for the first-order grouping, as the constraint attunement hypothesis predicted.

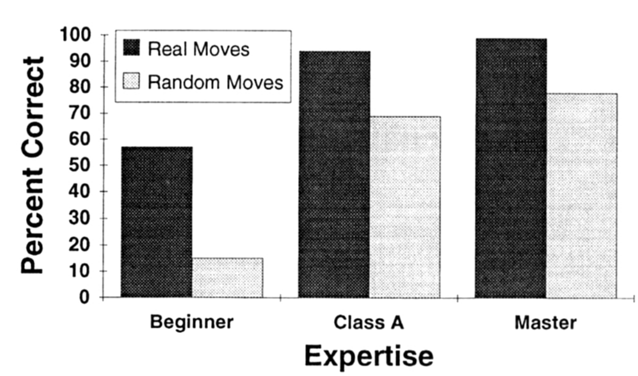

Figure 4. Mean number of chess pieces recalled from sequences of real moves and random moves as a function of level of expertise in Chase and Simon (1973b).

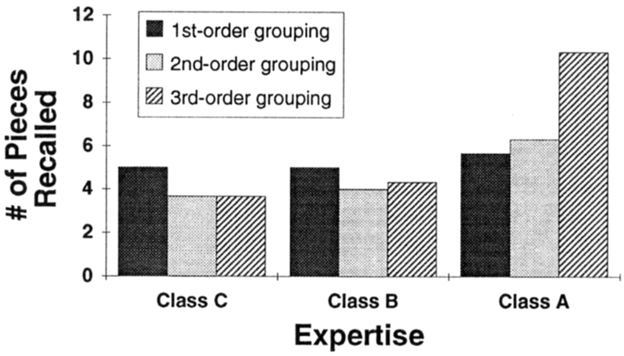

Figure 5. Mean number of chess pieces recalled for three different-order groupings as a function of level of expertise in Reynolds (1982).

Myles-Worsley et al. (1988)

A study conducted by Myles-Worsley et al. (1988) investigating the effects of expertise on recognition memory for X-rays provides another opportunity to evaluate the constraint attunement hypothesis. The layers of goal-relevant constraint in this domain may include the following: the purpose (detect abnormalities), physiological symptoms, and anatomical symptoms. Myles-Worsley et al. compared the ability of participants with various levels of expertise to recognize abnormal X-rays and normal X-rays, some of which had been presented earlier in the experiment. As in the case of the Vicente (1992) experiment, the crucial issue here is which of these two types of stimuli is the analog of a randomized chess position. Ericsson and Staszewski's (1989) criterion of familiarity for identifying nonrandom stimuli is not very helpful in this case because experts are very familiar with both normal and abnormal X-rays. Gobet and Simon's (1996b) criterion of "infrequently observed features" would suggest that abnormal X-rays are more similar to a random board than a normal X-ray because presumably abnormal X-rays are not encountered as frequently as normal X-rays. From this perspective, then, one might think that there would be an expertise advantage for normal X-rays but not for abnormal X-rays. Instead, Myles-Worsley et al. found that recognition increased on abnormal X-rays as a function of expertise and actually decreased on normal X-rays with expertise (see Figure 6). This result was unexpected because it did not fit with the theoretical suppositions motivating the study.

It might be tempting to attribute this result to the fact that the study involved a recognition task. However, there is a variety of studies that has replicated the classical findings of the memory recall paradigm by using a recognition task instead of a recall task (e.g., Beal, 1985; Chiesi et al., 1979; Goldin, 1979). Thus, the unexpected results observed by Myles-Worsley et al. (1988) are unlikely to be due to the fact that a recognition task was adopted instead of a recall task.

The constraint attunement hypothesis suggests that the interpretation that random and meaningful chess positions are equivalent to abnormal and normal X-rays, respectively, is a shallow one. The goal in reading an X-ray is to detect any existing abnormality. Thus, the goal-relevant constraints are those features that signify an abnormal X-ray, not those that characterize a normal X-ray. From this perspective, the results of Myles-Worsley et al. (1988) have a straightforward interpretation. Experts were more attuned to the goal-relevant information in the X-rays, whereas novices were more attuned to irrelevant information.

Barfield (1986)

Another set of findings that has not been theoretically accounted for is that of Barfield (1986), who conducted a memory recall study in the domain of computer programming. In this domain, the layers of goal-relevant constraint include the following: the purpose that the code is intended to achieve, the laws of logic, the logical functions required to satisfy the purpose, the syntax of the programming language, and the visual appearance of the code (e.g., indentation). Barfield's study included four expertise groups (naive, novice, intermediate, and expert) and three types of program organization (executable order [EO], random chunks [RC], and random lines [RL]). The EO condition, which is the equivalent of a meaningful chess position, presented participants with a complete program with all 25 lines of code in the proper order and with indentation intact. The remaining two conditions, RC and RL, were both randomized, although in a different manner. The RL condition used the same lines of code as the EO condition, but the order of those lines was randomized. The form of each individual line of code, however, was left intact, including the indentation. The RC condition also used the same lines of code, but the code was first divided into meaningful chunks, and then these chunks were presented in a randomized order.

Because the EO condition contained the program in its proper order, all of the goal-relevant constraints were retained in the stimulus. Thus, the magnitude of the expertise effect in this condition should be the greatest overall. The RL condition had the least amount of structure and so would be expected to lead to the smallest expertise advantage. The RC condition is an interesting one because it retained fewer constraints than the full EO condition (the correct order of the lines was not retained), but it also contained many more constraints than the RL condition. Even though the chunks in the RC condition were

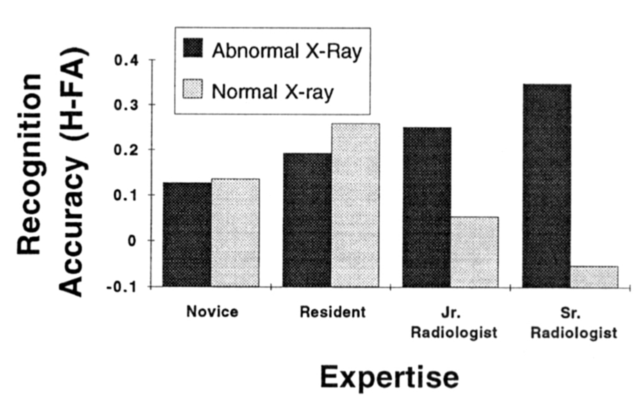

Figure 6. Recognition accuracy for normal and abnormal X-rays as a function of level of expertise in Myles-Worsley, Johnston, and Simons (1988). H = hits; FA = false alarms; Jr. = junior; Sr = senior.

presented in a randomized order, the stimulus still contained many goal-relevant constraints because the chunks were defined by meaningful program functions (e.g., initialization, arithmetic calculations, and for-next loops). This condition retained many of the higher order constraints that define a meaningful computer program. According to the constraint attunement hypothesis, the RC condition should lead to an intermediate expertise effect that is closer in magnitude to that for the EO condition than for the RL.

A remaining question is, Should there still be an expertise advantage in the RL condition? This condition might be considered the equivalent of a randomized chess board, in which each chess piece is placed on a randomly chosen square. If this comparison were accurate, one would expect equivalent performance from all expertise groups in this condition. However, an intact line of computer code is not equivalent to a chess piece because it still contains some goal-relevant constraints (e.g., the rules of syntax). Thus, an individual line of code has a function or meaning of its own, albeit highly limited, despite being separated from the overall structure of the program. Though the majority of the goal-relevant constraints was disrupted, the RL condition still retained some structure, though not nearly as much as the EO or RC conditions. Therefore, the constraint attunement hypothesis predicts a small expertise effect for the RL condition.

The results of this experiment are presented in Figure 7. Unfortunately, Barfield (1986) did not report between-participant, group statistical comparisons at various levels of organization. As a result, the best that can be done is to rely on the means in Figure 7 to make these comparisons. These informal comparisons support all of the predictions made by the constraint attunement hypothesis. First, the relative magnitude of expertise effects for the different program organization conditions was, in decreasing order, EO, RC, and RL, as predicted. Second, for the least structured, RL condition, there seemed to be an expertise effect with the expert and intermediate groups outperforming the other two. Third, there also seemed to be a strong expertise effect for the RC condition, despite the fact that the lines were random. In fact, there was only a slight difference between the RC and the EO conditions. The fact that all of the predictions of the constraint attunement hypothesis were confirmed is noteworthy if one considers that not only did Barfield not predict these results, but he also failed to offer a theoretical account of them.

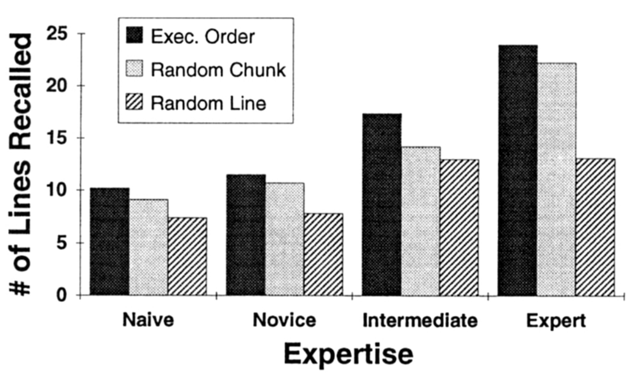

Figure 7. Mean number of lines of computer code recalled correctly as a function of code structure and level of expertise in Barfield (1986). Exec. = executable.

Coughlin and Patel (1987)

A memory recall experiment conducted by Coughlin and Patel (1987) in the domain of medical diagnosis also provides direct empirical support for the constraint attunement hypothesis. The layers of goal-relevant constraint in this domain include the following: the purpose, which is to identify the patient's ailment; the laws of physiology (e.g., pressure–volume relations); various pathological symptoms; and the form in which those symptoms are usually presented to doctors. Coughlin and Patel's experiment included two levels of expertise (novices and experts), two different clinical cases to be examined (endocarditis and arteritis), and two different structures of each case (structured and random). One case was intended to be uncommon but familiar (endocarditis), and the other was intended to be common but less familiar (arteritis). The randomized forms of the clinical cases contained exactly the same information as the structured forms, but the information was randomly ordered in the texts. Participants were asked to recall in writing the text of a random version of one case and a structured version of the other case, as well as provide a diagnosis for each case.

What predictions would the constraint attunement hypothesis make for this experiment? Coughlin and Patel (1987), like virtually all other researchers in this area, did not conduct a thorough analysis of the goal-relevant constraints in the domain before they conducted their experiment. However, in the results section, they introduced an important piece of information: The symptoms in the endocarditis case have a specific intrinsic temporal order associated with them, whereas those in the arteritis case do not. Thus, order is a goal-relevant constraint for the symptoms of endocarditis but not for those of arteritis.

The constraint attunement hypothesis would therefore predict the usual expertise effect for the normally ordered version of both cases. However, the predictions for the randomized cases differ for the two diseases. For endocarditis, randomizing the order of the information should disrupt the temporal order of the processes inherent in the disease. From the perspective of the constraint attunement hypothesis, a critical goal-relevant constraint was eliminated, and one would therefore expect a significant reduction in the magnitude of the expertise effect as a result. For arteritis, the temporal order is not a goal-relevant constraint, and so one would still expect to see an expertise effect of the same magnitude as in the normally ordered case, despite the fact that the information is randomly ordered.

The proportions of critical propositions recalled for the endocarditis and arteritis cases are presented in Figures 8a and 8b, respectively. Post hoc statistical comparisons of means were not reported by Coughlin and Patel (1987), so again, we can only rely on an informal comparison of the means. Figure 8a replicates the classic findings from chess: There is an expertise effect for the normally ordered case but not for the randomly ordered version. This result is not surprising. The results in Figure 8b, however, are more difficult to explain from the perspective of previous theories. For the arteritis case, there was a main effect of expertise across both the normally ordered and the randomly

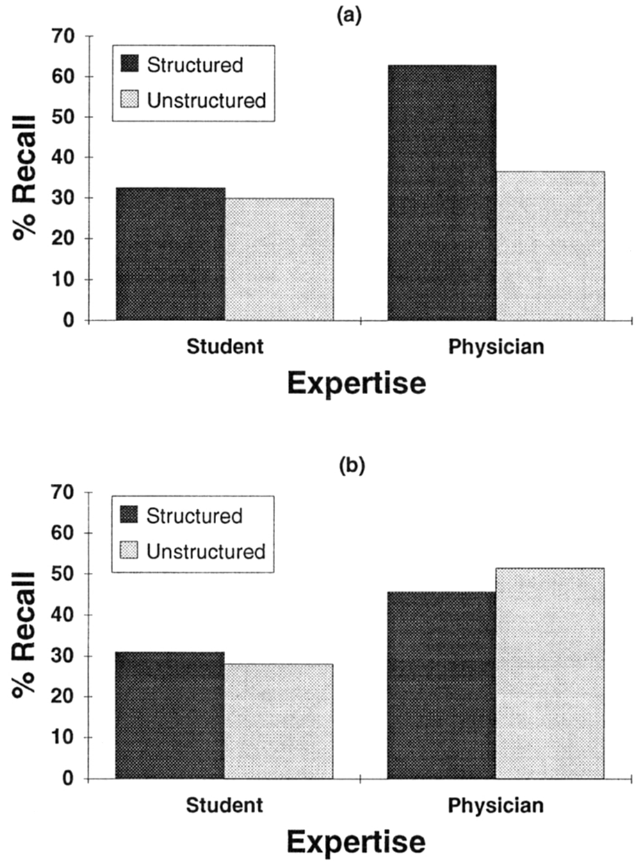

Figure 8. Mean percent of critical propositions recalled correctly by students and physicians from normally ordered and randomized texts of (a) an endocarditis case and (b) a temporal arteritis case, in Coughlin and Patel (1987).

ordered versions of the case. Not only was this result not anticipated by Coughlin and Patel, but it cannot be explained by existing theories of expertise effects in memory recall. Furthermore, as is discussed in more detail shortly, it explicitly contradicts the recent LTWM theory of Ericsson and Kintsch (1995). However, all of the results from this study are consistent with the predictions made by the constraint attunement hypothesis.

The Coughlin and Patel (1987) study also brings up a conceptual difficulty in the memory recall literature that is addressed by the constraint attunement hypothesis. As researchers have adapted the memory recall paradigm to domains other than chess, the question of what is the psychological equivalent of a random chess board arises in each domain. The answer to this question is not at all straightforward. In the domain of medicine for instance, Coughlin and Patel pointed out that several alternative approaches have been adopted (see also Norman et al., 1989, for additional approaches). What the constraint attunement hypothesis shows is that, in any domain, there can be varying degrees of randomness (or conversely, structure) that are defined by multiple levels of constraint. Thus, each domain should be analyzed to identify the levels of goal-relevant constraints that can productively be manipulated in memory recall research. Then, each stimulus should be analyzed to determine what subset of the available domain constraints is retained in the stimulus. This point is fundamental to the constraint attunement hypothesis, but it has not widely been acknowledged in the literature. The emphasis has instead been on investigating mental processes before understanding the constraints that the domain imposes on productive behavior. For example, the results of Coughlin and Patel cannot be accounted for unless one has a theory of how the structure of the environment constrains psychological processes. The constraint attunement hypothesis is the only theory of expertise effects in memory recall that directly addresses this critical issue.

Gobet and Simon (1994, 1996b): A Challenge?

There is a pair of recent empirical results in the literature that may seem to contradict the constraint attunement hypothesis. Gobet and Simon (1996b) reviewed 13 memory recall studies in chess and found that in almost all cases, skilled players' performance on random positions was slightly better than that of less skilled players in recall tasks with brief exposure times. These differences usually did not attain statistical significance, but Gobet and Simon hypothesized that this was because there were usually very few participants in each condition. They then analyzed the means from these 13 studies and found that the consistent advantage of experts on random trials in these experiments was statistically reliable. This is an important result because it contradicts the general view in the literature that there is no expertise effect on random positions with a brief presentation time.

Gobet and Simon (1994) conducted an experiment of their own to test this result in a more direct fashion. They included 21 participants in their study to ensure that the statistical power of the test was great enough to detect a small effect of expertise. The results show a slight, but nevertheless statistically significant, expertise advantage on random boards with brief exposure times, confirming the results obtained from the analysis of the 13 earlier studies. Gobet and Simon then provided a theoretical explanation for this result on the basis of an extension of the Chase and Simon (1973b) chunking theory that incorporates retrieval structures.

On first glance, this result may seem to contradict directly the constraint attunement hypothesis. After all, if there are no goal-relevant constraints in the stimulus (i.e., if it is completely random), then why should there be an expertise advantage? Interestingly, the constraint attunement hypothesis can account for this result. Gobet and Simon (1994) generated their random positions by using an accepted and frequently used method; namely, by randomly assigning the pieces from a game position to squares on the chessboard. However, the game positions were selected according to the following criteria (Gobet & Simon, 1994, p. 8): (a) it was White's turn to move, (b) the position was not in the middle of a sequence of exchanges, (c) the positions were reached after approximately 20 moves, and (d) the game was played by Masters or Grandmasters. Although the first two criteria are not critical for the present purposes, the latter two are because they constrain not just the positions of the pieces on the board but also what pieces are still on the board. The former are destroyed by the randomization manipulation, but the latter are not, meaning that the "random" positions inherit the selection criteria used to identify the game positions. To take but one example, there will always be a King on each side in the random positions, and it is well-known that the King provides an important source of information in chess (e.g., Gobet & Simon, 1996a). Thus, such stimuli are not really random because the pieces remaining on the board are constrained by the strategies governing midgame positions between Masters or Grandmasters. As far as we know, this observation has never before been made in the literature.

Although this may seem to be a minute detail, it has several important theoretical implications. First, it means that the results obtained by Gobet and Simon (1994, 1996b) are consistent with the constraint attunement hypothesis. There is still some, albeit very little, constraint in the random positions, and so a significant but very small expertise advantage would be expected. Second, this observation also invalidates one of the conclusions made by Gobet and Simon. They inferred that the slight expertise effects on random positions rule out other theories of chess expertise.

> It is hard to see how theories of chess skill based mainly on level of processing (Lane & Robertson, 1979) or high-level knowledge (Cooke et al., 1993) can account for this observed, if modest, effect with random positions. (Gobet & Simon, 1996b, p. 162)

The logic of this inference is correct, but the premise on which it is based (i.e., that the positions are truly random) is not. Thus, contrary to what Gobet and Simon claimed, the observed expertise effect is consistent with these other theories because experts' conceptual knowledge can be used to extract the little structure that exists in these stimuli. Third, our observation leads to a new experiment that contrasts our theory with that of Gobet and Simon (1994, 1996b). The two theories make different predictions for a truly random chess position, in which both the position and the selection of the pieces are randomly determined (as far as we know, this procedure has not been adopted before). Gobet and Simon (1996b) predicted a small but significant expertise effect for completely random positions, whereas our theory would predict no expertise effect.

Goldin (1978): An Anomaly?

A study by Goldin (1978, Experiment 1) provides another challenge to the constraint attunement hypothesis. The main focus here was comparing the memory recall performance of expert and less skilled participants on typical and atypical chess positions, the hypothesis being that typical positions would be recalled more accurately by all participants. The information theoretic analysis of the redundancy in chess conducted by de Groot (1966) suggests a close connection between typicality and constraint. His results indicate that prototypical chess positions are in fact more highly constrained than less typical positions. Thus, the constraint attunement hypothesis predicts a Typicality × Expertise interaction, with the expertise effect being greater on typical than atypical positions.

Goldin (1978) observed a significant main effect of typicality, with performance on the typical positions being better than that on the atypical positions. However, the Typicality × Expertise interaction predicted by the constraint attunement hypothesis was not statistically significant ($p > .10$). There are several reasons to reconsider this apparent negative result. First, the statistical power of the experiment seems to have been weak. There were only 4 participants in each expertise group. Also, only two levels of expertise were included in the study, neither of them representing novices; both groups were rated competent chess players. Moreover, the main effect for expertise was not statistically significant, a rarity in this area of research. Second, the recall data were averaged across five chess positions, consisting of opening positions, middle-game positions, and end-game positions. The average number of pieces in each type of position differed substantially (29.8 for openings, 25.4 for middle games, and 17.5 for end games). There are data to suggest that expertise effects would not be uniform across these positions (Cooke et al., 1993; Saariluoma, 1984). Therefore, averaging the data over these positions may have masked the expertise effects. Finally, the recall tests in this study were preceded by a 15-min period during which participants studied a chess game. This study period may have affected recall performance.

Therefore, while the results from Goldin (1978) contradict the constraint attunement hypothesis, there are several reasons to suggest that the experiment was not a sensitive, unconfounded test of expertise effects in memory recall performance on prototypical chess positions. Only further experimentation can determine whether the Goldin study was flawed or whether the constraint attunement hypothesis is in need of revision.

Lane and Robertson (1979)

So far, in this review we have focused on the presence of goal-relevant constraints as a prerequisite for an expertise effect. However, this is a necessary but not sufficient condition. As the name suggests, the constraint attunement hypothesis also predicts that participants must be attuned to those constraints for an expertise effect to be observed. Just because information is present in the stimulus does not mean that it will be picked up. There is at least one memory recall study that can be used to evaluate this prediction.

Lane and Robertson (1979) conducted a recall study in chess on the basis of the levels of processing theory in which they varied the instructions given to participants. Half of the participants received one trial with semantic-orienting instructions (decide which side had the advantage and determine the best move), and half received one trial with formal-orienting instructions (count the number of pieces on black squares and those on white squares). Each of these trials measured incidental learning because participants were not told that their memory recall performance was going to be tested after they performed the required task. In addition, both participant groups also performed an intentional learning task in which they were told to memorize the position of the pieces on the board. Memory performance on this intentional-learning trial served as a control condition for the two types of incidental learning. The participants in Experiment 1 were nonrated chess players. Experiment 2 consisted of a replication with more expert, rated chess players.

The results are presented in Figure 9. In both experiments, the formal-orienting manipulation led to a statistically significant decrease in recall performance compared with either the seman-

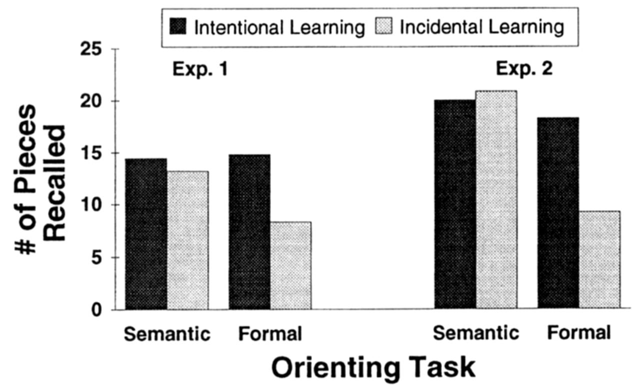

Figure 9. Number of chess pieces recalled correctly during intentional and incidental learning as a function of orienting task for Experiments 1 and 2 in Lane and Robertson (1979). Exp. = experiment.

tic-orienting condition or the intentional-learning condition (the latter two were not significantly different). This finding is consistent with the constraint attunement hypothesis because the formal-orienting instructions caused participants to divert their attention away from the goal-relevant constraints in the stimulus. Unfortunately, no statistical comparisons were made between expertise groups. However, the data in Figure 9 suggest that, under formal-orienting instructions, the effect of expertise is minimal. This result is interesting because it suggests that, when not attuned to goal-relevant constraints, there is very little difference between the recall performance of experts and lesser players, as predicted by the constraint attunement hypothesis.

Conclusion

All but one of the results reviewed in this section were consistent with predictions generated by the constraint attunement hypothesis, thereby providing a significant amount of empirical support for the theory. The close degree of agreement between theory and data is particularly notable because many of the results that were reviewed were anomalous in two senses: First, they were surprising because they were not predicted by the authors, and, second, they still could not be accounted for, a posteriori, by other competing theories. This consistency suggests that the constraint attunement hypothesis represents a unique contribution to our understanding of the relationship between domain expertise and memory recall.

A Remaining Theoretical Challenge

Another challenge to theories of expertise effects in memory recall is posed by the set of studies that has failed to show an expertise effect for overall memory recall of meaningful stimuli. One of these studies is in the domain of process control (Moray et al., 1993), but the remainder are all in the medical domain (Claesen & Boshuizen, 1985; Hassebrock et al., 1993; Patel & Frederiksen, 1984; Patel et al., 1986, 1990). Why have these studies failed to replicate the findings observed in chess and in so many other domains?

In the study conducted by Moray et al. (1993), participants were required to observe a dynamic scenario showing the behavior of a nuclear power plant under various fault conditions on a computer display. There were three groups of participants: undergraduate mechanical engineering students who were familiar enough with the domain to be able to interpret the stimuli, graduate mechanical and nuclear engineering students who had more domain expertise in the form of coursework, and licensed nuclear power plant operators who had more expertise in the form of hands-on experience. In one experimental condition, participants were required to recall the relative values of all of the displayed variables (35 in total) at the end of the trial. In addition, participants were asked to diagnose the nature of the abnormality that generated the symptoms observed during the trial. Overall memory recall performance failed to show a significant effect of expertise. However, in another experimental condition, participants' memory was assessed by having them answer a number of yes–no questions in addition to diagnosing the trial. These questions were at a higher level of abstraction; instead of having to remember the relative value of individual variables (e.g., steam generator level), the questions were directed at more general, functional issues (e.g., "Did the steam generator ever fail to act as a heat sink for the primary?"). Interestingly, the qualitative memory measures did show a significant expertise effect.

It seems that experts were able to remember the functional state of the system but not the detailed variable values. From the viewpoint of reconstructive memory (Bartlett, 1932), experts were able to recall the "gist" of the stimulus but were not able to reconstruct the lower level details. In contrast, in chess, experts were able to use their memory of gist to reconstruct the position of the individual pieces. Why were participants in the Moray et al. (1993) study not able to do this as well?

This question has not been addressed in the literature, but the difference can be accounted for by differences in the structure of knowledge between the two domains. In process control, it is generally not possible to reconstruct lower level (physical) details from knowledge of high-order (functional) relations because there is a very loose coupling between physical and functional levels of representation. More specifically, there is a bounded but infinite number of specific variable values that can account for a given functional state. For example, if one remembered that there was a leak in a reservoir after viewing a scenario of a process control system's behavior, this knowledge alone would not allow one to reconstruct what the reservoir input and output flow rates were. It is important to note that this is an analytically demonstrable fact and is not open to empirical falsification. In chess, there is a tighter coupling between physical and functional levels of representation in that the precise physical location of a constellation of pieces has a very important role in determining the meaning of a position (Gobet & Simon, 1996a). As a result, memory of the functional state of a chess board (e.g., Queen's Gambit) provides a great deal of help in reconstructing the specific position of the pieces. This allows experts to reconstruct many lower level details merely from knowledge of high-order relations (Cooke et al., 1993; de Groot & Gobet, 1996).

Can the lack of expertise effects in several memory recall studies in the medical domain also be explained in the same way? There is not enough evidence currently available to answer this question. However, there are a few empirical results that suggest that the answer may be yes. For example, Patel and Frederiksen (1984) found no expertise effect for memory for the stimulus material but did find an expertise effect for the number of propositions that were inferred from the stimulus, with experts inferring more than less skilled participants. Similar findings were obtained by Patel et al. (1986). These findings can be viewed as analogous to those of Moray et al. (1993), if one considers inferred information to be equivalent to (or to result from) higher order knowledge. With this interpretation, the experts in these medical studies seemed to have remembered the gist of the case but not the lower level details. However, in another study of recall in medicine, Patel et al. (1990) failed to replicate the expertise advantage for higher order information. Thus, the influence of the degree of domain coupling on memory recall performance is uncertain. Studies explicitly designed to address this issue need to be conducted.

To conclude, the lack of expertise effects in several recall studies presents a challenge to existing theories of expertise effects in memory recall. Why would chunks, or retrieval structures, lead to expertise advantages in some domains but not in others? This question cannot solely be answered by digging deeper into the psychological mechanisms inside the human. Only a detailed investigation of the nature of the constraints in the environment beyond the human skin (cf. Bentley, 1941) can hope to explain these differences. Because it has a psychologically relevant theory of the characteristics of various domains of expertise, the constraint attunement hypothesis can account for some of these findings. A comparison of the degree of coupling between constraints in the domains of chess and process control suggests that these anomalous results may be explained by further investigating differences in the structure of different domains of expertise.

Discussion

Relation to Other Research on Human Memory

Throughout this article, we have noted that the constraint attunement hypothesis is a novel theory of expertise effects in memory recall. However, what we did not mention is that it has many similarities with some existing work on human memory. In particular, David Rubin and colleagues have conducted a number of studies showing the impact of multiple constraints on LTM.

An early example of this line of research is Rubin and Wallace (1989), who looked at participants' ability to retrieve rhymes from LTM. They found that performance could be explained by looking at the role of multiple constraints in the stimulus (viz., meaning, rhyme, and other sound and orthographic properties). These multiple constraints work together in an interrelated fashion, increasing the number of cues for recall, limiting the response choices that are available to participants, and thereby reducing their memory load. For example,

> If you are looking for a word to end a line that has a particular rhyme and a particular meaning, chances are that it does not exist; but if it does exist, chances are that it is the only solution. (Rubin, 1995, p. 94)

The situation is similar to Vignette 1 in Figure 1, in which there are very few degrees of freedom available.

Hyman and Rubin (1990) in a follow-up study investigated participants' recall of Beatles song lyrics from LTM. Again, multiple constraints were identified in the stimulus material: poetics, rhythm, and meaning. These work together to reduce the possibilities available, thereby enhancing the accuracy of reconstructive recall from LTM. Importantly, Hyman and Rubin observed that when errors were made, they were usually consistent with the aforementioned constraints.

The role of multiple constraints has also been explored in the context of recall of ballads from LTM (Rubin, 1995; Rubin, Wallace, & Houston, 1993; Wallace & Rubin, 1988). The general findings obtained in earlier studies have been observed here as well. In addition, several other findings consistent with the constraint attunement hypothesis have also been obtained. For example, Wallace and Rubin (1988) found that with less constraints in the stimulus, there is greater variability in recall. Moreover, Rubin et al. (1993) examined the beginning of expertise in memory for ballads and found that, with experience, participants became more sensitive to the rhyme and line organization constraints in the ballads.

There is a great deal of conceptual similarity between the work of Rubin and colleagues and the theory presented in this article. Nevertheless, there are some important differences as well. First, the two lines of research, although both dealing with human memory, have been applied to different phenomena. The work of Rubin has focused on representative situations in which people are required to recall material from LTM, whereas the constraint attunement hypothesis focuses on explaining recall performance after a very brief presentation of the stimulus material. Moreover, in the latter case the recall task is usually contrived, whereas in Rubin's case it is usually intrinsic. Second, with one exception (Rubin et al., 1993), Rubin's work has not focused on comparisons between participants of different levels of expertise. In contrast, the effects of expertise play a central role in the constraint attunement hypothesis. Third, Rubin has identified multiple sources of constraint in the stimulus material in each of his studies, but he has not developed a framework for systematically organizing multiple constraints in the same fashion across multiple domains. The abstraction hierarchy fulfills this role in the constraint attunement hypothesis. Therefore, while similar to Rubin's work, the constraint attunement hypothesis leads to significant independent insights.

Constraints on a Process Theory

The constraint attunement hypothesis is best viewed as a stepping stone to developing viable process theories of expertise effects in memory recall. It can be used to rule out, or at least revise, process theories (or cognitive models) that do not generate expertise effects under the conditions, or of the magnitude, predicted by the theory. This role is valuable because it can make the search for a viable process theory more efficient than it has been.

LTWM. This function can be illustrated by comparing the constraint attunement hypothesis with LTWM (Ericsson & Kintsch, 1995). We have already mentioned that it is difficult to see how retrieval structures could account for expertise effects in domains in which memory recall is a contrived, rather than an intrinsic, task. However, LTWM suffers from other difficulties as well. First, it claims that the magnitude of expertise effects is "related to the level of attained skill and to the amount of relevant prior practice" (Ericsson & Kintsch, 1995, p. 238). The data reviewed earlier show that this claim is, at best, incomplete. Expertise effects in memory recall are also determined by the amount of structure in the domain (and by active attunement to that structure). The omission of this important factor probably results from the fact that the influence of the environment is not explicitly accounted for in LTWM. Second, LTWM cannot account for the lack of expertise effects in several memory recall studies in the domain of medicine. The constraint attunement hypothesis suggests a possible explanation for this by examining differences in the tightness of the coupling between constraints across domains. Third, LTWM cannot account for Coughlin and Patel's (1987) findings. The constraint attunement hypothesis explains why very different results were obtained for the endocarditis and arteritis cases. Randomly ordering the propositions in the endocarditis case disrupted the temporal processes inherent in the disease, and so an expertise advantage was not observed. In contrast, for the arteritis case, order is not a goal-relevant constraint, so an expertise effect was observed on both the normally ordered and randomly ordered versions of the case.

These results directly conflict with LTWM because Ericsson and Kintsch (1995) predicted that "experts should be relatively insensitive to the order in which information is presented" (p. 236). Although this lack of agreement between theory and data may seem insignificant, given all the other data accounted for by LTWM, the anomaly is crucial when one considers that "the strongest evidence for retrieval structures concern the ability of experts to independently store pieces of information when they are presented out of their normal context *in scrambled order* [italics added]" (Ericsson & Kintsch, 1995, p. 238). As mentioned earlier, the Coughlin and Patel (1987) data are of great theoretical interest because they cannot be accounted for without an explicit theory of how the goal-relevant constraints in the environment shape memory recall performance. Because the LTWM theory does not deal with the environment, its predictions are not supported by the Coughlin and Patel data. This analysis shows how a product theory that explicitly accounts for the environment's influence on behavior can make a unique contribution by guiding the search for a viable process theory of the same phenomenon.

Template theory. The relationship between our theory and the *template* theory of chess skill recently proposed by Gobet and Simon (1996a, 1996c) is also worthy of investigation. A template is a mental construct that is intended to represent a pattern of chess pieces found in classes of familiar openings and lines of play (e.g., Ruy Lopez opening). It contains more pieces than the chunks identified in the original Chase and Simon (1973b) theory. Templates play the same psychological role in memory as do retrieval structures, except that the former are acquired implicitly in the course of chess study, whereas the latter must be acquired deliberately and consciously to aid memorization. Template theory tries to overcome the problems that skilled-memory theory and EPAM-based models have in accounting for expertise effects in memory recall tasks in chess.

Comparing the constraint attunement hypothesis and the recent template theory would help address two questions. First, what is the origin of the mental construct of a template? If the constraint attunement hypothesis is correct, then templates correspond to a layer of constraint (probably at the strategic level) that is contained in the stimulus material. If so, then templates would have their origin in the structure of the environment. Second, can template theory be generalized beyond the domain of chess? If so, it could provide a basis for a process theory that accounts for the breadth of data currently explained by the constraint attunement hypothesis product theory. Both of these questions show how the constraint attunement hypothesis can help to constrain process theories of expertise in recall.

Relationship to a Theory of Expertise

The ultimate goal of theorizing should be to develop a theory of expertise in domain-relevant problems, not just expertise effects in memory recall. As several researchers have pointed out, this is a point that has gone relatively ignored for some time (Hassebrock et al., 1993; Holding, 1985). For instance, no computer program has ever been developed that can both play chess at an expert level and also exhibit expert-level performance in a memory recall task (Koedinger & Anderson, 1990). This fact calls into question how these two activities, chess performance and memory performance, are related. In this subsection we try to make some modest progress in this area by sketching some connections between the constraint attunement hypothesis and existing research on skilled performance.

One important finding from the expertise literature is that experts' information search activities are more economic than those of novices (Shanteau, 1992b). Novices cannot distinguish between relevant and irrelevant information and, consequently, search through a great deal of information. Experts, on the other hand, are able to evaluate effectively what is relevant in specific contexts and are thereby able to focus primarily, or exclusively, on the information that is of value for the task at hand. As Shanteau pointed out, the question of how experts know what kind of information to use has not been effectively dealt with in the literature.

The constraint attunement hypothesis may be of use here because it suggests that one of the hallmarks of expertise is attunement to goal-relevant constraints. Perhaps it is the focus provided by this educated attention that accounts for experts' economic information search behaviors. Other research has illustrated the role that high-level constraints can play in the processing of domain knowledge by experts (Cooke et al., 1993). If this view turns out to be correct, then skilled memory is actually only a by-product of chess skill, rather than the other way around. The fact that skilled memory is not a causal prerequisite for skilled domain performance is a view that has gained empirical support in recent years (Cooke et al., 1993; Ericsson & Harris, 1990; Holding, 1985).

Another important characteristic of expertise that has emerged from the literature (e.g., Koedinger & Anderson, 1990; Patel & Groen, 1991) is experts' ability to engage productively in forward reasoning (i.e., to reason from the evidence given to the goal state). The relationship between the constraint attunement hypothesis and forward reasoning is explained by Patel and Groen (1991):

> Forward reasoning . . . is highly error prone in the absence of adequate domain knowledge because there are no built-in checks in the legitimacy of the inferences. Therefore, success in using forward reasoning is *constrained by the environment* [italics added] because a great deal of relevant knowledge is necessary. (p. 94)

Thus, attunement to goal-relevant constraints seems to be a prerequisite for enjoying the advantages of forward reasoning.

As for chess in particular, the constraint attunement hypothesis seems to have strong relationships to two theories of skill: the SEEK theory of Holding (1992) and the apperception–restructuring theory of Saariluoma (1992), both of which place a great deal of importance on information search. For example, Holding's theory suggests that chess skill rests to a great extent on the ability to anticipate the consequences of moves. Attunement to goal-relevant constraints would seem to be an important part of such an ability. Saariluoma's theory suggests that the selection of moves is based on the apperception of critical paths that allow one to move from the initial position to a goal position. Again, it seems that constraint attunement would be a prerequisite for such behavior. The possibilities for integrating the constraint attunement hypothesis with existing theories of chess skill seem good. However, neither Holding's nor Saariluoma's theory specifically addresses in detail how chess skill accounts for expertise effects in recall tasks. Relating these theories to the constraint attunement hypothesis should lead to productive insights that could lead to an integrated account of chess skill and expertise effects in memory recall.

Ecological Theories of Skill Acquisition

The product theory proposed here also makes a contribution by extending ecological theories of skill acquisition to new ground. There are several reasons why this is worthy of note. First, it shows that ecological theories are relevant to cognitive phenomena, not just to perceptual-motor behavior. The broad relevance of ecological theories of skill acquisition opens the door for a cumulative theory of skill acquisition that can account for learning in a wide variety of tasks, perhaps taking the place of a large number of microtheories of very specific, isolated phenomena (cf. Newell, 1973). Second, the work presented here provides further evidence of the power of developing product theories before process theories. This approach is standard in ecological psychology, but as mentioned earlier, it has not widely been adopted by other experimental psychologists. Third, the constraint attunement hypothesis also shows the theoretical importance of studying the environmental constraints on productive behavior. Some, although by no means all, cognitive phenomena have their origins in constraints in the environment. As a result, to have a viable theory of psychology, one also needs a psychologically relevant theory of the environment (Brunswik, 1956; J. J. Gibson, 1979). The value of this metatheoretical assumption extends well beyond memory recall.

Future Research

There are several research issues that are motivated by the constraint attunement hypothesis. First, the relationship between prototypical stimuli and expertise effects should be investigated. Only one study has addressed this issue (Goldin, 1978), and its results are inconsistent with the constraint attunement hypothesis. However, there are strong reasons to question the sensitivity and internal validity of that experiment with respect to isolating memory recall effects.

Second, there is a set of results that has not been accounted for by any theory of expertise effects in memory recall. An analysis of these results has led to the suggestion that investigation of the differences between problem domains in terms of the degree of coupling between physical and functional levels of representation may explain these anomalous results. A rigorous evaluation of this conjecture would require one to thoroughly investigate the constraint structure of various problem domains and then compare the degree of coupling between levels.

Third, the utility of the abstraction hierarchy to broader research issues in expertise should be investigated. Many studies in the expertise literature, conducted in quite disparate domains, have observed that the knowledge representation of experts is organized in a functional hierarchy (Glaser & Chi, 1988). However, the precise structure of the hierarchical organization observed in these studies has typically not been well defined, neither within nor across domains (e.g., Chi, Feltovich, & Glaser, 1981). The abstraction hierarchy, defined as a stratified hierarchy with a means–end relations between levels, is an attempt at making more explicit exactly what such a functional hierarchy might look like for various domains. It makes the search for identifying and organizing domain constraints directed, rather than haphazard. Having a common framework for representing different problem domains would also facilitate the process of assessing the generalization of research results on expertise. Several authors have pointed to the importance of examining differences between domains and how such differences can mediate the nature of expertise (e.g., Hammond, Hamm, Grassia, & Pearson, 1987; Shanteau, 1992a; Stewart, Roebber, & Bosart, 1997). This, in turn, would make an integrated, cumulative theory of expertise a viable pursuit.

Fourth, the extension of the constraint attunement hypothesis to a process theory of expertise effects in memory recall should be pursued. Two specific paths seem promising. First, the relationship between the constraint attunement hypothesis and template theory (Gobet & Simon, 1996a, 1996c) should be investigated. Templates may be the knowledge structures that are responsible for the attunement to goal-relevant constraints predicted by the constraint attunement hypothesis. To see if this is so, researchers must examine the constraints in the problem domain to see if they correspond with the mental structures posited by the theory. The latter have received far more attention than the former.

A second important question to consider in developing a process theory is what are the psychological resources experts use to attune themselves to domain constraints? There are strong reasons to suggest that the expertise effects in memory recall can be driven be either of two completely different psychological processes: one being based on perceptual pattern recognition and another being based on conceptual knowledge (Cooke et al., 1993; Ericsson & Staszewski, 1989; Norman et al., 1989; Vicente, 1988). The original Chase and Simon (1973b) theory emphasizes the perceptual view, but more recent studies have begun to pay more attention to the conceptual view (e.g., Cooke et al., 1993; Egan & Schwartz, 1979; Vicente, 1992). A detailed process theory should explain how each of these types of processes can lead to expertise effects in memory recall.

Fifth, and perhaps most important of all, the possibility of linking the constraint attunement hypothesis to theories of expert performance in domain-relevant tasks deserves serious consideration. Otherwise, a unified theory of expertise will continue to remain elusive. A sketch of how to proceed in this direction has been provided here.

Conclusions

The study of expertise effects in memory recall performance has a long history in psychology. During this time, a large number of different domains of expertise have been experimentally investigated, and a very large amount of data has been generated. Most attempts to organize these data have taken the form of process theories—the perceptual chunking hypothesis of Chase and Simon (1973a, 1973b) being the best known contender. There is a general consensus that Chase and Simon's theory is seriously limited, yet no alternative theory has taken its place, despite several attempts. In the meantime, a number of anomalous findings have accumulated in the literature, many awaiting a theoretical explanation. The suggestion made by Gobet (1994) in the context of chess is generally applicable: A new theory is required to account for the wealth of data already existing in this area.

As a precursor to developing a viable process theory of these data, this article has proposed a product theory of expertise effects in memory recall performance on the basis of an ecological approach to skill acquisition. The constraint attunement hypothesis was suggested as a candidate for predicting under what conditions there will be an expertise advantage and what factors determine how large that advantage will be. It is the only theory of expertise effects in memory recall that has a framework for both identifying and organizing the multiple levels of goal-relevant constraint in a domain. As a result, it is the only such theory that systematically accounts for the contribution of the structure of the environment to behavior. The essence of the constraint attunement hypothesis is that experts become adapted to the constraints inherent in their domain of expertise (cf. Holyoak, 1991). This view of skill acquisition already has a great deal of empirical support, primarily from the perceptual learning literature (e.g., E. J. Gibson, 1991). More recently, the view of expertise as extreme adaptation to domain-specific constraints has also received more attention in the mainstream cognitive psychology literature (Ericsson & Lehmann, 1996), but very little attention has been given to describing these domain-specific constraints (for an exception, see Stewart et al., 1997). This observation is crucial because if one only has a loose, intuitive description of the object of adaptation, then the claim of adaptation is very difficult to defend.

A thorough review of the literature reveals that the constraint attunement hypothesis also accounts for a large number of results in the memory recall literature. Interestingly, most of these results were anomalies in that they had not been predicted by the researchers that generated them. Furthermore, few of these individual data can be explained, a posteriori, by other competing theories of expertise effects in memory recall. Even more important, no other competing theory can account for all of these data. Therefore, the constraint attunement hypothesis makes a novel and significant contribution to the literature.

Because it is a product theory, the constraint attunement hypothesis defines a path for future attempts at developing a process account of expertise effects in memory recall; any process theory that does not satisfy the predictions of the constraint attunement hypothesis is unlikely to be a viable candidate. Furthermore, the constraint attunement hypothesis also ties into contemporary accounts of expertise, thereby linking exceptional memory recall performance with exceptional performance at domain-relevant problems. These contributions show, specifically, the value of defining product theories before attempting to construct process theories and, more generally, the value of an ecological approach to psychology.

References

Adelson, B. (1981). Problem solving and the development of abstract categories in programming languages. *Memory & Cognition, 9,* 422–433.

Allard, F., Graham, S., & Paarsalu, M. E. (1980). Perception in sport: Basketball. *Journal of Sport Psychology, 2,* 14–21.

Anderson, J. R. (1990). *The adaptive character of thought.* Hillsdale, NJ: Erlbaum.

Barfield, W. (1986). Expert–novice differences for software: Implications for problem solving and knowledge acquisition. *Behaviour and Information Technology, 5,* 15–29.

Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology.* New York: Macmillan.

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

Beal, A. L. (1985). The skill of recognizing musical structures. *Memory & Cognition, 13,* 405–412.

Bentley, A. F. (1941). The human skin: Philosophy's last line of defense. *Philosophy of Science, 8,* 1–19.

Billman, D., & Shaman, D. (1990). Strategy knowledge and strategy change in skilled performance: A study of the game Othello. *American Journal of Psychology, 103,* 1–166.

Bisantz, A. M., & Vicente, K. J. (1994). Making the abstraction hierarchy concrete. *International Journal of Human–Computer Studies, 40,* 83–117.

Brunswik, E. (1956). *Perception and the representative design of psychological experiments* (2nd ed.). Berkeley: University of California Press.

Chang, K.-T., Lenzen, T., & Antes, J. (1985). The effect of experience on reading topographic relief information: Analyses of performance and eye movements. *The Cartographic Journal, 22,* 88–94.

Charness, N. (1976). Memory for chess positions: Resistance to interference. *Journal of Experimental Psychology: Human Learning and Memory, 2,* 641–653.

Charness, N. (1979). Components of skill in bridge. *Canadian Journal of Psychology, 33,* 1–16.

Charness, N. (1981). Aging and skilled problem solving. *Journal of Experimental Psychology: General, 110,* 21–38.

Chase, W. G., & Ericsson, K. A. (1981). Skilled memory. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 141–189). Hillsdale, NJ: Erlbaum.

Chase, W. G., & Ericsson, K. A. (1982). Skill and working memory. *The Psychology of Learning and Motivation, 16,* 1–58.

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

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

Chi, M. T. H. (1978). Knowledge structures and memory development. In R. S. Siegler (Ed.), *Children's thinking: What develops?* (pp. 73–96). Hillsdale, NJ: Erlbaum.

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

Chiesi, H. L., Spilich, G. J., & Voss, J. F. (1979). Acquisition of domain-related information in relation to high and low domain knowledge. *Journal of Verbal Learning and Verbal Behavior, 18,* 257–273.

Claesen, H. F. A., & Boshuizen, H. P. A. (1985). Recall of medical information by students and doctors. *Medical Education, 19,* 61–67.

Cooke, N. J., Atlas, R. E., Lane, D. M., & Berger, R. C. (1993). Role of high-level knowledge in memory for chess positions. *American Journal of Psychology, 106,* 321–351.

Coughlin, L. D., & Patel, V. L. (1987). Processing of critical information by physicians and medical students. *Journal of Medical Education, 62,* 818–828.

Deakin, J. M., & Allard, F. (1991). Skilled memory in expert figure skaters. *Memory & Cognition, 19,* 79–86.

de Groot, A. D. (1965). *Thought and choice in chess.* The Hague, the Netherlands: Mouton. (Original work published 1946)

de Groot, A. D. (1966). Perception and memory versus thought: Some old ideas and recent findings. In B. Kleinmuntz (Ed.), *Problem solving: Research, method, and theory* (pp. 19–50). New York: Wiley.

de Groot, A. D., & Gobet, F. (1996). *Perception and memory in chess: Studies in the heuristics of the professional eye.* Assen, the Netherlands: Van Gorcum.

Di Persio, T., Isbister, D., & Shneiderman, B. (1980). An experiment using memorization/reconstruction as a measure of programmer ability. *International Journal of Man–Machine Studies, 13,* 339–354.

Duncker, K. (1945). On problem solving. *Psychological Monographs, 58,* Whole No. 270.

Egan, D. E., & Schwartz, B. J. (1979). Chunking in recall of symbolic drawings. *Memory & Cognition, 7,* 149–158.

Eisenstadt, M., & Kareev, Y. (1975). Aspects of human problem solving. In D. A. Norman & D. E. Rumelhart (Eds.), *Explorations in cognition* (pp. 308–346). San Francisco: Freeman.

Engle, R. W., & Bukstel, L. (1978). Memory processes among bridge players of differing expertise. *American Journal of Psychology, 91,* 673–689.

Ericsson, K. A., & Harris, M. S. (1990, November). *Expert chess memory without chess knowledge: A training study.* Paper presented at the 31st Annual Meeting of the Psychonomic Society, New Orleans, LA.

Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review, 102,* 211–245.

Ericsson, K. A., & Lehmann, A. C. (1996). Expert and exceptional performance: Evidence of maximal adaptation to task constraints. *Annual Review of Psychology, 47,* 273–305.

Ericsson, K. A., & Polson, P. G. (1988). An experimental analysis of the mechanisms of a memory skill. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 305–316.

Ericsson, K. A., & Staszewski, J. J. (1989). Skilled memory and expertise: Mechanisms of exceptional performance. In D. Klahr & K. Kotovsky (Eds.), *Complex information processing: The impact of Herbert A. Simon* (pp. 235–267). Hillsdale, NJ: Erlbaum.

Frey, P. W., & Adesman, P. (1976). Recall memory for visually presented chess positions. *Memory & Cognition, 4,* 541–547.

Frijda, N. H., & de Groot, A. D. (1981). *Otto Selz: His contribution to psychology.* The Hague, the Netherlands: Mouton.

Garland, D. J., & Barry, J. R. (1991). Cognitive advantage in sport: The nature of perceptual structures. *American Journal of Psychology, 104,* 211–228.

Gibson, E. J. (1969). *Principles of perceptual learning and development.* New York: Appleton-Century-Crofts.

Gibson, E. J. (1991). *An odyssey in learning and perception.* Cambridge, MA: MIT Press.

Gibson, J. J. (1966). *The senses considered as perceptual systems.* Boston: Houghton-Mifflin.

Gibson, J. J. (1979). *The ecological approach to visual perception.* Boston: Houghton-Mifflin.

Gilhooly, K. J., Wood, M., Kinnear, P. R., & Green C. (1988). Skill in map reading and memory for maps. *The Quarterly Journal of Experimental Psychology, 40A,* 87–107.

Glaser, R., & Chi, M. T. H. (1988). Overview. In M. T. H. Chi, R. Glaser, & M. J. Farr (Eds.), *The nature of expertise* (pp. xv–xxviii). Hillsdale, NJ: Erlbaum.

Gobet, F. (1993). A computer model of chess memory. In *Proceedings of the 15th Annual Conference of the Cognitive Science Society* (pp. 463–468). Hillsdale, NJ: Erlbaum.

Gobet, F. (1994). *Memory in chess players: Chunks, schemata, or both?* (Complex Information Processing Working Paper No. 517). Pittsburgh, PA: Department of Psychology, Carnegie Mellon University.

Gobet, F., & Simon, H. A. (1994). *Role of presentation time in recall of game and random chess positions* (Complex Information Processing Working Paper No. 524). Pittsburgh, PA: Department of Psychology, Carnegie Mellon University.

Gobet, F., & Simon, H. A. (1996a). Recall of random and distorted chess positions: Implications for the theory of expertise. *Memory & Cognition, 24,* 493–503.

Gobet, F., & Simon, H. A. (1996b). Recall of rapidly presented random chess positions is a function of skill. *Psychonomic Bulletin and Review, 3,* 159–163.

Gobet, F., & Simon, H. A. (1996c). Templates in chess memory: A mechanism for recalling several boards. *Cognitive Psychology, 31,* 1–40.

Goldin, S. E. (1978). Memory for the ordinary: Typicality effects in chess memory. *Journal of Experimental Psychology: Human Learning and Memory, 4,* 605–616.

Goldin, S. E. (1979). Recognition memory for chess positions. *American Journal of Psychology, 92,* 19–31.

Hammond, K. R., Hamm, R. M., Grassia, J., & Pearson, T. (1987). Direct comparison of intuitive and analytical cognition in expert judgement. *IEEE Transactions on Systems, Man, and Cybernetics, SMC-17,* 753–770.

Hassebrock, F., Johnson, P. E., Bullemer, P., Fox, P. W., & Moller, J. H. (1993). When less is more: Representation and selective memory in expert problem solving. *American Journal of Psychology, 106,* 155–189.

Hoffman, R. R., Shadbolt, N. R., Burton, A. M., & Klein, G. (1995). Eliciting knowledge from experts: A methodological analysis. *Organizational Behavior and Human Decision Processes, 62,* 129–158.

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

Holding, D. H. (1992). Theories of chess skill. *Psychological Research, 54,* 10–16.

Holding, D. H., & Reynolds, R. I. (1982). Recall or evaluation of chess positions as determinants of chess skill. *Memory & Cognition, 10,* 237–242.

Holyoak, K. J. (1991). Symbolic connectionism: Toward third-generation theories of expertise. In K. A. Ericsson & J. Smith (Eds.), *Toward a general theory of expertise* (pp. 301–336). Cambridge, MA: Cambridge University Press.

Hyman, I. E., Jr., & Rubin, D. C. (1990). Memorabeatlia: A naturalistic study of long-term memory. *Memory & Cognition, 18,* 205–214.

Itoh, J., Yoshimura, S., Ohtsuka, T., & Masuda, F. (1990). Cognitive task analysis of nuclear power plant operators for man–machine interface design. In *Proceedings of the ANS Topical Meeting on Advances in Human Factors Research on Man–Computer Interactions: Nuclear and Beyond* (pp. 96–102). La Grange Park, IL: American Nuclear Society.

Jongman, R. W. (1968). *Het oog van de meester* [The eye of the master]. Amsterdam: Van Gorcum.

Koedinger, K. R., & Anderson, J. R. (1990). Abstract planning and perceptual chunks: Elements of expertise in geometry. *Cognitive Science, 14,* 511–550.

Korf, R. E. (1987). Planning as search: A quantitative approach. *Artificial Intelligence, 33,* 65–88.

Lane, D. M., & Robertson, L. (1979). The generality of the levels of processing hypothesis: An application to memory for chess positions. *Memory & Cognition, 7,* 253–256.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information.* San Francisco: Freeman.

McKeithen, K. B., Reitman, J., Rueter, H. H., & Hirtle, S. C. (1981). Knowledge organization and skill differences in computer programmers. *Cognitive Psychology, 13,* 307–325.

Mesarovic, M. D., Macko, D., & Takahara, Y. (1970). *Theory of hierarchical, multilevel systems.* New York: Academic Press.

Moray, N., Jones, B. J., Rasmussen, J., Lee, J. D., Vicente, K. J., Brock, R., & Djemil, T. (1993). *A performance indicator of the effectiveness of human–machine interfaces for nuclear power plants* (Rep. No. NUREG/CR-5977). Washington, DC: United States Nuclear Regulatory Commission.

Myles-Worsley, M., Johnston, W. A., & Simons, M. A. (1988). The influence of expertise on X-ray image processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 553–557.

Neisser, U. (1987). From direct perception to conceptual structure. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 11–24). Cambridge, MA: Cambridge University Press.

Newell, A. (1973). You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In W. G. Chase (Ed.), *Visual information processing* (pp. 283–308). New York: Academic Press.

Norman, G. R., Brooks, L. R., & Allen, S. W. (1989). Recall by expert medical practitioners and novices as a record of processing of attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 1166–1174.

Patel, V. L., & Frederiksen, C. L. (1984). Cognitive processes in comprehension and knowledge acquisition by medical students and physicians. In H. G. Schmidt & M. L. de Volder (Eds.), *Tutorials in problem-based learning: New directions in training for the health professions* (pp. 143–157). Assen, the Netherlands: Van Gorcum.

Patel, V. L., & Groen, G. J. (1991). The general and specific nature of medical expertise: A critical look. In K. A. Ericsson & J. Smith (Eds.), *Toward a general theory of expertise* (pp. 93–125). Cambridge, England: Cambridge University Press.

Patel, V. L., Groen, G. J., & Arocha, J. F. (1990). Medical expertise as a function of task difficulty. *Memory & Cognition, 18,* 394–406.

Patel, V. L., Groen, G. J., & Frederiksen, C. H. (1986). Differences between medical students and doctors in memory for clinical cases. *Medical Education, 20,* 3–9.

Pattee, H. H. (1972). The nature of hierarchical controls in living matter. In R. Rosen (Ed.), *Foundations of mathematical biology: Vol. 1. Subcellular systems* (pp. 1–21). New York: Academic Press.

Pfau, H. D., & Murphy, M. D. (1988). Role of verbal knowledge in chess skill. *American Journal of Psychology, 101,* 73–86.

Rasmussen, J. (1985). The role of hierarchical knowledge representation in decisionmaking and system management. *IEEE Transactions on Systems, Man, and Cybernetics, SMC-15,* 234–243.

Rasmussen, J. (1986). *Information processing and human–machine interaction: An approach to cognitive engineering.* Amsterdam: North-Holland.

Reed, E. S. (1988). *James J. Gibson and the psychology of perception.* New Haven, CT: Yale University Press.

Reitman, J. (1976). Skilled perception in Go: Deducing memory structures from inter-response times. *Cognitive Psychology, 8,* 336–356.

Reynolds, R. I. (1982). Search heuristics of chess players of different calibers. *American Journal of Psychology, 95,* 383–392.

Richman, H. B., Staszewski, J. J., & Simon, H. A. (1995). Simulation of expert memory using EPAM IV. *Psychological Review, 102,* 305–330.

Rubin, D. C. (1995). *Memory in oral traditions: The cognitive psychology of counting-out rhymes, ballads, and epic.* New York: Oxford University Press.

Rubin, D. C., & Wallace, W. T. (1989). Rhyme and reason: Analyses of dual retrieval cues. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 698–709.

Rubin, D. C., Wallace, W. T., & Houston, B. C. (1993). The beginnings of expertise for ballads. *Cognitive Science, 17,* 435–462.

Saariluoma, P. (1984). *Coding problem spaces in chess: A psychological study.* Helsinki, Finland: Societas Scientiarum Fennica.

Saariluoma, P. (1985). Chess players' intake of task-relevant cues. *Memory & Cognition, 13,* 385–391.

Saariluoma, P. (1992). Error in chess: The apperception–restructuring view. *Psychological Research, 54,* 17–26.

Selz, O. (1922). *Zur psychologie des produktiven denkens und des irrtums* [On the psychology of productive thinking and of error]. Bonn, Germany: Freidrich Cohen.

Shanteau, J. (1992a). Competence in experts: The role of task characteristics. *Organizational Behavior and Human Decision Processes, 53,* 252–266.

Shanteau, J. (1992b). How much relevant information does an expert use? Is it relevant? *Acta Pscyhologica, 81,* 75–86.

Shneiderman, B. (1976). Exploratory experiments in programmer behavior. *International Journal of Computer and Information Sciences, 5,* 123–143.

Simon, H. A. (1981). *The sciences of the artificial* (2nd ed.). Cambridge, MA: MIT Press.

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

Sloboda, J. A. (1976). Visual perception of musical notation: Registering pitch symbols in memory. *Quarterly Journal of Experimental Psychology, 28,* 1–16.

Spilich, G. J., Vesonder, G. T., Chiesi, H. L., & 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, 18,* 275–290.

Starkes, J. L. (1987). Skill in field hockey: The nature of the cognitive advantage. *Journal of Sport Psychology, 9,* 146–160.

Starkes, J. L., Deakin, J. M., Lindley, S., & Crisp, F. (1987). Motor versus verbal recall of ballet sequences by young expert dancers. *Journal of Sport Psychology, 9,* 222–230.

Stewart, T. R., Roebber, P. J., & Bosart, L. F. (1997). The importance of the task in analyzing expert judgement. *Organizational Behavior and Human Decision Processes, 69,* 205–219.

Sweller, J., & Cooper, G. A. (1985). The use of worked examples as a substitute for problem solving in learning algebra. *Cognition and Instruction, 2,* 59–89.

Vessey, I. (1987). On matching programmer's chunks with program structures: An empirical investigation. *International Journal of Man–Machine Studies, 27,* 65–89.

Vicente, K. J. (1988). Adapting the memory recall paradigm to evaluate interfaces. *Acta Psychologica, 69,* 249–278.

Vicente, K. J. (1992). Memory recall in a process control system: A measure of expertise and display effectiveness. *Memory & Cognition, 20,* 356–373.

Vicente, K. J., & Brewer, W. B. (1993). Reconstructive remembering of the scientific literature. *Cognition, 46,* 101–128.

Vicente, K. J., Christoffersen, K., & Pereklita, A. (1995). Supporting operator problem solving through ecological interface design. *IEEE Transactions on Systems, Man, and Cybernetics, SMC-25,* 529–545.

Vicente, K. J., & de Groot, A. D. (1990). The memory recall paradigm: Straightening out the historical record. *American Psychologist, 45,* 285–287.

Vicente, K. J., & Rasmussen, J. (1990). The ecology of human–machine systems. II: Mediating "direct perception" in complex work domains. *Ecological Psychology, 2,* 207–250.

Wallace, W. T., & Rubin, D. C. (1988). "The Wreck of the Old 97": A real event remembered in song. In U. Neisser & E. Winograd (Eds.), *Remembering reconsidered: Ecological and traditional approaches to the study of memory* (pp. 283–310). Cambridge, England: Cambridge University Press.

Williams, M., Davids, K., Burwitz, L., & Williams, J. (1993). Cognitive knowledge and soccer performance. *Perceptual and Motor Skills, 76,* 579–593.

Wolff, A. S., Mitchell, D. H., & Frey, P. W. (1984). Perceptual skill in the game of Othello. *The Journal of Psychology, 118,* 7–16.

Appendix

Two Examples of an Abstraction Hierarchy

This appendix provides two examples showing how the abstraction hierarchy can be applied to identify the goal-relevant constraints in a particular problem domain. It is important to distinguish between two usages of the abstraction hierarchy: (a) as a problem space representation in which verbal protocols from problem-solving activities can be mapped, and (b) as a framework for identifying the multiple levels of goal-relevant constraint that can be presented in a stimulus in memory recall experiments. The focus here is on the latter (for the former, see Vicente et al., 1995, and Bisantz & Vicente, 1994).

There are at least two different ways of developing an abstraction hierarchy representation. For domains that are well understood (e.g., physical systems), analytical models identify the various levels of constraint in the domain. For more ill-structured domains in which such models are not available (e.g., chess), one can use a combination of empirical and logical methods to identify the relevant constraints. An example for each of these classes is presented below.

Example 1: A Thermal–Hydraulic Process

Physical systems are typically described by using differential and algebraic equations, not an abstraction hierarchy. However, the constraints embedded in physical models can be transformed in such a manner that they are more purpose oriented and, therefore, more useful to psychology. This section provides an example of this analytical method of constructing an abstraction hierarchy representation for a thermal–hydraulic process.

The problem domain to be analyzed, DURESS (DUal REservoir System Simulation) II, is illustrated in Figure A1. The system consists of two redundant feedwater streams that can be configured to supply water to either of the two reservoirs. Each reservoir has associated with it an externally determined demand for water that can change over time. The system purposes are twofold: to keep each of the reservoirs at a prescribed temperature (40 °C and 20 °C) and to maintain enough water in each reservoir to satisfy the current water output demand. To accomplish these goals, the participant has control over eight valves (VA, VA1, VA2, VO1, VB, VB1, VB2, and VO2), two pumps (PA and PB), and two heaters (HTR1 and HTR2). The input temperature (T0), output temperatures (T1 and T2), and the volumes for both reservoirs (V1 and V2) are also displayed in Figure A1.

As shown in Table A1, the system can completely be described by 34 process variables. However, these variables are not independent and should not take on arbitrary variables. When DURESS II is operating as intended, these variables are governed by a number of constraints that, in turn, affect the behavior of the system. For example, the temperature for Reservoir 1 (T1) should be at 40 °C (±2°C), if the system is to function as intended. This simple goal-relevant constraint can be defined as follows: 38 °C < T1 < 42 °C. Analogous constraints can be defined for the temperature goal on Reservoir 2 and the two demand goals as well. As a second example, when the system is operating normally, it also obeys the following constraints that are defined by mass and energy conservation laws, respectively:

$$\frac{d\mathrm{V1}(t)}{dt} = \frac{\mathrm{FA1}(t) + \mathrm{FB1}(t) - \mathrm{D1}(t)}{\rho}$$

$$\frac{d\mathrm{E1}(t)}{dt} = \mathrm{EI1}(t) - \mathrm{EO1}(t),$$

where ρ is density of water. Analogous constraints can be defined for Reservoir 2. Moreover, there are other constraints describing the behavior of system variables, such as flow rates, heat transfer rates, and component settings (see Bisantz & Vicente, 1994).

In summary, the goal-relevant constraints that govern DURESS II are given by analytical engineering models. The next task is to transform this knowledge into an abstraction hierarchy representation that organizes these constraints into a form that is more psychologically useful.

A formal abstraction hierarchy representation for DURESS II implemented in LISP was developed by Bisantz and Vicente (1994). Figure A2 shows the multiple layers of goal-relevant constraint in the system corresponding to the various levels in the abstraction hierarchy (only one reservoir is represented). Five levels of description were adopted (Rasmussen, 1985). For reasons that are later explained, only four levels are illustrated in Figure A2.

The top level of *functional purpose* describes the constraints arising from system purposes. There are four goals in this system: keep the water at the set-point temperature for each reservoir (two goals) and keep enough water in each reservoir to keep up with the current demand flow rate (two goals).

The level of *abstract function* identifies the goal-relevant constraints arising from the laws of conservation of mass and energy for each reservoir subsystem. As shown in Figure A2, each subsystem has one mass and energy inventory (the reservoirs), one mass input (the incom-

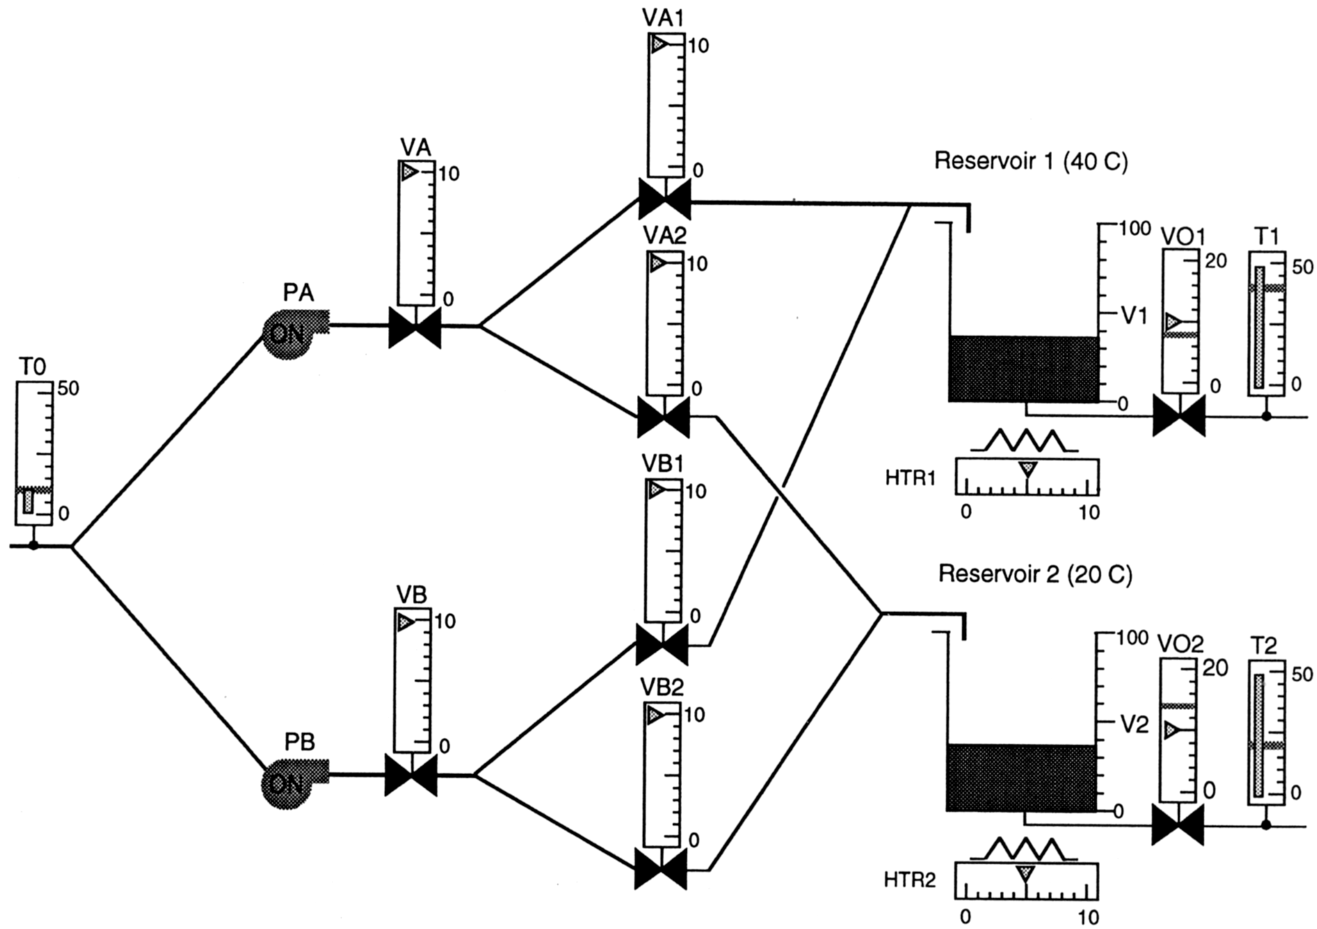

Figure A1. Mimic diagram representation showing the physical structure of the DURESS II system. VA, VA1, VA2, VO1, VO2, VB, VB1, and VB2 = eight valves. TO = input temperature; T1 and T2 = output temperatures; PA and PB = two pumps; HTR1 and HTR2 = two heaters. From "The Ecology of Human–Machine Systems. II: Mediating 'Direct Perception' in Complex Work Domains," by K. J. Vicente and J. Rasmussen, 1990, *Ecological Psychology, 2,* p. 233. Copyright 1990 by Lawrence Erlbaum Associates, Inc. Adapted with permission.

ing water), two energy inputs (the incoming water and the heater), and one mass and energy output (the demand). Topological links at this level, shown in Figure A2, indicate the flows of mass and energy through the subsystems.

At the level of *generalized function,* the system constraints pertinent to flows and storage of heat and water are represented. The rate of flow of water and heat transfer from the input stream, rate of heat transfer from the heating system, storage of heat and storage of water in the reservoirs, and rate of removal of heat and water due to demand are described at this level. The topological links, shown in Figure A2, indicate the flows of water and heat through the components.

At the level of *physical function,* the constraints pertinent to system components are described. Because of space restrictions, the components in each feedwater stream (one pump and three valves) have been grouped and labeled as a whole in Figure A2. At this level, the settings of valves, pumps, and heaters are described, along with the volume and temperature in the reservoir. Topological links, shown in Figure A2, indicate physical connections between components.

At the level of *physical form,* the appearance and location of the components are described. This level is not represented in Figure A2 because the units of analysis are essentially the same as those of the previous level. The main difference is that this level would describe the location and appearance of each component (e.g., via a video image) rather than the physical state of those components.

It is important to point out that, with the exception of the bottom level of physical form, there are quantitative equations and variables behind the structural relationships identified by the abstraction hierarchy in Figure A2 (several examples were given earlier).

Greater insight into the psychological value of the abstraction hierarchy representation can be obtained by examining Figure A3, which shows the means–ends relations that connect the various levels of abstraction. These relations reveal some important system properties to which an expert can become attuned. First, the mappings between functional purpose and abstract function indicate that the temperature goal is connected to both the mass and energy balances. The reason for this is that temperature is defined as energy per unit mass. In contrast, the demand goal is only connected to the mass balance, which means that changing the mass flow to affect demand may have the unintended side effect of changing temperature, whereas the reverse need not be true. A second property of the system revealed by the means–ends links shown in Figure A3 is the many-to-many mapping between abstract function and generalized function. There are several instances wherein a single

Table A1
Complete List of Process Variables in DURESS II

Variable	Description
Temperature	
T1	Temperature of Reservoir 1
T2	Temperature of Reservoir 2
Mass	
D1	Demand (output) flow rate for Reservoir 1
D2	Demand (output) flow rate for Reservoir 2
MI1	Mass input flow rate for Reservoir 1
MI2	Mass input flow rate for Reservoir 2
V1	Volume of Reservoir 1
V2	Volume of Reservoir 2
Energy	
E1	Total energy stored in Reservoir 1
E2	Total energy stored in Reservoir 2
EI1	Energy input flow rate for Reservoir 1
EI2	Energy input flow rate for Reservoir 2
EO1	Energy output flow rate for Reservoir 1
EO2	Energy output flow rate for Reservoir 2
Heat transfer	
FH1	Flow from Heater HTR1
FH2	Flow from Heater HTR2
Flow rates	
FA1	Flow rate from Valve VA1
FB1	Flow rate from Vavle VB1
FA2	Flow rate from Valve VA2
FB2	Flow rate from Valve VB2
FPA	Flow rate from Pump PA
FPB	Flow rate from Pump PB
FVA	Flow rate from Valve VA
FVB	Flow rate from Valve VB
Heaters	
HTR1	Setting for heater of Reservoir 1
HTR2	Setting for heater of Reservoir 2
Pumps	
PA	Setting of pump in Flow rate A
PB	Setting of pump in Flow rate B
Valves	
VA	Setting of initial valve in Flow rate A
VB	Setting of initial valve in Flow rate B
VA1	Setting of Valve 1 in Flow rate A
VB1	Setting of Valve 1 in Flow rate B
VA2	Setting of Valve 2 in Flow rate A
VB2	Setting of Valve 2 in Flow rate B
VO1	Setting output valve in Reservoir 1
VO2	Setting of output Valve 2 in Reservoir 2

Note. From "Memory Recall in a Process Control System: A Measure of Expertise and Display Effectiveness," by Kim J. Vicente, 1992, *Memory & Cognition, 20,* p. 362. Copyright 1992 by the Psychonomic Society. Adapted with permission.

subsystem is a means for controlling both mass and energy. For example, each feedwater stream serves both as a mass input and an energy input. Therefore, there is a structural coupling between the mass and energy configurations. A third critical property of the system is the many-to-one mapping between physical function and generalized function. This property results from the fact that either feedwater stream can be used to supply either reservoir. As a result, changing one of the components in one feedwater stream to control the flow to one reservoir could actually unintentionally change the flow to the other reservoir.

Each of the examples just described identifies a constraint that experts can exploit while reconstructing the state of the system in a memory recall task (or while problem solving). For example, if one remembers that the valves in one of the feedwater subsystems are in a certain configuration, then this fact can be used to reconstruct the water flow into the reservoir, as shown by the means–ends links between physical function and generalized function in Figure A3 (see Bisantz & Vicente, 1994, for the equations describing this constraint).

Example 2: Chess

In chess, instead of relying on analytical models, one must use empirical and logical methods to conduct an abstraction hierarchy analysis to identify goal-relevant domain constraints. A comprehensive analysis of this type cannot be provided here because little work has been done on representing chess positions (Gobet, 1994), so the knowledge required to conduct a comprehensive analysis is not available. Fortunately, research conducted by de Groot (1966) provides enough insight to suggest how such an analysis could be conducted.

A sketch of an abstraction hierarchy for chess, suggesting plausible layers of goal-relevant constraint, is presented in Table A2. The labels are similar to those for baseball (Figure 1), but the detailed contents of the levels differ. At the lowest level, there are the number and type of pieces and the physical properties of the board. Given the size of the board and the number of pieces available, there are certain physical configurations that are simply not possible (e.g., having more than 16 pieces of the same color). Thus, the board level puts a minimum amount of constraint on what constitutes a legal chess position. At the second level, paths, the constraints arising from the legal paths of movement

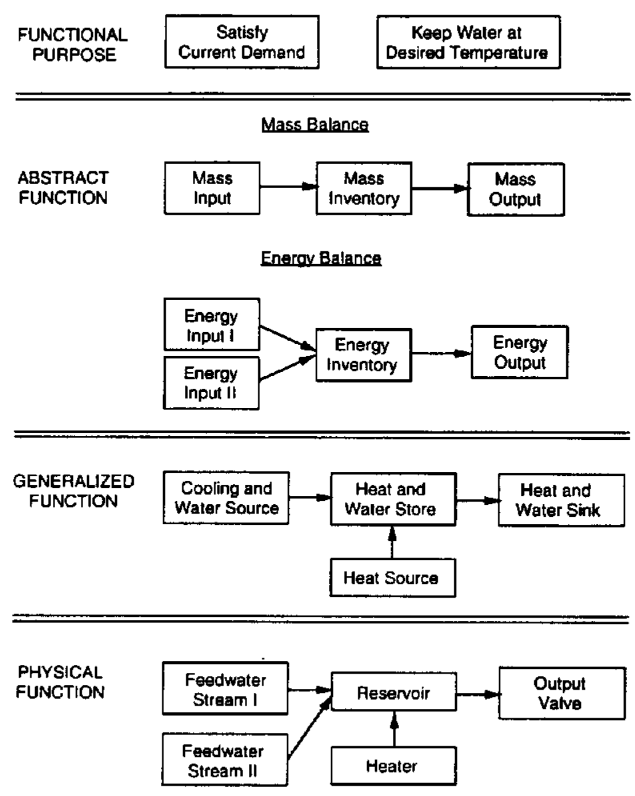

Figure A2. Abstraction hierarchy representation of DURESS II showing the various layers of goal-relevant constraint in the system. Note that the bottom level of constraint, physical form, is not depicted (see text). From "The Ecology of Human–Machine Systems. II: Mediating 'Direct Perception' in Complex Work Domains," by K. J. Vicente and J. Rasmussen, 1990, *Ecological Psychology, 2,* p. 234. Copyright 1990 by Lawrence Erlbaum Associates, Inc. Adapted with permission.

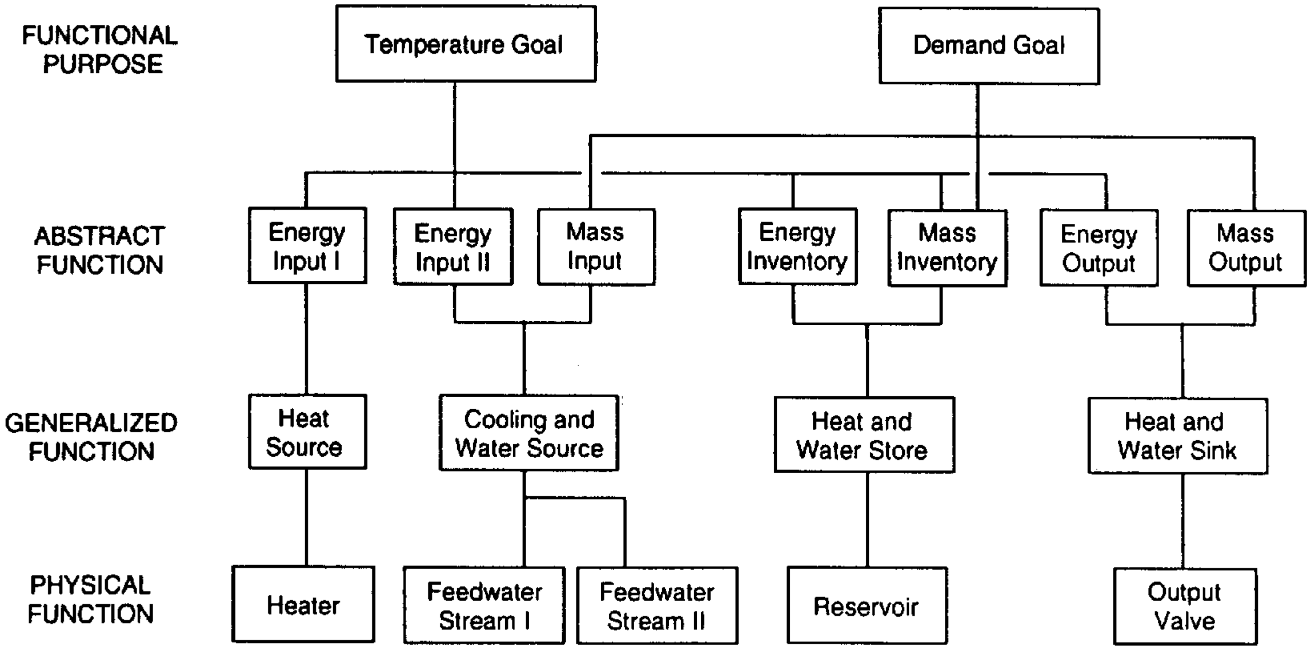

Figure A3. Abstraction hierarchy representation of DURESS II showing the means–ends relations between the various levels. Again, the bottom level of constraint, physical form, is not depicted (see text). From "The Ecology of Human–Machine Systems. II: Mediating 'Direct Perception' in Complex Work Domains," by K. J. Vicente and J. Rasmussen, 1990, *Ecological Psychology, 2,* p. 236. Copyright 1990 by Lawrence Erlbaum Associates, Inc. Adapted with permission.

that are available are represented. Because of the rules of the game, there are certain movement paths that are not allowed (e.g., moving pawns into the first row). At the next level, tactics, another layer of structure is added. The tactics level is constrained by the meaningful and effective ways in which moves can implement goal-relevant strategies. Note that this is an added layer of constraint that exists over and above the constraints specified at the lower two levels. The next level of constraint is that of strategies. Master chess players do not constrain their moves by the rules of the game or tactics alone. Rather, their actions are also highly constrained by the higher order strategic plans that they have adopted to achieve a win (cf. Holding, 1985). This layer of constraint is represented at the strategies level. Finally, at the top level of purpose, the player's goals will in turn constrain the meaningful strategies that can be selected. For instance, certain strategies are much more likely to lead to victory than others. Furthermore, the choice to play a certain style of game (e.g., risky or conservative) will also constrain the meaningful strategies that are available.

Table A2
Abstraction Hierarchy Representation for the Domain of Chess

Level	Description
1. Purpose	Player's goals (e.g., to win, to take risks, and to play an elegant game).
2. Strategies	Higher level lines of development that can achieve purpose.
3. Tactics	Lower level combinations that can implement a strategy.
4. Paths	Available paths of movement that can be selected to execute a tactic.
5. Board	Spatial location and appearance of pieces of board.

A combination of logical and empirical methods can be used to add substance to the skeleton structure in Table A2. Information theory serves as an effective tool because it allows one to identify quantitatively the degree of constraint in a stimulus. For instance, de Groot (1966) calculated that the maximum amount of information in a chess position with at least two Kings on the board is 143 bits. This calculation corresponds to the degree of constraint imposed by the board level of the abstraction hierarchy in Table A2. As the abstraction hierarchy analysis makes clear, however, there is much more structure in meaningful chess positions. De Groot's analysis confirms this. For example, one of the rules of chess is that pawns cannot be moved onto the first row. This constraint, represented in the paths level in Table A2, leads to a reduction of uncertainty of 5.6 bits (de Groot, 1966). This example illustrates how a logical analysis using information theory can be used to identify the goal-relevant constraints in an abstraction hierarchy for a domain in which analytical models are not available.

De Groot (1966) also provides several examples of how constraints can be identified in an empirical manner. To take but one example, he conducted an information theoretic analysis on a number of chess positions sampled from master-level games. This analysis suggests that there are approximately 60 bits of information in such positions, a very large reduction in uncertainty compared with the full 143 bits that are physically possible. A great deal of this constraint results from the strategic level in Table A2. As de Groot (1966) observed, chess is a highly stereotyped game, at least at the level of master play.

What Table A2 does not show is the set of means–ends links between the different levels in the abstraction hierarchy (cf. Figures 1 and A3). Each level of the abstraction hierarchy is nested within the context defined by the level above, and it is this nesting that provides the abstraction hierarchy with its psychological value. For example, identifying a player's purpose (e.g., to play a conservative game) will constrain the set of strategies that can be used as means to achieve this end. Some strategies simply need not be considered because they do not fulfill the desired end. These strategies would not be linked to this particular purpose. In contrast, conservative chess strategies would all have a link to the aforementioned goal. Similarly, the choice of a particular strategy adds even more constraint in that only a finite set of tactics can serve as means to achieve the end of implementing the desired strategy. These tactics would therefore be connected across levels of abstraction by a means–end relation to the strategies that they can effectively fulfill. Similar logic applies between any two levels of abstraction. The links between levels define what is goal relevant for that domain, thereby allowing people who are aware of these relations to reduce the number of meaningful alternatives that need to be considered in reconstructive recall.

The behavior of a novice, who is not familiar with these different levels of abstraction, would be very different. All he or she could focus on is the easily perceived, bottom level of abstraction (i.e., the spatial location of each piece on the board). Reconstructive recall based on this level is extremely difficult because, in the absence of higher order constraints, there seem to be an incredibly large number of possible positions. Without taking into account higher order functional constraints, it is psychologically infeasible to determine effectively which of those positions corresponds to the one that is supposed to be recalled.

Received March 30, 1995
Revision received December 23, 1996
Accepted February 13, 1997 ■

Sternberg Appointed Editor of *Contemporary Psychology (APA Review of Books)*, 1999–2004

The Publications and Communications (P&C) Board of the American Psychological Association announces the appointment of Robert J. Sternberg, Yale University, as editor of *Contemporary Psychology (APA Review of Books)* for a 6-year term beginning in 1999.

Contemporary Psychology has been in existence for 42 years and, for most of the time, has been operating under the same coverage model. The model is a good one, as the current issues edited by John H. Harvey reflect, and the journal has long met the needs of individuals and libraries. The pace of change has increased during the past few years, however, and the P&C Board recently decided that it was time for a new model, one that would reflect the 21st century reader's needs for information about books.

Sternberg, at the request of the P&C Board, will be embarking on a program to make the journal even more timely and interesting during his editor-elect year in 1998. Some of the changes envisioned include fewer but longer and more thoughtful reviews of books, reviews only of "new" books (with a few noteworthy exceptions), comparative textbook reviews at strategic times of the year, and changes in publication frequency and pricing. Sternberg welcomes suggestions for improving the journal and serving reader needs.

E-mail: robert.sternberg@yale.edu

Please note that all reviews are written by invitation. Publishers should note that books should not be sent to Sternberg. *Publishers should continue to send two copies of books to be considered for review plus any notices of publication to*

PsycINFO Services Department, APA
Attn: *Contemporary Psychology* Processing
750 First Street, NE
Washington, DC 20002-4242

As the editorial term of John H. Harvey comes to a close, the P&C Board wishes to express its appreciation for his hard work and dedication as well as that of his staff at the University of Iowa.