EXPERIENCE IS A DOUBLE-EDGED SWORD: A COMPUTATIONAL MODEL OF THE ENCODING/RETRIEVAL TRADE-OFF WITH FAMILIARITY

Lynne M. Reder, Christopher Paynter, Rachel A. Diana, Jiquan Ngiam, and Daniel Dickison

I. Introduction

Although many aspects of memory are not well understood, there are other aspects on which there is little debate. For example, one of the most basic laws of memory is that *practice benefits retention*. Indeed, the conventional wisdom that "practice makes perfect" is applicable whether the practice involves learning a skill (e.g., how to drive a car) or learning a fact (e.g., the name of the first American president). One need not to be a memory researcher to appreciate that the more experience one has with something, the easier it is to process. On the other hand, it is less appreciated that this same experience comes with costs. That is, familiarity with an item sometimes benefits and sometimes hurts performance, depending on the nature of the task.

One area in which this familiarity trade-off is increasingly evident is the domain of memory retrieval. Two decades ago, in this same *Psychology of Learning and Motivation* series, Reder (1988) wrote a chapter about the "strategic control of retrieval strategies" arguing against the (then) conventional wisdom that we always try to search our memory for an answer before

THE PSYCHOLOGY OF LEARNING
AND MOTIVATION VOL. 48
DOI: 10.1016/S0079-7421(07)48007-0

attempting to reason the answer by using other strategies That chapter highlighted the various factors that can make one strategy more useful than another, and also proposed that people *unconsciously* adapt their strategy use to optimize their performance (see also Cary & Reder, 2002; Koriat, 2000; Reder, Weber, Shang, & Vanyukov, 2003; Sun, 2000) A decade later, Schunn and Reder (1998) also wrote a chapter for this series, proposing that there are individual differences in the ability to rapidly adapt strategies to optimize performance Both chapters dealt with the notion that people do not behave in a monolithic fashion, but rather alter their strategies adaptively based on the contingencies of the environment, their own cognitive capacities, and the contents of their memory.

It is now generally understood and accepted that people use different strategies in different situations (Anderson & Betz, 2001; Reder, 1987; Shrager & Siegler, 1998) and that people vary in how quickly they adapt to how well a strategy is working (Schunn, Lovett, & Reder, 2001). In this chapter, we want to examine the variables that affect performance from the bottom up, rather than the top down That is, we will examine what aspects of the cognitive architecture make the same information an advantage or a liability depending on the task. Our focus is on the *trade-offs* that are inherent with experience and why these trade-offs occur from a mechanistic standpoint.

The first section of this chapter reviews the evidence that experience can be a liability when *retrieving* information and also explains the conditions when experience does not hurt performance at retrieval. In the second part of the chapter, we focus on how experience generally facilitates *encoding*, although we point out trade-offs here as well, such that familiarity can sometimes be a liability at encoding. As a part of these explanations, we describe a model that we have developed that can explain retrieval deficits with experience. The SAC model, which stands for source of activation confusion, has had success predicting many results, including some that were not intuitive. However, some additions to the model seem warranted in order to make it more complete and allow it to account for an even wider range of the data. We introduce a revised but more psychologically accurate model[1] that can explain how experience positively affects encoding.

[1] We will still call it SAC and like most computational models it undergoes additions and modifications to its assumptions. It is conventionally more parsimonious to keep the same name rather than to introduce a new name every time a change is made to a model If the changes were fundamental to the axiomatic assumptions of the model, then it would make sense to reject it and start over with something totally different That is not the case here

II. When and Why Experience Adversely Affects Memory Retrieval

If a person on the street were asked, "Do you think it is easier to answer a question about something if you know a lot about it?" the answer would almost certainly be, "Of course." Yet if the question was phrased, "If you were searching for a particular key would it be more difficult if there were many keys on the key ring or if there were only a few keys?" the answer would clearly be that discriminating a single key from many keys would be more difficult. This common intuition about physical search is just as applicable for memory search, that it is more difficult to find a specific fact if there are many contenders available. Below we review some of the evidence for the assertion that knowing more about a concept can hurt subsequent retrieval of any particular fact about the concept. We explain why that occurs from a mechanistic standpoint and why it does not always adversely affect performance.

A. The Fan Effect

Anderson and Bower (1973) demonstrated that when more statements had been previously studied that shared concepts with a given test probe, subjects were slower and less accurate to recognize that the test probe had been seen before. For instance, subjects were slower and less accurate to verify a studied sentence such as "The hippie touched the debutante" if more sentences had also been studied that shared the same terms (e.g., *hippie*, *touch*, or *debutante*). They dubbed this phenomenon the "fan effect" because they assumed a representation in which concepts were represented as nodes and associations connected the concepts such that the more concepts that "fanned" out of a node, the less activation could spread to any other associated node. Speed and accuracy are related to the amount of activation that reaches another node to make it available.

These types of effects have been demonstrated in many paradigms with many types of stimuli (Anderson & Paulson, 1978; Lewis & Anderson, 1976; Reder, Donavos, & Erickson, 2002; Zbrodoff, 1995), although there are some who have questioned the generality of these effects (Radvansky, 1999; Smith, Adams, & Schorr, 1978). The fan effect shows that having more information about a topic does not necessarily decrease memory retrieval time for probes of that topic and might increase it. Nevertheless, one might question whether fan effects observed in the laboratory are relevant to attempts to retrieve information in the real world.

1. *The Paradox of the Expert*

Smith et al. (1978) noted that a logical conclusion of the claim that fan effects are ubiquitous is that experts should be too slow to answer any questions posed to them and should always be lost in thought. Although anecdotal evidence seems to suggest that experts often cannot give a "straight answer," the authors' point is well taken, as it certainly does not seem experts are unable to give responses. Smith et al. demonstrated that when the facts used in a fan experiment belonged to a theme such as a ship christening (e.g., *Marty broke the bottle*), knowing more facts about an item (Marty) that were all consistent with the theme did not produce a fan effect. They suggested that thematically related information is organized into schemas that are represented in a qualitatively different way than a semantic network such as the one proposed by Anderson and Bower (1973). Moreover, they suggested that only when the materials were unrelated and unintegrated (and presumably, unnatural), the fan effect would occur. This seems to suggest that increasing experience may not decrease memory performance in most cases.

2 *Strategy Variability and Strategy Selection*

An alternative explanation that we ultimately put forward is that whether the fan effect hurts an expert (or anyone else) depends on the nature of the task requirements. Specifically, in some situations (e.g., memory tasks), people are obliged to use a "direct retrieval" strategy that is adversely affected by fan. In other situations, question answering can occur without using direct retrieval.

A few decades ago, the conventional wisdom concerning strategy use in question answering was that people first used a direct retrieval strategy wherein they searched for the answer to a question and only used an inference strategy if that initial direct retrieval attempt failed (Anderson, 1976; Kintsch, 1974; Norman, Rumelhart, & the LNR research group, 1975). Reder (1979, 1982) discovered that this conventional wisdom was erroneous. That is, people do not necessarily search for the answer to a question (direct retrieval) before adopting an inference strategy (plausible reasoning) to answer a question even when they are expressly told to search for a specific fact. Conceivably, the subjects in the Smith et al. (1978) paradigm were frequently opting to use a type of plausible reasoning or consistency strategy to answer the questions in their experiment, and the foils being used in their experiment did not preclude this behavior.[2] The hypothesis that Reder and Anderson (1980) tested was that depending on the type of foil, different strategies for question-answering would be selected.

[2] Smith et al. tested Reder's explanation (provided in a personal communication) by inserting a novel lexical item into the test probes, for example, "Marty broke the *champagne* bottle," and did not find that the fan effect reappeared. Reder discounted Smith et al.'s finding because the low-frequency novel lexical item provided an additional means of rejecting the probe as unstudied. Reder felt that it was important that the experiment control the familiarity of foils which motivated the study by Reder and Anderson (1980).

In that study, subjects produced fan effects, but only in certain trial blocks, depending on the nature of the foils in that block. In blocks in which the foils were not thematically related to study items, subjects could use a consistency or plausibility strategy (Reder, 1982, 1987; Reder, Wible, & Martin, 1986), and Reder and Anderson (1980) obtained the same null fan effect observed by Smith et al. (1978). However, in blocks in which a consistency strategy would not work because foils were thematically related, the fan effect re-emerged, suggesting that a direct retrieval strategy was used. The notion that subjects can adapt their strategy choice from one block to another has subsequently been demonstrated many times (Cary & Reder, 2002; Lemaire & Reder, 1999; Lovett & Schunn, 1999; Reder, 1982, 1987; Reder & Ross, 1983; Reder et al., 1986; Schunn & Reder, 1998).

Reder and Ross (1983) went on to show that the flat or null fan effect that emerged when subjects could get away with a consistency strategy actually resulted from a mixture of two processes: on some trials, subjects actually searched for the specific fact using the effortful retrieval process, while on other trials a subject would adopt the faster consistency judgment strategy (the fact retrieved is consistent with the probe statement). In the former case, the more related facts studied, the slower the verification; however, Reder and Ross also demonstrated that when subjects used the consistency strategy, the more relevant facts studied, the faster subjects were to verify the statement. They added a third type of test block in which subjects were specifically told to make their decision based on consistency. In the blocks that forced specific search because the foils were thematically related, the fan effect was found. In recognition blocks in which the foils were not thematically related and subjects could get away with using plausibility, the fan effect was flat or null. Importantly, in those blocks in which subjects were specifically instructed to base their judgments on the consistency of the probe to the studied statements regardless of whether that specific statement had been studied, verification was faster when more relevant facts had been studied. In other words, Reder and Ross (1983) found a *negative fan effect* when the appropriate strategy was plausibility or consistency rather than retrieving a specific statement from memory. The paradox of the expert was solved.

3 *Fan Effects with Real-World Knowledge*

Although the paradox of the expert was "solved" in that experts did not really search for an exact fact in memory, one could still wonder whether these manipulations only had effects on material learned in the laboratory. That is, the original demonstrations of the fan effect involved contrived laboratory statements that no undergraduate would ever believe was true, motivating the research by Smith et al. (1978) discussed above. Conceivably real semantic facts stored in memory would not be affected by this fan manipulation.

That question motivated several laboratory investigations of whether real-world knowledge could be affected by laboratory fan manipulations (Lewis & Anderson, 1976; Peterson & Potts, 1982). In those experiments, subjects learned fantasy facts (Lewis & Anderson) or esoteric (unknown) but true facts (Peterson & Potts) about famous individuals (e.g., George Washington, Napoleon Bonaparte) and later had to verify which newly learned statements had been studied about the famous character. The number of novel facts learned about a famous person was randomly determined for each subject. The time to verify a specific new fact increased monotonically with the number of studied facts, replicating the typical fan effect. The more interesting result was the effect that fan manipulation had on the time to verify previously known facts about a famous person. These real-world facts were also adversely affected by the number of new facts that had been learned about an individual. In other words, both episodic and semantic (real-world knowledge) memory were shown to be vulnerable to the fan effect.

4. *A Mechanistic Account of Retrieval Effects*

The original fan effects of Anderson and Bower (1973) were modeled with mathematical equations that produced excellent fits to the data. The response times were derived from the estimated time to activate the memory structure due to activation spread from the content words (source nodes) in the test probe to the connected representation in memory. The amount of activation spread[3] depended on the number of competitors sharing the activation of each of the probes.

Reder and Ross (1983) suggested that consistency judgments were based on the amount of activation that accrues at a given theme (e.g., lawyer) due to its relationship with a particular character (e.g., Marty). This activation accrual is affected by the number of themes associated with the character. The more themes associated with a person, the slower the response times for consistency judgments; however, the more facts associated with a given thematic node, the faster to make a consistency judgment. Reder and Ross (1983) presented a verbal description that is consistent with recent modeling implementations. Specifically, they suggested that the theme node and the link between it and the character node would become stronger with each additional thematic fact studied.

[3] When first proposed, the description involved time for activation to spread. In revisions of the theory, the assumptions changed to the amount of activation available to spread. Latency is an inverse function of activation.

Neither of these mathematical models was implemented as a computational model. However, Anderson in recent decades has developed a sophisticated cognitive architecture, ACT-R (Anderson & Lebiere, 1998) that can easily account for these types of fan effects (Anderson & Reder, 1999). Reder developed a related, but simpler model of memory called SAC that does not address skill learning, but that has been used to account for a wide variety of memory phenomena (some not easily accommodated by ACT-R). These include feeling of knowing effects (Reder & Schunn, 1996; Schunn, Reder, Nhouyvanisvong, Richards, & Stroffolino, 1997), word frequency mirror effects (Reder et al., 2000), perceptual match effects (Diana, Peterson, & Reder, 2004, Reder et al., 2002), paired associate learning and cued recall (Reder, Park, & Kieffaber, 2007a), and aging effects on memory (Buchler & Reder, 2007). The ACT-R mechanism for spread of activation was included in SAC assumptions, so the explanation for the fan effect is the same.

Although many of the assumptions of SAC were imported from ACT-R, other assumptions of SAC are not part of the ACT-R architecture. For example, SAC allows phenomenological judgments to be made based on activation values of nodes (chunks) while ACT-R does not allow activation levels to be "read"[4] in this way. It is worth emphasizing that the fan effect, which plays an important role in both SAC and ACT-R, is concerned only with retrieval, not encoding. At this time, ACT-R does not make any assumptions about differential probability of encoding. In the second half of this chapter, we will describe modifications to SAC that posit differential probability of encoding information. These modifications allow the model to account for various effects demonstrating both the advantages and disadvantages of familiarity in memory.

B. The SAC Memory Model: The Role of Experience in Recognition Memory

The SAC model was initially developed to account for a series of feeling of knowing experiments (Reder & Ritter, 1992; Reder & Schunn, 1996; Schunn et al., 1997).[5] However, SAC also makes very strong predictions concerning the role of experience on memory performance, and these basic assumptions

[4] It seems likely that ACT-R could be modified to make the same predictions as SAC. In our view, some of the SAC assumptions provide a better account of certain phenomena; however, it is probably not practical for ACT-R to import those assumptions now. Since all theories are only approximations to the truth, hopefully the better assumptions of theories will be adopted by other theories and ultimately become one and the same.

[5] The motivation for those experiments was to test the assumption that people could quickly evaluate whether to search for an answer or use a reasoning strategy (Reder, 1987; Reder et al., 1986).

and necessary predictions seemed inconsistent with findings in the literature. Specifically, others had claimed manipulating word frequency in a recognition memory task produced a *dissociation* such that recollection judgments are affected by word frequency but familiarity judgments are not (Gardiner & Java, 1990). It is a central assumption of SAC that high- and low-frequency words should differ in their inherent familiarity because they differ in how often they have been previously experienced. This apparent contradiction of a basic axiom of the model motivated further exploration of this claimed dissociation. Further research made it clear that the conventional wisdom was incorrect.

Before recounting those experiments, a description of the assumptions of SAC is in order. These are the original assumptions of the simpler version of the model. The recent elaborations to SAC that incorporate assumptions about working memory (WM) and how experience affects encoding will be introduced later in the chapter.

SAC is an experience/history sensitive model that represents information as a set of interconnected concepts (we refer to them as nodes). Concept nodes are linked to semantically related nodes as well as nodes representing the constituent features of the concept (e.g., phonemic and lexical features, semantic features).[6] There also exist episode nodes that are linked to the concept nodes and which provide information about having seen a concept in a particular context. Any idiosyncratic features of the experience will be individually bound to the episode node, which is connected through memory linkages to both conceptual and perceptual aspects of the experience. There is also a node for the general experimental context in the model that has features of the experiment bound to it and which is also linked to the episode nodes. An illustration of these representational assumptions is shown in Fig. 1. A central assumption is that all aspects of a memory experience follow the same principles, regardless of whether the information is conceptual or perceptual. In other words, all nodes in the network strengthen and decay according to the same rules. Although this model uses a localist, rather than a distributed representation such as the PDP framework of McClelland and Rumelhart (1985), each concept is associated with a wide variety of features, a subset of which can activate the episode node. It is the detailed specification

[6] The representation is necessarily schematic and not all features of the experience are represented such as the language that the word is presented in; however, we believe that the perceptual and lexical features are often part of the representation, depending on the attention given to various aspects of the experience. For simplicity, we do not represent features that are probably part of the mental representation and do not affect our account of the phenomena.

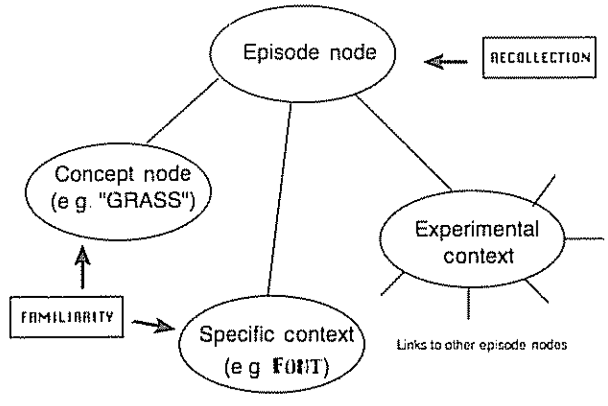

Fig 1 A schematic representation of the structure of the SAC model

of how representations change with experience and how activation values are interpreted in particular situations that allows SAC to make specific, quantifiable predictions for many types of tasks

1 Node Strength

The strength of a concept (node in our theory) represents the history of exposure to that concept, with more exposure producing greater strengthening Strength can also be thought of as the baseline or resting level of activation of a node Increases and decreases in this baseline strength change according to a power function:

$$B = c \sum t_i^{-d} \qquad (1)$$

where B is the base-level activation, c and d are constants, and t_i is the time since the ith presentation This function captures both power law decay of memories with time and power law learning of memories with practice Very strong regularities have been found wherever these issues have been studied (Anderson & Schooler, 1991) The central feature of power law decay is that initially memories decay quickly and then much more slowly at increasing delays Similarly, the central feature of power law learning is that first exposures to an item contribute more than subsequent exposures That is, the incremental contribution of each new exposure decreases with increasing numbers of exposures

2. *Link Strength*

Links connect nodes that have been associated together by being thought of or experienced at the same time. The strength of these links will vary as a function of how many times the concepts had been associated together and the time delay between exposures. Specifically, we assume a power function given by:

$$S_{s,r} = \sum t_i^{-d_L} \tag{2}$$

where $S_{s,r}$ is the strength of the link from the node s to node r, t_i is the time since the ith coexposure, and d_L is the decay constant for links.

3. *Spread of Activation*

The current activation level of a node can increase by receiving environmental stimulation directly or by receiving activation that has "spilled over" from another node in the network to which it is linked. The increase in activation of some node r, which is receiving activation from other nodes, is computed by summing the activation it is receiving from all (source) nodes. However, the amount of activation each source node sends depends on (a) that source node's strength and (b) how much competition the connection from the source to node r has from other links associated with that source. The change in activation of some node r is computed by summing the spread of activation from all source nodes s connected to node r according to the equation:

$$\Delta A_r = \sum\left(\frac{A_s S_{s,r}}{\Sigma S_{s,i}}\right) \tag{3}$$

where ΔA_r is the change in activation of the receiving node r, A_s is the activation of each source node s, $S_{s,r}$ is strength of the link between nodes s and r, and $\Sigma S_{s,i}$ is the sum of the strengths of all links emanating from node s. The effect of the ratio $S_{s,r}/\Sigma S_{s,i}$ is to limit the total spread from a node s to all connected nodes such that it is equal to the node's current activation A_s. This feature gives the model the ability to simulate the *fan effects* (Anderson, 1974; Reder & Ross, 1983) we have discussed. For example, if a node had three connections emanating from it with link strengths of 1, 2, and 3, then the activation spread along those links would be, respectively, 1/6, 1/3 (i.e., 2/6), and 1/2 (i.e., 3/6) of the node's current activation level.

4 *Current Activation of a Node*

The base or resting level of activation of a node should be distinguished from the current activation value of a node. The current level of a node will be higher than its baseline whenever it receives stimulation from the environment, that is, when the concept is mentioned or perceived, or when the concept receives activation from other nodes. While baseline strength decays according to a power function (i.e., first quickly and then slowly), current activation decays rapidly and exponentially toward its base level. Let A represent the current level of activation and B represent the base level of activation. Then, the decrease in *current* activation will be:

$$\Delta A = -\rho(A - B) \tag{4}$$

such that, after each unit of time, the current activation will decrease for every node by the proportion ρ multiplied by that node's current distance from its base-level activation.

C The SAC Model of Word Recognition and the Word Frequency Mirror Effect

Researchers have found that differential experience with words has profound effects both in ease of reading (making lexical decisions, naming times) and in memory for the words. One of the conundrums of memory research is the problem of the word frequency mirror effect in recognition memory (Glanzer & Adams, 1985; Glanzer & Bowles, 1976; Gorman, 1961; Greene & Thapar, 1994; Hintzman, Caulton, & Curran, 1994; Hockley, 1994). Normative word frequency attempts to measure the extent of previous every day experience with each word (although the estimates are usually derived from books). The word frequency mirror effect is given its name because the pattern of hit rates is a mirror image of the pattern of false alarm rates: Low frequency words produce more hits and fewer false alarms than high-frequency words. In other words, people are more likely both to recognize a previously seen low-frequency word compared with a high-frequency word and to correctly reject a low-frequency foil compared to a high-frequency foil. This effect has been seen as counterintuitive because it provides a case in which familiarity with a concept produces poorer memory performance.

The SAC architecture posits a *dual-process* account of recognition, and the word frequency mirror effect follows naturally from the original SAC assumptions (Reder et al., 2000; Reder, Angstadt, Cary, Erickson, & Ayers, 2002).

The SAC representation of words studied in an experiment is shown in Fig. 1. By dual-process, we mean that when a subject is asked whether a test probe had been studied as part of a list of words presented earlier, the subject has two routes through which he/she may recognize the probe word. Recognition can occur because (a) the subject *recollects* having studied the word on the list, which means retrieving specific episodic details of the appropriate previous encounter, or (b) the test probe seems so *familiar* that the inference is drawn that the familiarity must be the result of a recent previous exposure. The dual-process theory of recognition is becoming increasingly accepted among memory researchers (Jacoby, 1991; Joordens & Hockley, 2000; Mandler, 1980; Reder et al., 2000; Yonelinas, 1994), but what sets the SAC dual-process theory apart from the others is that it is computationally implemented (see Diana, Reder, Arndt, & Park, 2006 for a review).[7]

The Remember/Know paradigm is often used as an assessment of recollection and familiarity-based processes (Tulving, 1985). In this paradigm, participants are asked to make a Remember response when they recognize an item and can recall some detail about the context in which they studied the item. Know responses are made when the participant feels the item is familiar, but is unable to recall any details about the context in which he/she studied the item. Remember responses index the recollection process and Know responses index the familiarity process. We have used the terms know and familiar interchangeably for the same judgment.

Figure 2 illustrates how the role of normative word frequency affects recognition memory, especially Remember versus Know judgments. Using the assumptions described above, SAC can predict the percentage of recollection-based and familiarity-based responses that will be produced under the various conditions of a recognition task. These predicted response percentages are based on the current activation values of memory traces within the model. The percentage of recollection and familiarity responses can be combined to predict old/new responses.

When real words are used in an experiment, SAC assumes that the concept nodes already exist in memory and their base-level activation is determined by their history of previous exposure (frequency and recency of exposure). In order to approximate a given word's base-level activation value, we use its word frequency value in standard norms (Kucera & Francis, 1967).[8]

[7] An important part of the debate between single- and dual-process models is the value and diagnosticity of the phenomenological judgments of recollection. In our view, the cumulative evidence is too compelling to reject the dual-process account (see Diana et al., 2006 for a further discussion of this point).

[8] We raised that word frequency value to the power 0.7 for base level activation and 0.4 for the amount of preexperimental fan. We have used those values in all experiments in which we modeled effects of normative word frequency.

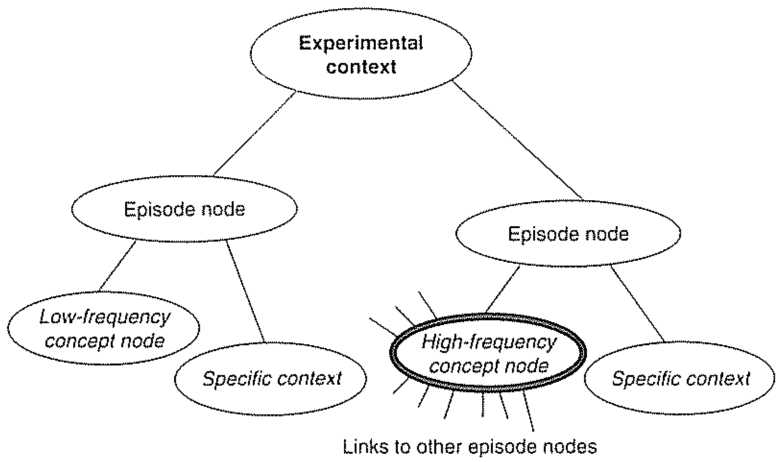

Fig 2 SAC's representation of high- and low-frequency words studied in an experiment

At study, we assume that the to-be-remembered word is activated and linked to the context in which it occurred. This context can include those characteristics of the environment that the subject experiences during the experiment, such as the lighting, equipment in the room, and the participant's mood during the task. Features that are general to the entire experiment are bound together as a general experimental context node. A specific context node also may be created during a study trial to capture a novel element of context that differs from the general experimental context. This might include the presentation of a word in a unique font, a sound occurring outside the room, or the participant's response to the stimulus. These three types of information: the concept node, specific context node, and experimental context node, are bound together by an episode node, which represents the experience of studying the word in the experiment.

When a probe word is presented at test, its concept node is activated along with the experimental context node. The contextual features of the test probe will also be activated. If the word is presented in the same specific context that was linked to the episode node during study, the specific context will be a relevant source of activation that can spread to the episode node. The activation from the concept and context nodes may intersect at the same episode node (depending on whether the probe is a target item or a foil and whether the specific context is similar). Recollection responses are based on the activation of the episode node, where activation accrues due to spread from associated concept nodes, specific context nodes, and experimental context nodes. Familiarity responses are based on the activation of the concept node and sometimes spuriously from the specific context node.

Activation spreads from each node in the structure that is activated by the environment (including concept nodes, specific context nodes, and experimental context nodes) according to the number and relative strength of the links connected to the node. The more links there are emanating from a node, the less activation spreads along any one of the node's individual links. See Eq. (3) above or consult Reder et al. (2000) for more details.

The probability of a Remember response depends on the current activation of the episode node and the subject's individual threshold for giving a "Remember" response. We assume the same parameters for strengthening, decay, spread of activation, and so on, but we assume that each individual has his or her own threshold for giving a Remember and a Know response. The probability of a Know response is the probability of not responding Remember multiplied by the probability of the concept or specific context node's activation being above threshold.[9] It is important to note that the Remember and Know judgments are not assumed to be independent. The proportion of Remember responses affects Know responses, but not the converse because participants are instructed to respond Remember if any recollected information is available, even when the item is familiar.

We assume that when the node binding the episodic details to the conceptual information is not sufficiently strong to pass threshold, the subject will rely on the less accurate process of familiarity. The familiarity-based (*Know*) response is based on the activation of the concept node. Given that the entire history of experience influences the node's strength or activation value, this judgment is less accurate for episodic tasks that require context-specific judgments of familiarity.

SAC got its name, *Source of Activation Confusion*, because of the assumption that people are unable to distinguish between activation due to recent exposure and activation due to a buildup of prior exposures. This principle is central to the SAC explanation of the word frequency mirror effect. The strength of the word concept node is affected by whether the word has been recently seen and how often it has been seen previously. High-frequency words have higher concept node strength due to prior exposure, and thus high-frequency lures would be more likely to produce familiarity-based false alarms than low-frequency lures.

[9] The effect of the activation from the specific context node on the probability of making a Know response is important when the specific context can be varied between study and test (see Diana et al., 2004 for more details).

As described earlier, another principle of SAC is that activation spreads along links between nodes according to the number and relative strength of the links. Therefore, less activation spreads along any one link from a node that has a greater number of links. A high-frequency word has more pre-experimental contextual associations than a low-frequency word and thus can be expected to have more links emanating from its word concept node.[9] This makes it less likely that a sufficient amount of activation will spread from a high-frequency word concept node to its episode node than that sufficient activation will spread from a low-frequency word concept node to its episode node. Recollection-based responses are made when the activation of an episode node surpasses threshold. Therefore, SAC predicts more hits to low-frequency words than high-frequency words, but also predicts that this difference should be seen in the Remember responses (Fig. 2).

According to SAC, the familiarity of a word is affected by whether or not the word has recently been seen and how frequently it has been seen overall such that both normative word frequency and recent exposure affect a word's familiarity. Because familiarity can arise from multiple causes, an accurate recognition judgment is based on the retrieval of the study event node (i.e., a true recollection), while responses based on the word node (i.e., familiarity-based responses) are error prone. There are more false alarms for high-frequency words than low-frequency words because high-frequency words are more familiar (have a higher base-level activation), and hence are more likely to seem old when a response is made based on the word node.

The SAC model of the word frequency mirror effect was formally implemented in Reder et al. (2000). It was shown to successfully fit the empirical data. However, the predictions and data obtained by Reder et al. were inconsistent with the findings obtained by Gardiner and Java (1990). Similar to the Reder et al. (2000) finding, Gardiner and Java (1990) found that for the hit portion of the mirror effect, there were more Remember responses to low-frequency targets than high-frequency targets. This led the authors to conclude that retrieval is responsible for the mirror effect. SAC also predicts that there will be more Know responses to high frequency than low-frequency words, but Gardiner and Java found no evidence of this. In order to confirm their finding of a difference in Know responses, Reder et al. (2000) analyzed the results of five previous papers testing the word frequency mirror effect with Remember/Know judgments. They found that high-frequency words produced a significantly higher proportion of Know responses compared with low-frequency words, confirming the SAC prediction. Figure 3 shows the model fits to the empirical data for Remember and Know judgments as a function of the experimental and preexperimental frequency of the stimuli.

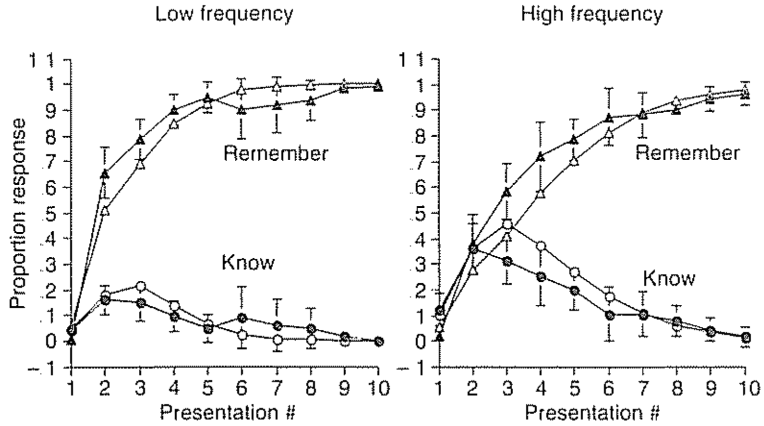

Fig. 3 The proportion of Remember and Know responses for words as a function of word frequency. Triangles represent Remember responses, circles represent Know responses. Closed symbols with solid lines represent the actual data. Open symbols represent the model predictions. The error bars represent 95% confidence intervals. From Reder et al., 2000, p. 310. Copyright 2000 by the American Psychological Association. Reprinted with permission of the author.

D. Converging Evidence for SAC Explanation Using Other Types of Stimuli

Recently we have tested our explanation of the effect of prior experience on retrieval in studies that manipulated exposure to perceptual (as opposed to conceptual) information. This involved presenting words in unusual fonts during study and then measuring word recognition as a function of whether the font at test matched the encoding font and as a function of the number of other words studied in that unusual font (Diana et al., 2004; Reder, Donavos et al., 2002). We represent the unusual font as an idiosyncratic contextual cue associated with the episode node for the studied word. If the word is tested in the same font used during encoding, then there is an extra source of activation that can spread to the episode node, and there should be a greater chance for a recollection (Fig. 4). However, if the font was used with many other words, then the fan of the font node will diminish the amount of activation that will get to any one of the associated episode nodes. As predicted, there were more hits and more Remember responses when the font matched and, most importantly, the advantage of the font matching was modulated by the fan of the font, such that the greater the font fan, the smaller the advantage of matching font.

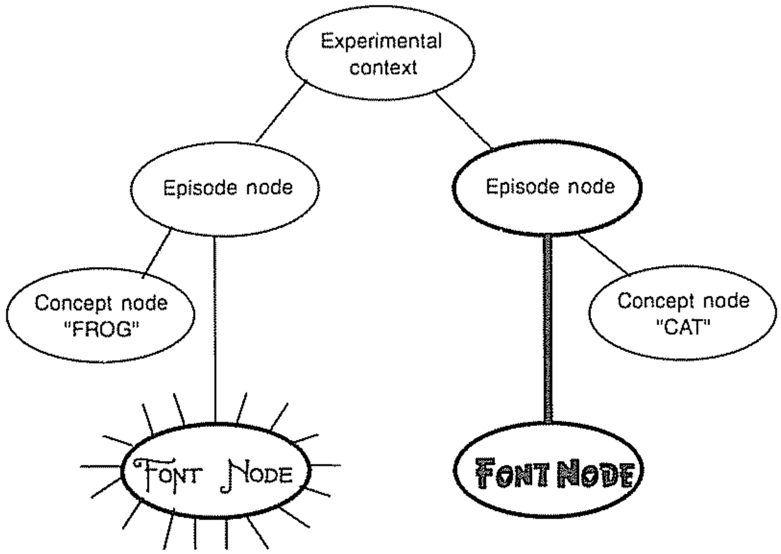

Fig 4 SAC's representation of high fan and low fan fonts reinstated at test

Further evidence for this explanation comes from a study by Park, Arndt, and Reder (2006). In order to test our hypothesis that these effects were driven by the fan of the contextual cue reinstated at test, subjects were asked to study a series of words presented individually on a screen in one of a number of unusual fonts while simultaneously hearing the word pronounced through a pair of headphones in one of a set of unfamiliar voices. A given word was presented in either a high fan font (seen with many words) or a low fan font (seen with only a few words). If the font was high fan, the voice would be low fan and vice versa. Assignment of voices and fonts to fan condition and to words was randomly determined for each subject. At test, when a probe was presented it was only presented in one modality, either font or voice (for both new and studied words). The context provided always matched the encoding features.

As predicted, recognition was more accurate when the feature that was reinstated was low fan. Not only do these findings provide additional evidence that the fan effects found for word frequency apply to perceptual information, but they also imply that these effects occur at retrieval rather than encoding. Subjects studied all words for the same amount of time, regardless of fan condition, and it was the fan of the reinstated feature that mattered at test.

Note that this explanation for more Remember responses with a low-frequency font is analogous to the explanation for more Remember hits for low-frequency words. Also as predicted, there were more false alarms to foils that were tested in high-frequency fonts than low-frequency fonts. In other words, we obtained a mirror effect for font frequency, just as one sees for word frequency. Since the assignment of fonts to be either high or low frequency (seen with one or many words) was randomly determined for each subject, the font frequency mirror effect does not suffer the interpretation problems of a quasi-experimental design that typically plague studies of the word frequency mirror effect.

Indeed, Maddox and Estes (1997) proposed that word frequency, per se, was not the real cause of the mirror effect. They manipulated exposure to artificial words (pseudowords) and found a concordant pattern of hits and false alarms such that high-frequency pseudowords produced more hits and false alarms. However, we suspected that their frequency manipulation was too weak, and that they were replicating a finding that rare words produce fewer hits (Schulman, 1976). Reder, Angstadt et al. (2002) exposed subjects to these pseudowords for an entire semester. Early in the training, they replicated the results of Maddox and Estes. However, by the end of training, they produced the standard mirror effect, including more Remember responses for low-frequency pseudowords. More recently, Nelson and Shiffrin (2006) have replicated our result of a mirror effect for differentially experienced stimuli, in this case Chinese characters.

In summary, given that differential exposure to fonts, pseudowords, or Chinese characters all produce the mirror effect and that the assignment of stimuli to frequency category was randomly determined for each subject, this effect must be due to the previous exposure to the stimuli and not something inherent in the stimuli, per se. This finding supports the claim that familiarity alone can be the source of a reduction in memory performance.

1. Converging Evidence Using Synthetic Amnesia

Although word frequency manipulations in tests of recognition memory almost always produce a mirror effect, there are situations where this regularity does not occur, such as in studies with amnesiacs or participants under the influence of midazolam. It is often proposed that patients with Alzheimer's disease and other forms of anterograde amnesia have damage to the recollection capability in memory, but that their familiarity capabilities remain largely intact (Balota & Ferraro, 1996). Hirshman, Fisher, Henthorn, Arndt, and Passannante (2002) induced temporary anterograde amnesia using the drug midazolam and showed that when participants were under the influence of midazolam, the hit rate portion of the mirror effect did not occur. A concordant pattern emerged such that there were more hits and false alarms to high-frequency words than low-frequency words. However, participants in the control condition, who received an injection of saline, did show the typical word frequency mirror effect. It is thought that midazolam affects people's ability to recollect information from study, but that it does not impair familiarity processes (Hirshman et al., 2002).

Dual-process models like SAC can explain these data: the hit rate portion of the mirror effect is due to a recollection process which is disturbed by the drug (or organic amnesia), but the false-alarm portion results from a familiarity process that is not affected by the drug. According to SAC, high-frequency words have a higher base-level familiarity that results in more hits (and false alarms) when retrieval of contextual associations cannot be used.

2. *Source Memory Studies Provide Further Support*

Evidence from the source memory literature further supports the SAC account. Low-frequency words are more likely to be associated with correct source judgments than high-frequency words (Guttentag & Carroll, 1997; Rugg, Cox, Doyle, & Wells, 1995). Source judgments ask participants to report a contextual detail from the study phase that was varied systematically when they recognize a test word. This type of task is thought to use recollection-based processing (Quamme, Frederick, Kroll, Yonelinas, & Dobbins, 2002). The research found that low-frequency words were more likely to be correctly judged old and to be assigned to the correct study context than were high-frequency words. This indicates that participants could more easily recollect the specific context for low-frequency items and thus were more able to use recollection processing for low-frequency words. This is consistent with a dual-process account claiming that the increased hit rate for low-frequency words is based on better recollection. These findings provide supporting evidence for our model that the hit portion of the mirror effect is driven by recollection-based responses while the false-alarm portion is driven by familiarity-based responses.

3. *The Costs of Lifelong Experience on Retrieval*

An interesting implication of the theory we have presented to explain the word frequency mirror effect and other phenomena is that the base-level activation and contextual fan of words should continue to increase over a person's lifetime because the words continue to be experienced. We propose that some of the memory deficits associated with advancing age can be explained with these same assumptions (Buchler & Reder, 2007). Although there has been a great deal of research done on the biological and physiological bases of age-related memory problems, there has been surprisingly little attention devoted to the potential effects of experience itself.

SAC predicts that familiarity processes should be relatively unaffected in that familiarity is enhanced with continued experience (base-level activation goes up). However, the fan out of each word also accumulates with age making the recollection process more difficult. Many studies support our position that age-related deficits are found in the recollection-based component rather than the familiarity component (Balota, Burgess, Cortese, & Adams, 2002; Burke & Light, 1981; Castel & Craik, 2003; Chalfonte & Johnson, 1996; Kliegl & Lindenberger, 1993; Light, Healy, Patterson, & Chung, 2005; Naveh-Benjamin, 2000; Simons, Dodson, Bell, & Schacter, 2004; Spencer & Raz, 1995). Buchler and Reder (2007) used a two-parameter model of aging to successfully account for a number of previous results that compared young and old memory performance. The older adults were assumed to differ from the younger on only two parameters, one representing the extra increase in baseline activation and another representing the increased fan. The fit to the published data was quite good (generally with an r^2 of .98 or better using only these two parameters, and sometimes only one, to fit the data).

Despite the excellent fits to five different published data sets, we recognize that other factors besides these two parameters affect differences in performance between young and older adults. We will discuss those in the second part of this chapter. For one thing, there is evidence that older adults use different cognitive strategies, presumably to try to compensate for whatever detrimental effects do arise from aging. Reder et al. (1986) explored whether the tendency to use "direct retrieval" as opposed to a plausibility strategy differed with age. Some subjects of both age groups (young versus old) were explicitly asked to judge whether a sentence was consistent with what they had read before while the two other groups were explicitly asked to determine whether a specific sentence had been read earlier (direct retrieval). Although older subjects were slower to respond in all cases, they were actually better than their younger counterparts at the plausibility task in terms of accuracy. However, as predicted, they were much worse when direct retrieval was required.

E. Summary of How Experience Hurts Retrieval

In this section, we have reviewed a number of experiments that report that knowing more about a concept hurts one's ability to retrieve specific information associated with that concept. We have used the explanation of the "fan effect" to account for various aspects of the word frequency mirror effect, as well as reviewing the larger literature on the fan effect that shows accuracy and latency are adversely affected by knowing more about a concept. We showed that this effect is not limited to experimental material generated in the laboratory, but applies to prior knowledge about famous individuals. We also showed that our computational model could account for effects of fan on perceptual information such as font during encoding and showed that it is the fan of the contextual features reinstated at test that matters, rather than the fan of the features used during encoding.

We also explained how it is that people avoid the "paradox of the expert" by using strategies other than "direct retrieval." Not only can individuals be manipulated to use direct retrieval or plausibility as the preferred strategy by manipulating prior history of success, or cues in the question (Reder, 1987, 1988; Reder & Ritter, 1992; Reder & Schunn, 1996), people's appreciation of their general ability to use retrieval, as a function of age, also influences tendencies to use one question-answering strategy or another (Reder et al., 1986).

Despite all the evidence showing how detrimental prior experience can be to the retrieval process, there is also evidence that prior experience can be a benefit during encoding. The rest of this chapter is devoted to presenting the evidence for this point of view and the additions to SAC to explain these effects.

III. When and Why Experience Facilitates Memory Encoding

It is generally accepted that novel stimuli attract attention (Johnston, Hawley, Plewe, Elliott, & DeWitt, 1990; Sokolov, 1963) even for infants (Fagan, 1970). That observation has been used by some theorists to explain the word frequency mirror effect (Glanzer & Adams, 1990). Rao and Proctor (1984) demonstrated that when encoding is self-paced, participants study low-frequency words longer than high-frequency words. Conceivably, the longer study times for low-frequency words arises as a result of people preferring novel stimuli and therefore allocating more attention to them. This leads to better recollection for low-frequency words. In the previous section, we offered a different explanation for the word frequency mirror effect; we think that the longer study time for low-frequency words results from the fact that less familiar stimuli are actually *more difficult to encode* and, as a result, *require more attention* in order to be processed.

The arguments put forward in the first half of this chapter concerned the adverse effects of experience, when attempting to retrieve associations to frequently experienced concepts. Now we want to examine the other side of the coin and argue that frequently experienced concepts are actually easier to encode. This encoding advantage occurs despite the novelty bias in attention, which we speculate may occur in part as a compensation for the encoding disadvantage. In this section, we will review some of the evidence that has led us to this conclusion and describe our modifications to SAC in order to account for the encoding advantage. We also provide model fits to a number of the phenomena that we intend to explain with the revised model.

Some aspects of an apparent encoding advantage, such as faster naming times and faster reading times for high-frequency words are consistent with the SAC assumption that high-frequency words have a higher base level of activation and are therefore more accessible. What was missing from SAC was the assumption that there is a finite pool of WM resources and that the ability to encode a stimulus depends on both the familiarity of the stimulus and the amount of WM resources available. Before providing the details of the change in the SAC architecture, we will review some of the findings that motivated the modifications to the model.

In an unpublished paper, Spehn and Reder (2000) (available on the web at http://www.memory.psy.cmu.edu/unpublished/SpehnLMR.pdf) found that subjects were better at learning novel first names to famous names such as Einstein or Travolta than to unfamiliar last names such as Kounkel. When tested on their memory for just the last names of the studied first–last name pairs, famous last names were recognized best, rare names intermediate, and common names such as Smith were worst. In contrast, when the recognition test required judging whether the first name was studied with the last name, common last names did exceptionally well.

In our view, this result is analogous to the finding that although high-frequency words are not well recognized, they do better in word-pair recognition than low-frequency word pairs (Clark, 1992). Like high-frequency words, common names have greater fan (many first names already associated with them), so it is harder to retrieve the pairing if only given the last name as the test probe. That is one reason why common last names were recognized worst when tested in isolation. The other reason is that basing the recognition judgment on familiarity (when retrieval of the first name failed) will be error prone just as it is for high-frequency words.

On the other hand, if the task is name-pair recognition the first name is provided at test as well. In that case, there are two sources of activation to send to the episode node that binds the names together. With two sources of activation, the effects of fan should be reduced, enabling the encoding advantage of common names to be observed. In other words, we believe that it is easier to link an arbitrary first name in memory to a common name than to a rare name like Kounkel or Nhouyvanisvong because those names are quite unfamiliar and take up considerable resources just to encode those names.[10]

Another study conducted by Diana and Reder (2006) supports the role of familiarity at encoding. Subjects were presented with high- or low-frequency words that were superimposed on pictures of common objects and instructed to try to remember both the pictures and the words. Assignment of words to pictures was randomized for each subject. For example, a picture of a basketball might have a high-frequency word (e.g., tree) or a low-frequency word (e.g., aspirin) superimposed on it. At test, pictures were presented without any words and subjects were asked to recognize the studied pictures. Recognition memory for the pictures was better when the superimposed word at study was high frequency rather than low frequency. Not only was recognition accuracy better when the picture was studied with a high-frequency word, but the proportion of "Remember" judgments was greater when the encoding word was of high frequency. This latter point is important because the binding operation that we believe requires WM is manifest in Remember responses. "Familiar" (or "Know") responses do not depend on this binding process because they reflect only the activation of the concept node.

Although picture memory was better when high-frequency words were superimposed, recognition memory for the words themselves (tested separately from the pictures) showed the typical pattern whereby low-frequency words were recognized better than high-frequency words. In our view, recognition is better for low-frequency words despite their encoding disadvantage because the retrieval advantage masks the encoding disadvantage unless there are increased WM demands at encoding. Another study by Diana and Reder (2006) found that when two words are presented for study simultaneously, both high- and low-frequency words are more easily recollected later if the word it was paired with was of high frequency. That is, pairing a word with a low-frequency word at study makes recollection more difficult.

An alternative explanation for the picture encoding advantage with high-frequency words is that there is a tacit trade-off in attention between the word and the picture such that low-frequency words grab more of the attention than high-frequency words and the total amount of attention is limited. That is, more novel words attract more attention leaving less for the pictures. High-frequency words are less unusual and therefore more attention is allocated to the picture, increasing its chances of being recognized later. This alternative account cannot explain the findings with lists of pure high frequency or pure low-frequency words in paired associate recognition or recall. In those cases, high-frequency words are at an advantage (Clark, 1992;

[10] Ngiam also modeled some of these data successfully. For reasons of space and time considerations (he did not have time to model all of the results), we are not reporting those efforts.

Deese, 1960). We will describe these patterns in more detail when we fit SAC to the empirical results.

Participants may tacitly appreciate this trade-off between encoding and retrieval for word frequency. When queried before the experiment begins "how difficult will each item be to recognize", they predict that high-frequency words will be easier to recognize. However, when asked the same question during the test phase, participants make the correct judgment, noticing that low-frequency words are easier to recognize (Benjamin, 2003). This suggests that participants may experience high- and low-frequency words differently during encoding as well as supporting the idea that low-frequency words are more likely to produce a recollection-based response, which would lead participants to feel that such words are particularly memorable.

Recognition memory tests also show list composition effects whereby the low-frequency word advantage is augmented in mixed lists of predominantly high-frequency words (Dewhurst, Hitch, & Barry, 1998; Malmberg & Murnane, 2002). Also, rare words (e.g., "iatrogenic") do not show the normal hit rate advantage in standard recognition memory experiments that low-frequency words enjoy (Schulman, 1976). This may be because the rare words are so difficult to parse or comprehend that it becomes difficult to form any associative link to them whatsoever. Thus, the postulation of a low-frequency encoding disadvantage can explain a range of phenomena in the literature on memory for words.

High-frequency words also show an advantage in associative recognition tasks. Associative recognition requires the formation of associations between items. In these tasks, participants study pairs of words and at test are asked to discriminate between words that were presented as pairs at study (that should be judged as old) and those that are recombinations of studied items from different pairs (that should be judged as new). Unlike item recognition, associative recognition shows a mirror effect for high-frequency words: previously seen high-frequency word pairs produce more hits while high-frequency recombined pairs produce fewer false alarms than low-frequency pairs (Clark, 1992). These findings from associative recognition and recall provide evidence that the formation of associative links between items in memory, such as between arbitrary word pairs presented in associative tasks or from word to word in serial recall tasks, may be easier for high-frequency words than low-frequency words.

A. Augmentation of SAC: How WM and Prior Experience Interact to Affect Ease of Encoding

We have previously implemented SAC models that vary the probability of encoding an event to explain aging effects (Reder et al., 2007a) and to simulate the effects of midazolam (Reder et al., 2007b). We accomplished these effects by merely positing different probabilities of forming a link. Although those modifications worked well, they were ad hoc. The addition of a WM component to the SAC architecture enables the probability of encoding to vary in a more principled fashion (i.e., without merely fitting a parameter that varies the success of the binding).

We assume that there is a finite amount of WM resources that can be used to encode stimuli, build associations, perform tasks, and so on and that this pool returns to its full capacity over the time. Resources are drawn from this pool of WM to activate a stimulus so that it can be encoded in a way that enables the construction of a link between two elements. For example, this could be the binding of a word to an experimental context or forming an association between two words. Importantly, how much activation must be drawn from the pool of WM resources *depends on the resting level of activation of the concept* such that the weaker the base-level activation of the concept, the more activation that is required to build a new association. As such, familiar concepts (e.g., words with higher normative frequency) make fewer demands on the WM pool when attempting to bind an item to context or to another concept. This implies that the more elements that need to be encoded and processed, the greater the demand on this pool of WM (Anderson, Reder, & Lebiere, 1996). The amount of WM expended in encoding one concept is:

$$\mathrm{WM}_{\mathrm{encode}} = \tau - B \tag{5}$$

where τ is the threshold and B is the node's base-level activation [Eq. (1)]. The WM pool replenishes at a linear rate, r, such that the pool at time t is given by:

$$\mathrm{WM}_t = \min(\mathrm{WM}_{\max}, \mathrm{WM}_{t-1} + r) \tag{6}$$

Thus, the WM extensions to SAC involve 2 new parameters: the maximum WM pool quantity, $WM_{\max}$, and the WM recovery rate, r.

We also assume that if there is sufficient WM to get a concept over threshold, the amount of activation that is sent from a source node is unaffected by the base-level activation, although it remains proportional to the relative link strength. Familiarity judgments are now a function of the amount of WM resources required to get the word up to threshold (much like "perceptual fluency," see Whittlesea, Jacoby, & Girard, 1990) such that the fewer WM resources needed to reach threshold, the more perceptually fluent and the more familiar the concept appears.[11]

[11] These different assumptions do not change the behavioral predictions of the model for the datasets already fit. Familiarity judgment calculations are isomorphic. The spread of activation values are almost the same as well.

These assumptions mean that a person is less likely to be able to bind a concept to a context if (a) the concept is unfamiliar, (b) there are many other stimuli to encode at the same time, (c) the stimulus is perceptually degraded, or (d) the WM pool is small, either because it has not finished being replenished or because the person has a smaller pool to begin with. We assume that the amount of WM varies among individuals (Daily, Lovett, & Reder, 2001; Lovett, Daily, & Reder, 2000; Lovett, Reder, & Lebiere, 1997), as well for a particular individual as a function of fatigue, and so on.[12]

It is important to note that these assumptions concerning encoding also apply at test when the probe(s) need to be encoded. When there are more stimuli as part of the test probe that need to be encoded (word pair vs a single item) or when the stimuli are less familiar (low-frequency words, words presented in unusual fonts), more WM resources are depleted in the effort to get each concept of the test probe up to threshold. If there are sufficient resources to get a concept up to threshold, then activation can spread to its associated nodes.

1. *A Limit on Concept Strengthening*

We have also added the assumption that a node is not strengthened when its current activation is above a specific level. This assumption could be viewed as a proxy for habituation such that when the same information is experienced over and over it no longer attracts as much attention and does not gain strength indefinitely; however, we are not claiming that the links are not formed or strengthened when the item is repeated at threshold; therefore, it should not be taken as a complete analogue to habituation.

2. *Partial Match and Spurious Recollection*

In order to model false alarms that are reported as "recollections", a spurious recollection mechanism has been introduced to SAC. Previously, SAC only accounted for false alarms as familiarity-based "Know" false alarms and did not allow any "Remember" false alarms by spuriously activating the wrong episode node. That simplifying assumption seems odd in hindsight because the original SAC model of feeling of knowing (Reder & Schunn, 1996; Schunn et al., 1997) accounted for spurious feelings of knowing that were generated from partial matching. Specifically, we modeled that a spurious

[12] Reder's previous work on individual differences in working memory capacity used the ACT-R framework. In ACT-R, working memory differences are assumed to only affect retrieval, not encoding. There are currently no assumptions about differential probability of encoding or binding in ACT-R.

feeling of knowing would occur if sufficient activation accumulated at the problem node even if an element of the problem (such as the operator) did not match. We now appreciate that the same assumptions should have remained in SAC when we modeled recognition.

We now allow for an analogous mechanism in recognition to occur by letting the model attempt to retrieve the episode node with the highest activation regardless of whether or not that episode node corresponds to the concept in the probe. If a spurious episode node is retrieved, the participant may still be able to recall the original concept that the episode had been linked to and reject it on that basis (recall to reject). For more information on spurious recollection, see Cook, Reder, Buchler, Hashemi, and Dickison (in preparation).

B. Illustrations of Model Fits with the New Encoding Assumptions

Earlier in this chapter we described a study by Diana and Reder (2006) in which words were superimposed on pictures and subjects were responsible for remembering both aspects of the stimulus. In this model, pictures and words are represented by concept nodes, with an attempt to link each concept node to an episode node at study. The concept nodes for the pictures were given base activation levels approximated from medium-frequency words. During each study trial, consisting of a superimposed word and a picture, two links needed to be formed from the picture concept node and the word concept node to their respective episode nodes. This link is only formed when sufficient resources exist in the WM pool. Therefore, when a low-frequency word is presented with a picture, fewer resources remain to allow encoding of the picture than when a high-frequency word is presented with a picture. Model fits to this experiment, comparing the SAC predictions to the actual data, were quite good, with Pearson's $r^2 = .95$. These fits are shown in Fig. 5.

As described earlier, Diana and Reder (2006) found that low-frequency words were better recognized if they were encoded with a high-frequency word while high-frequency words were recognized worse if encoded with a low-frequency word even though participants were instructed to remember each word separately and were not tested on their memory for which words were paired together. The results and model fit are shown in Fig. 6. Here too the fit was quite good, $r^2 = .96$. As in other models, words are represented by concept nodes, with each concept node linked to its own episode node. Because participants were instructed to remember each word separately, we did not include a link between the episode nodes for words studied at the same time.

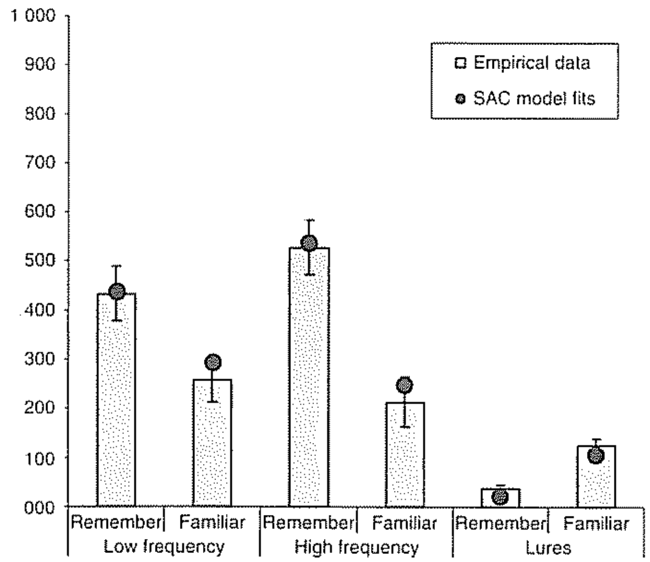

Fig. 5 SAC model fits for Remember-Familiar responses during the word test in the picture-word interference experiment

In each study trial consisting of two words (to be encoded separately), recollection requires that the word concept node be bound to the experimental context node by creating an episode node that links them. The formation of each episode node requires resources to be drawn from the pool of WM resources. Because high-frequency words have a greater base-level activation, fewer WM resources are required to create an episode node linking the concept and context nodes, while low-frequency words require relatively more WM resources in order to form an episode node. In the event of a link formation failure, the concept node will not be linked to the episode node at all. This reflects a failure in binding and the item cannot be retrieved using recollection. In the case of a link formation failure, no resources are subtracted from the WM pool.

Word frequency manipulations produce different effects depending on the composition of the study lists. When items are encoded on lists of either purely high-frequency or purely low-frequency words, high-frequency items produce better performance on cued recall and associative recognition tests (Clark & Burchett, 1994). On lists with both high- and low-frequency words, the high-frequency advantage in cued recall does not occur. Simple recall also shows a

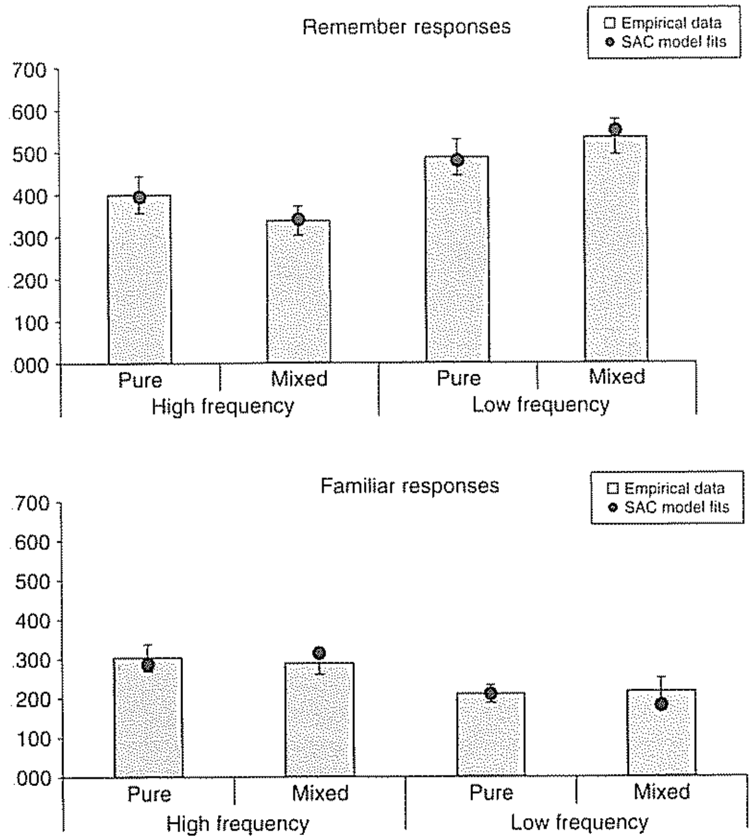

Fig 6. SAC model fits for Remember-Familiar responses during the word-word interference experiment

high-frequency advantage only for pure lists (MacLeod & Kampe, 1996; Watkins, LeCompte, & Kim, 2000). Even in recognition, the ubiquitous low-frequency advantage is affected by list composition. There is some evidence that high-frequency words show an advantage when items are presented on pure lists (Dewhurst et al., 1998). Also, when the proportion of high-frequency words on a list is increased, the low-frequency advantage increases (Malmberg & Murnane, 2002).

If low-frequency words in fact use more WM capacity during encoding, the presence of more low-frequency words on a list may reduce the processing resources that are available to encode all words on this list. This is because

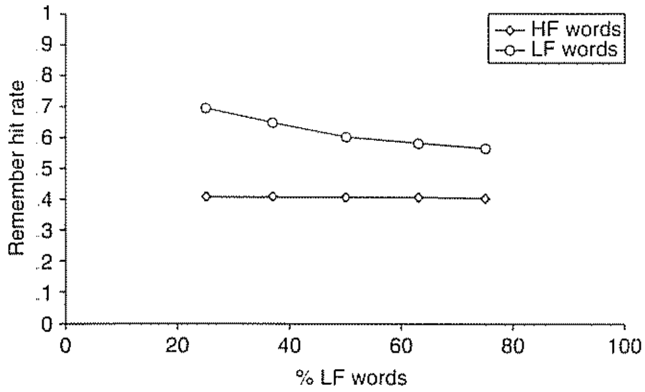

Fig. 7 SAC simulation of Remember hit rate to high- and low-frequency words across different proportions of low-frequency words included in the study list.

the low-frequency words may recruit WM capacity from high- or low-frequency words presented on subsequent trials. That is, encoding of a previous low-frequency word may still be occurring during later study trials. In this case, we would expect better encoding of low-frequency words on a randomized list that contained fewer low-frequency words and better encoding of high-frequency words on a randomized list that contained only high-frequency words. To test whether our explanation of this pattern could actually be simulated, we developed a SAC simulation of learning a study list that varied in the proportion of low- and high-frequency words and tested its ability to retrieve the episode node. Figure 7 shows the results of that simulation. Note that this pattern is consistent with the findings of Malmberg and Murnane: As the proportion of low-frequency words on the list increases, there is a reduction in the proportion of low-frequency Remember hits while high-frequency word Remember hits were largely unaffected by this manipulation.

C. The Consequences of Minimal "Lifelong" Experience on Encoding

In the previous section, we discussed how and why experience hurts the elderly when it comes to using prior knowledge in a fact retrieval situation. The other side of this coin is the demonstration that young children are less able to encode information because of their limited experience with the stimuli. Whitehouse, Maybery, and Durkin (2006) found that the picture superiority effect (over words) in free-recall tests increases from middle childhood to adolescence.

Given that word reading does not decline with age, and pictures should be more important for younger children, the explanation cannot be due to simple identification of the stimuli. Indeed word recall did not improve from grades 2/3 to 10/11, while picture recall improved substantially. The interpretation of Whitehouse et al. is that the picture superiority effect is "contingent on the encoding of pictorial information through two different routes." While many would have predicted that the picture superiority effect would decrease from elementary school to secondary school, Whitehouse et al. speculate that the converse finding results from the development of inner speech with age, and that inner speech allows for the dual-code advantage postulated by Paivio (1971).

Our interpretation is similar but is based on lower familiarity of concepts for young children. Concepts that have a lower level of activation are more difficult to bind to an episode, making recall more difficult. The picture task uses more WM resources because the picture has to be translated into a word to get the second code. Within SAC, the recovery rate of the WM pool takes time and is affected by the amount of depletion. We would argue that in grades 2/3, fewer of the pictures benefit from the secondary code, but as each of the concepts gets stronger, the number of concepts that can be bound to the episode node increases. In other words, it is the tacit secondary task of converting pictures to words that creates the dual codes but also taxes WM, meaning that more concepts fail to be bound.

D. Extremely Low-Frequency Stimuli: Experience Enables Unitization (Chunking)

Although low-frequency words typically show an advantage in tests of recognition memory over high-frequency words, this effect is reversed when rare words (e.g., "iatrogenic") are used (Schulman, 1976). We believe this is because the rare words are so unusual that they are not *chunks*. Stimuli that are not chunks have a weak node binding the components together and WM resources are used to bind together the constituents of the rare stimulus rather than binding it to the experimental context.

Another study from our lab (Reder et al., 2006a) provided additional support for the notion that unfamiliar stimuli are difficult to encode and therefore bind to context, despite their unusual status. Subjects studied words, photographs, and abstract pictures for a subsequent recognition test on the same day. Each subject participated in two sessions with two separate lists of stimuli. In one session they received an injection of the drug midazolam, a benzodiazepine that creates temporary anterograde amnesia, before studying the list of items that they would then have to recognize. In the other session they studied different items from the same stimulus classes, but after an injection of saline. Neither the

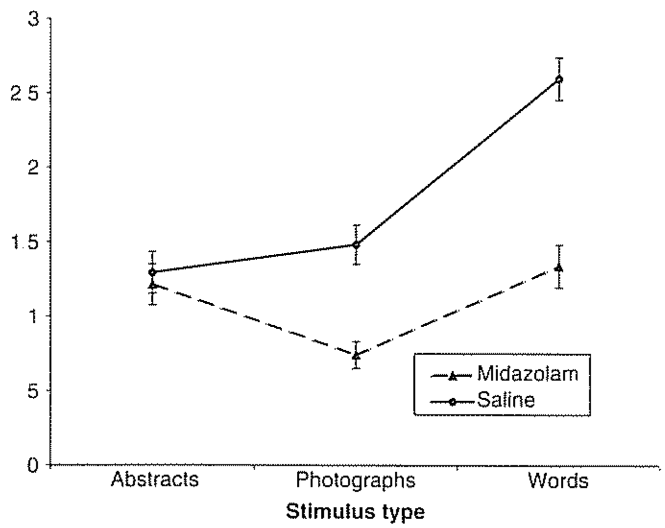

Fig. 8 Recognition memory measured in *d'* as a function of stimulus type (words, photographs, and abstract pictures) and drug condition (saline vs midazolam). From Reder et al., 2006a, p. 565. Copyright 2006 by the Association for Psychological Science. Reprinted with permission of the author

participant, nor the nurse, nor the experimenter knew which day a particular subject was given saline or midazolam, (i.e., testing conditions were double blind). The striking result was that midazolam affected recognition memory for words most and affected memory for abstract pictures least (Fig. 8).

Our explanation for this result is that (a) midazolam only affects the ability to create new bindings (Park, Quinlan, Thornton, & Reder, 2004; Reder et al., 2006b) and (b) only a *unitized chunk* can be bound to an experimental context. The abstract pictures could not be bound to the experimental context even in the saline condition, and therefore the effect of the drug was minimized for that stimulus class. Another finding by Dobbins and Kroll (2005) can be interpreted as supporting our hypothesis. They found that recognition memory was superior for scenes and faces that were known, but that the advantage for those stimulus types was eliminated when subjects were forced to respond quickly or when testing was delayed for one week. Our interpretation is that binding concepts to experimental context is much more likely for known faces and scenes; however, if responding must be rapid, judgments are based on familiarity and so there is no advantage to having formed an episode node. With a one week delay the episode node and link will have decayed substantially making reliance on familiarity the dominant process.

The notion that unitization requires prior experience is not a new idea Hayes-Roth (1977) and Servan-Schreiber (1991) have hypothesized something similar; however, no one has thus far suggested that the strength of a chunk predicts the probability of encoding it and binding it to other chunks. Our explanation is that an item with no prior representation must be encoded in terms of the component features that are strongly activated. With repeated exposure, the node that binds the constituents together becomes a chunk in its own right, forming a new, higher-level chunk involving the grouping of these features At that point the higher-level chunk is sufficiently strong (i.e., has strong enough base-level activation) to be bound with other co-occuring stimuli or bound to the experimental context to make an episodic event The abstract pictures had not been experienced before and recognition could only be based on the familiarity of the elements that were primed from exposure.[13]

Further support for the notion that chunks are constructed as their constituent elements become more familiar comes from studies with chess masters (Chase & Simon, 1973; Simon & Gilmartin, 1973) who have acquired thousands of hours of experience with various chess patterns. Although chess masters are much better than novices at reproducing a chessboard configuration when it was displayed tachistoscopically (very briefly), they are not better than a novice if the configuration of chess pieces on the board is random (de Groot, 1965). In addition, the latency between chess pieces that were put down on the board to reproduce the flashed display mirrored the chunks that one would expect. That is, subjects had shorter pauses when putting down pieces within a chunk (e.g., a Sicilian defense), but longer pauses when switching to recall of another chunk

IV. General Discussion

Sometimes psychologists will say with a wry smile, "Psychology is the science penetrating the obvious." Whether or not that adage is valid, it seems obvious (with hindsight) that experience should facilitate encoding. However, it has also been demonstrated that novel stimuli attract far more attention, and it has often been claimed that the disadvantage of high-frequency words in recognition results from poorer encoding. In this chapter, we have argued that high-frequency words are encoded more easily than low-frequency words, but that their deficit in recognition occurs despite their encoding advantage.

[13] There is also the possibility of recollection from a subset of the features, that is binding some of the features to context. The danger with that strategy is that the features that are strong enough to bind to context could also be shared with foil pictures

The important contribution of this chapter is not articulating what some might consider the obvious (at least in hindsight), but rather articulating a mechanistic account of when and why familiarity helps encoding. That is, the familiarity advantage at encoding matters more when there is a demand on WM resources. We also offered an explanation of how and why familiarity enables the binding of context to concepts. Finally, we reviewed the evidence that knowing more about a concept means that retrieving any one fact about it is slower or less accurate. This seems obvious when reframed as "it is harder to find a specific strand of hay in a haystack than on a clean floor." If the details of the retrieved information are unimportant, then the effect of fan goes away or even reverses.

This chapter went beyond verbal explanations to account for classes of phenomena. We offered a computationally implemented model that accounts for both the costs of experience at retrieval and the benefits of experience at encoding within the same framework. We went beyond demonstrations of qualitative fits to the empirical data and provided excellent quantitative fits that involved estimating few new free parameters (i.e., most parameter values have remained the same across all SAC models). We did not attempt to fit all the data we reported that provides converging evidence for our point of view, but we are confident that these phenomena could also be modeled within our framework. We have also fit some phenomena that we did not describe such as differential effects of word frequency as a function of the presentation rate.

A. Explaining Related Phenomena with Our Model

All of the phenomena that we have modeled have either involved simple numerical problems or words or word pairs and perceptual contexts (e.g., font or voice). These domains have the property that individual differences in semantic memory are not too relevant to performance (unless one gets into free-recall tasks) and we do not need to model language parsing. In order to model phenomena that involve the semantics of the stimuli we would need to speculate on the semantic content of people's memories, a complex task that we do not feel equipped to undertake. Nonetheless, a number of ideas described here apply to other phenomena that have not been modeled in SAC but seem consistent with the architectural principles.

For example, we reviewed the findings that new information about famous people can produce fan effects (interference) with real-world knowledge about them when the task requires retrieval of specific facts rather than consistency judgments about these people. The explanation that chess masters have acquired higher-level chunks from the experience of building up constituent (smaller) chunks with experience is something that is predicted by the model. A prediction of our model is that if chess masters were presented with chess configurations at the same time as orally presented words, then recognition for the chess patterns presented would show a "mirror effect" such that the very common patterns would have fewer hits and more false alarms than the somewhat less common chess configurations; however, we would also predict that the words presented with the common chess patterns would produce more "Remember" responses than those words studied with the less common (lower frequency) chess patterns, analogous to what we have seen with words and pictures (except that here chess patterns are mapped to the words in terms of our predictions).

Our explanation of why high-frequency words are easier to encode involves the assumption that they have a higher resting level of activation, which we have also used to explain the misattribution of activation that creates spurious familiarity judgments This assumption follows from the architectural principles of strengthening chunks with repeated exposures and also explains a number of other phenomena associated with words of different frequency. For example, word naming tasks, used primarily in the study of semantic memory, show a high-frequency advantage such that high-frequency words produce faster responses than low-frequency words (Frost & Katz, 1989). Also consistent with our framework, when a secondary task is added to the word naming task, effects of secondary task difficulty are larger for low-frequency words than high-frequency words (Becker, 1976; Goldinger, Azuma, Abramson, & Jain, 1997). When longer delays between word presentation and response are used, the high-frequency advantage disappears (Becker, 1976; Connine, Mullenix, Shernoff, & Yelen, 1990). Seidenberg (1985) argued that higher frequency words are more visually familiar and this visual familiarity allows lexical access without generation of phonology With regard to the current question of the effects of frequency at encoding, this idea could be simplified to the view that access to memory representations of high-frequency words is faster than access to representations of low-frequency words.

V. Summary and Conclusions

In this chapter, we have proposed that experience can facilitate cognition, but that it also carries costs. We have provided both empirical evidence to support these claims and a computational mechanism to show how these processes interact with other aspects of the mind. Our cognitive architecture also has neurophysiological support for its assumptions. For example, there is evidence that repetition priming produces a reduction in the BOLD response (see Henson, 2003 for a review), consistent with the idea that a node with a stronger base-level activation (from recent boosts in activation) requires less processing to get to threshold. Likewise, there is evidence that high-frequency words produce a reduced signal both in fMRI (de Zubicaray, McMahon, Eastburn, Finnigan, & Humphreys, 2005) and EEG (Hauk & Pulvermuller, 2004) compared with low-frequency words, which is also consistent with our assumptions. Likewise, there is neuroimaging evidence that increased fan creates a greater BOLD response, which supports the view that it is more difficult to retrieve something for which there are more associations (D'Arcy, Ryner, Richter, Service, & Connolly, 2004).

The first half of this chapter reviewed the evidence for the important role of experience at retrieval. We argued that greater experience makes retrieval of specific facts more difficult, but that it facilitates judgments based on inference (familiarity based, consistency based, and so on). As we age, we have more wisdom, and more knowledge and more experience, so it is natural that we rely more on this experience and make more inferential judgments. The second half of this chapter extended our implemented mechanistic account of implicit and explicit memory effects that can account for the mirror effect of word frequency among many other phenomena. In the augmentation of SAC, we provided insights as to how familiarity can provide an advantage in cognitive processing by facilitating encoding. The value of a computational model such as SAC is that it can be integrated to explain many phenomena with the same set of assumptions. As Herb Simon said, "If the goal of psychology is to prove a theory wrong, we can all go home now because all theories are wrong" (personal communication, 2000). Yet Herb Simon was one of the strongest advocates for developing computational models and frameworks or architectures. The goal is to move toward closer and closer approximations to the truth by building models that can account for more and more phenomena.

Acknowledgments

Preparation of this chapter and the work described herein was supported by National Institute of Mental Health grants 5R01MH052808 and 5T32MH019983. The latter training grant on "Combined Computational and Behavioral Approaches to the Study of Cognition" supported R.A.D. and C.P. The authors wish to thank Joyce Oates, Paul Kieffaber, and Alexandra Tsarenko for their comments on an earlier version of the chapter.

References

Anderson, J. R. (1974). Retrieval of propositional information from long-term memory. *Cognitive Psychology, 5*, 451–474.

Anderson, J. R. (1976). *Language, memory, and thought*. Hillsdale, NJ: Erlbaum.

Anderson, J. R., & Betz, J. (2001). A hybrid model of categorization. *Psychonomic Bulletin & Review, 8*(4), 629–647.

Anderson, J. R., & Bower, G. H. (1973). *Human associative memory*. Washington: Winston and Sons.

Anderson, J. R., & Lebiere, C. (1998). *Atomic components of thought*. Mahwah, NJ: Erlbaum.

Anderson, J. R., & Paulson, R. (1978). Interference in memory for pictorial information. *Cognitive Psychology, 30*, 221–256.

Anderson, J. R., & Reder, L. M. (1999). The fan effect: New results and new theories. *Journal of Experimental Psychology: General, 128*, 186–197.

Anderson, J. R., Reder, L. M., & Lebiere, C. (1996). Working memory: Activation limitations on retrieval. *Cognitive Psychology, 30*, 221–256.

Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science, 2*, 396–408.

Balota, D. A., Burgess, G. C., Cortese, M. J., & Adams, D. R. (2002). The word-frequency mirror effect in young, old, and early-stage Alzheimer's disease: Evidence for two processes in episodic recognition performance. *Journal of Memory and Language, 46*, 199–226.

Balota, D. A., & Ferraro, F. R. (1996). Lexical, sublexical, and implicit memory processes in healthy young and healthy older adults and in individuals with dementia of the Alzheimer type. *Neuropsychology, 10*(1), 82–95.

Becker, C. A. (1976). Allocation of attention during visual word recognition. *Journal of Experimental Psychology: Human Perception & Performance, 2*(4), 556–566.

Benjamin, A. S. (2003). Predicting and postdicting the effects of word frequency on memory. *Memory & Cognition, 31*(2), 297–305.

Buchler, N. E. G., & Reder, L. M. (2007). Modeling age-related memory deficits: A two-parameter solution. *Psychology & Aging, 22*(1), 104–121.

Burke, D. M., & Light, L. L. (1981). Memory and aging: The role of retrieval processes. *Psychological Bulletin, 90*, 513–546.

Cary, M., & Reder, L. M. (2002). Metacognition in strategy selection: Giving consciousness too much credit. In P. Chambres, M. Izaute, and P. J. Marescaux (Eds.), *Metacognition: Process, function, and use* (pp. 63–78). New York, NY: Kluwer.

Castel, A. D., & Craik, F. I. M. (2003). The effects of aging and divided attention on memory for item and associative information. *Psychology and Aging, 18*, 873–885.

Chalfonte, B. L., & Johnson, M. K. (1996). Feature memory and binding in young and older adults. *Memory & Cognition, 24*, 403–416.

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

Clark, S. E. (1992). Word frequency effects in associative and item recognition. *Memory & Cognition, 20*(3), 231–243.

Clark, S. E., & Burchett, R. E. R. (1994). Word frequency and list composition effects in associative recognition and recall. *Memory & Cognition, 22*(1), 55–62.

Connine, C. M., Mullenix, J., Shernoff, E., & Yelen, J. (1990). Word familiarity and frequency in visual and auditory word recognition. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 16*(6), 1084–1096.

Cook, S., Reder, L. M., Buchler, N., Hashemi, S., & Dickison, D. A mechanistic account of spurious recollection. Manuscript in preparation.

Daily, L. Z., Lovett, M. C., & Reder, L. M. (2001). Modeling individual differences in working memory performance: A source activation account. *Cognitive Science, 25*, 315–353.

D'Arcy, R. C. N., Ryner, L., Richter, W., Service, E., & Connolly, J. F. (2004). The fan effect in fMRI: Left hemisphere specialization in verbal working memory. *Neuroreport: For Rapid Communication of Neuroscience Research, 15*(12), 1851–1855.

Deese, J. (1960). Frequency of usage and number of words in free recall: The role of association. *Psychological Reports, 7*, 337–344.

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

de Zubicaray, G. I., McMahon, K. L., Eastburn, M. M., Finnigan, S., & Humphreys, M. (2005). fMRI evidence of word frequency and strength effects during episodic memory encoding. *Cognitive Brain Research, 22*, 439–450.

Dewhurst, S. A., Hitch, G. J., & Barry, C. (1998). Separate effects of word frequency and age of acquisition in recognition and recall. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 24*(2), 284–298.

Diana, R. A., Peterson, M. J., & Reder, L. M. (2004). The role of spurious feature familiarity in recognition memory. *Psychonomic Bulletin & Review, 11*(1), 150–156.

Diana, R. A., & Reder, L. M. (2006). The low-frequency encoding disadvantage: Word frequency affects processing demands. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 34*(4), 805–815.

Diana, R. A., Reder, L. M., Arndt, J., & Park, H. (2006). Models of recognition: A review of arguments in favor of a dual-process account. *Psychonomic Bulletin & Review, 13*, 1–21.

Dobbins, I. G., & Kroll, N. E. A. (2005). Distinctiveness and the recognition mirror effect: Evidence for an item-based criterion placement heuristic. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 31*(6), 1186–1198.

Fagan, J. F., III (1970). Memory in the infant. *Journal of Experimental Child Psychology, 9*, 217–226.

Frost, R., & Katz, L. (1989). Orthographic depth and the interaction of visual and auditory processing in word recognition. *Memory & Cognition, 17*(3), 302–310.

Gardiner, J. M., & Java, R. I. (1990). Recollective experience in word and non-word recognition. *Memory & Cognition, 18*, 23–30.

Glanzer, M., & Adams, J. K. (1985). The mirror effect in recognition memory. *Memory & Cognition, 13*(1), 8–20.

Glanzer, M., & Adams, J. K. (1990). The mirror effect in recognition memory: Data and theory. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 16*, 5–16.

Glanzer, M., & Bowles, N. (1976). Analysis of the word-frequency effect in recognition memory. *Journal of Experimental Psychology: Human Learning and Memory, 2*, 21–31.

Goldinger, S. D., Azuma, T., Abramson, M., & Jain, P. (1997). Open wide and say "Blah!" Attentional dynamics of delayed naming. *Journal of Memory and Language, 37*, 190–216.

Gorman, A. M. (1961). Recognition memory for nouns as a function of abstractness and frequency. *Journal of Experimental Psychology, 61*(1), 23–29.

Greene, R. L., & Thapar, A. (1994). Mirror effect in frequency discrimination. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*(4), 946–952.

Guttentag, R. E., & Carroll, D. (1997). Recency judgments as a function of word frequency: A framing effect and frequency misattributions. *Psychonomic Bulletin & Review, 4*(3), 411–415.

Hauk, O., & Pulvermuller, F. (2004). Effects of word length and frequency on the human event-related potential. *Clinical Neurophysiology, 115*(5), 1090–1103.

Hayes-Roth, B. (1977). Evolution of cognitive structures and processes. *Psychological Review, 84*(3), 260–278.

Henson, R. N. A. (2003). Neuroimaging studies of priming. *Progress in Neurobiology, 70*, 53–81.

Hintzman, D. L., Caulton, D. A., & Curran, T. (1994). Retrieval constraints and the mirror effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*(2), 275–289.

Hirshman, E., Fisher, J., Henthorn, T., Arndt, J., & Passannante, A. (2002). Midazolam amnesia and dual-process models of the word-frequency mirror effect. *Journal of Memory and Language, 47*(4), 499–516.

Hockley, W. E. (1994). Reflections of the mirror effect for item and associative recognition. *Memory & Cognition, 22*, 713–722.

Jacoby, L. L. (1991). A process-dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language, 30*, 513–541.

Johnston, W. A., Hawley, K. J., Plewe, S. H., Elliott, J. M., & DeWitt, M. J. (1990). Attention capture by novel stimuli. *Journal of Experimental Psychology: General, 119*(4), 397–411.

Joordens, S., & Hockley, W. E. (2000). Recollection and familiarity through the looking glass: When old does not mirror new. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*(6), 1534–1555.

Kintsch, W. (1974). *The representation of meaning in memory*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Kliegl, R., & Lindenberger, U. (1993). Modeling intrusions and correct recall in episodic memory: Adult age differences in encoding of list context. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19*, 617–637.

Koriat, A. (2000). The feeling of knowing: Some metatheoretical implications for consciousness and control. *Consciousness and Cognition, 9*, 149–171.

Kucera, H., & Francis, N. W. (1967). *Computational analysis of present-day American English*. Providence, RI: Brown University Press.

Lemaire, P., & Reder, L. M. (1999). What affects strategy selection in arithmetic? An examination of parity and five effects on product verification. *Memory & Cognition, 27*(2), 364–382.

Lewis, C. H., & Anderson, J. R. (1976). Interference with real world knowledge. *Cognitive Psychology, 7*, 311–335.

Light, L. L., Healy, M. R., Patterson, M. M., & Chung, C. (2005). Dual-process models of memory in young and older adults: Evidence from associative recognition. In L. Taconnat, D. Clarys, S. Vanneste, and M. Isingrini (Eds.), *Manifestations cognitives du vieillissement psychologique. Actes des VIIemes journees du vieillessement cognitif* (pp. 95–116). Paris, France: Publibook.

Lovett, M. C., Daily, L. Z., & Reder, L. M. (2000). A source activation theory of working memory: Cross-task prediction of performance in ACT-R. *Journal of Cognitive Systems Research, 1*, 99–118.

Lovett, M. C., Reder, L. M., & Lebiere, C. (1997). Modeling individual differences in a digit working memory task. *Proceedings of the Nineteenth Annual Cognitive Science Conference*, 460–465. Mahwah, NJ: Erlbaum.

Lovett, M. C., & Schunn, C. D. (1999). Task representations, strategy variability, and base-rate neglect. *Journal of Experimental Psychology: General, 128*(2), 107–130.

MacLeod, C. M., & Kampe, K. E. (1996). Word frequency effects on recall, recognition, and word fragment completion tests. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 22*(1), 132–142.

Maddox, W. T., & Estes, W. K. (1997). Direct and indirect stimulus-frequency effects in recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23*, 539–559.

Malmberg, K. J., & Murnane, K. (2002). List composition and the word-frequency effect for recognition memory. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 28*(4), 616–630.

Mandler G (1980) Recognizing: The judgment of previous occurrence *Psychological Review 87*(3), 252–271

McClelland J L & Rumelhart D E (1985) Distributed memory and the representation of general and specific information *Journal of Experimental Psychology General 114*(2) 159–188

Naveh-Benjamin M (2000) Adult age differences in memory performance: Tests of an associative deficit hypothesis *Journal of Experimental Psychology Learning Memory and Cognition 26*, 1170–1187

Nelson, A B, & Shiffrin, R M (2006) Modeling the effect of differential experience on perception and memory [abstract] *Proceedings of the Psychonomic Society USA* Vol 11, p 17

Norman D A, Rumelhart, D E, & the LNR Research Group (1975) *Explorations in cognition* San Francisco CA: Freeman

Paivio, A (1971) *Imagery and verbal processes* New York NY: Holt, Rinehart, Winston

Park H, Arndt J D & Reder, L M (2006) A contextual interference account of distinctiveness effects in recognition *Memory & Cognition, 34*(4), 743–751

Park, H, Quinlan, J J, Thornton E R, & Reder L M (2004) The effect of midazolam on visual search: Implications for understanding amnesia *Proceedings of the National Academy of Sciences 101*(51), 17879–17883

Peterson S B, & Potts, G R (1982) Global and specific components of information integration *Journal of Verbal Learning & Verbal Behavior 21*(4) 403–420

Quamme, J R, Frederick C, Kroll N E, Yonelinas, A P & Dobbins I G (2002) Recognition memory for source and occurrence: The importance of recollection *Memory & Cognition, 30*(6), 893–907

Radvansky, G A (1999) The fan effect: A tale of two theories *Journal of Experimental Psychology Genera, 128*(2), 198–206

Rao K V, & Proctor, R W (1984) Study-phase processing and the word frequency effect in recognition memory *Journal of Experimental Psychology Learning Memory and Cognition, 10*(3), 386–394

Reder, L M (1979) The role of elaborations in memory for prose *Cognitive Psychology 11*, 221–234

Reder, L. M (1982) Plausibility judgments vs fact retrieval: Alternative strategies for sentence verification *Psychological Review, 89*, 250–280

Reder, L M (1987) Strategy selection in question answering *Cognitive Psychology 19*(1), 90–138

Reder, L M (1988) Strategic control of retrieval strategies In G Bower (Ed), *The psychology of learning and motivation* (Vol 22, pp 227–259) New York, NY: Academic Press

Reder L M, & Anderson J R (1980) A partial resolution of the paradox of interference: The role of integrating knowledge *Cognitive Psychology, 12*, 447–472

Reder, L M, Angstadt, P, Cary, M, Erickson, M A, & Ayers M A (2002) A reexamination of stimulus-frequency effects in recognition: Two mirrors for low- and high-frequency pseudowords *Journal of Experimental Psychology Learning, Memory and Cognition, 28*, 138–152

Reder, L M, Donavos D K, & Erickson, M A (2002) Perceptual match effects in direct tests of memory: The role of contextual fan *Memory & Cognition, 30*(2) 312–323

Reder, L M, Park, H, & Kieffaber, P D (2007a) Memory systems do not divide on consciousness Under review

Reder, L M, Oates, J M, Dickison, D, Anderson, J R, Gyulai, F, Quinlan, J J, et al (2007b) Retrograde facilitation under midazolam: The role of general and specific interference *Psychonomic Bulletin & Review, 14*(2) 261–269

Reder, L. M., Nhouyvansivong, A., Schunn, C. D., Ayers, M. S., Angstadt, P., & Hiraki, K. (2000). A mechanistic account of the mirror effect for word frequency: A computational model of remember/know judgments in a continuous recognition paradigm. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 26*(2), 294–320

Reder, L. M., Oates, J. M., Thornton, E. R., Quinlan, J. J., Kaufer, A., & Sauer, J. (2006a). Drug-induced amnesia hurts recognition, but only for memories that can be unitized. *Psychological Science, 17*(7), 562–567

Reder, L. M., Proctor, I., Anderson, J. R., Gyulai, F., Quinlan, J. J., & Oates, J. M. (2006b). Midazolam does not inhibit association formation, just its storage and strengthening. *Psychopharmacology, 188*(4), 462–471

Reder, L. M., & Ritter, F. (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

Reder, L. M., & Ross, B. H. (1983). Integrated knowledge in different tasks: The role of retrieval strategy on fan effects. *Journal of Experimental Psychology: Learning, Memory and Cognition, 9*, 55–72

Reder, L. M., & Schunn, C. D. (1996). Metacognition does not imply awareness: Strategy choice is governed by implicit learning and memory. In L. M. Reder (Ed.), *Implicit memory and metacognition* (pp. 45–77). Mahwah, NJ: L. Erlbaum

Reder, L. M., Weber, K., Shang, Y., & Vanyukov, P. (2003). The adaptive character of the attentional system: Statistical sensitivity in a target localization task. *Journal of Experimental Psychology: Human Perception and Performance, 29*(3), 631–649

Reder, L. M., Wible, C., & Martin, J. (1986). Differential memory changes with age: Exact retrieval versus plausible inference. *Journal of Experimental Psychology: Learning, Memory and Cognition, 12*(1), 72–81

Rugg, M. D., Cox, C., Doyle, M. C., & Wells, T. (1995). Event-related potentials and the recollection of low and high frequency words. *Neuropsychologia, 33*(4), 471–484

Schulman, A. I. (1976). Memory for rare words previously rated for familiarity. *Journal of Experimental Psychology: Human Learning & Memory, 2*(3), 301–307

Schunn, C. D., Lovett, M. C., & Reder, L. M. (2001). Awareness and working memory in strategy adaptivity. *Memory & Cognition, 29*(2), 254–266

Schunn, C. D., & Reder, L. M. (1998). Strategy adaptivity and individual differences. In D. L. Medin (Ed.), *The psychology of learning and motivation* (pp. 115–154). Academic Press, York: PA.

Schunn, C. D., Reder, L. M., Nhouyvanisvong, 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*, 1–27

Seidenberg, M. S. (1985). The time course of phonological code activation in two writing systems. *Cognition, 19*, 1–30

Servan-Schreiber, E. (1991). The competitive chunking theory: Models of perception, learning and memory. *Dissertation Abstracts International, 52*(6-B), 3279

Shrager, J., & Siegler, R. S. (1998). SCADS: A model of children's strategy choices and strategy discoveries. *Psychological Science, 9*(5), 405–410

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

Simons, J. S., Dodson, C. S., Bell, D., & Schacter, D. L. (2004). Specific- and partial- source memory: Effects of aging. *Psychology and Aging, 19*, 689–694

Smith, E. E., Adams, N., & Schorr, D. (1978). Fact retrieval and the paradox of interference. *Cognitive Psychology, 10*, 438–464

Sokolov, E. N. (1963). Higher nervous functions: The orienting reflex. *Annual Review of Physiology, 25*, 545–580.

Spehn, M. K., & Reder, L. M. (2000). The Role of Familiarity and Associative Competition in Building Novel Structures. *Unpublished manuscript.*

Spencer, W. D., & Raz, N. (1995). Differential effects of aging on memory for content and context: A meta-analysis. *Psychology and Aging, 10*, 527–539.

Sun, R. (2000). *Explicit and implicit processes of metacognition* (Tech. Rep. No. TR-00-001). Columbia, MO: University of Missouri, CECS Department.

Tulving, E. (1985). Memory and consciousness. *Canadian Psychology, 26*, 1–12.

Watkins, M. J., LeCompte, D. C., & Kim, K. (2000). Role of study strategy in recall of mixed lists of common and rare words. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 26*(1), 239–245.

Whitehouse, A. J. O., Maybery, M. T., & Durkin, K. (2006). The development of the picture-superiority effect. *British Journal of Developmental Psychology, 24*(4), 767–773.

Whittlesea, B. W., Jacoby, L. L., & Girard, K. (1990). Illusions of immediate memory: Evidence of an attributional basis for feelings of familiarity and perceptual quality. *Journal of Memory and Language, 29*(6), 716–732.

Yonelinas, A. P. (1994). Receiver-operating characteristics in recognition memory: Evidence for a dual-process model. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*(6), 1341–1354.

Zbrodoff, N. J. (1995). Why is 9 + 7 harder than 2 + 3? Strength and interference as explanations of the problem-size effect. *Memory & Cognition, 23*, 689–700.