7

Human Memory: Retention and Retrieval

Popular fiction sometimes includes a protagonist who is unable to recall some critical information — either because of a head injury or because of repression due to a traumatic experience, or just because the passage of time has seemed to erase the memory. The critical event in the story occurs when the protagonist is able to recover the memory — perhaps because of hypnosis, clinical treatment, returning to an old context, or (particularly improbable) sustaining another head injury. Although our everyday struggles with our memory are seldom so dramatic, we all have had experiences with memories that are just on the edge of availability. For instance, try remembering the name of someone who sat beside you in class in grade school or the name of a grade school teacher. Often, we can picture the person but struggle to retrieve the person's name — a struggle at which we may or may not succeed. This chapter will answer the following questions:

- How does memory for information fade with the passage of time?
- How do other memories interfere with the retrieval of a desired memory?
- How can other memories support the retrieval of a desired memory?
- How does a person's internal and external context influence the retrieval of a memory?
- How can our past experiences influence our behavior without our being able to recall these experiences?
- What role do temporal (and particularly hippocampal) regions of the brain play in memory?

Are Memories Really Forgotten?

Figure 7.1 identifies the prefrontal and temporal structures that have proved important in studies of memory (see Figure 6.1 for a different view of these regions). An early study on the role of the temporal cortex in memory seemed to provide evidence that forgotten memories are still there even though we cannot retrieve them. As part of a neurosurgical procedure, Penfield (1959) electrically stimulated portions of patients' brains and asked them to report what they experienced (patients were conscious during the surgery, but the stimulation was painless). In this way, Penfield determined the functions of various portions of the brain. Stimulation of the temporal lobes led to reports of memories that patients were unable to report in normal recall, such as events from childhood. This seemed to provide evidence that much of what seems forgotten is still stored in memory. Unfortunately, it is hard to know whether the patients' memory reports were accurate — that is, whether the reported events actually occurred. Therefore, although suggestive, the Penfield experiments are generally discounted by memory researchers.

A better experiment, conducted by Nelson (1971), also indicated that forgotten memories still exist. He had participants learn a list of 20 paired associates, each consisting of a number for which the participant had to recall a noun (e.g., *43-dog*). The subjects studied the list and were tested on it until they could recall all the items without error. Participants returned for a retest 2 weeks later and were able to recall 75% of the associated nouns when cued with the numbers. However, the research question concerned the 25% that they could no longer recall — were these items really forgotten? Participants were given new learning trials on the 20 paired associates. The paired associates they had missed were either kept the same or changed. For example, participants who had learned *43-dog* but failed to recall the response *dog* to *43* might now be trained on either *43-dog* (unchanged) or *43-house* (changed).

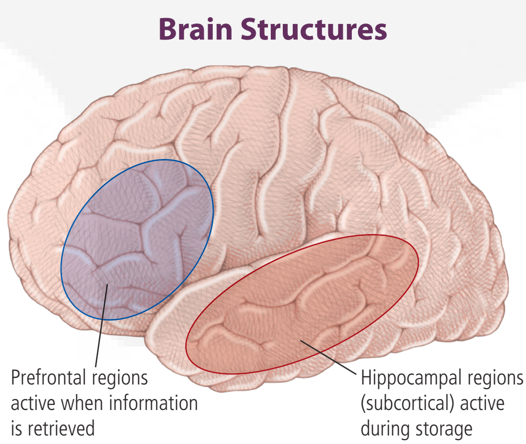

Figure 7.1 Brain Structures Involved in the Creation and Storage of Memories. Prefrontal regions are responsible for the creation of memories. The hippocampus and surrounding structures in the temporal cortex are responsible for the permanent storage of these memories.

Participants were tested after studying the new list once. If the participants had lost all memory for the forgotten pairs, there should have been no difference between recall of changed and unchanged pairs. However, participants correctly recalled 78% of the unchanged items formerly missed, but only 43% of the changed items. This large advantage for unchanged items indicates that participants had retained some memory of the original paired associates, even though they had been unable to recall them initially.

Johnson, McDuff, Rugg, and Norman (2009) report a brain-imaging study that also shows that there are records of experiences in our brain that we can no longer remember. Participants saw a list of words and, for each word, were asked either to imagine how an artist would draw the object denoted by the word or to imagine uses for the object. The researchers trained a pattern classifier (a program for analyzing patterns of brain activity, as discussed in Implications 4.1) to distinguish between words assigned to the artist task and words assigned to the uses task, based on differences in brain activity during the two tasks. Later, the classifier was applied to participants' brain activation patterns while they were shown the words again and asked to recall the type of task they had assigned to each word. The classifier was able to recognize from these patterns what task the word had been assigned to with better than chance accuracy. It was successful at recognition both for words that participants could recall studying and for words they could not remember, although the accuracy was somewhat lower for the words they could not remember. This indicates that even though we may have no conscious memory of something, aspects of how we experienced it may be retained in our brains.

These experiments do not prove that everything is remembered. They show only that appropriately sensitive tests can find evidence for remnants of some memories that appear to have been forgotten. In this chapter, we will discuss first how memories become less available with time, then some of the factors that determine our success in retrieving these memories.

> *Even when people appear to have forgotten memories, there is evidence that they still have some of these memories stored.*

The Retention Function

The processes by which memories become less available are extremely regular, and psychologists have studied their mathematical form. Wickelgren did early, systematic research on memory retention functions, and his data are still used today. In one recognition experiment (Wickelgren, 1975), he presented participants with a sequence of words to study and then examined the probability of their recognizing the words after delays ranging from 1 min to 14 days. **Figure 7.2** shows performance as a function of delay. The performance measure Wickelgren used is called d' (pronounced d-prime), which is derived from the probability of recognition. Wickelgren interpreted d' as a measure of memory strength.

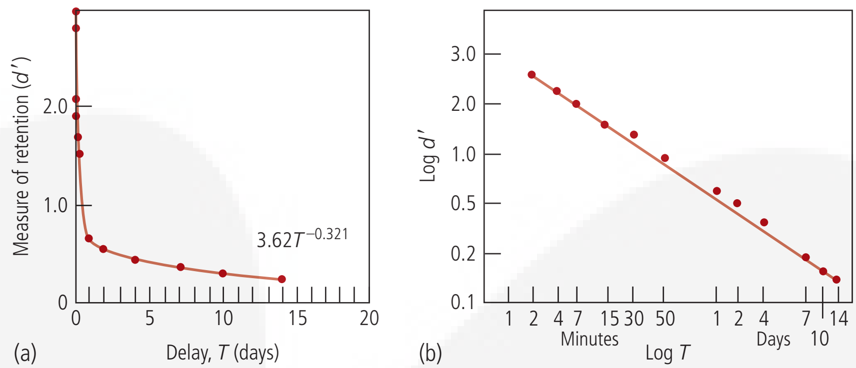

Figure 7.2 Results from Wickelgren's (1975) Experiment on Memory Retention. (a) Success at word recognition, as measured by d', as a function of delay T. (b) The data in (a) replotted on logarithmic scales. *(Data from Wickelgren, 1975.)*

Figure 7.2 shows that this measure of memory systematically deteriorates with delay. However, the memory loss is *negatively accelerated* — that is, the rate of change gets smaller and smaller as the delay increases. Figure 7.2b replots the data as the logarithm d' versus the logarithm of delay. Marvelously, the function becomes linear:

$$\log d' = A - b \log T$$

where A is the value of the function at 1 min $[\log(1) = 0]$ and b is the slope of the function in Figure 7.2b (= 0.321 in this case).

This equation can be transformed to

$$d' = cT^{-b}$$

where $c = 10^A$ (= 3.62 in this case). Such a functional relationship is called a *power function* because the independent variable (the delay T in this case) is raised to a power ($-b$ in this case) to produce the performance measure (d' in this case). In a review of research on forgetting, Wixted and Ebbesen (1991) concluded that retention functions are generally power functions. This relationship is called the **power law of forgetting.** Recall from Chapter 6 that there is also a power law of learning: Both power functions are negatively accelerated, but with an important difference: Whereas the practice functions associated with the power law of learning show diminishing improvement with practice (see Figure 6.12), the retention functions associated with the power law of forgetting show diminishing loss with delay (see Figure 7.2).

A very extensive investigation of the negative acceleration in retention functions was produced by Bahrick (1984), who looked at participants' retention of English–Spanish vocabulary items anywhere from immediately to 50 years after

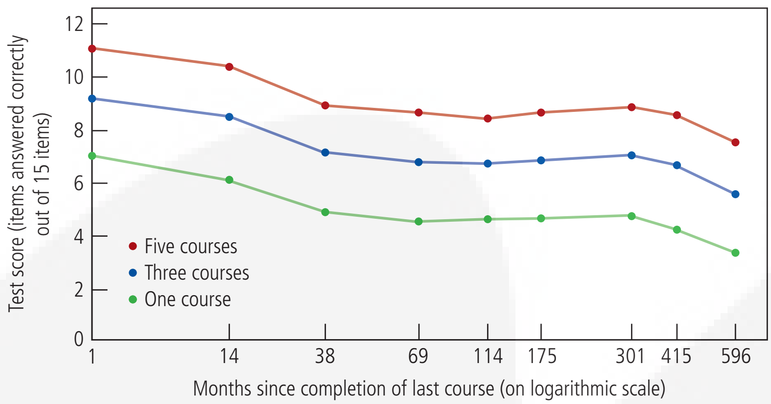

Figure 7.3 Results from Bahrick's (1984) Experiment Measuring Retention of English–Spanish Vocabulary Items over Various Time Periods. The number of items correctly recalled out of a total of 15 items is plotted as a function of the logarithm of the time since course completion. *(Data from Bahrick, 1984.)*

they had completed courses in high school and college. **Figure 7.3** plots the number of items correctly recalled out of a total of 15 items as a function of the logarithm of the time since course completion. Separate functions are plotted for students who had one, three, or five courses. The data show a slow decay of memory combined with a substantial practice effect (the greater the number of courses, the better the recall, regardless of time since completion). In Bahrick's data, the retention functions are nearly flat between 3 and 25 years (as would be predicted by a power function), with some further drop-off from 25 to 50 years (which is more rapid than would be predicted by a power function). Bahrick (personal communication, circa 1993) suspects that this final drop-off is probably related to memory deterioration associated with old age.

There is some evidence that the explanation for these retention functions may be found in the associated neural processes. In Chapter 6, we saw that long-term potentiation (LTP, an increase in neural responsiveness that occurs as a reaction to prior electrical stimulation) mirrors the power law of learning (see Figure 6.14). Raymond and Redman (2006) found a decrease in LTP in the rat hippocampus with delay after electrical stimulation. **Figure 7.4** shows their

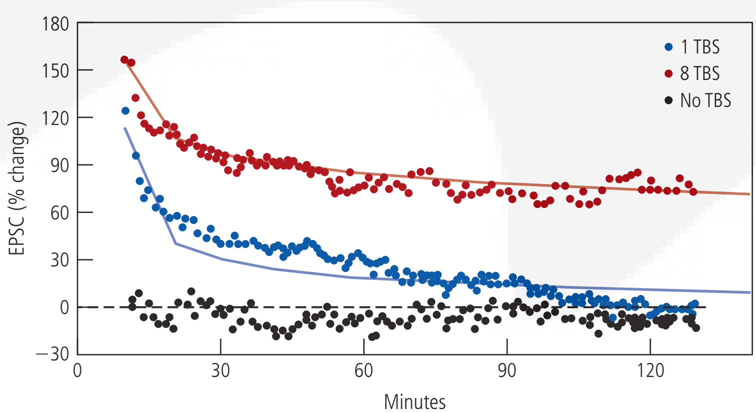

Figure 7.4 Neural Forgetting Mirrors Behavioral Forgetting. At 10 minutes in this experiment, Raymond and Redman (2006) stimulated rats' hippocampus with 1 or 8 theta-burst stimulations (TBS). The decrease in ESPC (excitatory postsynaptic current—a measure of LTP) with delay (the two lines represent the best-fitting power functions) mirrors the decline in memory with delay (e.g., the power function in Figure 7.2a). The "No TBS" results represent a control condition that received no stimulation.

results from three conditions: no stimulation (a control condition, with no resulting LTP), a single stimulation to induce LTP, and eight such stimulations. The level of LTP is greater in the condition with eight stimulations than in the condition with one (a learning effect), but both conditions show a drop-off with delay, and the best-fitting curves for those two conditions (the smooth lines in Figure 7.4) show that maintenance of LTP has the form of a power function. Thus, the time course of this neural forgetting mirrors the time course of behavioral forgetting, just as the neural learning function mirrors the behavioral learning function. In terms of the strength concept introduced in Chapter 6, the assumption is that the strength of a memory decays with time. The data on LTP suggest that this strength decay involves changes in the strength of connections between synapses. Thus, there may be a direct relationship between the concept of strength defined at the behavioral level and strength defined at the neural level.

The idea that memories simply decay in strength with time is one of the common explanations of forgetting; it is called the **decay theory** of forgetting. In the next section, we will review one of the major competitors of this theory.

> *The strength of a memory decays as a power function of the period of time over which it is retained.*

How Interference Affects Memory

The discussion to this point might lead one to infer that the only factor affecting loss of memories is the passage of time. However, it turns out that retention is strongly impacted by another factor, interfering material — this is the **interference theory** of forgetting. Much of the original research on interference investigated how learning a list of paired associates would affect memory for a previously learned list. **Table 7.1** illustrates paired-associates lists made up by associating nouns as stimuli to two-digit numbers as responses. While various experiments will involve many sorts of pairings besides nouns and numbers, such items are typical of the rather arbitrary associates participants are asked to learn. As in the table, there are two critical groups, experimental and control. The experimental group learns two lists of paired associates, first the list designated A–B and then the list designated A–D. These lists are so designated because they share common stimuli (the A terms — e.g., *cat* and *house* in Table 7.1) but different responses (the B and D terms — e.g., *43* and *82* in Table 7.1). The control group also first studies the A–B list but then studies a completely different second list, designated C–D, which does not share any stimuli with the first list (e.g., none of the C terms in Table 7.1 matches any of the A terms). After learning their respective second lists, both groups are

TABLE 7.1 Sample Paired-Associates Lists in a Typical Interference Experiment

Experimental Group	*Control Group*
Learn A–B	*Learn A–B*
cat-43	cat-43
house-61	house-61
apple-29	apple-29
etc.	etc.
Learn A–D	*Learn C–D*
cat-82	bone-82
house-37	cup-37
apple-45	chair-45
etc.	etc.

retested for memory of their first list (the same A–B list in both cases). Often, this retention test is administered after a considerable delay, such as 24 hours or a week. In general, the experimental group, which learns the A–D list, does not do as well as the control group, which learns the C–D list: The experimental group shows both a slower rate of learning of the second list and poorer retention of the original A–B list (see Keppel, 1968, for a review). Such experiments provide evidence that learning the A–D list interferes with retention of the A–B list, causing it to be forgotten more rapidly.

More generally, research has shown that it is difficult to maintain multiple associations to the same items. It is harder both to learn new associations to old items and to retain the old ones if new associations are learned. These results might seem to have rather a dismal implication for our ability to remember information: that it would become increasingly difficult to learn new information about a concept — for example, every time we learned a new fact about a friend, we would be in danger of forgetting an old fact about that person. Fortunately, there are important additional factors that counteract such interference. Before discussing these factors, however, we need to examine in more detail the basis for interference effects. It turns out that a rather different experimental paradigm has been helpful in identifying the cause of those effects.

> *Learning additional associations to an item can cause old ones to be forgotten.*

The Fan Effect: Networks of Associations

The interference effects discussed above can be understood in terms of how much activation spreads to stimulate a memory structure (refer back to the activation equation in Chapter 6). The basic idea is that when participants are presented with a stimulus such as *cat,* activation will spread from this source stimulus to all of its associated memory structures. However, the total amount of activation that can spread from a source is limited; thus, the greater the number of associated memory structures, the less the amount of activation that will spread to any one structure.

In one of my dissertation studies illustrating these ideas (Anderson, 1974), I asked participants to memorize 26 sentences of the form a-person-is-in-a-location, like the four example sentences listed below. As you can see from these examples, some persons were paired with only one location (1–1 and 1–2) and some locations with only one person (1–1 and 2–1), whereas other persons were paired with two locations (2–1 and 2–2) and other locations with two persons (1–2 and 2–2):

1. The doctor is in the bank. (1–1)
2. The fireman is in the park. (1–2)
3. The lawyer is in the church. (2–1)
4. The lawyer is in the park. (2–2)

TABLE 7.2 Recognition of Studied Sentences When Presented with Re-paired Sentences

Sentence Type	*Recognition Time (s)*
1–1 (person appears in 1 studied sentence, location appears in 1 studied sentence)	1.11
1–2 (person appears in 1 studied sentence, location appears in 2 studied sentences)	1.17
2–1 (person appears in 2 studied sentences, location appears in 1 studied sentence)	1.17
2–2 (person appears in 2 studied sentences, location appears in 2 studied sentences)	1.22

(Data from Anderson, 1974.)

Participants were drilled on 26 sentences like these until they knew the material well. Then participants were presented with a set of test sentences that consisted of studied sentences mixed in with new sentences created by re-pairing people and locations from the study set, and participants had to recognize the sentences from the study set.

The recognition times displayed in **Table 7.2** show that recognition time increases as a function of the sum of the two numbers used to classify the example sentences above — that is, sentences that could be labeled 1–1 are fastest to be recognized (sum of associations = 2), sentences that could be labeled 1–2 or 2–1 are next fastest (sum of associations = 3), and sentences that could be labeled 2–2 are slowest (sum of associations = 4). The increase in recognition time from fastest to slowest is not much more than a hundred milliseconds ($1.22\ s - 1.11\ s = 0.11\ s = 110$ milliseconds), but such effects can add up in situations like taking a test under time pressure: Taking a little more time to answer each question could mean not finishing the test.

These interference effects — that is, the increases in recognition time — can be explained in terms of activation spreading through network structures, like the network structure in **Figure 7.5**, which represents the four example sentences above. According to the spreading-activation theory, recognizing a sentence (i.e., retrieving the memory of that sentence) would involve the following discrete steps:

1. Presentation of a sentence activates the representations of the concepts in the sentence. In Figure 7.5, the concepts are *doctor, lawyer, fireman, bank, church,* and *park,* which are each associated with either one or two of the four sentences.
2. Activation spreads from these source concepts to memory structures representing the associated sentences. In Figure 7.5, the ovals represent these memory structures, and the arrows represent the activation pathways from the concepts. However, as noted above, the total amount of activation that can spread from a source is limited. This means,

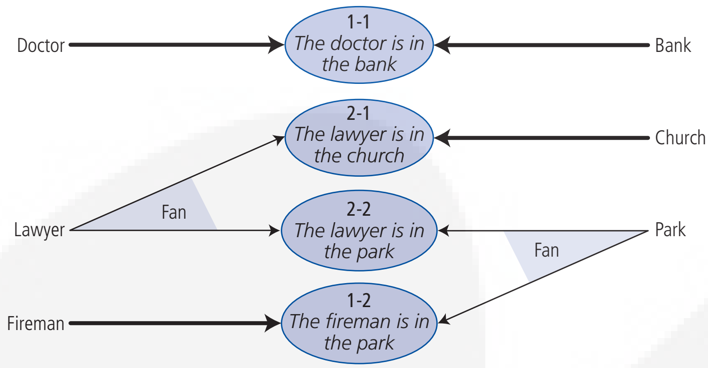

Figure 7.5 How Spreading Activation Works. The ovals are the memory structures of the sentences to be remembered (from Anderson, 1974). Each memory structure is labeled with the number of associations of the person and location in the sentence — e.g., in the 2–1 sentence, the person (*lawyer*) has 2 associations and the place (*church*) has 1 association. The sources of activation are the concepts *doctor, lawyer, fireman, bank, church,* and *park;* the arrows represent the activation pathways; and the thickness of each arrow represents the amount of activation along that pathway.

for example, that each of the two pathways from *lawyer* carries less activation than the single pathway from *doctor*. In Figure 7.5, the thickness of each arrow represents the amount of activation it carries.

3. As activation spreading down the pathways converges on the memory structures, the memory structures are activated to various levels. These activations sum to produce an overall level of activation of the memory structure. Because of the limitation on the total activation from any one source, a memory structure's activation level is inversely related to the sum of associations of the source concepts.
4. A sentence is recognized in an amount of time that is inversely related to the activation level of its memory structure — that is, the greater the activation level, the less time required to retrieve the memory and recognize the sentence. Or, to put it in terms of associations, the greater the number of associations of the source concepts, the more time required to recognize the sentence.

So, given a structure like that shown in Figure 7.5, participants should be slower to recognize the sentence involving *lawyer* and *park* than the one involving *doctor* and *bank* (as is the case in Table 7.2) because more paths emanate from the first set of concepts. That is, in the *lawyer* and *park* case, two paths point from each of the concepts to the two sentences in which each was studied, whereas only one path leads from each of the *doctor* and *bank* concepts. The increase in reaction time related to an increase in the number of memory structures associated with a concept is called the **fan effect.** It is so named because the increase in reaction time is related to an increase in the fan of activations emanating from the network representation of the concept (see the fans in Figure 7.5).

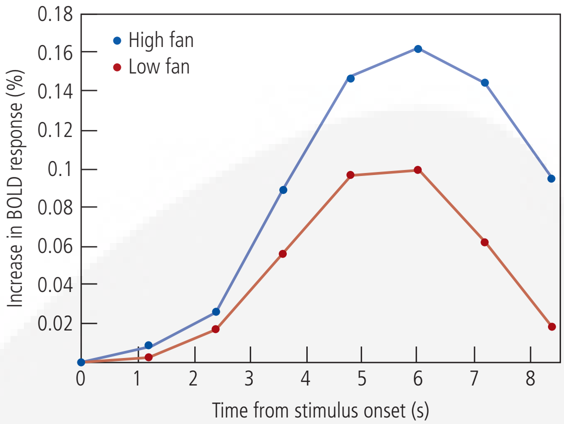

Figure 7.6 Hemodynamic Response in the Prefrontal Cortex During Retrieval of Low-Fan vs. High-Fan Sentences. The increase in BOLD response is plotted against the time from stimulus onset. The greater response for high-fan sentences reflects the greater amount of mental work required to retrieve them from memory. *(Data from Sohn et al., 2003.)*

In an fMRI brain-imaging study, Sohn, Goode, Stenger, Carter, and Anderson (2003) looked at the response in the prefrontal cortex during the recognition of such sentences. They contrasted recognition of high-fan sentences (composed of concepts that appeared in many other sentences) with low-fan sentences (composed of concepts that appeared in few sentences). **Figure 7.6** compares the hemodynamic response in the two conditions and shows that there is a greater hemodynamic response for the high-fan sentences, which have lower activation. One might have expected lower activation to map onto weakened hemodynamic response. However, the prefrontal structures must work harder to retrieve memories with lower activation. As we will see throughout the later chapters of this text, in which we look at higher mental processes like problem solving, more difficult conditions are associated with higher metabolic expenditures, reflecting the greater mental work required under more difficult conditions.

The more facts associated with a concept, the slower the retrieval of any one of the facts.

The Interfering Effect of Preexisting Memories

Do such interference effects occur with material learned outside of the laboratory? As one way to address this question, Lewis and Anderson (1976) investigated whether the fan effect could be obtained with material the participant knew before the experiment. We had participants learn fantasy "facts" about public figures, in the form of statements such as *Napoleon Bonaparte was from India*. Participants studied from zero to four such fantasy facts about each public figure. After learning these "facts," they proceeded to a recognition test phase, in which they saw three types of sentences: (1) statements they had studied in the experiment; (2) true statements about the public figures (e.g., *Napoleon Bonaparte was an emperor*); and (3) statements about the public figures that were false in the real world and had not been studied. Participants had to respond to the first two types of statements as true and to the last type as false.

Figure 7.7 presents participants' times in making these judgments as a function of the number (or fan) of the fantasy facts studied about the person. Note that the recognition time increased with the fan for all types of statements. Also note that participants responded much faster to actual true statements than to the fantasy facts they had studied. The advantage of true statements can be explained by the observation that these facts would be much more strongly encoded in memory than the fantasy facts studied. The most important result to note in Figure 7.7 is that the more fantasy facts participants learned about an individual such as Napoleon Bonaparte, the longer they took to recognize a fact that they already knew about the individual (e.g., *Napoleon Bonaparte was an emperor*). This shows that we can produce interference with preexperimental material. (For further research on this topic, see Peterson & Potts, 1982.)

> *Material learned in the laboratory can interfere with material learned outside of the laboratory.*

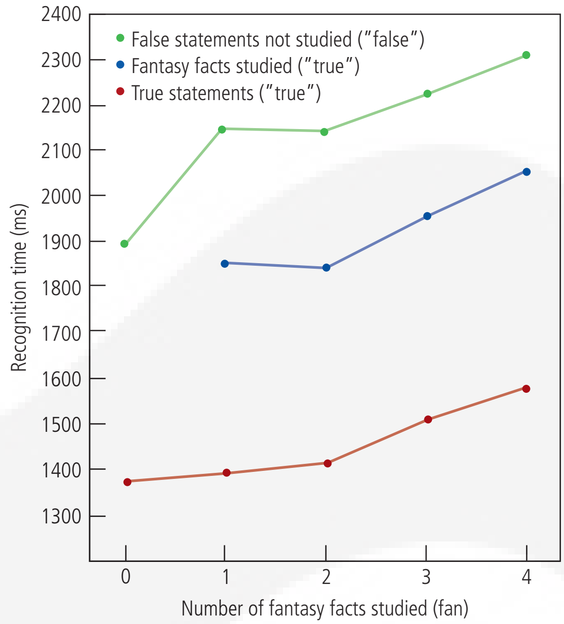

Figure 7.7 Can the Fan Effect Be Obtained Using Already Known Material? In Lewis and Anderson (1976), participants studied 0–4 fantasy "facts" about public figures and then saw a mix of those statements, true statements about the public figures, and false statements about them that they had not studied. The task was to respond "true" to the first two types of statements and "false" to the last type. Participants' recognition times are plotted as a function of the number (or fan) of the fantasy facts studied. The time participants took to make all three judgments increased as they learned more fantasy facts. *(Data from Lewis and Anderson, 1976.)*

Both Decay and Interference?

We have discussed two mechanisms that can produce forgetting: decay of memory strength and interference from other memories. There has been some speculation among researchers that what appears to be decay may really reflect interference. That is, the reason memories appear to decay over a retention interval is that they are interfered with by additional memories that the participants have stored. Objections have been raised to decay theories because they do not identify the psychological factors that produce the forgetting but rather just assert that forgetting occurs spontaneously with time. It is possible, however, that there is no explanation of decay at the purely psychological level. The explanation may be physiological, as we saw with respect to the LTP data (see Figure 7.4). Thus, it seems that the best conclusion, given the available data, is that both interference and decay contribute to forgetting (Sadeh, Ozubko, Winocur, & Moscovitch, 2016).

> *Forgetting results both from decay in memory strength and from interference from other memories.*

An Inhibitory Explanation of Forgetting?

A more recent controversy in psychology concerns the issue of whether interference effects are due to an inhibition process that actively suppresses the competing memories rather than being a passive side effect of storing and strengthening memories. The inhibition account has been championed by Michael Anderson (e.g., Anderson, 2003). Evidence for this account comes from a variety of paradigms that show that trying to retrieve certain items causes others to be forgotten. For instance, participants might learn a list of category–exemplar pairs where there are multiple instances of the same category. After the initial study, participants are given practice on only some of the pairs studied; then they are given a recall test in which they see the category names and have to recall all the instances they studied. For example, pairs such as the four below might be studied (among others), followed by practice on *Red–Blood* but not on the other three pairs, followed by a recall test with the results shown in parentheses (these results are from one of the early experiments — Anderson & Spellman, 1995):

Red–Blood (practiced)	(74% correct recall)
Red–Tomato	(22% correct recall)
Food–Strawberry	(22% correct recall)
Food–Cracker	(36% correct recall)

Not surprisingly, participants showed the highest recall for *Red–Blood,* which they practiced. Interest focuses on recall of the other pairs, which were not practiced. Note that recall was lower for both *Red–Tomato* and *Food–Strawberry* than for *Food–Cracker*. Michael Anderson argues that while practicing *Red–Blood,* participants were inhibiting all other red things, including *Strawberry,* which they did not even study as a red thing. The lower recall for *Red–Tomato* can be explained by other types of interference, such as competition from the strengthened *Red–Blood* association, but the lower recall of *Food–Strawberry* is considered evidence for the inhibition account.

Another source of evidence for the inhibition account comes from what is called the *think/no-think paradigm* (Anderson & Green, 2001). Participants study pairs like *Ordeal–Roach*. Then they are presented with the first item in a pair (e.g., *Ordeal*) and asked either to think about the response or to suppress thinking about the other item in the pair (e.g., *Roach*). Participants are then tested with a probe like *Insect–R,* where their task is to produce a word from the pairs studied that is associated to the first term (*Insect*) and that begins with the given first letter (*Roach* in this example). Participants are less likely to recall the target word (*Roach*) if they have been suppressing it.

Unfortunately for the purposes of presenting firm conclusions, there have been a number of recent critiques of this research (e.g., Raaijmakers & Jakab, 2013; Verde, 2012). Other researchers sometimes can replicate these results but oftentimes cannot. Much effort has been put into trying to elucidate the cause of this mixed empirical picture. One idea that has emerged is that the occurrence of these "inhibition" effects may be the result of unobserved strategies used by participants. For instance, in the think/no-think paradigm, participants may think of some other insect to prevent themselves from thinking of *Roach*. In the first experiment we discussed, when subjects are given the stimulus *Food,* they might be tempted to respond with items associated with the category *Red,* because some of the food items were red. Thus, what appears to be the result of inhibition of a response item like *Roach* or *Strawberry* may actually be the result of competition from other, implicit stimuli generated by the participant's strategy. Such strategies could vary with many factors, and this strategy variation could explain the inconsistent results. There is some evidence for the existence of such covert strategies (e.g., Camp, Pecher, & Schmidt, 2005), although the evidence has been disputed (see Huddleston & Anderson, 2012).

In some ways retrieval-induced suppression is not a new idea. It harkens back to Freud, who argued that we suppress unpleasant memories.

Freud's hypothesis was thought to apply only to highly emotional memories, and even there it is controversial (see the later section of this chapter on the false memory controversy). Freud's original account of the mechanisms that produced suppressed memories is not generally accepted. One of the criticisms of current inhibition ideas is that proponents have not described mechanisms that might produce such inhibition. This is similar to the criticisms of decay theory for not describing mechanisms that might produce the decay. (Regardless of the exact mechanisms behind forgetting, there is evidence that forgetting itself may be adaptive from an evolutionary standpoint — see **Implications 7.1** for discussion of this idea.)

It has been argued that forgetting may also be produced by active inhibition (suppression) of memories, but the evidence is inconclusive.

Implications 7.1

Is Forgetting Adaptive?

We tend to view the fact we forget things as an annoying feature of our brains. However, Schooler and Anderson (2017) argue that forgetting is actually quite adaptive. That argument begins with the observation that it is impossible for any physical system to store everything. Libraries throughout history have had only so much space for books (or, before books, for scrolls), meaning that hard decisions had to be made about what to keep and what to discard (Agee & Naper, 2007). Likewise, there is only so much storage available in the human brain. Even in the current digital age, with the immense storage capacity available, there are limits. For instance, in response to the search query "cognitive psychology," Google reports that it has millions of related results. However, if you explore those results, you will find that only a few hundred pages are available. The strength of search engines is that they make the best results available — users do not go beyond the first page of results for 94% of all searches (Buddenbrock, 2016). Likewise, human memory makes some information much more available than other information. While we might struggle to remember the name of our third grade teacher, we do not have that problem with the name of our significant other, parent, or best friend.

Schooler and Anderson observed that human memory tends to forget what is less useful and tends to make most available what is most useful. For instance, consider the power law of forgetting (see Figure 7.2), describing how the probability of remembering some piece of information decreases with the delay since the information was last recalled. Across a wide range of environments, the longer it has been since we have encountered something, the less likely we are to encounter it again and hence it is less costly to lose memory access to it. The striking observation is that the probability of encountering some person or thing declines as a power function, just like human forgetting. Interestingly, **Figure 7.8** shows that this observation is true of both humans and chimpanzees. Figure 7.8a (data from Pachur, Schooler, & Stevens, 2014) shows how the probability of meeting someone declines as a function of how long it has been since the last face-to-face meeting. Figure 7.8b (data from Stevens, Marewski, Schooler, & Gilby, 2016) shows that a similar relationship holds for chimpanzee meetings in the wild. In both the human and chimpanzee cases, the rate of decline closely fits a power function like the function for the human law of forgetting. This match between species suggests that the human forgetting function may have been shaped through a long history of evolutionary adaptations, consistent with what we see with respect to long-term potentiation (Figure 7.4).

As another example of the adaptiveness of human memory, Anderson and Schooler review evidence that memories typically need to be accessed in particular contexts. For instance, in the gym the thing I am most likely to need to remember is the combination of the lock on my locker. Fortunately memories do tend to be more available in the context in which they are experienced

Implications 7.1 (*Continued*)

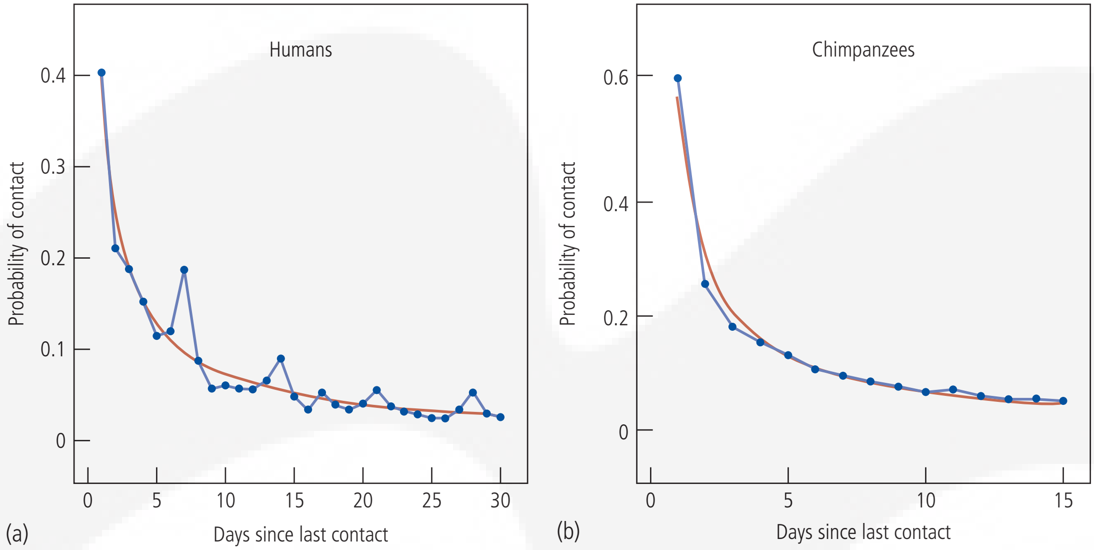

Figure 7.8 Probability of Contact as a Function of Time Since Last Contact. (a) For people going about their daily lives, the probability of a face-to-face encounter with a specific other person declines at a rate that is proportional to the time since the last encounter. *(Data from Pachur et al., 2014, and Stevens et al., 2016.)* (b) The same is true for chimpanzees in the wild encountering a specific other chimpanzee. *(Data from Stevens et al., 2016.)* For both species, the rate of decline in this probability is a power function, like the power function that describes the power law of forgetting. (The smooth curves are the best-fitting power functions.)

(see Figure 7.12). Some cues tend to be associated with particular facts. For instance, my lock is nearly uniquely associated with its combination whereas other cues (like my clothes) are associated with many things. Correspondingly, human memory gives more reliable and more rapid access to memories when cued with items that are uniquely associated with the memory (see Tables 7.1 and 7.2). This kind of context-dependent access is similar to having a web page show up high in the results for a search query when items in the query are found only on that web page.

Relatedness Protects Against Interference

There is a major qualification about the situations in which interference effects are seen: Interference occurs only when one is learning multiple pieces of information that have no intrinsic relationship to one another. In contrast, interference does not occur when the pieces of information are meaningfully related. An experiment by Bradshaw and Anderson (1982) illustrates the contrasting effects of related versus unrelated information. These researchers looked at participants' ability to learn some little-known information about famous people, under three conditions: single, unrelated, and related:

- In the single condition, they had participants study just one fact:

  Newton became emotionally unstable and insecure as a child.

- In the unrelated condition, they had participants learn a target fact about the individual:

  Locke was unhappy as a student at Westminster.

  plus two unrelated facts:

  Locke felt fruits were unwholesome for children.
  Locke had a long history of back trouble.

TABLE 7.3 Recall of Related and Unrelated Information

	Recall of Target Facts (% of participants)	
Condition	*Immediately*	*After 1 Week*
Single	92	62
Unrelated	80	45
Related	94	73

(Data from Bradshaw & Anderson, 1982.)

- In the related condition, participants learned a target fact:

  Mozart made a long journey from Munich to Paris.

  plus two additional facts that were causally related to the target fact:

  Mozart wanted to leave Munich to avoid a romantic entanglement.
  Mozart was intrigued by musical developments coming out of Paris.

Participants were then tested for their ability to recall the target facts immediately after studying them and after a week's delay. In each test, they were presented with names such as Newton, Mozart, and Locke and asked to recall what they had studied. **Table 7.3** shows the results in terms of the percentage of participants who recalled the target facts. Comparing the unrelated condition with the single condition, we see the standard interference effect: Recall was worse when there were more facts to be learned about an item. However, the result is quite different when we compare the related condition to the single condition. Here, particularly after a week's delay, recall was better when there were more facts to be learned, presumably because the additional facts were causally related to the target facts.

To understand why the effects of interference are eliminated or even reversed when there is relatedness among the materials to be learned requires that we move on to discussing the retrieval process and, in particular, the role of inferential processes in retrieval.

Learning related material does not interfere with retrieval of a target memory and may even facilitate retrieval.

Retrieval and Inference

Often, when people cannot remember a particular fact, they are able to retrieve related facts and so infer the target fact on the basis of the related facts. For example, in the case of the Mozart facts just discussed, if the participants who could not recall the target fact (that Mozart made a long journey from Munich to Paris) were able to retrieve the other two facts, they might have then been able to infer the target fact. There is considerable evidence that people make such inferences at the time of recall. They seem unaware that they are making inferences but rather think that they are directly recalling what they studied.

Bransford, Barclay, and Franks (1972) reported an experiment that demonstrates how inference can lead to incorrect recall. They had participants study one of the following sentences:

1. Three turtles rested beside a floating log, and a fish swam beneath them.
2. Three turtles rested on a floating log, and a fish swam beneath them.

Participants who had studied sentence 1 were later asked whether they had studied this sentence:

3. Three turtles rested beside a floating log, and a fish swam beneath it.

Not many participants thought they had studied this sentence. Participants who had studied sentence 2 were tested with:

4. Three turtles rested on a floating log, and a fish swam beneath it.

The participants in this group judged that they had studied sentence 4 much more often than participants in the other group judged that they had studied sentence 3. Sentence 4 is implied by sentence 2, whereas sentence 3 is not implied by sentence 1. Thus, participants thought that they had actually studied what was implied by the studied material.

A study by Sulin and Dooling (1974) illustrates how inference can bias participants' memory for a text. They asked participants to read the following passage:

> **Gerald Martin's Seizure of Power**
> Gerald Martin strove to undermine the existing government to satisfy his political ambitions. Many of the people of his country supported his efforts. Current political problems made it relatively easy for Martin to take over. Certain groups remained loyal to the old government and caused Martin problems. He confronted these groups directly and so silenced them. He became a ruthless, uncontrollable dictator. The ultimate effect of his rule was the downfall of his country.

A second group of participants read the same passage, except that the name *Adolf Hitler* was substituted for *Gerald Martin*. A week after reading the passage, participants were given a recognition test in which they were presented with a sentence and asked to judge whether it had occurred in the passage they read originally. One of the critical test sentences was *He hated the Jews particularly and so persecuted them*. Only 5% of participants who read the Gerald Martin passage accepted this sentence, but a full 50% of the participants who read the Adolf Hitler version thought they had read the sentence. The second group of participants had elaborated the story with facts they knew about Adolf Hitler. Thus, it seemed reasonable to them at test that this sentence had appeared in the studied material, but in this case their inference was wrong.

We might wonder whether an inference such as *He hated the Jews particularly and so persecuted them* was made while the participant was studying the passage or only at the time of the test. This is a subtle issue, and participants certainly do not have reliable intuitions about it. However, a couple of techniques seem to yield evidence that some inferences are being made at test. One method is to determine whether the inferences increase in frequency with delay. With delay, participants' memory for the studied passage should deteriorate, and if they are making inferences at test, they will have to do more reconstruction via inference, which in turn will lead to more inferential errors. Both Dooling and Christiaansen (1977) and Spiro (1977) found evidence for increased inferential intrusions with increased delay of testing. Dooling and Christiaansen used another technique with the Gerald Martin passage to show that inferences were being made at test. They had the participants study the passage and then told them a week later, just before test, that Gerald Martin really was Adolf Hitler. In this situation, participants also made many inferential errors, accepting such sentences as *He hated the Jews particularly and so persecuted them.* Because they did not know that Gerald Martin was Adolf Hitler until test, they must have made the inferences at test. Thus, it seems that participants do make such reconstructive inferences at time of test.

> *In trying to remember material, people will use what they can remember to infer what else they might have studied.*

Plausible Retrieval

In the foregoing analysis, we spoke of participants as making errors when they recalled or recognized facts that were not explicitly presented. In real life, however, such acts of recall often would be regarded not as errors but as intelligent inferences. Reder (1982) has argued that much of recall in real life involves plausible inference rather than exact recall. For instance, in deciding that Darth Vader was evil in *Star Wars,* a person does not search memory for the specific proposition that Darth Vader was evil, although it may have been directly asserted in the movie. The person infers that Darth Vader was evil from memories about the *Stars Wars* movies.

Reder demonstrated that people will display very different behavior, depending on whether they are asked to engage in exact retrieval or plausible retrieval. She had participants study passages such as the following:

> The heir to a large hamburger chain was in trouble. He had married a lovely young woman who had seemed to love him. Now he worried that she had been after his money after all. He sensed that she was not attracted to him. Perhaps he consumed too much beer and French fries. No, he couldn't give up the fries. Not only were they delicious, he got them for free.

Then she had participants judge sentences such as

1. The heir married a lovely young woman who had seemed to love him.
2. The heir got his French fries from his family's hamburger chain.
3. The heir was very careful to eat only healthy food.

The first sentence was studied; the second was not studied, but is plausible; and the third neither was studied nor is plausible. Participants in the exact condition were asked to make exact recognition judgments, in which case the correct responses would have been to accept the first sentence and reject the second two. Participants in the plausible condition were asked to judge whether the sentence was plausible given the story, in which case the correct responses would have been to accept the first two and reject the last. Reder tested participants immediately after studying the story, 20 min later, or 2 days later.

Reder was interested in the response time for participants in the two conditions, exact versus plausible. **Figure 7.9** shows the results from her experiment, plotted as the average response times as a function of delay. As might be expected, participants' response times increased with delay in the exact condition. However, the response times actually decreased in the plausible condition. They started out slower in the plausible condition than in the exact condition, but this trend was reversed after 2 days. Reder argued that participants respond more slowly in the exact condition after a lengthy delay because the memories of exact wording are getting weaker. A plausibility judgment, however, does not depend on any particular memory and so is not similarly vulnerable to forgetting. Participants respond faster in the plausible condition with delay because they no longer try to retrieve facts, which often have been forgotten, but instead immediately begin making inferences on which to base plausibility judgments.

In another experiment, Reder and Ross (1983) compared exact versus plausible judgments after participants had studied sentences such as

Alan bought a ticket for the 10:00 a.m. train.
Alan heard the conductor call, "All aboard."
Alan read a newspaper on the train.
Alan arrived at Grand Central Station.

Figure 7.9 Exact Retrieval vs. Plausible Retrieval. In Reder (1982), participants in the exact condition had to recognize whether presented information exactly matched the information studied, whereas participants in the plausible condition had to judge whether presented information was plausible based on the information studied. Exact judgments were made faster than plausible judgments immediately after study, but slower after a delay of 2 days. *(Data from Reder, 1982.)*

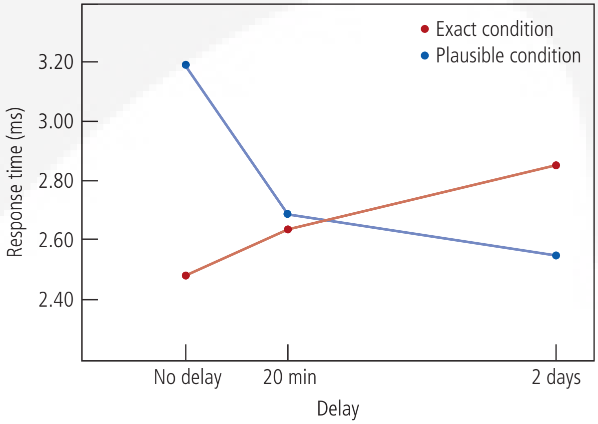

They manipulated the number of sentences that participants had to study about a particular person such as Alan. Then they looked at the times participants took to recognize sentences such as

1. Alan heard the conductor call, "All aboard."
2. Alan watched the approaching train from the platform.
3. Alan sorted his clothes into colors and whites.

In the exact condition, participants had to judge whether the sentence had been studied. So, given the foregoing material, participants would accept sentence 1 and reject sentences 2 and 3. In the plausible condition, participants had to judge whether it was plausible that Alan was involved in the activity, given what they had studied. Thus, participants would accept sentences 1 and 2 and reject sentence 3.

In the exact condition, Reder and Ross found that participants' response times increased when they had studied more facts about Alan. This is basically a replication of the fan effect discussed earlier in this chapter. In the plausible condition, however, participants' response times decreased when they had learned more facts about Alan. The more facts they knew about Alan, the more ways there were to judge a particular fact to be plausible. Thus, plausibility judgments did not have to depend on retrieval of a particular fact.

> *People will often judge what plausibly might be true rather than try to retrieve exact facts.*

The Interaction of Elaboration and Inferential Reconstruction

In Chapter 6, we discussed how people tend to display better memories if they elaborate the material being studied. We also discussed how semantic elaborations are particularly beneficial. Such semantic elaborations should facilitate the process of inference by providing more material from which to infer. Thus, we expect elaborative processing to lead to both an increased recall of what was studied and an increase in the number of inferences recalled. An experiment by Owens, Bower, and Black (1979) confirms this prediction. Participants (college students divided into two conditions, a theme condition and a neutral condition) studied a story that followed the principal character, a college student, through a day in her life: making a cup of coffee in the morning, visiting a doctor, attending a lecture, shopping for groceries, and attending a party. The following is a passage from the story:

> Nancy went to see the doctor. She arrived at the office and checked in with the receptionist. She went to see the nurse, who went through the usual procedures. Then Nancy stepped on the scale and the nurse recorded her weight. The doctor entered the room and examined the results. He smiled at Nancy and said, "Well, it seems my expectations have been confirmed." When the examination was finished, Nancy left the office.

The only difference between the two groups of participants was that those in the theme condition had read the following additional information at the beginning:

> Nancy woke up feeling sick again and she wondered if she really were pregnant. How would she tell the professor she had been seeing? And the money was another problem.

Participants in the theme condition characterized Nancy as an unmarried student who is afraid she is pregnant as a result of an affair with a college professor. Participants in the neutral condition, who had not read this opening passage,

TABLE 7.4 The Interactive Effects of Elaboration and Inference

	Number of Propositions Recalled	
	Theme Condition	Neutral Condition
Studied propositions	29.2	20.3
Inferred propositions	15.2	3.7

(Data from Owens, Bower, & Black, 1979.)

had no reason to suspect that there was anything special about Nancy. We would expect participants in the theme condition to make many more theme-related elaborations of the story than participants in the neutral condition.

Participants were asked to recall the story 24 hours after studying it. Those in the theme condition introduced a great many more inferences that had not actually been studied. For instance, many participants reported that the doctor told Nancy she was pregnant. Intrusions of this variety are expected if participants reconstruct a story on the basis of their elaborations. **Table 7.4** reports some of the results from the study. As can be seen, many more inferences were added in recall for the theme condition than for the neutral condition. A second important observation, however, is that participants in the theme condition also recalled more of the propositions they had actually studied. This indicates that the additional elaborations these participants made enabled them to recall more of the story.

We might question whether participants really benefited from their elaborations, because they also mistakenly recalled many things that did not occur in the story. However, it is wrong to characterize these inferences as errors. Given the theme information, participants were perfectly right to make inferences. In a nonexperimental setting, such as recalling information for an exam, we would expect these participants to recall such inferences as easily as they recalled material they had actually read.

> *When participants elaborate on material while studying it, they tend to recall more of what they studied and also tend to recall inferences that they did not study but made themselves.*

Eyewitness Testimony and the False-Memory Controversy

The ability to elaborate on and make inferences from information, both while it is being studied and when our recall is being tested, is essential to using our memory successfully in everyday life. Inferences made while studying material allow us to extrapolate from what we actually heard and saw to what is probably true. When we hear that someone found out she was pregnant during a visit to a doctor, it is a reasonable inference that the doctor told her. So such inferences usually lead to a much more coherent and accurate understanding of the world. There are circumstances, however, in which we need to be able to separate what we actually saw and heard from our inferences. The difficulty of doing so can lead to false memories that may have negative consequences; the Gargoil example in **Implications 7.2** represents only the tip of the iceberg.

Implications 7.2

How Have Advertisers Used Knowledge of Cognitive Psychology?

Advertisers often capitalize on our tendency to embellish what we hear with plausible inferences. Consider the following description of a portion of an old Listerine commercial:

> "Wouldn't it be great," asks the mother, "if you could make him cold proof? Well, you can't. Nothing can do that." [Boy sneezes.] "But there is something that you can do that may help. Have him gargle with Listerine Antiseptic. Listerine can't promise to keep him cold free, but it may help him fight off colds. During the cold-catching season, have him gargle twice a day with full-strength Listerine. Watch his diet, see he gets plenty of sleep, and there's a good chance he'll have fewer colds, milder colds this year."

A verbatim text of this commercial, with the product name changed to "Gargoil," was used in an experiment conducted by Harris (1977). After hearing this commercial, all 15 of his participants recalled that "gargling with Gargoil Antiseptic helps prevent colds," although this assertion was clearly not made in the commercial. The Federal Trade Commission explicitly forbids advertisers from making false claims, but does the Listerine ad make a false claim? In a landmark case, the courts ruled against Warner-Lambert, makers of Listerine, for implying false claims in this commercial. As a corrective action the court ordered Warner-Lambert to include in future advertisements the disclaimer "contrary to prior advertising, Listerine will not help prevent colds or sore throats or lessen their severity." They were required to continue this disclaimer until they had expended an amount of money equivalent to their prior 10 years of advertisements.

ballyscanlon/Getty Images

One situation in which it is critical to separate inference from actual experience is in eyewitness testimony. It has been shown that eyewitnesses are often inaccurate in their testimony, even though jurors accord it high weight. One reason for the low accuracy is that people confuse what they actually observed about an incident with what they learned from other sources. Loftus (Loftus & Zanni, 1975; Loftus, Miller, & Burns, 1978) showed that subsequent information can change a person's memory of an observed event. In one study, for instance, Loftus asked participants who had witnessed a traffic accident about the car's speed when it passed a Yield sign. Although there was no Yield sign, many participants subsequently remembered having seen one, confusing the implication of the question they were asked with what they had actually seen.

The Loftus study provided an early example of errors in memory reports arising from confusion about the source of information, about whether information came from what was observed or from some other source. Identifying the source of one's memory is crucial for avoiding errors in eyewitness testimony (Davis & Loftus, 2017). Such errors, called *source monitoring errors* (Johnson et al., 1993), can affect eyewitness testimony in many ways. For instance, Brown, Deffenbacher, and Sturgill (1977) found that seeing a person's face when reviewing a series of mugshots can make participants in an experiment more likely to identify that person in a lineup as the person they saw committing a crime. Such misidentifications happen in reality as well as in experiments. For instance, a man named Walter Snyder was wrongly convicted of rape in 1985 and later exonerated on the basis of DNA evidence. The victim had seen his face in a mugshot and did not identify him at that time as the perpetrator, but upon actually seeing him months later she identified him as her rapist (Scheck, Neufeld, & Dwyer, 2000). Source monitoring errors can also arise when keeping straight what one actually observed versus what was reported in the media or by other witnesses, or even versus what was just imagined. For instance, over 40% of participants in a U.K. study (Ost, Vrij, Costall, & Bull, 2002) claimed to have seen a video of the car crash in which Diana, Princess of Wales, died, even though there is no such video.

Another kind of memory confusion — the so-called **false-memory syndrome** — has produced a great deal of controversy in cases where individuals claim to have recovered previously suppressed memories of childhood sexual abuse (Schacter, 2001). Many of these memories are recovered in the process of therapy, and some memory researchers have questioned whether these recovered memories reflect what really happened and have hypothesized that such memories might have been implanted as a result of therapists' suggestions. For instance, one therapist said to patients, "You know, in my experience, a lot of people who are struggling with many of the same problems you are, have often had some kind of really painful things happen to them as kids — maybe they were beaten or molested. And I wonder if anything like that ever happened to you?" (Forward & Buck, 1988, p. 161). Given the evidence we have reviewed about how people will put information together to make inferences that they then "remember," one could wonder whether patients hearing such suggestions from their therapist might remember what did not happen.

A number of researchers have shown that it is indeed possible to create false memories by use of suggestive interview techniques. For instance, Loftus and Pickerall (1995) had adult participants read four stories from their childhood written by an older relative — three were true, but one was a false story about being lost in a mall at age 5. After reading the story, about 25% of participants claimed to remember the event of being lost in a mall. In another study, Wade, Garry, Read, and Lindsay (2002) inserted an actual photo from the participants' childhood into a picture of a hot-air balloon ride that never happened (see **Figure 7.10**). Fifty percent of their participants then reported false memories about the experience. There is a lot of room for error in the process by which we distinguish between memory and imagination, and it is easy to become confused about the source of information. Of course, it would not be ethical to intentionally try to implant false memories about something as traumatic as childhood sexual abuse, and there are questions about whether it is even possible to create false memories about such awful events (e.g., Pope, 1996).

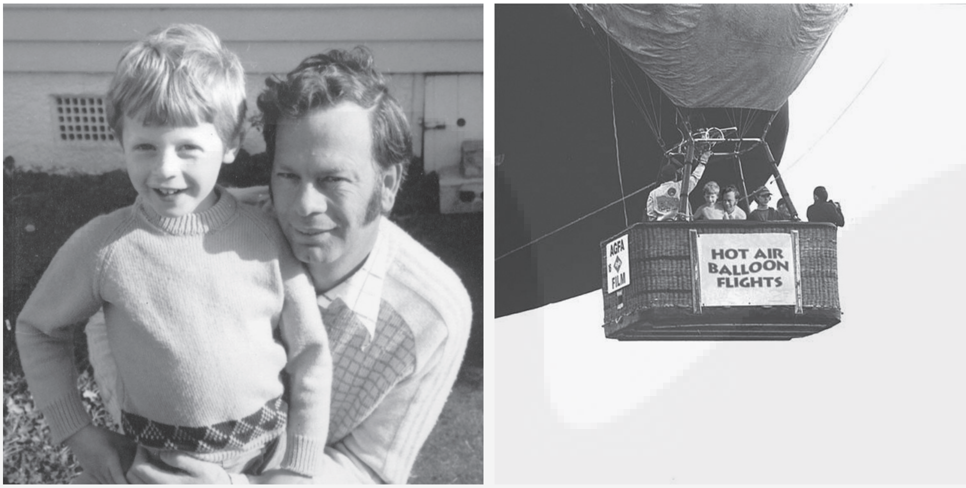

Figure 7.10 Remembering What Did Not Happen. The actual childhood photo on the left was embedded into the picture on the right to help create a false childhood memory of a balloon ride. *(Reprinted with permission from Kimberly Wade.)*

The intense debate about how much credibility should be given to recovered memories of childhood abuse might tempt one to conclude either that all reports of recovered memories of abuse should be believed or that all should be discounted, but the truth does not appear to be so simple. In some cases, recovered memories of abuse seem to be strongly supported by documentation (Sivers, Schooler, & Freyd, 2002), whereas in other cases people who claimed to have recovered such memories have subsequently retracted their claims and said they were misled in their memories (Schacter, 2001).

A similarly nuanced conclusion seems appropriate with regard to eyewitness testimony. While there are certainly cases where wrongful convictions were based on inaccurate eyewitness testimony, Wixted, Mickes, and Fisher (2018) argue that eyewitness testimony is often quite accurate. They argue that, just as major precautions need to be taken not to contaminate DNA evidence, similar efforts should be taken not to contaminate the memories of eyewitnesses. Eyewitnesses who have not had contaminating experiences, such as seeing photos in mugshots before seeing suspects in person, are typically very confident and very accurate in their identifications (Wixted & Wells, 2017). **Table 7.5** lists recommendations from Wixted and Wells for best practices in conducting a lineup, and some jurisdictions have begun to adopt such practices.

> *Serious errors of memory can occur when people fail to separate what they actually experienced from what they inferred, imagined, or were told. In cases with serious personal or legal consequences, care must be taken to avoid such contamination of memory.*

TABLE 7.5 Best Practices in Conducting a Lineup
1. Include only one actual suspect per lineup. 2. Make sure that the suspect does not stand out in the lineup. 3. Caution the eyewitness that the offender might not be in the lineup. 4. Use double-blind testing (i.e., neither the eyewitness nor those conducting the lineup should know who the suspect is in the lineup). 5. If the eyewitness makes an identification at the lineup, collect a confidence statement at that time (i.e., a statement by the eyewitness of how confident he or she is in the identification).
(Data from Wixted & Wells, 2017.)

False Memories and the Brain

Using less exotic paradigms than the one illustrated in Figure 7.10, researchers have developed the ability to explore the neural basis of false memories. In experiments based on the **Deese-Roediger-McDermott paradigm** — originally developed by Deese (1959) and elaborated by Roediger and McDermott (1995) — participants study lists of words and then have their brain activity monitored as they are tested on their recall of which words they studied. One such list of words might contain *thread, pin, eye, sewing, sharp, point, prick, thimble, haystack, thorn, hurt, injection, syringe, cloth, knitting;* a second list might contain *bed, rest, awake, tired, dream, wake, snooze, blanket, doze, slumber, snore, nap, peace, yawn, drowsy.* In the later test, participants are shown a series of words and must decide whether they have studied each word. There are three types of words:

True (e.g., *sewing, awake*) — words that were in the lists studied

False (e.g., *needle, sleep*) — words that are strongly associated with words in the lists studied but were not in those lists

New (e.g., *door, candy*) — words that were not in the lists studied and are unrelated to any of the words in those lists

Typically, participants say they have studied most of the true words and reject most of the new words, but they have difficulty rejecting the false words. For example, Cabeza, Rao, Wagner, Mayer, and Schacter (2001) found that 88% of the true words and only 12% of the new words were accepted, but 80% of the false words were also accepted — almost as high a percentage as the true words.

Cabeza et al. examined the activation patterns that these different types of words produced in the cortex. **Figure 7.11** illustrates such activation profiles in hippocampal and parahippocampal structures. In the hippocampus proper, true words and false words produced almost identical fMRI responses, which were stronger than the responses produced by the new words. Thus, these hemodynamic responses appear to match up pretty well with the behavioral data, where participants cannot discriminate between true words and

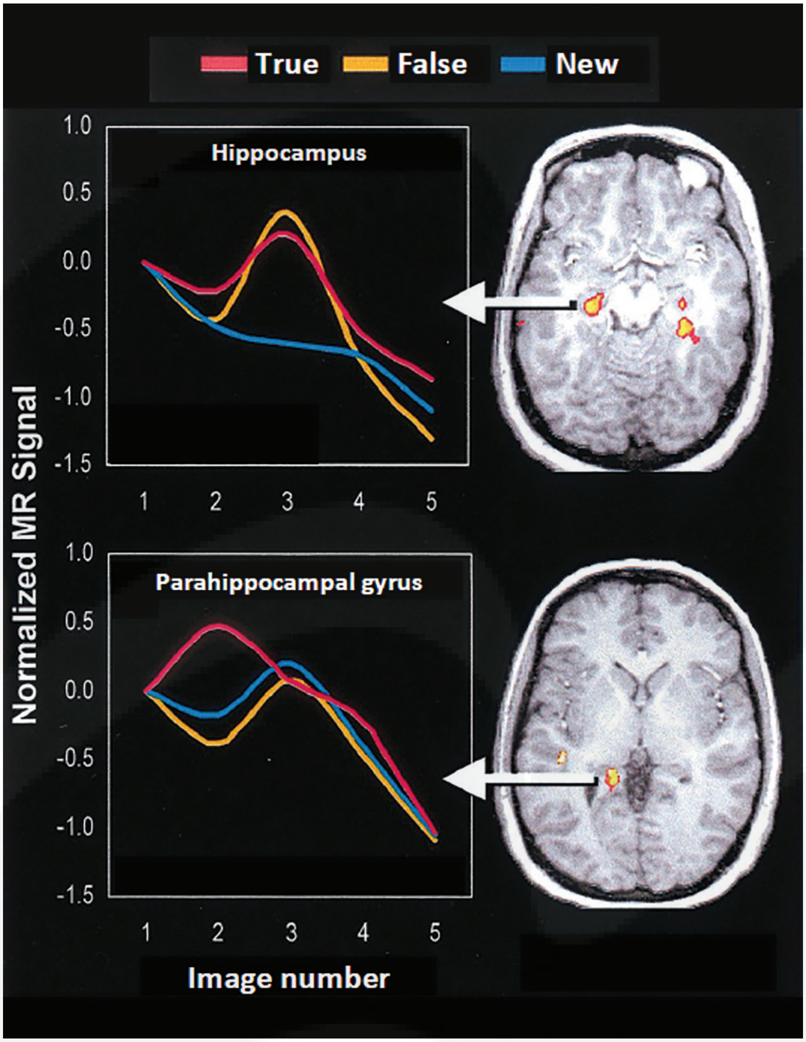

Figure 7.11 Brain Activation Patterns While Recalling Words. "True words" were in lists of words that participants studied; "false" words were not in the lists studied but were strongly associated with words in the lists; "new" words were not in the lists studied and were unrelated to any words in those lists. In the hippocampus (top), hemodynamic responses align with behavioral data: participants are almost equally likely to accept true words and false words as being in the lists studied, but reject new words. In contrast, in the parahippocampal gyrus (bottom), hemodynamic responses for both false words and new words were weaker than for true words, perhaps reflecting the parahippocampal area's closer connection to sensory areas. *(Cabeza et al., Can medial temporal lobe regions distinguish true from false?* Proceedings of the National Academy of Sciences Apr 2001, 98*(8) 4805–4810; Copyright (2001) National Academy of Sciences, U.S.A.)*

false words. However, in the parahippocampal gyrus, an area just adjacent to the hippocampus, both false and new words produced weaker responses than the true words. The parahippocampus is more closely connected to sensory regions of the brain than is the hippocampus, and Cabeza et al. suggested that the parahippocampus retains the original sensory experience of seeing the word, whereas the hippocampus maintains a more abstract representation, which would explain why true words produce a larger hemodynamic response in the parahippocampus. Schacter (e.g., Dodson & Schacter, 2002a, 2000b) has suggested that people can be trained to pay more attention to these distinctive sensory features and so improve their resistance to false memories. For instance, distinctiveness training could be used to help elderly patients who find it hard to remember whether they have seen something or just imagined it (Henkel, Johnson, & DeLeonardis, 1998).

> *Activation in the hippocampus is about the same for false memories and true memories, reflecting our difficulties in discriminating between what was experienced, what was inferred, and what was imagined.*

Associative Structure and Retrieval

The spreading-activation theory described in Chapter 6 implies that we can improve our ability to retrieve particular memories by prompting ourselves with closely associated memories. You may find yourself practicing this technique when you try to remember the name of an old classmate. You may prompt your memory with names of other classmates or memories of things you did with that classmate. Often, the name does seem to come to mind as a result of such efforts. An experiment by Tulving and Pearlstone (1966) provides one demonstration of this technique. They had participants learn lists of 48 words that included words from particular categories — for example, a list might include the words *dog, cat, horse,* and *cow,* which belong to the category of domesticated mammals. Participants were asked to try to recall all the words in the list, and they displayed better memory when they were given category names as prompts — for example, a prompt such as *mammal* would cue their memory for the words denoting members of that category.

The Effects of Encoding Context

Among the cues that can become associated with a memory are those from the context in which the memory was formed. This section will review some of the ways that such contextual cues influence memory.

Smith, Glenberg, and Bjork (1978) performed an experiment that showed the importance of physical context. In their experiment, participants learned two lists of paired associates on different days and in different physical settings. On day 1, participants learned the paired associates in a windowless room in a building near the University of Michigan campus. The experimenter was neatly groomed, dressed in a coat and a tie, and the paired associates were shown on slides. On day 2, participants learned the paired associates in a tiny room with windows on the main campus. The experimenter was dressed sloppily in a flannel shirt and jeans (it was the same experimenter, but some participants did not recognize him), and the paired associates were presented via a tape recorder. A day later, participants were tested for their recall of half the paired associates in one setting and half in the other setting. They could recall 59% of the list learned in the same setting as the one in which they were tested, but only 46% of the list learned in the other setting. Thus, it seems that recall is better if the context during test is the same as the context during study.

Perhaps the most dramatic manipulation of context was performed by Godden and Baddeley (1975). They had divers learn a list of 40 unrelated words either on the shore or 20 feet under the sea. The divers were then asked to recall the list either in the same environment or in the other environment. **Figure 7.12** displays the results of this study. Participants clearly showed superior memory when they were asked to recall the list in the same environment in which they studied it. So, it seems that contextual elements do get associated with memories and that memory is better when participants are provided with these contextual elements when being tested. This result actually has serious implications for diver instruction, because most of the instructions are given on dry land but must be recalled underwater.

The degree to which such contextual effects are obtained has proved to be quite variable from experiment to experiment (Roediger & Guynn, 1996). Fernandez and Glenberg (1985) reported a number of failures to find any context dependence; and Saufley, Otaka, and Bavaresco (1985) reported a failure to find such effects in a classroom situation. Eich (1985) argued that the magnitude of such contextual effects depends on the degree to which the participant integrates the context with the memories. In his experiment, he read lists of nouns to two groups of participants. In one condition, participants were instructed to imagine the referents of the nouns alone (e.g., imagine a *kite*); in the other, they were asked to imagine the referents integrated with the experimental context (e.g., imagine a *kite* on the table in the corner of the room). Participants were then tested for their recall of the nouns in one of two contexts, either in the same room where they had heard the nouns or in a different room. Eich found that participants had much better memory for the nouns that they had imagined as integrated with the experimental context, but only when participants were tested in that same context.

Bower, Monteiro, and Gilligan (1978) showed that emotional context can have the same effect as physical context. They instructed participants to learn two lists.

Figure 7.12 Context Effects on Recall of Words. Divers learned lists of words either on shore or underwater and were then tested for recall either on shore or underwater. Participants' recall was better when the environment at test matched the environment at study. *(Data from Godden & Baddeley, 1975.)*

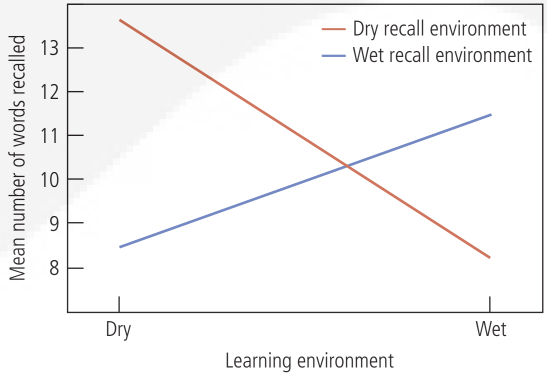

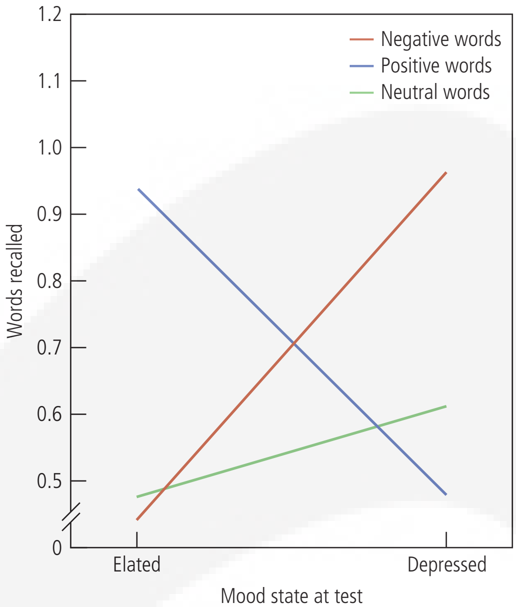

Figure 7.13 Mood Congruence. At study, participants in a neutral mood learned a list of positive, negative, and neutral words. Then, at test, participants were induced into either an elated or a depressed mood. Participants recalled more of the words that matched their mood at test. *(Data from Teasdale & Russell, 1983.)*

For one list, they hypnotically induced a positive state by having participants review a pleasant episode in their lives; for the other, they hypnotically induced a negative state by having participants review a traumatic event. A later recall test was given under either a positive or a negative emotional state (again hypnotically induced). Better memory was obtained when the emotional state at test matched the emotional state at study.[1]

Not all research shows such mood-dependent effects. For instance, Bower and Mayer (1985) failed to replicate the Bower et al. (1978) result. Eich and Metcalfe (1989) found that mood-dependent effects tend to be obtained only when participants integrate what they are studying with mood information. Thus, like the effects of physical context, mood-dependent effects occur only in special study situations.

While an effect of match between study mood and test mood is only sometimes found, there is a more robust effect called **mood congruence:** Memories are easier to recall when the emotional content of the memory matches the mood at recall. In other words, memories of events that made one happy are easier to recall when one is in a happy mood, and memories of events that made one sad are easier to recall when one is in a sad mood. Thus, mood congruence is an effect of the content of the memories rather than the emotional state of the participant during study. For instance, Teasdale and Russell (1983) had participants learn a list of positive, negative, and neutral words in a neutral mood. Then, at test, they induced either an elated or a depressed mood. Their results, illustrated in **Figure 7.13**, show that participants recalled more of the words that matched their mood at test. When a particular mood is created at test, elements of that mood will prime memories that share these elements. Thus, mood elements can prime both memories whose content matches the mood, as in the Teasdale and Russell experiment, and memories that have such mood elements integrated as part of the study procedure (as in Eich & Metcalfe, 1989).

A related phenomenon is **state-dependent learning.** People find it easier to recall information if they can return to the same emotional and physical state they were in when they learned the information. For instance, it is

[1] As an aside, it is worth commenting that, despite popular reports, the best evidence is that hypnosis per se does nothing to improve memory (see Hilgard, 1968; Lynn, Lock, Myers, & Payne, 1997; Smith, 1982), although it can help memory to the extent that it can be used to re-create the contextual factors at the time of test. However, much of a learning context can also be re-created by nonhypnotic means, such as through free association about the circumstances of the event to be remembered (e.g., Geiselman, Fisher, Mackinnon, & Holland, 1985).

TABLE 7.6 State-Dependent Learning: The Effects of Drugged State at Study and at Test

	Words Correctly Recalled at Test (%)		
At Study	*Ordinary Cigarette*	*Marijuana Cigarette*	*Average*
Ordinary cigarette	25	20	23
Marijuana cigarette	12	23	18

(Data from Eich et al., 1975.)

often casually claimed that when heavy drinkers are sober, they are unable to remember where they hid their alcohol when drunk, and when drunk, they are unable to remember where they hid their money when sober. In fact, some experimental evidence does exist for this state dependency of memory with respect to alcohol, but the more important factor seems to be that alcohol has a general debilitating effect on the acquisition of information (Parker, Birnbaum, & Noble, 1976). Marijuana has been shown to have similar state-dependent effects. In one experiment (Eich, Weingartner, Stillman, & Gillin, 1975), participants studied a list of words after smoking either a marijuana cigarette or an ordinary cigarette. Participants were tested 4 hours later — again after smoking either a marijuana cigarette or a regular cigarette. **Table 7.6** shows the results from this study. Two effects were seen, both of which are typical of research on the effects of psychoactive drugs on memory. First, there is a state-dependent effect reflected by better recall when the state at test matched the state at study. Second, there is an overall higher level of recall when the material was studied in a nonintoxicated state.

People show better memory if their external context and their internal states are the same at the time of study and the time of test.

The Encoding-Specificity Principle

Memory for studied material can also depend heavily on the context of the other material in which it is embedded at study and at test. A series of experiments (e.g., Tulving & Thomson, 1973; Watkins & Tulving, 1975) illustrate how memory for a word can depend on how well the test context matches the original study context. There were three phases to a typical experiment in this series:

1. *Study.* Participants learned pairs of words, such as *train–black*, and were told that they were responsible for remembering only the second word, referred to as the "to-be-remembered word" (in this example, *black*).

2. *Generate and recognize*. Participants were presented with words not studied and asked to generate four free associates to each word. The presented words were chosen to have a high probability of eliciting to-be-remembered words. For instance, participants might be presented with *white* and might generate the free associates *snow, black, wool,* and *pure* (*white* has a high probability of eliciting *black*). For each presented word, participants were asked to indicate which of the four associates they generated was a to-be-remembered word. In cases where a to-be-remembered word was generated, participants correctly chose it only 54% of the time. Because participants were always forced to indicate a choice, some of these correct choices must have been lucky guesses. Thus, true recognition was lower than 54%.
3. *Cued recall.* Participants were presented with the original context words (the first word in each pair studied — e.g., *train*) and asked to recall the to-be-remembered words (in this case, *black*). Participants recalled 61% of the to-be-remembered words — better than their 54% recognition rate without any correction for guessing. Moreover, Watkins and Tulving found that 42% of the words recalled in this phase had not been recognized earlier when the participants gave them as free associates.[2]

Recognition is usually superior to recall. Thus, in the experiment just described, we would expect participants' recognition rate in phase 2 to be better than their recall rate in phase 3; in particular, we would expect that participants who could not recognize a word would be unable to recall it (just as we expect to do better on a multiple-choice test than on a recall-the-answer test). Experiments such as this provided dramatic reversals of such standard expectations. The results can be understood in terms of the similarity of the test context to the study context. The test context in the generate-and-recognize phase was quite different from the context in which the words had originally been studied. The test context in the cued-recall phase, by contrast, matched the study context. This shows that, if contextual factors are sufficiently weighted in favor of recall, recall can be superior to recognition. Tulving interprets these results as illustrating what he calls the **encoding-specificity principle:** The probability of recalling an item at test depends on the similarity of the context during encoding at test to the context during encoding at study.

> *People show better word memory if the words are tested in the context of the same words with which they were studied.*

[2] A great deal of research has been done on this phenomenon (for a review, see Nilsson & Gardiner, 1993.)

The Hippocampal Formation and Amnesia

In Chapter 6, we discussed the fictional character Leonard from the movie *Memento*, who suffered amnesia resulting from hippocampal damage. A large amount of evidence points to the great importance of the hippocampal formation, a structure embedded within the temporal cortex, for the establishment of permanent memories. In animal studies (typically rats or primates; for a review, see Eichenbaum, Dudchenko, Wood, Shapiro, & Tanila, 1999; Squire, 1992), lesions in the hippocampal formation produce severe impairments to the learning of new associations, particularly those that require remembering combinations or configurations of elements. Damage to the hippocampal area also produces severe **amnesia** (memory loss) in humans. One of the most studied amnesic patients is known in the literature as H.M.[3] In 1953, when he was 27 years old, large parts of his temporal lobes were surgically removed to cure epilepsy. His epilepsy after the surgery was much less severe, but he suffered one of the most profound amnesias ever recorded and was studied for decades. He had normal memories of his life up to the age of 16 but forgot most of the subsequent 11 years preceding the surgery. Moreover, he was almost totally unable to remember new information and events. He appeared in many ways as a normal person with a clear self-identity, but his identity was largely as the person he was when he was 16, when his memories stopped (although he realized he was older, and he had learned some general facts about the world). His surgical operation involved complete removal of the hippocampus and surrounding structures, and this is considered the reason for his profound memory deficits (Squire, 1992).

Only rarely is there a reason for surgically removing the hippocampal formation from humans. However, humans can suffer severe damage to this structure and the surrounding parts of the temporal lobe. One common cause is a severe blow to the head, but other frequent causes include brain infections (such as encephalitis) and chronic alcoholism, which can result in a condition called **Korsakoff syndrome.** Such damage (including the damage associated with Korsakoff syndrome) can result in two types of amnesia: **retrograde amnesia** (loss of memory for events that occurred before an injury) and **anterograde amnesia** (an inability to store new things in memory).

In the case of a blow to the head, the amnesia often is not permanent but displays a particular pattern of recovery. **Figure 7.14** displays the pattern of recovery for a patient who was in a coma for 7 weeks following a head injury that did not require surgery. Tested 5 months after the injury, the patient showed total anterograde amnesia — he could not remember anything that had happened since the injury. He also displayed total retrograde amnesia

[3] Henry Gustav Molaison died in 2008 at the age of 82. There is an interesting discussion of him in the article "The Man Who Forgot Everything," by Steven Shapin (*New Yorker,* October 14, 2013).

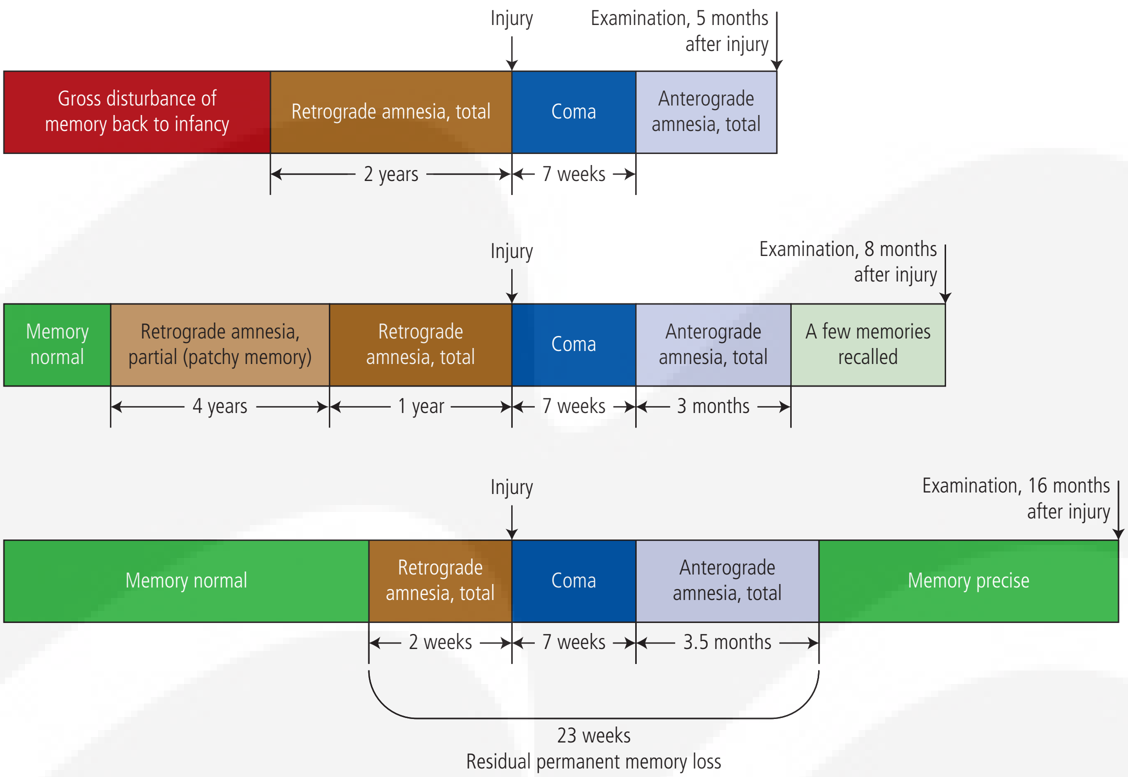

Figure 7.14 Recovery from Amnesia Caused by a Head Injury. This patient's time course of anterograde and retrograde amnesia over a period of 16 months following a head injury is typical. The severity of amnesia symptoms can vary among patients, but memories for events close in time to the injury (both before and after) are generally never recovered. *(Data from Barbizet, 1970.)*

for the 2 years preceding the injury and substantial disturbance of memory beyond that. When tested 8 months after the injury, the patient showed some ability to remember new experiences (but still had total anterograde amnesia for the 3 months following the injury), and the period of total retrograde amnesia had shrunk to 1 year. When tested 16 months after injury, the patient had full ability to remember new events and had only what proved to be a permanent 2-week period of total retrograde amnesia and a permanent 3.5-month period of total anterograde amnesia preceding and following the time of the injury (and of course total amnesia for the time in coma). It is characteristic that retrograde amnesia is for events close in time to the injury and that events just before the injury are never recovered. In general, anterograde and retrograde amnesia show this pattern of occurring and recovering together, although the severity of the retrograde and the anterograde symptoms can differ in different patients.

A number of striking features characterize cases of amnesia. The first is that anterograde amnesia can occur along with some preservation of long-term memories. This was particularly the case for H.M., who remembered many things from his youth but was unable to learn new things. The existence of such cases indicates that the neural structures involved in forming new memories are distinct from those involved in maintaining old ones. It is thought that the hippocampal formation is particularly important in creating new memories and that old memories are maintained in the cerebral cortex. It is also thought that events just prior to an injury are particularly susceptible to retrograde amnesia because they still require the hippocampus for support. A second striking feature of these amnesia cases is that the memory deficit is not complete: the patient can still acquire certain kinds of memories (this feature will be discussed in the next section, on implicit and explicit memory). A third striking feature of amnesia is that patients might remember things for short periods but then forget them. Thus, H.M. would be introduced to someone and told the person's name, would use that name for a short time, and would then forget it after a half minute. Thus, the problem in anterograde amnesia is retaining memories for more than 5 or 10 seconds.

> *Patients with damage to the hippocampal formation show both retrograde amnesia and anterograde amnesia.*

Implicit Versus Explicit Memory

Another famous case of amnesia involves the British musicologist Clive Wearing, who suffered herpesviral encephalitis that attacked his brain, particularly the hippocampus. His case is documented by his wife (Wearing, 2011) in *Forever Today: A Memoir of Love and Amnesia* and in the ITV documentary "The Man with a 7 Second Memory" (you can probably find videos by searching the Internet for "Clive Wearing"). He has nearly no memory for his past at all, and yet he remains a proficient pianist. Thus, while he cannot recall facts, he has perfect knowledge of all that is needed to play a piano. This illustrates the distinction between **explicit memory,** what we can consciously recall, and **implicit memory,** what our actions imply we remember in the absence of conscious memory.

While Clive Wearing is an extreme example, we all have implicit memories for things that we cannot consciously recall. However, because there is no conscious involvement, we are not aware of the extent of such memories. One example that some people can relate to is memory for the location of the keys of a computer keyboard. Many proficient typists cannot recall the arrangement of the keys except by imagining themselves typing (Snyder, Ashitaka, Shimada, Ulrich, & Logan, 2014). Clearly, their fingers know where the keys are, but they have no conscious access to this knowledge. Such implicit memory demonstrations highlight the significance of retrieval conditions in

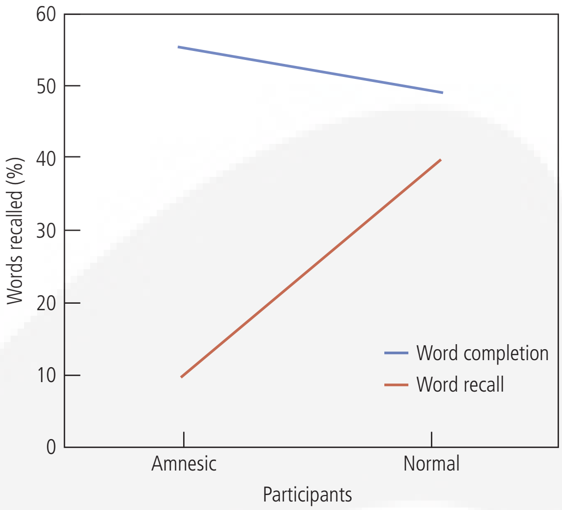

Figure 7.15 Word Recall and Word Completion in Amnesic versus Normal Participants. Both groups of participants studied lists of words and were then tested on their recall of the studied words. Then they were given the task of completing words when presented with the first three letters of studied words. Amnesic participants did much worse than normal participants on the word-recall task, but amnesic and normal participants were about equally likely to produce studied words in the word-completion task. *(Data from Graf, Squire, & Mandler, 1984.)*

assessing memory. If we asked typists to tell us where the keys are, we would conclude they had no knowledge of the keyboard. If we tested their typing, we would conclude that they had perfect knowledge. This section discusses such contrasts, or **dissociations,** between explicit and implicit memory. In the keyboard example above, explicit memory shows no knowledge, while implicit memory shows total knowledge.

A considerable amount of research has been done on implicit memory in amnesic patients. For instance, Graf, Squire, and Mandler (1984) compared amnesic versus normal participants with respect to their memories for a list of words. After studying the list, participants were asked to recall the words. As shown in **Figure 7.15**, amnesic participants did much worse than normal participants. Then participants were given a word-completion task. They were shown the first three letters of a word they had studied (e.g., *banana*) and were asked to complete it to make an English word (e.g., they might be asked to complete *ban-*). There is less than a 10% probability that participants will generate a word such as *banana* just given the prompt *ban-* without having studied the word, but the results show that participants in both groups came up with the studied word more than 50% of the time. There was no significant difference between the amnesic and normal participants in the word-completion task. Clearly, then, the amnesic participants did have memory for the word list, although they could not gain conscious access to that memory in a free-recall task. Rather, they displayed implicit memory in the word-completion task. The patient H.M. was also capable of implicit learning. For example, he was able to improve on various perceptual-motor tasks from one day to the next, although each day he had no memory of the task from the previous day (Milner, 1962).

Amnesic patients often cannot consciously recall a particular event but will show in implicit ways that they have some memory for the event.

Implicit Versus Explicit Memory in Normal Participants

A great deal of research (for reviews, see Richardson-Klavehn & Bjork, 1988; Schacter, 1987) has also looked at dissociations between implicit and explicit memory in normal individuals. It is often impossible with normal participants to obtain the dramatic dissociations we see in amnesic individuals, who can show no explicit memory but have normal implicit memory (as in the Graf et al., 1984, experiment discussed above). It has been possible, however, to demonstrate that certain variables have different effects on tests of explicit memory than on tests of implicit memory. For instance, Jacoby (1983) had participants respond to words under three different conditions: In the no-context condition, participants studied each word alone (e.g., *woman*); in the context condition, participants studied each word in the presence of an antonym (e.g., *man–woman*); and in the generate condition, participants would see a word and have to generate an antonym (e.g., see *man* and have to say *woman*). Participants were then tested in two ways, designed to tap either explicit memory or implicit memory. In the explicit memory test, participants were presented with a list of words, some studied or generated and some not, and asked to recognize the studied or generated words. In the implicit memory test, participants were presented with a studied or generated word for a very brief period (40 ms) and asked to identify the word. **Figure 7.16** shows the results.

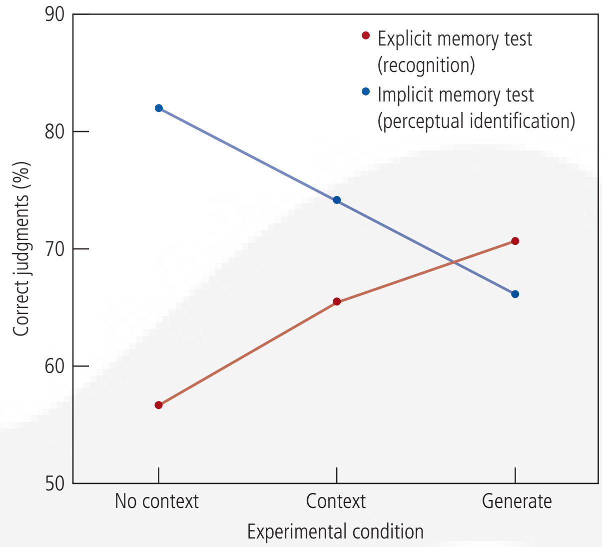

Figure 7.16 Explicit Memory versus Implicit Memory in Normal Participants. In the explicit memory test, participants were presented with a list of words and had to recognize which words had been studied (in either the context or no-context condition) or generated (in the generate condition). In the implicit memory test, participants had to identify studied or generated words on the basis of a very brief presentation (too brief for conscious recognition). These results indicate that explicit memory improves as the amount of elaborative processing increases (from no context to context to generate) but that implicit memory is best in the no-context condition, where perceptual processing is most prominent. *(Data from Jacoby, 1983.)*

Performance on the explicit memory test was best in the generate condition, which involved more semantic and generative processing — consistent with the research we have reviewed on elaborative processing. In contrast, performance on the implicit perceptual identification was worst in the generate condition and best in the no-context condition. All three conditions showed better perceptual identification than would have been expected if the participants had not studied the word at all (there is 60% correct perceptual identification for non-studied words). This enhancement of perceptual identification is referred to as **priming,** and Jacoby argues that participants show greatest priming in the no-context condition because that is the condition in which they had to rely most on perceptual encoding to recognize or identify the word.[4] In the generate condition, participants did not even have the word to read. Similar contrasts have been shown in memory for pictures: Elaborative processing of a picture will improve explicit memory for the picture but not affect implicit perceptual identification (e.g., Schacter, Cooper, Delaney, Peterson, & Tharan, 1991).

In another experiment, Jacoby and Witherspoon (1982) tested whether participants would display more priming for words they could recognize than for words they could not. Participants first studied a set of words. Then, in one phase of the experiment (involving explicit memory), participants had to recognize the words they had studied. In another phase (involving implicit memory), participants had simply to say what word they had seen after a very brief presentation. Participants showed better ability to identify the briefly presented words that they had studied than words

[4] Not all research has found better implicit memory in the no-context condition. However, all research finds an interaction between study condition and type of memory test. (For further discussion, see Masson & MacLeod, 1992.)

they had not studied. However, their identification success was no different for words they had studied and could recognize than for words they had studied but could not recognize. Thus, exposure to a word improves normal participants' ability to perceive that word (i.e., improves implicit memory), even when they cannot recall having studied the word (i.e., in the absence of explicit memory).

Research comparing implicit and explicit memory suggests that the two types of memory are supported rather differently in the brain. We have already noted that amnesics with hippocampal damage and normal individuals perform similarly in studies of priming, whereas such amnesics can show dramatic deficits in explicit memory. Research has shown that the drug midazolam — used for sedation in patients undergoing surgery — can produce similar deficits in normal patients (Victoria & Reder, 2010). For example, midazolam produces severe anterograde amnesia for the period of time it is in a patient's system, although the patient functions normally during that period (Polster, McCarthy, O'Sullivan, Gray, & Park, 1993). Participants given the drug just before studying a list of words showed greatly impaired explicit memory for the studied words but intact priming for these words (Hirshman, Passannante, & Arndt, 2001). Midazolam has its effect on neurotransmitters that are found throughout the brain but that are particularly abundant in the hippocampus and prefrontal cortex. Thus, the explicit memory deficits it produces are consistent with the association of the hippocampus and the prefrontal cortex with explicit memory. The drug's lack of implicit memory effects suggests that implicit memories are supported elsewhere in the brain.

Neuroimaging studies suggest that implicit memories are supported in the cortex. As we have discussed, there is increased hippocampal activity when memories are explicitly retrieved (Schacter & Badgaiyan, 2001). During priming, in contrast, there is often decreased activity in various cortical regions. For instance, in one fMRI study (Koutstaal et al., 2001), priming produced decreased activation in visual areas responsible for the recognition of pictures. (As discussed earlier in relation to fMRI studies, decreased activation reflects the fact that the brain regions responsible for the processing have to work less and so produce a weaker fMRI response).

A general interpretation of these results would seem to be that new explicit memories are formed in the hippocampus and prefrontal cortex; but with experience, this information is transferred elsewhere in the cortex. That is why hippocampal damage does not eliminate old memories formed before the damage. The permanent knowledge deposited in the cortex includes such information as word spelling and what things look like. These cortical memories are strengthened when they are primed and become more available in a later retest.

> *New explicit memories are built in hippocampal regions, but old knowledge can be implicitly primed in cortical structures.*

Procedural Memory

Implicit memory is defined as memory without conscious awareness. By this definition, rather different things can be considered implicit memories. Sometimes, implicit memories involve perceptual information relevant to recognizing words. These types of memories result in the priming effects we saw in experiments such as the one illustrated in Figure 7.16. In other cases, implicit memories involve knowledge about how to perform tasks. An important type of implicit memory involves **procedural knowledge,** the knowledge of how to perform a task such as riding a bike. Most of us have learned to ride a bike but, if asked, would have little ability to say explicitly what it is we have learned. Memory for such procedural knowledge is spared in amnesic individuals.

An experiment by Berry and Broadbent (1984) involved a procedural learning task with a more cognitive character than riding a bike. They asked participants to try to control the monthly output of a hypothetical sugar factory by manipulating the size of the workforce (all simulated by a computer program). Participants would see the factory's sugar output for the current month (in thousands of tons — e.g., 6,000 tons) and current workforce (in hundreds of workers — e.g., 500 workers) and then have to choose the next month's workforce (e.g., 700). They would then see the next month's sugar output with that new workforce (e.g., 8,000 tons) and have to pick the workforce for the following month (e.g., keep it at 700 or reduce to 600). The goal was to keep sugar production within the range of 8,000–10,000 tons. **Table 7.7** shows a series of participant interactions with the hypothetical sugar factory.

The relationship of sugar output to workforce is rule governed, but the series of interactions in Table 7.7 does not make the rule particularly obvious. The computer program used the following formula to relate the next month's sugar output in thousands of tons (S) to the size of the next month's workforce in hundreds of workers (W), given the current month's sugar output in thousands of tons (S_1):

$$S = 2W - S_1$$

In addition, a fluctuation of 1,000 tons of sugar was randomly added to or subtracted from S, and S and W were constrained to stay within the bounds of 1 to 12. Oxford undergraduates given 60 trials at trying to control the factory output got quite proficient at it. However, they were unable to state what the rule was and claimed they made their responses on the basis of "some sort of intuition" or because it "felt right." Thus, participants were able to acquire implicit knowledge of how to operate such a factory without acquiring the corresponding explicit knowledge. Amnesic participants have also been shown to be capable of acquiring this type of implicit knowledge (Phelps, 1989).

TABLE 7.7 Procedural Memory: A Series of Choices for Workforce Size in a Hypothetical Sugar Factory When Trying to Produce a Desired Output

Workforce Size (number of workers)	*Output (tons of sugar)*
700	8,000
900	10,000
800	7,000
1,000	12,000
900	6,000
1,000	12,000
1,000	8,000

Sequence learning (Curran, 1995) has also been used to study the nature of procedural memory, including its realization in the brain. There are a number of sequence-learning models, but in the basic procedure, a participant observes a sequence of lights flash and must press corresponding buttons. For instance, there may be four lights with a button under each, and the task is to press the buttons in the same order as the lights flash. The typical manipulation is to introduce a repeating sequence of lights and contrast how much faster participants can press the keys in this sequence than when the sequence is random. For instance, in the original Nissen and Bullemer (1987) study, the repeating sequence might be 4-2-3-1-3-2-4-3-2-1. After a sufficient number of repeats, people are faster with such a repeating sequence than when the lights come up in a random order. There has been much interest in the question of whether participants are aware that there is a repeating sequence. In some experiments, they are aware of the repetition; but in many others, they are not. They tend not to notice the repeating sequence when the experimental pace is fast or when they are performing some secondary task. However, participants are faster at the repeated sequence whether they are aware of it or not.

It does not appear that the hippocampus is critical to developing proficiency with the repeated sequence, because amnesics show an advantage for the repeated sequence, as do normal patients with pharmacologically induced amnesia. Rather, a set of subcortical structures, referred to as the *basal ganglia* (see Figure 1.8), does appear to be critical for sequence learning. It has long been known that the basal ganglia are critical to motor control, because damage to these structures produces the deficits associated with Huntington's and Parkinson's diseases, which are characterized by uncontrolled movements. However, there are rich connections between the basal ganglia and the prefrontal cortex, and it is now known that the basal ganglia are also important in cognitive functions. They have been shown to be active during the learning of a number of skills, including sequence learning (Middleton & Strick, 1994). Nonhuman primates are capable of mastering sequence tasks and have been used to study the neural basis of sequence learning. Such primate studies have shown that the basal ganglia are critical to early learning of a sequence. For instance, Miyachi, Hikosaka, Miyashita, Karadi, and Rand (1997) were able to impair early sequential learning in monkeys by injecting their basal ganglia with a chemical that temporally inactivated it. Other neural structures appear to be involved in sequence learning as well. For instance, similar chemical inactivation of structures in the cerebellum impairs later learning of a sequence. All in all, the evidence is pretty compelling that procedural learning involves structures different from those involved in explicit learning.

> *Procedural learning is another type of implicit learning and is supported by the basal ganglia.*

Conclusions: The Many Varieties of Memory in the Brain

Figure 7.17 shows the different types and subtypes of memory proposed by Squire (1987). The major distinction in his classification is between explicit memory and implicit memory, which he calls *declarative memory* and *nondeclarative memory*. It appears that the hippocampus is particularly important for the establishment of declarative memories. Within the declarative memory system, there is a distinction between episodic and semantic memory. Episodic memories include information about where and when they were acquired. For example, a memory of a particular newscast can be considered an episodic memory. This chapter and Chapter 6 have discussed these kinds of memories. Semantic memories, discussed in Chapter 5, reflect general knowledge of the world, such as what a dog is or what a restaurant is.

Figure 7.17 makes it clear that there are many kinds of nondeclarative, or implicit, memories. We have just completed a discussion of procedural memories and the critical role of the basal ganglia and cerebellum in their formation. We also talked about priming and the fact that priming seems to entail changes to cortical regions directly responsible for processing the information involved. There are other kinds of learning that we have not discussed but that are particularly important in studies of animal learning. These include conditioning and the two types of memory labeled *nonassociative* in Figure 7.17, habituation and sensitization, all of which have been demonstrated in species ranging from sea slugs to humans. Evidence suggests that conditioning in mammals involves many different brain structures (Anderson, 2000), as do many other types of learning, with different brain structures supporting different kinds of learning.

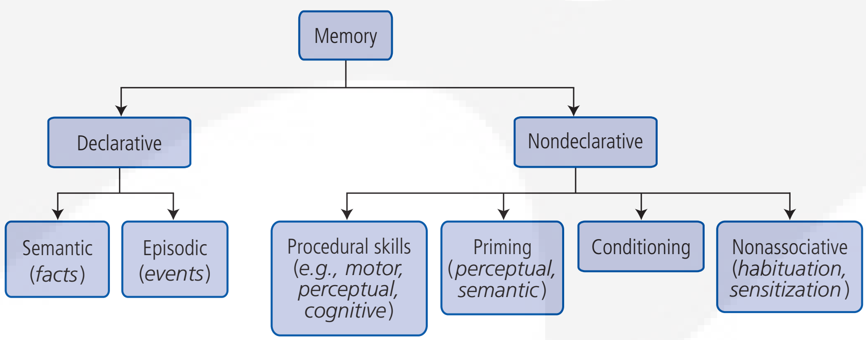

Figure 7.17 The Varieties of Memory Proposed by Squire. This classification system suggests that there are two memory systems (declarative and nondeclarative, corresponding to explicit memory and implicit memory), presumably supported by different brain structures. (*Research from Squire, 1987.*)

Questions for Thought

1. One of the exceptions to the decay of memories with time is the "reminiscence bump" (Berntsen & Rubin, 2002) — people show better memory for events that occurred in their late teens and early 20s than for earlier or later events. What might be the explanation of this effect?
2. The story is told about David Starr Jordan, an ichthyologist (someone who studies fish), who was the first president of Stanford University. He tried to remember the names of all the students but found that whenever he learned the name of a student, he forgot the name of a fish. Does this seem a plausible example of interference in memory?
3. Do the false memories created in the Deese-Roediger-McDermott paradigm reflect the same sort of underlying processes as false memories of childhood events?
4. It is sometimes recommended that students study for an exam in the same room that they will be tested in. According to Eich (1985), discussed earlier in this chapter, how would one have to study to make this an effective procedure? Would this be a reasonable way to study for an exam?
5. Squire's classification in Figure 7.17 would seem to imply that explicit and implicit memories involve different memory systems (declarative and nondeclarative) and, presumably, different brain structures. However, Reder, Park, and Keiffaber (2009) argue that the same memory system and the same brain structures are sometimes involved in both types of memories (ones of which we are consciously aware and others of which we are not). How could one determine whether implicit memory and explicit memory correspond to different memory systems?

Key Terms

amnesia
anterograde amnesia
decay theory
Deese-Roediger-McDermott paradigm
dissociations
encoding-specificity principle
explicit memory
false-memory syndrome
fan effect
implicit memory
interference theory
Korsakoff syndrome
mood congruence
power law of forgetting
priming
procedural knowledge
retrograde amnesia
state-dependent learning