
Knowledge and Expertise

Stephan Lewandowsky and Daniel Little

University of Western Australia, Australia

and

Michael L. Kalish

University of Louisiana, USA

KNOWLEDGE IN APPLIED SETTINGS

What is knowledge? In fact, it may be simpler to ask, what is *not* knowledge? There is perhaps no single, more all-encompassing concept in cognitive psychology than knowledge. Knowledge contributes to simple perceptual tasks such as object recognition, when people identify an ambiguous stimulus on the basis of prior knowledge (e.g., Bar & Ullman 1996). Knowledge contributes to memory performance in myriad ways, for example when people reconstruct events according to a schema or script (e.g., Roediger *et al.* 2001; Tuckey & Brewer 2003). Finally, knowledge is the fundamental ingredient of human cognition at its best, namely expert performance. Accordingly, the literature on knowledge is vast, and its sheer size prevents a summary by a few simple assertions. To keep our task manageable we have therefore imposed some strong constraints on this chapter.

Charness and Schultetus (1999) defined knowledge as “acquired information that can be activated in a timely fashion in order to generate an appropriate response” (p. 61). We accept this as our working definition, but restrict consideration to manifestations of knowledge that have been variously called declarative or explicit (Reber & Squire 1994; Shanks & Johnstone 1998). These forms of knowledge are characterized by being accessible to awareness and verbal report, for example in response to a query such as “What is the capital of France?” We do not give much consideration to issues of training and knowledge acquisition, which are the domain of Chapter 21. Finally, we use the applied focus of this handbook to guide which topics to foreground and which to downplay. Accordingly, we omit discussion of computational models of knowledge and its acquisition and transfer (e.g., ACT; Anderson 1990) as extensive treatments of those models can be found elsewhere (e.g., Singley & Anderson 1989). Instead, we foreground research on expertise and expert performance; we focus on the fractionation and encapsulation of knowledge; and we examine the success or failure of the transfer of that knowledge.

We proceed as follows. In the first major section we discuss the nature of expert behavior. In particular, we suggest that expertise is the result of specific learned adaptations to cognitive processing constraints. In consequence, expertise turns out to be very specific and “brittle”; that is, experts may encounter difficulties when tasks are altered or when transfer to new problems is expected. We conclude the section on expertise by examining three shortcomings and sources of error that experts frequently encounter. In the second major section, we consider more conventional, non-expert manifestations of knowledge. We begin by considering the widespread view that knowledge is integrated and coherent, exemplified by knowledge space theories as well as the mental model approach. We then consider the alternative position; namely, that knowledge is often fragmented or partitioned, and that multiple alternative pieces of knowledge are often held simultaneously. In the final major section, we consider the mechanisms underlying the transfer of learned knowledge to novel situations. We suggest that transfer succeeds only if people perceive the similarity between their existing knowledge and a novel problem, and we then review the factors that affect the perception of similarity. We conclude the section by examining additional factors that may lead to the failure of transfer. Throughout the chapter, we place particular emphasis on the problems and shortcomings associated with those processes, because they are of major relevance to the practitioner.

EXPERTISE AND ITS LIMITATIONS

We begin by considering the performance of the most skilled of individuals – the experts. Analysis of expertise can illustrate the essential characteristics of human knowledge; indeed, some have gone as far as to argue that expertise is an indicator of consciousness (Rossano 2003). Our intention in this section is to provide a fairly atheoretical summary of the performance characteristics and shortcomings of experts. The subsequent sections provide a more theoretical discussion of the properties of knowledge in general, and in so doing also provide another, more theoretical perspective on expertise.

Although the many definitions of an expert include anecdotal descriptions such as “anyone who is holding a briefcase and is more than 50 miles from home” (Salthouse 1991, p. 286) or “someone who continually learns more and more about less and less” (Salthouse 1991, p. 286), there is common agreement that an expert is characterized by reproducible superior performance in a particular domain. Any coherent set of tasks and problems that is amenable to objective performance measurement (Ericsson 1996) can constitute a domain of expertise. Accordingly, researchers have examined domains as diverse as the linking of car crime series by expert investigators (Santtila *et al.* 2004), the ability to predict the spread of bush fires by expert firefighters (Lewandowsky & Kirsner 2000), medical diagnosis (e.g., Patel *et al.* 1996), seemingly mundane but highly sophisticated activities such as driving a car (see Chapter 15 for more details), or the performance of chess masters (e.g., Charness *et al.* 1996). In all cases, expert performance has been consistently and reliably found to be outstanding and superior to that of novices.¹

In chess, for example, expertise is associated with an extraordinary ability to remember the location of pieces on a board after a few seconds of viewing, and the ability to play several games at the same time (e.g., de Groot 1965). In medical diagnosis, experienced radiologists reliably outperform residents when inspecting x-rays (Norman *et al.* 1992). A contemporary mnemonist, Rajan, has memorized the digits of π to over 30,000 places

and can reproduce sequences of up to 75 digits with great ease (e.g., Thompson *et al.* 1993). Even relatively mundane tasks such as waiting tables (Ericsson & Polson 1988) and transcription typing (Salthouse 1991) can involve astonishing levels of knowledge and cognitive performance. (Chapter 3, this volume, provides more information about some of those feats.) Notwithstanding the generally high level of performance, expertise is characterized by several attributes which, in addition to supporting exceptional performance, engender intriguing limitations and create the potential for serious error.

Characteristics of Expertise

Circumventing Known Processing Limitations

Expert performance often seems to defy known human processing limitations. For example, it is known that people cannot tap a finger repetitively more than about six times a second, even if they do not have to respond to specific stimuli (Freund 1983). In conjunction with the known lower limit on response latency to successive stimuli (around 550 ms; Salthouse 1984), these constraints seem to dictate a maximum transcription typing speed of somewhere between 20 and 75 words per minute. Yet, expert typists can enter text at a rate exceeding 75 words per minute. Salthouse (1984) showed that typists achieve this high level of performance by developing specific strategies to circumvent these processing constraints. Specifically, the maximum typing speed a typist can achieve is correlated with the number of characters that must be simultaneously visible for the typist to maintain their maximum speed. This correlation indicates that growing expertise is associated with increased parallelism of processing, such as that used to pre-plan keystrokes involving opposite hands. One index of this planning is the strong negative correlation between expertise and the delay between keystrokes involving alternate hands, as when “w” is followed by “o.” That is, coordination between the two hands increases with the expertise of a typist. The further fact that the correlation between expertise and inter-key intervals is substantially smaller for repetitions of the same letter – which necessarily involves repeated tapping of the same finger – indicates that expertise often involves the acquisition of skills to circumvent “hard” constraints, rather than a relaxation of those biological or cognitive constraints.

Similarly, outstanding memorial abilities appear to be based on acquired strategies and techniques. To illustrate, consider individuals who gradually raised their digit span by deliberate acquisition of mnemonic techniques. In some particularly dramatic instances, a person’s span increased from the standard 7 ± 2 to 80 or even higher (e.g., Ericsson *et al.* 1980; Staszewski 1993). These remarkable abilities relied on the acquisition of increasingly larger, richly integrated hierarchical retrieval structures (e.g., Staszewski 1993), an observation supported by computer simulation (Richman *et al.* 1995). Ericsson *et al.* (2004) recently confirmed that a similar account can capture the immediate memory abilities of the mnemonist Rajan mentioned earlier, notwithstanding earlier opinions to the contrary (Thompson *et al.* 1993).

The view that expertise represents a learned adaptation to task constraints – as opposed to being the result of innate “talent” – has found a theoretical focus in the work by Anders Ericsson and colleagues (e.g., Ericsson 2003, 2005). The principal tenet of Ericsson’s view is that expertise arises not from mere prolonged exposure to a task, but from extensive

Table 4.1 Commonalities between experts in different domains identified by Holyoak (1991)

1	Experts perform complex tasks in their domains much more accurately than do novices
2	Experts solve problems in their domains with greater ease than do novices
3	Expertise develops from knowledge initially acquired by weak methods
4	Expertise is based on the automatic evocations of actions by conditions
5	Experts have superior memory for information related to their domains
6	Experts are better at perceiving patterns among task-related cues
7	Expert problem-solvers search forward from given information rather than backward from goals
8	One's degree of expertise increases steadily with practice
9	Learning requires specific goals and clear feedback
10	Expertise is highly domain-specific
11	Teaching expert rules results in expertise
12	Performances of experts can be predicted accurately from knowledge of the rules they claim to use

“deliberate practice.” Deliberate practice differs from mere exposure and repetition in several important ways. First, deliberate practice involves a well-defined specific task that the learner seeks to master. Second, task performance is followed by immediate feedback. Third, there is opportunity for repetition. Ericsson *et al.* (1993) provided a very extensive characterization of deliberate practice, and Ericsson (2005) surveys specific instances in which the role of deliberate practice has been established in a variety of expert domains.

There is now considerable consensus that, irrespective of the domain of expertise, ten years of deliberate practice are required to attain outstanding levels of performance (e.g., Ericsson 1996). Moreover, experts exhibit some notable commonalities across domains. Table 4.1 lists some of the commonalities that were identified by Holyoak (1991). The bold-faced entries correspond to issues that we take up in this chapter because we consider them to be particularly critical; the reader is referred to Holyoak (1991) for a discussion of the remainder.

The fact that expertise is the result of very specific adaptations to task demands and processing constraints entails two related consequences: First, experience is typically very specific and limited to the trained domain. Second, expertise is often quite brittle, and seemingly small deviations from a routine task can be associated with surprisingly large performance decrements.

Specificity of Expertise

It should come as no surprise that expert archaeologists are not necessarily also outstanding oceanographers, and that expert psychologists are unlikely also to be world-class ornithologists. However, the extent of that specificity may exceed the expectations and intuitions of most practitioners. For example, individuals who acquire a phenomenally large digit span after extended training (e.g., Ericsson *et al.* 1980), somewhat soberingly

retain the standard limited capacity for other information – approximately seven symbols (e.g., Chase & Ericsson 1981). That is, the same person may struggle to recall “C F G K L P Z” in the correct order while being able to reproduce the sequence “2 9 0 3 4 1 8 9 2 3 0 5 7 1 4 5 2 2 8 1 0” (or an even longer series of digits) flawlessly. Similarly, expert pianists’ acquired ability to tap fingers particularly rapidly (Ericsson *et al.* 1993) does not generalize to an ability to tap feet at a particularly rapid rate (Keele & Ivry 1987). Perhaps the most astounding demonstration of specificity is the finding that one year after learning to read text in an unfamiliar transformation (e.g., letters flipped upside down and mirror reversed), people can re-read pages from a year ago reliably more quickly than new text that is presented in the same transformed script (Kolers 1976).

Brittleness of Expertise

A corollary of the specificity of expertise is its “brittleness”; that is, the deterioration in performance that is observed when a domain-relevant task is altered slightly and thus becomes atypical. A classic example involves memory for chess configurations. Chase and Simon (1973) found that an expert chess player could recall the identity and location of pieces on a chessboard after fairly brief (5 seconds) exposure with remarkable accuracy. However, this striking ability was limited to plausible configurations that might arise during an actual game. When pieces were randomly arranged, and hence no longer formed a meaningful pattern, the performance of the expert deteriorated dramatically. The deterioration of expert memory when domain-relevant stimuli are rendered meaningless by randomization or some other disruption is a fundamental attribute of expertise that has been observed in many domains. A review by Ericsson and Lehmann (1996) cites areas as diverse as the games of bridge, GO, Othello, snooker, basketball, field hockey, volleyball, and football, and professional disciplines such as medicine, computer programming, and dance.

Another intriguing aspect of these results, in particular those involving chess, arises from detailed comparisons between experts and novices. For meaningful game positions, the reproduction skills of chess masters are indubitably far superior to those of novices. For random positions, it used to be a matter of consensus that the expert advantage was completely eliminated. The belief that experts and novices did not differ in their memorial abilities for random board configurations was sufficiently entrenched to be echoed in recent textbooks (e.g., Medin *et al.* 2001). However, when the evidence from numerous studies is considered jointly in a meta-analysis, increasing expertise is found to be associated with a small but clear memory advantage even for random board positions (Gobet & Simon 1996). This small advantage is most likely due to the experts’ ability to discover even small regularities in otherwise random positions by matching board positions against a repertoire of an estimated 50,000 or so chess patterns stored in long-term memory (Simon & Gilmartin 1973).

Accordingly, when the degree of randomness (defined by the extent to which basic game constraints are violated) is manipulated, players with greater expertise are better able to exploit any remaining regularities than players with lesser expertise (Gobet & Waters 2003). The specificity of expertise thus extends to highly subtle regularities indeed.

Expert Transfer

The characteristics of expertise reviewed in the foregoing should readily generate expectations about how expertise transfers from one task to another. It would seem safe to assume a fair degree of within-domain transfer, albeit perhaps bounded by the observed brittleness of expertise, combined with the likely absence of transfer to tasks outside the expert's domain.

1 Indeed, there is considerable support for within-domain transfer. For example, Novick and colleagues (Novick & Holyoak 1991) showed that mathematical expertise predicts the degree to which solution strategies are transferred from one algebra word problem to another that appears different at the surface but shares the same deep structure. In fact, transfer is observed even when the two problems are presented under two separate experimental cover stories. In one study, the amount of transfer among experts was found to be up to nine times greater than among novices (Novick 1988, Experiment 1).

Similarly, in the domain of accounting, Marchant *et al.* (1991) showed that experts in general exhibit significantly more transfer than novices between problems involving the application of taxation laws. An accompanying finding was that when the problems were "anomalous," that is, constituted exceptions to a general taxation principle, the experts' subsequent transfer was often reduced to the level shown by novices. Marchant *et al.* argued that processing of the first exceptional case "increased the salience of a highly proceduralized strategy that overrides transfer from the analogy in the more experienced group" (p. 283). Thus, while expertise generally facilitates within-domain transfer, it may not do so in cases involving exceptional problems, because experts cannot help but activate their general knowledge even when exceptions to that knowledge must be processed. In consequence, the strongly activated general knowledge may prevent the renewed recognition of an exception to the general knowledge. We revisit the theme of the inevitable activation of expert knowledge below.

Turning to the issue of transfer outside a problem domain, it is unsurprising that such transfer is often absent. What is perhaps more surprising is how little deviation from a routine task it takes in order to eliminate transfer. Sims and Mayer (2002) examined the spatial skills of expert "Tetris" players. "Tetris" is a computer game that requires the player mentally to rotate shapes presented on the screen in a limited amount of time. People who were experienced "Tetris" players (either pre-experimental experts or trained in the experiment) did not differ from novices on a whole battery of spatial tests, with the highly selective exception of mental rotation tests involving shapes used in "Tetris" or very similar ones. That is, even though "Tetris" relies almost entirely on mental rotation skills, and even though people improved those skills during training, this improvement was narrowly limited to a certain type of stimuli and did not transfer to other shapes.

The characteristics of expertise just reviewed can engender specific performance errors and shortcomings that are worthy of the practitioner's attention. We next review those errors and shortcomings before examining the knowledge structures that underlie skilled performance in general and expertise in particular.

Expert Errors and Shortcomings

There is growing recognition that the analysis of performance errors and limitations contributes in fundamental ways to our understanding of the nature of expert knowledge (e.g.,

Table 4.2 Expert Shortcomings (items 1–8 were identified by Holyoak 1991)

1	Inflexibility
2	Expediency
3	Mediocrity
4	Inefficiency
5	Poorer memory for cases outside domain
6	Poorer perception of patterns unrelated to expert performance
7	Asymptotic performance
8	Domain specificity
9	Subjectivity
10	Lack of knowledge integration
12	Knowledge Inaccessibility

Johnson *et al.* 1992). Holyoak (1991) provided a list of expert limitations that are reproduced in Table 4.2, together with others identified by ourselves. Three of those limitations and shortcomings – inflexibility, expediency, and mediocrity – are particularly relevant here.

Inflexibility

Inflexibility is revealed when experts are confronted with novel task demands that are inconsistent with their existing knowledge base. In those situations, the need for adaptation may prove to be more challenging to experts than to novices (Frensch & Sternberg 1989; Sternberg & Frensch 1992). Using the game of bridge as their domain of expertise, Sternberg and Frensch (1992) examined the effects of various arbitrary rule changes on the performance of expert and novice bridge players. In general, experts were found to suffer more than novices from any rule change, although the extent of their impairment differed with the type of change. When the rule change involved surface modifications, such as introducing new nonsense names for suits and honor cards, experts suffered less of a performance decrement than when the deep structure of the game was changed, for example, by altering the rule determining the opening of each play. The fact that expert disruption was maximal after a change to the deep structure suggests that experts, unlike novices, routinely processed the task at that deep level; a finding that is consonant with much prior research (e.g., Chi *et al.* 1981; Dunbar 1995). Highly skilled performance may thus entail the general cost of reduced flexibility in the face of novel task demands.

In a related vein, Wiley (1998) showed that experts cannot suppress the retrieval of domain-relevant knowledge, even when participants are warned that their knowledge may be inappropriate or misleading in the current task setting. Wiley used a remote association task, in which people have to generate a word that can form a familiar phrase with each one of three presented items. For example, given the stimuli *plate*, *broken*, and *rest*, the word *home* can be used to form the meaningful phrases *home plate*, *broken home*, and *rest home*. Readers with expertise in baseball may have found this example particularly easy because the target phrase *home plate* represents a crucial concept in baseball. But what if the stimuli had instead been *plate*, *broken*, and *shot*? The intended word here is *glass*, although the first two words are compatible with the baseball-consistent completion

home. Wiley found that baseball experts, unlike novices, had great difficulty with items that implied – but did not permit – a domain-consistent completion.

The experts' difficulty persisted even when they were warned beforehand that their domain knowledge would be misleading, suggesting that activation of expert knowledge is automatic and cannot be suppressed.

Expediency

Expediency, by contrast, concerns the acquisition phase of expertise, and refers to the fact that experts emphasize efficiency when acquiring a skill. They may, for example, trade knowledge for extended search where many cues could be considered (Johnson 1988; Charness 1991). Thus, accumulation of a large knowledge-base allows experts to select the key features of the problem, thereby reducing the number of variables chosen for consideration. An illustrative case of expert expediency was reported by Lewandowsky and Kirsner (2000), who asked experienced wildfire commanders to predict the spread of simulated wildfires. The spread of wildfires is primarily determined by two physical variables: fires tend to spread with the wind and uphill. It follows that with light downhill winds, the outcome depends on the relative strengths of the competing predictors. If the wind is sufficiently strong, the fire spreads downhill with the wind, whereas if the wind is too light, the fire spreads uphill against the wind. Lewandowsky and Kirsner found that (with an intriguing exception that we discuss in a later section) experts completely ignored the slope and based their predictions entirely on the wind. While this gave rise to correct predictions in most circumstances, any fire in which light winds were overridden by a strong slope was systematically mis-predicted.

Mediocrity

Finally, imperfect expert performance has also been associated with situations in which probabilistic cues must be used to predict uncertain outcomes. For example, when predicting the likely success of applicants to medical school from their prior record (e.g., grades, letters of recommendation), expert accuracy is often inferior to that achieved by simple linear regression models, and only slightly superior to that of novices (Johnson 1988; Camerer & Johnson 1991). Most reports of this expert “mediocrity” have relied on domains in which there are no unequivocally correct rules but only sets of more or less accurate heuristics (Johnson 1988), which human experts have difficulty applying and combining in the correct statistical manner. In consequence, performance in those domains can be optimized by forming weighted linear combinations of probabilistic cues, a process embodied in linear regression but apparently difficult to achieve by humans (Camerer & Johnson 1991). To circumvent those difficulties, human experts use a variety of alternative combinatorial strategies. One of them, known as configural reasoning, consists of considering predictor variables in a categorical manner rather than by weighted addition. For example, a configural rule in medical diagnosis might be: “If the patient experiences headaches that have a gradual onset, with no periods of remission, and has high levels of spinal fluid protein, then diagnose a brain tumor” (Schwartz & Griffin 1986, p. 94). Configural reasoning is often observed in experts, but unlike weighted linear regression its

all-or-none character renders it vulnerable to small variability in measurements (Camerer & Johnson 1991).

Adaptive Expertise

Thus far, we have limited our discussion to what some have described as “routine expertise,” in contrast to what is termed “adaptive expertise” (e.g., Gott *et al.* 1993; Kimball & Holyoak 2000). Adaptive expertise has been defined as “an advanced level of problem-solving performance . . . characterized by principled representations of knowledge . . . as opposed to representations dominated by surface features” (Gott *et al.* 1993, p. 259).

Although routine and adaptive expertise are often seen as two contrasting concepts (e.g., Kimball & Holyoak 2000), we are reluctant to accept this contrast for a variety of reasons. First, we are not aware of an independent criterion that identifies a particular expert, or a particular domain of expertise, as adaptive. Instead, expertise appears to be considered adaptive whenever it transfers well and it is considered routine whenever it does not. Second, empirical examinations of adaptive expertise converge on identification of the same, or similar, cognitive principles that are also involved in non-adaptive settings. For example, Barnett and Koslowski (2002) presented experienced restaurant managers and business consultants without any experience in the hospitality industry with problems relating to the management of hypothetical restaurants. Because the specific problems were novel to both groups of participants, Barnett and Koslowski considered them to represent “transfer” problems. Notwithstanding the lack of domain-specific expertise, the business consultants were found to outperform the restaurant managers, suggesting that the consultants were “adaptive” experts whereas the managers’ expertise was more “routine.” Further analysis identified the amount of prior consulting history (i.e., strategic business advisory experience) as the crucial variable underlying the performance difference. A crucial characteristic of business consulting, in turn, is the extreme breadth and variability of the problems that consultants tend to encounter. Barnett and Koslowski therefore conclude that “a possible explanation for the observed differences is . . . the wide variety of business problem-solving experience to which the consultants, but not the restaurant managers, have been exposed” (p. 260).

As we review below, variability among training instances is a known strong predictor of transfer in general. We therefore propose that adaptive expertise does not differ qualitatively from routine expertise, and that the observed differences in transfer ability are best explained within known principles of knowledge and expertise.

We now turn to an examination of those broader principles of knowledge in contexts other than expertise. This examination, in turn, provides another, more theoretical perspective on the nature of expertise.

STRUCTURE OF KNOWLEDGE

Overview

By discussing the structure of knowledge, we implicitly assume that knowledge can have a structure – that we can reasonably discuss constructs such as individual parcels of

knowledge (what might once have been called “ideas”), pathways connecting the parcels, domains that the parcels apply to, dominance relations between parcels, and so on. As befits this general review, we abstain from detailed analysis of the structure of any particular domain and instead concentrate on a central dichotomy: is knowledge generally integrated or is it generally fragmented?

The idea that knowledge is highly integrated is a central, if often tacit, tenet of views of expert knowledge (e.g., Bédard & Chi 1992; Ericsson & Lehmann 1996). Glaser (1996), for example, explicitly cited the centrality of the “acquisition of well-organized and integrated knowledge that provides a structure for representation that goes beyond surface features. For the development of expertise, knowledge must be acquired in such a way that it is highly connected and articulated, so that inference and reasoning are enabled as is access to procedural actions” (p. 306). Finally, it has been assumed that during training, experts notice inconsistencies in their current knowledge, “which in turn will serve as a stimulus for further analysis . . . until an acceptable reintegration . . . is attained” (Ericsson 1996, p. 38). To foreshadow our main point, in contrast to the prevailing views of expert knowledge, we will reach the conclusion that knowledge is often more fragmented and disconnected than experience would seem to indicate.

Discussions of knowledge integration are made difficult by the essentially intuitive nature of the term “integrated.” As diSessa and colleagues (2004) point out, it is possible to analyze the term (they use “coherent” for our “integrated”) as consisting of a position on dimensions of contextuality, relational structure, and specification. Integrated knowledge (such as that described by Gott *et al.* 1993) is largely decontextualized in that irrelevant surface features are not represented or considered; the color of a body in motion does not enter into the laws of thermodynamics. Integration also implies increased structure in the relations between elements, for example, in that some elements are logically deduced from others. Integration also entails the ability of an observer to specify knowledge in a more compact manner. It takes less verbiage to describe the ideal and derived gas laws than it does to explain a naïve theory of balloons, drinking straws, sweating lemonade glasses, and exploding soda cans.

The distinction between integration and fragmentation can also be expressed within a spatial metaphor; for example by noting the proximity of parcels to each other, or their size and relation to the size of the domain of interest. The spatial metaphor turns out to be at the heart of recent formal analyses of people’s knowledge.

Knowledge Spaces

The notion that knowledge may be organized spatially relates to the acknowledgment that similarity can be understood as a spatial relationship. In other words, the “space” in which one can represent knowledge elements is a space in which distances are defined by a metric of similarity. Shepard (1987) pointed out that the ability to generalize afforded by a spatial representation of similarity is a potential solution to one of the core problems in psychology; how new knowledge can be related to old. This relationship is critical in order for people to learn as rapidly as they do, given the type of data they are presented with. As Landauer and Dumais (1997) say, this leads to a potential solution of “Plato’s problem,” namely, induction.

Landauer and Dumais (1997) introduced a method known as latent semantic analysis, which can derive and represent the relationships between knowledge parcels as spatial

relations. Latent semantic analysis (LSA) assumes that words with similar meanings are likely to occur in the company of similar words; put another way, if words appear in similar contexts, they likely have similar meanings. LSA uses large corpora of text in which paragraph-length sections serve as context. Each word in the corpus is tallied for each context, and statistical dimension-reduction techniques are used to summarize the context–word relationships. Each word is then identified as a point in a space of lower dimensions than the number of contexts, but still of relatively high dimensionality (approximately 300 dimensions works well for many applications). The model assumes that all words can be represented in a single semantic space, and in this sense is clearly integrated rather than fragmented.

The hypothesis that a single semantic space is an effective way to represent general knowledge about word meanings is supported by Vigliocco *et al.*'s (2004) featural and unitary semantic space model (FUSS). FUSS differs from LSA in that its metric dimensions are interpretable; this is because the data for FUSS are not word–context relations in corpora but speakers' judgments of the semantic features of words. In a related vein, Griffiths and Steyvers (2002) used corpora to arrive at interpretable dimensions in their Topics model, although their statistical technique for dimension reduction sees similarity as a relation between probability distributions rather than spatial locations (a move similar to that taken by probabilistic latent semantic analysis, Hoffman 1999).

None of these models, in their current form, purports to be a fully adequate description of the sum of an individual's knowledge about the meanings of words in their language. However, the move to represent meanings in integrated formats, within a spatial metaphor of some form, is one that must be recognized.

Alternative Approaches

The theories of knowledge spaces just discussed stand in contrast to several alternative approaches that deserve mention. These alternative approaches have variously involved concepts such as feature lists (e.g., Barsalou & Hale 1993), marker networks (e.g., Berg 1992; Lange 1992; Delisle *et al.* 2003), schemata (e.g., Marshall 1995; Kintsch 2000), or mental models.

The concept of a mental model has a particularly long and venerable history in applied and industrial settings. Mental models are thought to consist of look-up tables, lists of formal declarative knowledge, or mental images that contain representations of objects, functions, and causal interconnections between objects necessary to perform a task (see, e.g., Moray 1999). Mental models offer a variety of functions such as providing rules for predicting future states from current information (Bellenkes *et al.* 1997). For example, expert pilots expend more gaze time on predictive instruments than novices, indicating that they utilize their mental model for predicting future states more than novices (Bellenkes *et al.* 1997).

Coexistence of Alternative Knowledge

We now examine evidence for the continued coexistence of alternative ways in which people use their knowledge at all levels of skill acquisition and expertise (e.g., Shrager &

Siegler 1998; Lovett & Schunn 1999). Reder and Ritter (1992) and Schunn *et al.* (1997) presented participants repeatedly with two-digit \times two-digit multiplication problems (e.g., 43×19). Before responding, participants had to rapidly indicate whether they could retrieve the correct answer from memory (which they then had to report immediately) or whether they would need to compute the answer (in which case extra time was allotted). Most relevant for present purposes is the finding that across repeated presentations of a given problem, people were found to switch strategies not just once but between two and three times, and switches were separated by up to 50 per cent of all learning trials (reported in Delaney *et al.* 1998). This suggests that both forms of knowledge – retrieval and computation – continued to coexist throughout the experiment.

Prolonged coexistence of alternative knowledge has also been observed at a much larger time-scale, namely across grades in primary school (e.g., Siegler 1987; Shrager & Siegler 1998). This research showed that children approach single-digit mental arithmetic with immense cognitive variability, and that some techniques – such as counting fingers vs. retrieving the answer from memory – may coexist for several years and compete for selection whenever a problem is presented. Correspondingly, even adult performance can be characterized by an interaction between memory retrieval and alternative strategies (Griffiths & Kalish 2002). Griffiths and Kalish showed that many, but not all, systematic aspects of the pattern of errors observed in simple multiplication problems (Campbell 1994) could be explained by the similarity (and thus confusability) of the problems as predicted by a retrieval-based response strategy.

The persistence of competing knowledge structures is consonant with the suggestion that people ordinarily maintain multiple parcels of knowledge that could apply to any given situation (diSessa 1988). This suggestion has two non-trivial implications. First, it presupposes that there is a selection process that can choose among plausible alternative parcels. This selection process is presumably based on the structural and perceptual similarities of the current situation with those stored in memory (Gentner 1989). Second, if the ordinary state of knowledge includes multiple overlapping parcels, then at least some of those parcels might contain mutually inconsistent and contradictory information. This, indeed, appears to be the case.

Knowledge Partitioning

Consider, first, an instance of contradictory knowledge that, though consolidated by experience, is at the lower end of the expertise spectrum. Tirosh and Tsamir (1996) reported inconsistencies in high school students' understanding of the concept of mathematical infinity. Depending on the surface structure of the problem presentation, the distribution of responses differed greatly: Whereas with one surface structure, 80 per cent of participants correctly identified two infinite sets as containing the same number of elements, the vast majority of the same respondents (70 per cent) gave the opposing, inconsistent answer with the other surface structure. In a related study, also involving mathematical knowledge, Even (2003) showed that few prospective secondary mathematics teachers spontaneously linked an expression to its isomorphic graphical representation, even though this linkage would have facilitated solution of the problem.

The reverse was also true: people had difficulty deriving an expression from a graphical representation of the same function. Given that subjects were highly conversant with both

representations of all functions used in the study, ~~the~~ Even's findings point to heterogeneity even in consolidated knowledge.

Contradictory elements of knowledge have also been revealed in another naturalistic domain known as "street mathematics." This research focused on people who lack formal schooling but are able to solve mathematical problems in everyday contexts, for example street vendors, fishermen, construction foremen, and cooks in Brazil (e.g., Carraher *et al.* 1985; Nunes *et al.* 1993). Notwithstanding their minimal formal schooling, the participants were highly competent at solving mathematical problems associated with their domain of expertise.

Of greatest interest here is a context manipulation involving expert cooks (Schliemann & Carraher 1993). Participants were presented with identical proportionality problems either in a pricing context ("If 2 kg of rice cost 5 cruzeiros, how much do you have to pay for 3 kg?"), or in a recipe context ("To make a cake with two cups of flour you need five spoonfuls of water; how many spoonfuls do you need for three cups of flour?"). Importantly, both problem contexts were familiar to participants and relevant to their domain of expertise. Schliemann and Carraher reasoned that social convention dictated accuracy in the pricing context, whereas estimation might be acceptable for recipes. Those expectations were confirmed. In the pricing context, subjects used a variety of identifiable mathematical strategies in preference to estimation, with the result that accuracy was in excess of 90 per cent. In the recipe context, by contrast, accuracy was dramatically lower (20 per cent) and half of the responses given were based on estimation.

In the preceding cases, contradictory performance arose between variants of problems that differed not only according to the context in which they were presented (e.g., their cover story), but also their surface structure. An even purer instance of contradiction, involving reasoning about materially identical problems that differed only in an irrelevant context, was observed in the study by Lewandowsky and Kirsner (2000) mentioned earlier. Lewandowsky and Kirsner asked experienced wildfire commanders to predict the spread of simulated wildfires. The experts' predictions were found to depend on an additional variable, the physically irrelevant problem context. When a fire was presented as one that had to be brought under control, experts nearly always expected it to spread with the wind. When an identical fire was presented as a "back burn," experts predicted the reverse, namely that the fire would spread uphill and into the wind. Back burns are fires that are lit by firefighters in the path of an advancing to-be-controlled fire to starve it of fuel; back burns obey the same laws of physics as any other fire, in the same way that apples and oranges both obey the laws of gravity.

Before presenting an explanatory framework for these results, it is essential to differentiate them from conventional context effects, such as those reviewed earlier which underscored the specificity of expertise. Four attributes of the Lewandowsky and Kirsner study are relevant in this regard: (1) The nature of the problem and its surface structure arguably did not differ between contexts. That is, unlike the conventional context effects in expertise, the problem was no more typical of the domain in one context than the other. (2) By implication, unlike the related study by Schliemann and Carraher (1993), the change in context was a minimal alteration of a verbal label that accompanied presentation of a problem. (3) Both domain-relevant contexts were part of the training regime of the experts and both regularly occurred in the field. (4) The context shift resulted not only in a reduction of performance – as, for example, observed with chess masters' memory of random board configurations – but in a qualitative reversal of the response. That is, the

same problem yielded two mutually exclusive and contradictory predictions, each of which was consistent with application of a domain-relevant predictor variable.

These attributes are sufficiently unique to warrant the assertion that knowledge, even within a well-learned domain, may exhibit little homogeneity. Indeed, it appears that experts may sometimes, perhaps often, have knowledge that is overlapping and contradictory. As we observed earlier, overlapping knowledge parcels are indicative of fragmented knowledge structures, and fragmentation has been assumed to be the norm for naïve theories (diSessa *et al.* 2004). Lewandowsky and Kirsner (2000) suggested that the fragmentation observed in experts be considered an example of knowledge partitioning, caused by associative learning and thus a natural consequence of acquiring expertise.

Lewandowsky *et al.* (2002) proposed that associative learning produces knowledge partitioning in essentially the following way. Early in learning, when few cases are known, the learner acquires information rapidly about the few available cases. As learning continues, the most effective strategy to deal with new problems that are similar to the learned cases is to use the initially learned information. Thus, it is effective for learners to protect their old knowledge and apply it whenever it is applicable. People achieve this protection through rapid shifts in attention (Kalish *et al.* 2004). This process of learning new cases when necessary and deflecting change from old cases creates knowledge parcels that may contain contradictory information. So long as the stored cases do not overlap, this is not a problem for the learner, and so the associative theory predicts that partitioning is only sustainable when such conflict does not routinely occur. In the firefighting example, this may indeed have been the case as wildfires tended (in the experts' experience) to occur during high-wind periods and back burns tended to be encountered (or set) primarily when winds were light.

Our discussion of knowledge "parcels" is not to give the impression that knowledge representations are necessarily static. On the contrary, there is evidence that knowledge, specifically conceptual knowledge, is not static and may be created or altered "on the fly." For example, knowledge assembly theory (Hayes-Roth 1977) posits that repeated activation of the same components leads to the unitization of these components into a configuration which is then activated as a single, integrated entity. As another example, Barsalou's (1983) work with ad hoc categories has demonstrated that highlighting or making salient a particular goal can alter one's judgment about an item's category membership, its typicality, and the activation of other related items (Barsalou 1982, 1983, 1985). More recently, Barsalou (1999) has taken the notion of "on the fly" recruitment of knowledge even further, by suggesting that knowledge is fundamentally linked to physical experience. According to his perceptual symbols theory, knowledge of a concept is represented as a modality-specific "simulation" of one's physical experience with that concept. For example, the sweetness of a strawberry is represented by "simulating" (or imagining) its taste. Knowledge is thus fragmented according to modalities, of which Barsalou identifies six: vision, audition, taste, smell, touch, and action (Pecher *et al.* 2004). Which of these is activated for simulation depends on the context in which the concept is encountered (Pecher *et al.* 2004), thus further underscoring the dynamic – and fractionated – nature of knowledge representations within this framework.

Given the ready occurrence of partitioning and fragmentation, the apparent integration of knowledge in the expert may now appear all the more remarkable. However, close examination of the way experts apply their knowledge suggests that this appearance is at least partially an illusion. We have suggested that knowledge is frequently accessible only

from an associated context, or, equivalently, that knowledge is often represented at a grain size that is smaller than the domain the knowledge ought to (in a normative sense) cover. One measure of this grain size is the ease of transfer; problems within a knowledge parcel's domain should see transfer, ~~the~~ knowledge should be used with more difficulty on problems outside the parcel's boundaries. In the next section, we take up this ~~measure~~ and evaluate the integration of knowledge with respect to transfer.

TRANSFER OF KNOWLEDGE


The use of existing knowledge in new situations, known as transfer, is perhaps the most important test of one's current knowledge structures. Transfer necessarily involves linking or mapping from what is known to a new or novel situation (Holland *et al.* 1986). This mapping entails a tradeoff between expediency, which requires the rapid application of knowledge, and efficiency, which requires the selective application of only those cognitive resources that are necessary for the task (Besnard & Cacitti 2005). Inherent in this tradeoff is the potential for transfer to fail.

4

Failure of transfer can have drastic consequences in an applied setting. Besnard and Cacitti (2005) described an industrial accident at a French steel factory, where the installation and use of new machinery among several older machines resulted in the death of a factory worker. The worker was operating a thread-drawing machine, a device used to reduce the diameter of a metal thread by gradually increasing its tension. The output of this machine is wound tightly onto a drum and held in place by a set of pressing wheels controlled by the operator. On the new machine, the two key buttons controlling the opening and closing of the pressing wheels were swapped with respect to the older machines. The experienced operator mistakenly opened the pressing wheels on the new machine at a time when the metal thread was tightly wound, causing the thread to uncoil violently and resulting in the death of the worker. Prior experience with the old machines led the worker to transfer an existing skill to a situation that required similar skills, but applied in different manner, with deadly results.

This is not to conclude that successful transfer is impossible or rare; we have already seen that experts are extremely adept at transferring their knowledge within their domain of expertise. There is also evidence for successful within-domain transfer (also called "near" transfer) among novices. For example, having learned a specific rule to categorize stimuli, people are able quickly to learn to categorize novel stimuli that share the same rule but are instantiated by different dimensions (Shepard *et al.* 1961).

5

However, as we show later, the extent of transfer between tasks often falls short of what intuition might lead one to expect. For example, in the domain of artificial grammar learning, transfer is much better if the surface structure of the training set remains the same during the test phase (Brooks & Vokey 1991; Gomez *et al.* 1994). Slight contextual changes (e.g., replacing colors with color names; Dienes and nann 1998) can reduce or eliminate transfer altogether.

We now examine the conditions that determine whether or not transfer is successful. We focus on cognitive factors and do not consider variables that are beyond the scope of this chapter, such as organizational factors (e.g., perceived support; Flint 2003), characteristics of the individual (e.g., IQ; Ceci & Ruiz 1993; Ceci *et al.* 1999); motivation

(Bereby-Meyer & Kaplan 2005), or social factors like mentoring or supervisor support (e.g., Cromwell & Kolb 2004).

Similarity and Transfer

For transfer to occur, people must necessarily *perceive* two tasks as being similar. The emphasis on perception is crucial, because transfer depends primarily on the perceiver's psychological processing rather than objective measurements of the tasks involved. We consider four factors that are known to affect the perception of similarity.

Perceived Similarity: Structure vs. Surface

Perhaps the most important differentiation between forms of similarity involves the distinction between "deep" structural similarity and "surface" similarity, which comprise two potentially independent means of describing the relations between two situations or objects. For example, two fables that involve completely different sets of characters (and hence share little surface similarity) may nonetheless make the same moral point (thus having identical deep structure). Conversely, two fairy tales may involve the same set of characters but provide completely different messages. The latter situation can be particularly harmful because when surface similarity lures people into attempting transfer between tasks that are structurally dissimilar, negative transfer may result (Hershey & Walsh 2000). For example, a novice attempting to understand the game of cricket might mistakenly apply his or her knowledge of American baseball because a bat and a ball are used in both games. This attempt at transfer is fatal because the deep structure of cricket – which is sufficiently grave and complex to be summarized not by mere rules but by "laws" – deviates considerably from the comparatively simple deep structure of baseball.

Conversely, a change in cover story or surface presentation can reduce transfer, notwithstanding deep structural identity between the two tasks. For example, in a now classic study, Gick and Holyoak (1980, 1983) taught participants to solve a problem involving the storming of a fortress surrounded by a minefield, in which the key to successful conquest was to send numerous platoons from all directions simultaneously which then converged onto the target. After learning this solution, people were unable to apply that knowledge to an isomorphic radiation convergence problem, in which removal of a tumor without damaging the surrounding tissue could be achieved only by applying weak intersecting radiation from all directions (Gick & Holyoak 1980, 1983). Hence, while transfer to similar problems is possible, a surprisingly small change in context or cover story can eliminate that transfer quite readily.

Similarity and Expertise

The distinction between surface and structural similarity is particularly relevant when comparing the transfer abilities of novices and experts. One of the primary differences between expert and novice problem-solving is that experts focus on the deep structure of

the task (e.g., Chi *et al.* 1981; Dunbar 1995). Accordingly, experts will attempt transfer if two tasks share structural similarities even if their surface similarity is low. As we have seen, mathematical expertise predicts the ease with which people transfer between superficially different word problems with the same structure (Novick 1988). Novices, by contrast, will attempt transfer only if there are salient surface similarities between the source and the target, despite the fact that the same solution process is needed (Cormier 1987). For example, in the earlier convergence problems, novices, after being trained in the radiation context, are more likely to attempt transfer to a problem that is similar at the surface, because it involves x-rays, rather than to a structurally similar problem that is less similar at the surface because it involves ultrasound (Holyoak & Koh 1987).

Conceptual vs. Structural Similarity

Dixon and colleagues (Dixon & Gabrys 1991; Dixon *et al.* 1997) differentiate between conceptual similarity, which is based on information about why a procedure works or how a device operates and allows the application of conceptually similar problem-solving steps, and structural similarity, which is similarity based on the steps that must be performed to solve a problem and allows for the application of identical procedural steps. In their studies, participants were initially trained to operate a complex device through a series of sub-goals comprised of a number of different steps (e.g., for an airplane device, the sub-goal “Engine Start Up” might consist of the steps “engine 1,” “engine 2,” followed by “ignition”). Consequently, conceptual similarity (which, in this example, is isomorphic to superficial similarity) could be manipulated independently of structural similarity by changing the order of the sub-goals but maintaining the same order of steps within each sub-goal.

Following initial training with one device, transfer to a second, superficially unrelated device (e.g., an alarm system) was impaired compared to transfer to a superficially related device (e.g., an airplane with differently labeled controls) when the order of sub-goals was changed. Transfer was poorer still when the order of steps within the sub-goals was changed, compared to when the steps were unchanged, regardless of whether the sub-goal order was also altered. Hence, the order of the steps comprising the deep structure was not as important to maintaining acceptable transfer as the order of the routines within these steps (Dixon *et al.* 1997).

Similarity and Context

The context in which a judgment is made can greatly alter perceived similarity. For example, changing the context during problem-solving can affect encoding of the problem, which in turn can either facilitate or deter successful transfer. In a “pass-along” task, in which blocks of various sizes are shifted within a frame from an initial configuration to a known goal-state, completion of a difficult problem becomes easier if an analogy can be identified between the difficult problem and an easier one (Zamani & Richard 2000). For instance, in the difficult problem, two rectangular blocks are encoded as either two halves of a square or as two separate rectangles depending on whether an easier problem with an identical solution, but with a square block instead of two rectangular blocks, is

shown first or not at all (Zamani & Richard 2000). Furthermore, for the difficult problem to be used as an analogue for an even harder problem, the problems must share the same goal-state. Recognition of similar goal-states allows for the application of the same solution strategies for both problems (Zamani & Richard 2000). If the two difficult problems have different goal-states, and hence different solution procedures, knowledge is not transferred from one to the other (Zamani & Richard 2000).

The context of training can also affect the judgment of similarity between problems. For example, in category learning, if people are trained to associate a context cue with a particular region of category space, then despite the fact that the context cue does not predict category membership, people will use context to gate their responses and will treat an *identical* stimulus differently in two different contexts (Yang & Lewandowsky 2003, 2004; see also the earlier discussion of Lewandowsky and Kirsner 2000).

Failures of Transfer

When people fail to perceive the similarity between what they know and a novel task, transfer does not occur. People may fail to note relevant isomorphisms for a variety of reasons.

Context Specificity

Transfer fails more readily if the first task is more context-specific than the second one. For example, if people are initially trained to answer physics problems and are then tested with more general algebra problems, which nonetheless involve the same concepts, transfer is poor (Bassok & Holyoak 1989). Conversely, if people are initially trained with algebra problems, transfer to physics problems remains intact (Bassok & Holyoak 1989).

Extent of Learning

Failures of transfer can stem from failures of learning, for example when training involves only a limited number of problems (see, e.g., Catrambone & Holyoak 1990; Loewenstein *et al.* 1999). Likewise, if training involves only prototypical examples, then transfer will only be possible for target problems that are suitably similar to the prototypes (Elio & Anderson 1984; see also Gick & Holyoak 1987). The inverse of this statement, that transfer is greater if training involves a broader range of problems, is also true and we have already noted that it may underlie apparent instances of “adaptive” expertise.

Accordingly, techniques that improve learning have also been shown to improve subsequent transfer (Aleven & Koedinger 2002). For instance, compared to rote learning, transfer is better after learning that required participants to generate solutions to problems (Flint 2003) or test hypotheses (Burns & Vollmeyer 2002). Similarly, training which emphasizes different objectives can facilitate the transfer of different skills (Bitan & Karni 2003). For example, when trained to recognize nonsense words (e.g., PON, LOP) composed of a Morse code-like script (e.g., PON = |^**|^|^, LOP = *|^*||^*), people were able

to transfer knowledge of the specific “letters” (i.e., they could recognize novel “words” composed of the letters used at training) only when initially instructed on how to decode the script. When people were trained on non-alphabetical words (i.e., the Morse code-like script did not consistently map to specific letters), they were unable to transfer any of the learning to novel words (Bitan & Karni 2003). Importantly, people who learned how to decode the script performed much worse on old words comprised of new symbols than people trained on non-alphabetical words. The demonstration of both positive and negative transfer within the same condition illustrates the differential effects of training with different objectives.

Negative Transfer

Negative transfer is said to occur when performance on a novel task following training is poorer than it would have been without any prior training. Although typically not accompanied by the fatal consequences that struck the unfortunate French steel worker, negative transfer can occur whenever surface similarities mask structural differences. For example, Woltz *et al.* (2000) trained subjects in a complex number reduction task in which the rule for reduction of a larger number to a smaller number was determined by the relationship between the first two digits of the larger number (e.g., if the two digits differ by a value of two, replace those two digits with the midpoint between the two digits, or if two digits are equal remove the first digit). Participants were initially trained solely on stimuli that required the application of one sequence of rules. For example, the numbers 3565 and 9767 both require application of the “midpoint” rule for the first and second reduction and the “equal” rule for the final reduction (e.g., 3565 becomes 465, 465 becomes 55, and 55 becomes 5). Subjects produced errors when new sequences were presented at transfer that initially resembled training sequences but that required the application of a different rule for the last reduction (e.g., 3567 requires application of the midpoint rule twice to reduce the number to 57, but the final reduction is different from the training stimuli). Hence, for these “garden path” stimuli, the initial similarity in rule application masked the necessity of a different final rule and thus resulted in negative transfer compared to stimuli which were either completely dissimilarity or highly similar to the training stimuli.

Moreover, when task complexity was increased (e.g., by increasing the number of possible rules), increased training led not only to increased positive transfer for novel sequences (i.e., faster response latencies), but also to an increase in undetected errors for new “garden path” sequences (Woltz *et al.* 2000, Experiment 2). In this case, as in the case of the French factory worker, increased training and expertise led to enhanced negative transfer.

These instances of negative transfer, which usually occur spontaneously, have often been referred to as “strong-but-wrong” errors (Norman 1981). Reason (1990) linked strong-but-wrong errors to a process of frequency gambling and similarity matching. That is, if a process has been successful in the past, people are more likely to continue applying that same process if it appears similar to the target task (see also Lovett & Anderson 1996). Negative transfer is distinct from other forms of error, such as simply computing an incorrect response from a correct algorithm, because the performance decrement involves the application of prior knowledge or training in a situation that does not require

it. In Woltz *et al.*'s (2000) number-reduction task, we can distinguish between negative transfer and calculation error because errors on "garden path" problems were committed with the same speed as correct responses to training problems. By contrast, latencies for novel regular transfer sequences were longer than for training sequences. One implication of negative transfer is that the resultant "strong-but-wrong" errors go unnoticed and thus escape the possibility of discovery and correction (Woltz *et al.* 2000).

Increasing Positive Transfer

On the basis of the preceding discussion, one might be tempted to assume that simply informing people about the similarity between two tasks might enhance transfer. Contrary to that intuition, it turns out that transfer is facilitated if the process of discovering and extracting similarities between tasks, particularly similarities between relational information, is self-initiated (Dixon & Dohn 2003). In their study, people were given problems consisting of different numbers of alternating, connected beams, each supported by a fulcrum, which acted in a seesaw fashion with the action of the first beam affecting the action of the second, and so forth. People were told in which direction the first beam was moving and were asked to predict the direction of movement of the final beam. Participants were either told to classify each beam as an *up* beam or a *down* beam in alternation or were given no instructions. Half of the people who did not receive instructions quickly discovered the up/down strategy, while all of the people given the up/down instructions used that strategy exclusively. When shown new problems involving gear systems, those people who received no instructions and nonetheless discovered the up/down strategy also quickly discovered an analogue of the up/down strategy and applied it to the gear problems. The people who received instructions, however, fell back to a less efficient tracing strategy. Hence, people demonstrated better transfer when allowed to discover more efficient strategies for themselves, a process labeled "redescription" by Dixon and Dohn (2003). It follows that training regimes which allow for self-discovery should lead to more effective transfer, although this notion has yet to be tested.

CONCLUSIONS

We have touched on a variety of issues in research on knowledge and expertise. At the risk of glib oversimplification, we propose to condense our review to the claim that knowledge is best understood by rejecting its existence as a coherent concept: Instead of talking about "knowledge," we prefer to talk about a *set of highly context-specific learned responses*. Adopting this perspective appears to be particularly useful for the practitioner because it automatically accommodates the following major limitations and shortcomings of knowledge and expertise:

1. Expertise is highly domain-specific and brittle. Accordingly, expert performance can suffer dramatically if the deep structure of a task is altered, even if only slightly. In an applied setting, any alteration of an expert's domain is likely to result in decreased performance. This decrease is likely to be more severe if the deep structure of the task,

- such as the number and sequence of steps involved in performing a task or a task rule, is altered. It follows that practitioners should ensure that the steps and rules in the current task closely match the steps and rules in the expert's domain of expertise.
2. While expertise transfers well within a domain, little or no transfer can be expected outside a domain. In general, transfer often falls short of what intuition might lead one to expect because it occurs only if people correctly *perceive* two tasks to be similar. In practice, it is crucial that practitioners wanting to ensure positive transfer maximize the likelihood of two tasks being perceived similarly. A pertinent example of this is the release of software updates. If the updated software changes the settings so that different labels are given to identical function, the user's performance will suffer, until he or she is well practiced with the new version.
 3. Negative transfer can result, with potentially grievous consequences, if people are misled by the surface similarity between two tasks with very different deep structures. The resulting "strong-but-wrong" errors often escape detection and correction. In practice, this means that tasks which require different operations should be given different surface features to minimize negative transfer. This is particularly true when two tasks with different operations are required in close temporal or spatial proximity.
 4. Knowledge frequently reveals itself to be fragmented and partitioned, with people exhibiting quite contradictory behavior on an otherwise identical problem in different contexts. Even experts, who are typically assumed to have a highly integrated knowledge base, may exhibit surprisingly contradictory behavior. In response, practitioners should ensure that training occurs in a variety of contexts, and should make trainees aware of the potential hazards of partitioned knowledge.
 5. Experts also often exhibit expediency; that is, the tendency to master a task by focusing on a few key variables at the expense of ignoring other, and sometimes crucial, information. One way to tackle this issue is to highlight the importance of considering all information during training. If training highlights the pitfalls of ignoring relevant information, perhaps by designing a task where ignoring relevant information leads to failure, then the learner will hopefully be more aware of this shortcoming of expertise. It has also been suggested that providing the opportunity for learning by self-discovery can result in more flexible knowledge that is readily transferred.

Although the list of limitations and shortcomings is perhaps sobering, it need not detract from the stunning achievements that people are capable of. Whether considered "knowledge" or a "set of context-specific responses," abilities such as the retention of 30000 digits of π or 50000 chess patterns, or the ability safely to operate a machine as complex as an Airbus A380, are remarkable by any criterion.

NOTE

- 1 By the same token, research has identified domains in which exceptional performance cannot be detected. For example, people who claim to be speed readers have been found to exhibit remarkable dexterity at turning pages without displaying any comprehension of the text (Homa 1983). Those "domains" are commonly excluded from consideration in research on expertise.

REFERENCES

- Aleven, V. A. W. M. M. & Koedinger, K. R. (2002). An effective metacognitive strategy: learning by doing and explaining with a computer-based Cognitive Tutor. *Cognitive Science*, 26, 147–179.
- Anderson, J. R. (1990). *The Adaptive Character of Thought*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Bar, M. & Ullman, S. (1996). Spatial context in recognition. *Perception*, 25, 343–352.
- Barnett, S. M. & Koslowski, B. (2002). Adaptive expertise: effects of type of experience and the level of theoretical understanding it generates. *Thinking and Reasoning*, 8, 237–267.
- Barsalou, L. W. (1982). Context-independent and context-dependent information in concepts. *Memory & Cognition*, 10, 82–93.
- Barsalou, L. W. (1983). Ad hoc categories. *Memory & Cognition*, 11, 211–227.
- Barsalou, L. W. (1985). Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11, 629–654.
- Barsalou, L. W. (1999). Perceptua. Symbol systems. *Behavioral and Brain Sciences*, 22, 577–660.
- Barsalou, L. W. & Hale, C. R. (1993). Components of conceptual representation: from feature lists to recursive frames. In I. Van Mechelen, J. Hampton, R. S. Michalski & P. Theuns (eds.), *Categories and Concepts: Theoretical Views and Inductive Data Analysis* (pp. 97–144). San Diego, CA: Academic Press.
- Bassok, M. & Holyoak, K. J. (1989). Interdomain transfer between isomorphic topics in algebra and physics. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 15, 153–166.
- Bassok, M. & Holyoak, K. J. (1993). Pragmatic knowledge and conceptual structure: determinants of transfer between quantitative domains. In D. K. Detterman & R. J. Sternberg (eds.), *Transfer on Trial: Intelligence, Cognition, and Instruction*. Norwood, NJ: Ablex Publishing Corporation.
- Bédard, J. & Chi, M. T. H. (1992). Expertise. *Current Directions in Psychological Science*, 1, 135–139.
- Bellenkes, A. H., Wickens, C. D. & Kramer, A. F. (1997). Visual scanning and pilot expertise: the role of attentional flexibility and mental model development. *Aviation, Space & Environmental Medicine*, 68, 569–579.
- Bereby-Meyer, Y. & Kaplan, A. (2005). Motivational influences on transfer of problem-solving strategies. *Contemporary Educational Psychology*, 30, 1–22.
- Berg, G. (1992). Representational adequacy and the case for a hybrid connectionist/marker-parsing model. In R. G. Reilly & N. E. Sharkey (eds.), *Connectionist Approaches to Natural Language Processing* (pp. 253–272). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Besnard, D. & Cacitti, L. (2005). Interface changes causing accidents: an empirical study of negative transfer. *International Journal of Human-Computer Studies*, 62, 105–125.
- Bitan, T. & Karni, A. (2003). Alphabetical knowledge from whole words training: effects of explicit instruction and implicit experience on learning script segmentation. *Cognitive Brain Research*, 16, 323–337.
- Brooks, L. R. & Vokey, J. R. (1991). Abstract analogies and abstracted grammars: comments on Reber (1989) and Mathews *et al.* (1989). *Journal of Experimental Psychology: General*, 120, 316–323.
- Burns, B. D. & Vollmeyer, R. (2002). Goal specificity effects on hypothesis testing in problem solving. *The Quarterly Journal of Experimental Psychology*, 55, 241–261.
- Camerer, C. F. & Johnson, E. J. (1991). The process-performance paradox in expert judgment: how can experts know so much and predict so badly? In K. A. Ericsson & J. Smith (eds.), *Towards a General Theory of Expertise: Prospects and Limits* (pp. 195–217). New York: Cambridge University Press.
- Campbell, J. I. D. (1994). Architectures for numerical cognition. *Cognition*, 53, 1–44.
- Carraher, T. N., Carraher, D. W. & Schliemann, A. D. (1985). Mathematics in the streets and in schools. *British Journal of Developmental Psychology*, 3, 21–29.

- Catrambone, R. & Holyoak, K. J. (1990). Learning subgoals and methods for solving probability problems. *Memory & Cognition*, 18, 593–603.
- Ceci, S. J., Rosenblum, T. B. & DeBruyn, E. (1999). Laboratory versus field approaches to cognition. In R. J. Sternberg (ed.), *The Nature of Cognition* (pp. 385–408). Cambridge, MA: MIT Press.
- Ceci, S. J. & Ruiz, A. (1993). Transfer, abstractness, and intelligence. In D. K. Detterman & R. J. Sternberg (eds.), *Transfer on Trial: Intelligence, Cognition, and Instruction* (pp. 168–191). Norwood, NJ: Ablex.
- Charness, N. (1991). Expertise in chess: the balance between knowledge and search. In K. A. Ericsson & J. Smith (eds.), *Toward a General Theory of Expertise: Prospects and Limits* (pp. 39–63). New York: Cambridge University Press.
- Charness, N., Krampe, R. T. & Mayr, U. (1996). The role of practice and coaching in entrepreneurial skill domains: an international comparison of life-span chess skill acquisition. In K. A. Ericsson (ed.), *The Road to Excellence: The Acquisition of Expert Performance in the Arts, Sciences, Sports, and Games* (pp. 51–80). Mahwah, NJ: Lawrence Erlbaum Associates.
- Charness, N. & Schultetus, R. S. (1999). Knowledge and expertise. In F. T. Durso, R. S. Nickerson, R. W. Schvaneveldt et al. (eds.), *Handbook of Applied Cognition* (pp. 57–81). Chichester: John Wiley & Sons.
- Chase, W. G. & Ericsson, K. A. (1981). Skilled memory. In J. R. Anderson (ed.), *Cognitive Skills and Their Acquisition* (pp. 141–189). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Chase, W. G. & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, 4, 55–81.
- 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.
- Cormier, S. M. (1987). The structural process underlying transfer of training. In S. M. Cormier & J. D. Hagman (eds.), *Transfer of Learning: Contemporary Research and Applications* (pp. 151–181). San Diego, CA: Academic Press, Inc.
- Cromwell, S. E. & Kolb, J. A. (2004). An examination of work-environment support factors affecting transfer of supervisory skills training to the workplace. *Human Resource Development Quarterly*, 15, 449–471.
- Delaney, P. F., Reder, L. M., Staszewski, J. J. & Ritter, F. E. (1998). The strategy specific nature of improvement: the power law applies by strategy within task. *Psychological Science*, 9(1), 1–8.
- Delisle, S., Moulin, B. & Copeck, T. (2003). Surface-marker-based dialog modeling: a progress report on the MAREDI project. *Natural Language Engineering*, 9, 325–363.
- Dienes, Z. & Altmann, G. (1998). Transfer of implicit knowledge across domains: how implicit and how abstract? In D. Berry (ed.), *How Implicit is Implicit Learning?* (pp. 107–123). London: Oxford University Press.
- diSessa, A. A. (1988). Knowledge in pieces. In G. Forman & P. B. Pufall (eds.), *Constructivism in the Computer Age. The Jean Piaget Symposium Series* (pp. 49–70). Hillsdale, NJ: Lawrence Erlbaum Associates.
- diSessa, A. A., Gillespie, N. M. & Esterly, J. B. (2004). Coherence versus fragmentation in the development of the concept of force. *Cognitive Science*, 28, 843–900.
- Dixon, J. A. & Dohn, M. C. (2003). Redescription disembeds relations: evidence from relational transfer and use in problem solving. *Memory & Cognition*, 31, 1082–1093.
- Dixon, P. & Gabrys, G. (1991). Learning to operate complex devices: effects of conceptual and operational similarity. *Human Factors*, 33, 103–120.
- Dixon, P., Zimmerman, C. & Nearey, S. (1997). Prior experience and complex procedures. *Memory & Cognition*, 25, 381–394.
- Dunbar, K. (1995). How scientists really reason: scientific reasoning in real-world laboratories. In R. J. Sternberg & J. E. Davidson (eds.), *The Nature of Insight* (pp. 365–395). Cambridge, MA: MIT Press.
- Elio, R. & Anderson, J. R. (1984). The effects of information order and learning mode on schema abstraction. *Memory & Cognition*, 12, 20–30.
- Ericsson, K. A. (1996). The acquisition of expert performance: an introduction to some of the issues. In K. A. Ericsson (ed.), *The Road to Excellence: The Acquisition of Expert Performance in the Arts and Sciences, Sports, and Games* (pp. 1–50). Mahwah, NJ: Lawrence Erlbaum Associates.

- Ericsson, K. A. (2003). Memorizers are made, not born. *Trends in Cognitive Science*, 7, 233–235.
- Ericsson, K. A. (2005). Recent advances in expertise research: a commentary on the contributions to the special issue. *Applied Cognitive Psychology*, 19, 233–241.
- Ericsson, K. A., Chase, W. G. & Faloon, S. (1980). Acquisition of a memory skill. *Science*, 208, 1181–1182.
- Ericsson, K. A., Delaney, P. F., Weaver, G. A. & Mahadevan, S. (2004). Uncovering the structure of a memorist's superior "basic" memory capacity. *Cognitive Psychology*, 49, 191–237.
- Ericsson, K. A., Krampe, R. T. & Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, 100, 363–406.
- Ericsson, K. A. & Lehmann, A. C. (1996). Expert and exceptional performance: evidence of maximal adaptation to task. *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.
- Even, R. (1988). Factors involved in linking representations of functions. *Journal of Mathematical Behavior*, 17, 105–121.
- Flint, W. (2003). Transfer of learning: evaluating workplace transfer of training techniques. In J. A. Chambers (ed.), *Selected Papers from the 14th International Conference on College Teaching and Learning* (pp. 93–101). Jacksonville, FL: Florida Community College.
- Frensch, P. A. & Sternberg, R. J. (1989). Expertise and intelligent thinking: when is it worse to know better? In R. J. Sternberg (ed.), *Advances in the Psychology of Human Intelligence* (Vol. 5, pp. 157–188). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Freund, H. J. (1983). Motor unit and muscle activity in voluntary motor control. *Physiological Review*, 63, 387–436.
- Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (eds.), *Similarity and Analogical Reasoning* (pp. 199–241). London: Cambridge University Press.
- Gick, M. L. & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology*, 12, 306–355.
- Gick, M. L. & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, 15, 1–38.
- Gick, M. L. & Holyoak, K. J. (1987). The cognitive basis of knowledge transfer. In S. M. Cormier & J. D. Hagman (eds.), *Transfer of Learning: Contemporary Research and Applications* (pp. 9–46). San Diego, CA: Academic Press.
- Glaser, R. (1996). Changing the agency for learning: acquiring expert performance. In K. A. Ericsson (ed.), *The Road to Excellence: The Acquisition of Expert Performance in the Arts and Sciences, Sports, and Games* (pp. 303–311). Mahwah, NJ: Lawrence Erlbaum Associates.
- Gobet, F. & Simon, H. A. (1996). The roles of recognition processes and look-ahead search in time-constrained expert problem solving: evidence from grandmaster level chess. *Psychological Science*, 7, 52–55.
- Gobet, F. & Waters, A. J. (2003). The role of constraints in xpert memory. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 29, 1082–1094.
- 6 Gomez, R. L., Gerken, L. & Schevaneveldt, R. W. (2000). The basis of transfer in artificial grammar learning. *Memory & Cognition*, 28, 253–263.
- Gott, S. P., Hall, E. P., Pokorny, R. A. *et al.* (1993). A naturalistic study of transfer: adaptive expertise in technical domains. In D. K. Detterman & R. J. Sternberg (eds.), *Transfer on Trial: Intelligence, Cognition and Instruction* (pp. 258–288). Norwood, NJ: Ablex.
- Griffiths, T. L. & Kalish, M. (2002). A multidimensional scaling approach to mental multiplication. *Memory and Cognition*, 30, 97–106.
- Griffiths, T. L. & Steyvers, M. (2002). A probabilistic approach to semantic representation. In *Proceedings of the Twenty-Fourth Annual Conference of Cognitive Science Society*. George Mason University, Fairfax, VA.
- de Groot, A. D. (1965). *Thought and Choice in Chess*. The Hague: Mouton.
- Hayes-Roth, B. (1977). Evolution of cognitive structures and processes. *Psychological Review*, 84, 260–278.
- Hershey, D. A. & Walsh, D. A. (2000). Knowledge versus experience in financial problem solving. *Current Psychology: Developmental, Learning & Personality*, 19, 261–291.

- Hoffman, T. (1999). Probabilistic latent semantic analysis. In K. B. Laskey & H. Prade (eds.), *UAI '99: Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, Stockholm, Sweden, July 30–August 1, 1999. Morgan Kaufmann.
- Holland, J. H., Holyoak, K. J., Nisbett, R. E. & Thagard, P. R. (1986). *Induction: Processes of Inference, Learning, and Discovery*. Cambridge, MA: MIT Press.
- Holyoak, K. J. (1991). Symbolic connectionism: toward third-generation theories of expertise. In A. Ericsson & J. Smith (eds.), *Toward a General Theory of Expertise: Prospects and Limits* (pp. 301–355). Cambridge: Cambridge University Press.
- Holyoak, K. J. & Koh, K. (1987). Surface and structural similarity in analogical transfer. *Memory & Cognition*, 15, 332–340.
- Homa, D. (1983). An assessment of two “extraordinary” speed-readers. *Bulletin of the Psychonomic Society*, 21, 123–126.
- Johnson, E. J. (1988). Expertise and decisions under uncertainty: performance and process. In M. L. H. Chi, R. Glaser & M. J. Farr (eds.), *The Nature of Expertise* (pp. 209–228). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Johnson, P. E., Grazioli, S., Jamal, K. & Zualkernan, I. A. (1992). Success and failure in expert reasoning. *Journal of Organizational Behavior and Human Decision Processes*, 53, 173–203.
- Kalish, M. L., Lewandowsky, S. & Kruschke, J. K. (2004). Population of linear experts: knowledge partitioning and function learning. *Psychological Review*, 111, 1072–1099.
- Keele, S. W. & Ivry, R. I. (1987). Modular analysis of timing in motor skill. In G. H. Bower (ed.), *The Psychology of Learning and Motivation* (Vol. 21, pp. 183–228). San Diego, CA: Academic Press.
- Kimball, D. R. & Holyoak, K. J. (2000). Transfer and expertise. In E. Tulving & F. I. M. Craik (eds.), *The Oxford Handbook of Memory* (pp. 109–122). London: Oxford University Press.
- Kintsch, W. (2000). The control of knowledge activation in discourse comprehension. In W. J. Perrig & A. Grob (eds.), *Control of Human Behavior, Mental Processes, and Consciousness: Essays in Honor of the 60th Birthday of August Flammer* (pp. 137–146). Mahwah, NJ: Lawrence Erlbaum Associates.
- Kolers, P. A. (1976). Pattern analyzing memory. *Science*, 191, 1280–1281.
- Landauer, T. K. & Dumais, S. T. (1997). A solution to Plato’s problem: the Latent Semantic Analysis theory of acquisition, induction and representation of knowledge. *Psychological Review*, 104, 211–240.
- Lange, T. E. (1992). Hybrid connectionist models: temporary bridges over the gap between the symbolic and the subsymbolic. In J. Dinsmore (ed.), *The Symbolic and Connectionist Paradigms: Closing the Gap* (pp. 237–289). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Lewandowsky, S., Kalish, M. & Ngang, S. K. (2002). Simplified learning in complex situations: knowledge partitioning in function learning. *Journal of Experimental Psychology: General*, 131, 163–193.
- Lewandowsky, S. & Kirsner, K. (2000). Knowledge partitioning: context-dependent use of expertise. *Memory & Cognition*, 28, 295–305.
- Loewenstein, J., Thompson, L. & Gentner, D. (1999). Analogical encoding facilitates knowledge transfer in negotiation. *Psychonomic Bulletin & Review*, 6, 586–597.
- Lovett, M. C. & Anderson, J. R. (1996). History of success and current context in problem solving: combined influences on operator selection. *Cognitive Psychology*, 31, 168–217.
- Lovett, M. C. & Schunn, C. D. (1999). Task representations, strategy variability, and base-rate neglect. *Journal of Experimental Psychology: General*, 128, 107–130.
- Marchant, G., Robinson, J., Anderson, U. & Schadewald, M. (1991). Analogical transfer and expertise in legal reasoning. *Organizational Behavior & Human Decision Processes*, 48, 272–290.
- Marshall, S. P. (1995). *Schemas in Problem Solving*. New York: Cambridge University Press.
- Medin, D., Ross, B. H. & Markman, A. B. (2001). *Cognitive Psychology* (3rd edn). Chichester: John Wiley & Sons.
- Moray, N. (1999). Mental models in theory and practice. In D. Gopher & A. Koriat (eds.), *Attention and Performance XVII: Cognitive Regulation of Performance: Interaction of Theory and Application. Attention and Performance* (pp. 223–258). Cambridge, MA: MIT Press.
- Norman, D. A. (1981). Categorization of action slips. *Psychology Review*, 88, 1–15.

- Norman, G. R., Brooks, L. R., Coblenz, C. K. & Babcock, C. J. (1992). The interdependence of feature and category identification in diagnostic radiology. *Memory & Cognition*, *20*, 344–355.
- Novick, L. R. (1988). Analogical transfer, problem similarity and expertise. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *14*, 510–520.
- Novick, L. R. & Holyoak, K. J. (1991). Mathematical problem solving by analogy. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 398–415.
- Nunes, T., Schliemann, A. D. & Carraher, D. W. (1993). *Street Mathematics and School Mathematics* (pp. 77–126). New York: Cambridge University Press.
- Patel, V. L., Kaufman, D. R. & Magder, S. (1996). The acquisition of medical expertise in complex dynamic decision-making environments. In A. Ericsson (ed.), *The Road to Excellence: The Acquisition of Expert Performance in the Arts and Sciences, Sports and Games* (pp. 127–165). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Pecher, D., Zeelenberg, R. & Barsalou, L. W. (2004). Sensorimotor simulations underlie conceptual representations: modality-specific effects of prior activation. *Psychonomic Bulletin & Review*, *11*, 164–167.
- Reason, J. T. (1990). *Human Error*. New York: Cambridge University Press.
- Reber, P. J. & Squire, L. R. (1994). Parallel brain systems for learning with and without awareness. *Learning and Memory*, *1*, 217–229.
- Reder, L. M. & Ritter, F. E. (1992). What determines initial feeling of knowing? Familiarity with question terms, not with the answer. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 435–451.
- Richman, H. B., Staszewski, J. J. & Simon, H. A. (1995). Simulation of expert memory using EPAM IV. *Psychological Review*, *102*, 305–330.
- Roediger, H. L. III, Meade, M. L. & Bergman, E. T. (2001). Social contagion of memory. *Psychonomic Bulletin & Review*, *8*, 365–371.
- Rossano, M. J. (2003). *Evolutionary Psychology: The Science of Human Behaviour and Evolution*. Chichester: John Wiley & Sons.
- Salthouse, T. A. (1984). Effects of age and skill in typing. *Journal of Experimental Psychology: General*, *13*, 345–371.
- Salthouse, T. A. (1991). Expertise as the circumvention of human processing limitations. In K. A. Ericsson & J. Smith (eds.), *Toward a General Theory of Expertise: Prospects and Limits* (pp. 286–300). New York: Cambridge University Press.
- Santtila, P., Korpela, S. & Häkkänen, H. (2004). Expertise and decision-making in linking car crime series. *Psychology, Crime & Law*, *10*, 97–112.
- Schliemann, A. D. & Carraher, D. W. (1993). Proportional reasoning in and out of school. In P. Light & G. Butterworth (eds.), *Context and Cognition: Ways of Learning and Knowing* (pp. 47–73). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Schunn, C. D., Reder, L. M., Nhoyvannisvong, A., Richards, D. R. & Stroffolino, P. J. (1997). To calculate or not calculate: a source activation confusion (SAC) model of problem-familiarity's role in strategy selection. *Journal of Experimental Psychology: Learning, Memory & Cognition*, *23*, 3–29.
- Schwartz, S. & Griffin, T. (1986). Evaluating medical information. In: *Medical Thinking. The Psychology of Medical Judgment and Decision Making* (pp. 56–96). New York: Springer-Verlag.
- Shanks, D. R. & Johnstone, T. (1998). Implicit knowledge in sequential learning tasks. In M. A. Stadler & P. A. Frensch (eds.), *Handbook of Implicit Learning* (pp. 533–572). Thousand Oaks, CA: Sage.
- Shepard, R. N. (1987). Towards a universal law of generalization for psychological science. *Science*, *237*, 1317–1323.
- 7 Shepard, R. N., Hovland, C. L. & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs*, *75*, ••.
- Shrager, J. & Siegler, R. S. (1998). SCADS: a model of children's strategy choices and strategy discoveries. *Psychological Science*, *9*, 405–410.
- Siegler, R. S. (1987). Some general conclusions about children's strategy choice procedures. *International Journal of Psychology*, *22*, 729–749.

- Simon, H. A. & Gilmarin, K. J. (1973). A simulation of memory for chess positions. *Cognitive Psychology*, 5, 29–46.
- Sims, V. K. & Mayer, R. E. (2002). Domain specificity of spatial expertise: the case of video game players. *Applied Cognitive Psychology*, 16, 97–115.
- Singley, M. K. & Anderson, J. R. (1989). *Transfer of Cognitive Skill*. Cambridge, MA: Harvard University Press.
- Staszewski, J. J. (1993). A theory of skilled memory. In *Proceedings of the 15th Annual Conference of the Cognitive Science Society* (pp. 971–975). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Sternberg, R. J. & Frensch, P. A. (1992). On being an expert: a cost–benefit analysis. In R. Hoffman (ed.), *The Cognition of Experts: Psychological Research and Empirical AI* (pp. 191–203). New York: Springer-Verlag.
- Thompson, C. P., Cowan, T. M. & Frieman, J. (1993). *Memory Search by a Memorist*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Tirosh, D. & Tsamir, P. (1996). The role of representations in students' intuitive thinking about infinity. *Journal of Mathematical Education in Science and Technology*, 27, 33–40.
- Tuckey, M. R. & Brewer, N. (2003). How schemas affect eyewitness memory over repeated retrieval attempts. *Applied Cognitive Psychology*, 17, 785–800.
- Vigliocco, G., Vinson, D. P., Lewis, W. & Garrett, M. F. (2004). Representing the meanings of object and action words: the Featural and Unitary Semantic System (FUSS) hypothesis. *Cognitive Psychology*, 48, 422–488.
- Wiley, J. (1998). Expertise as mental set: the effects of domain knowledge in creative problem solving. *Memory & Cognition*, 26, 716–730.
- Woltz, D. J., Gardner, M. K. & Bell, B. G. (2000). Negative transfer errors in sequential cognitive skills: strong-but-wrong sequence application. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 26, 601–625.
- Yang, L-X. & Lewandowsky, S. (2003). Context-gated knowledge partitioning in categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 663–679.
- Yang, L-X. & Lewandowsky, S. (2004). Knowledge partitioning in categorization: constraints on exemplar models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30, 1045–1064.
- Zamani, M. & Richard, J-F. (2000). Object encoding, goal similarity, and analogical transfer. *Memory & Cognition*, 28, 873–886.










AUTHOR QUERY-FORM

Dear Author,

During the preparation of your manuscript for publication, the questions listed below have arisen. Please attend to these matters and return this form with your proof.

Many thanks for your assistance.

Query References	Query	Remarks
1	Au: Novick and colleagues 1988 has not been found in the list.	
2.	Au: Griffiths and Kalish 2001 has been changed to Griffiths and Kalish 2002 so that this citation matches the list.	
3.	Au: Even 1998 has been changed to Even 1988 so that this citation matches the reference list. Please confirm that this is correct.	
4.	Au: Besnard & Cacitti 2004 has been changed to Besnard & Cacitti 2005 so that this citation matches the reference list. Please confirm that this is correct.	
5.	Au: Gomez et al. 1994 has not been found in the reference list.	
6.	Au: Gomez et al. 2000 has not been found in the text.	
7.	Au: Please supply the page range.	

HAC04

M