Chapter 7

Long-Term Memory for Meaning and Knowledge

"The dissolving power of modern research seems to have split Memory into a number of variously related functions. Remembering is not … entirely distinct from perceiving, imaging, or even from constructive thinking, but it has intimate relations with them all."

F. C. Bartlett, *Remembering,* 1932–1967, pp. 12–13.

PREVIEW QUESTIONS

- *Is human long-term memory a unitary structure, or are there multiple types of memories?*
- *How is meaning represented in memory?*
- *How are simple sentences represented and understood?*
- *Are there representations in memory for large structures like maps or stories?*
- *What brain processes are involved in storing and retrieving meaningful information?*

7.1 Introduction to Long-Term Memory

We have memories for words, faces, music, and mathematical operations. We know that robins are animals, dandelions are plants, and that Los Angeles is in California. We also know what we did yesterday, and most of us know how to hit a ball with a bat and ride a bicycle. We might even remember who won the 2003 World Series or where we were when we heard about the September 11, 2001, World Trade Center and Pentagon attacks. How are these different kinds of knowledge learned, stored, accessed, and used when we need them? Surely there must be multiple types of memories, only some of which we can recall to consciousness. Other memories might affect our behavior, but they are likely to remain unconscious if we cannot trace their origin or verbalize much about them, such as our memory for how to ride a bicycle.

Psychologists have argued that there are differences between general knowledge (semantic memory) and memory for events (episodic memory), between verbal and visuo-spatial memories, and between declarative (knowledge-of) and procedural (knowledge-how-to) memories. Some memories seem to be the result of explicit attempts to attend to and learn information, whereas others seem to be acquired implicitly, without awareness or intention to learn (see Figure 7.1). Studies of brain-damaged individuals support some of these distinctions. Some studies have even suggested that different kinds of information are stored in different parts of the brain, as specific types of memories can be affected in patients with damage to certain brain areas, but other memories can be spared. For example, damage to parts of the temporal lobes can result in a patient's inability to recognize the faces of familiar people, but objects might be readily named. Other patients with different injuries might be able to recognize family members, but common tools might be unrecognizable. Imaging studies of brain activity in healthy people have corroborated these findings, as different parts of the brain are active when we learn

FIGURE 7.1

Taxonomy of Long-Term Memory for Meaning and Knowledge

Taxonomy of memory by Larry Squire (From Gazzaniga, Ivry, & Mangun, 2002).

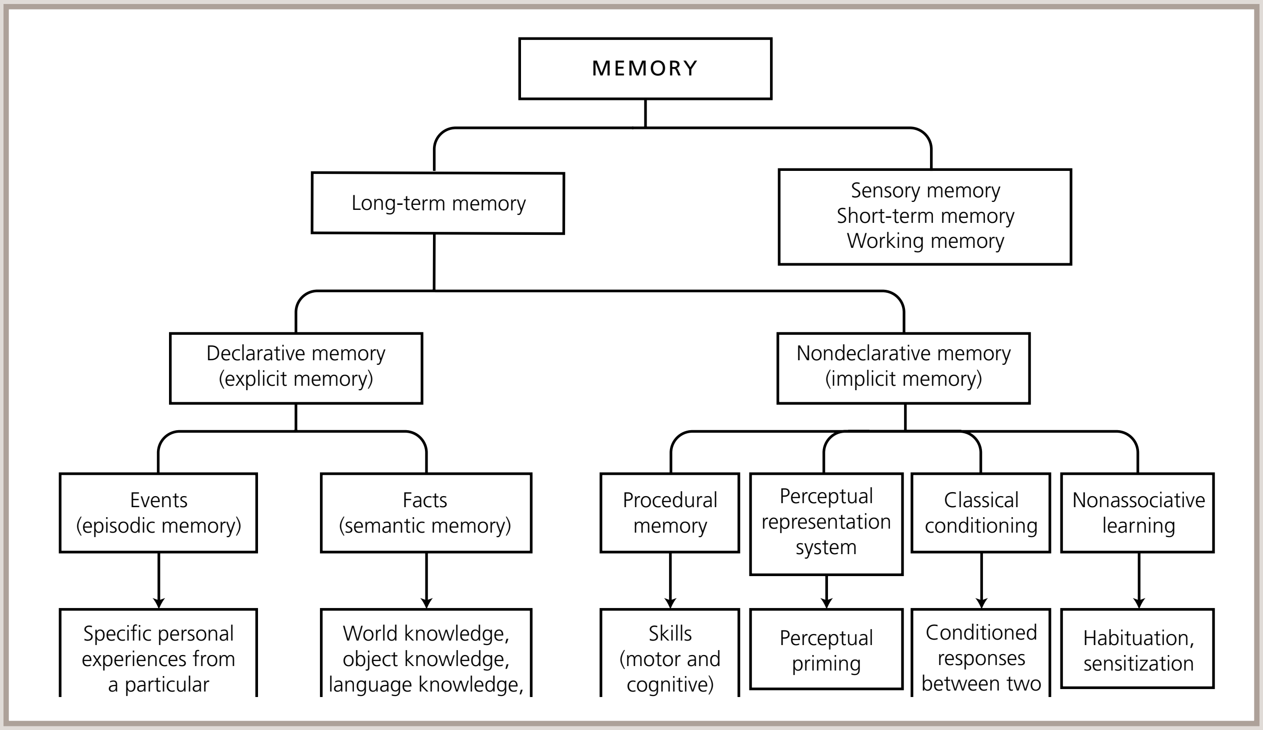

and remember different kinds of things. These results help us to develop and test theories of memory to explain why some things are easy to learn and other things are difficult, and why we remember some things well while others we seem to forget.

7.2 Semantic versus Episodic Memories

Over 40 years ago, Endel Tulving (1972) chaired a meeting of prominent memory researchers at the University of Pittsburgh and could not help noticing that they had aligned themselves into two groups. One group followed the tradition of Ebbinghaus (1885/1913) and studied how people learn and recall lists of verbal items, most commonly words that were specifically chosen to be unrelated to one another. The other group followed the tradition of Bartlett (1932/1967) and studied how people learn and remember meaningful sentences and stories. Tulving generalized beyond these differences to conclude that these two groups were in fact studying different kinds of memories, not just memory for meaningless versus meaningful collections of words. He was among the first to make the distinction between what he called "semantic memory" and "episodic memory."

According to Tulving's theory (1972), **semantic memory** is the general knowledge that we all have about the physical world and its symbolic representations, including language and mathematics. This common store of knowledge enables people to agree about how we perceive and identify objects and events and how we describe them to each other. We can witness a dog chasing a ball, and we will agree with one another about the basic perceptual facts of what a dog and a ball are. We are even likely to share inferences and attributions that are not directly perceptible, such that the ball belongs to the dog's owner and that the dog is enjoying the game of pursuit and retrieval. Both our memories of common events and our later discussions about them are likely to reveal the commonalties of our perceptions, interpretations, and recollections of everyday occurrences.

Semantic memory The permanent repository of all general knowledge in memory, according to Tulving's (1972) theory; the basis of understanding the meanings of objects, events, words, and sentences, as well as the basis of mathematical and encyclopedic knowledge

Episodic memory The personal memory of our experiences, according to Tulving's (1972) theory; it records events with date and setting information, and as such it is unique to each of us

Episodic memory, on the other hand, is the personal, autobiographical record of our daily experiences that is stored away in a type of individual memory journal. Rather than being general knowledge about dogs, balls, and what fun is, episodic memories are specific to particular experiences as witnessed by us. For example, we might remember the most recent time that we played fetch with a dog, whose dog it was, and where the event took place. Episodic memories tend to be unique to each one of us. However, despite the uniqueness of our individual experiences, we all live in a world with similar physical properties, similar animals and objects, and similar interactions among people, artifacts, and other living things. Therefore, semantic memories that store the general properties of our experiences and knowledge are likely to be similar across members of a common cultural and linguistic community, if not across all people generally.

If semantic memory is the part of long-term memory that contains all general knowledge, including mathematics, music, language, and familiarity with the three-dimensional physical world and the objects and living things that inhabit it, where does such knowledge come from? The obvious answer is that semantic memory is developed through some combination of biological predisposition and environmental experience, just as we learn to speak our native language and recognize members of our family. We cannot dismiss a large, inherited tendency to encode the world in a certain way, and this tendency also helps us to understand how we come to learn so much so fast. We will address some of these issues of nature versus nurture in Chapter 9 on language acquisition, and similar arguments can be raised to explain learning in other domains.

It is likely that semantic memory is built up through abstractions and generalizations from the set of experiences that make up episodic memory (Moscovitch et al., 2007). Indeed, nothing could be recognized at all if the currently encoded stimulus did not resemble a stored representation of a previous experience. Repeated experiences presumably result in stored traces, and a number of similar traces based on similar episodes are likely to be stored together in some sense. That is, the nervous system is hardwired to direct similar sensory experiences to similar parts of the brain. Repeated activations of brain structures due to repeated experiences are likely to produce changes in the brain itself (e.g., Hebb, 1949; Squire, 1987; captured in the phrase "neurons that fire together wire together"). These changes could include synthesis of new synapses and the development of large-scale interactions among different brain regions (e.g., Johnson, 2004). The repetition of familiar events promotes learning of general concepts that are central to the development of semantic memory. Any group of people living in a world with similar physical properties and filled with similar natural and artificial objects, who enjoy a common cultural and linguistic heritage, could well be expected to induce internal representations of their world that are more similar than different across individuals. That is, although the specific experiences that determine episodic memory will indeed differ across people, the abstracted generalities will be largely the same. This sameness is our shared semantic memory system.

7.3 Models of Semantic Memory

7.3.1 The Collins and Quillian Model

If semantic memory is an abstract representation of the commonalties of human experience, then it should include the processes and structures of human knowledge based on that experience. Any theoretical account of such memory structures and processes should detail the contents of the structures and how they are accessed, retrieved, and used in making judgments about real-world events. One of the first attempts to create such a theory, and also one of the first **network models** of human memory, was realized in Quillian's (1966) PhD dissertation, the "Teachable Language Comprehender." The model was later expanded (Collins & Quillian, 1969) and is today recognized as one of the first successful computer simulation models of human memory.

Network model A structural description of some domain (e.g., a part of semantic memory; Collins & Quillian, 1969) in which concepts are represented by nodes that are connected to other nodes via relational links; subsequent processing assumptions can include activation that spreads through the network from an input node along the links to other nodes in a diminishing wave of excitation

The structural components of the model included nodes, arranged into a hierarchy, that are accessed through a network of links connecting the nodes. Collins and Quillian chose to represent only a small part of semantic memory (i.e., some of the knowledge that we have about animals.) The nodes represent concepts, such as "animal" and "bird," and the links between them reflect relations among the concepts, such as subset/superset relations and property relations. Simple ideas can be represented by nodes and the associations between them. For example, propositions, such as "A bird is an animal," "A robin is a bird," and "A bird has feathers" are represented in a hierarchical network by two or more nodes and the relations that link them. There are two main types of relational links. Subset/superset relations define a concept's level in a hierarchy, to show, for example, that a robin is a type of bird and a bird is a type of animal. Other relations point to properties or features of the concept such as whether it can fly or has feathers. A schematic representation of such a model is shown in Figure 7.2.

The processing components of the model include search and decision processes. Search is simply an automatic spread of activation along the links from one node to others in the network that gradually diminishes as it extends out from its source like ripples in a pond around the splash of a stone. Activation can be regarded as a graded tendency for stored information to become available for further processing. Relatively inactive concepts remain below some threshold for influencing one's thoughts and behaviors, whereas highly activated concepts contribute to our current thoughts and perceptions. Activation is initiated either externally, as when we hear or see a word or encounter a familiar object, or internally, when we think of some concept. In either case, related concepts are activated to progressively lower levels as *distance* increases from the activated node. Distance is commonly equated with the number of links that have to be traversed between any two concepts in moving through the network. These mechanistic analogies should not be taken too literally, as the theory proposes only to represent a neural network in the brain, not to describe its actual structures and processes.

A decision is made whenever the search process is terminated or redirected. For example, Collins and Quillian (1969) originally tested their model by asking people to judge whether or not certain test sentences were true. They compared times taken by people to verify sentences of the form "A robin is a bird" versus "A robin is an animal." In either case, the search should terminate whenever the search process finds a link or links between the subject and predicate concept nodes. At that time, a decision could be made about whether or not the links existing in memory are consistent with the relation expressed in the test sentence. Since the sentence involving robins and birds is more direct than the one involving robins and animals, spreading

FIGURE 7.2

A Schematic Representation of Part of Semantic Memory as Modeled in the Theory of Collins and Quillian (1969)

Nodes representing concepts are arranged into a hierarchy with some links representing subset-superset relations (e.g., "a bird is an animal"). Other links identify concept properties (e.g., "a bird can fly" and "a fish has fins"). If the model is a good representation of human semantic memory, it should be able to predict how long people take to verify or refute simple sentences based on concepts and relations contained in the model.

From Goldstein, E. B. (2005). *Cognitive Psychology.* Belmont, CA: Thompson Wadsworth, p. 287.

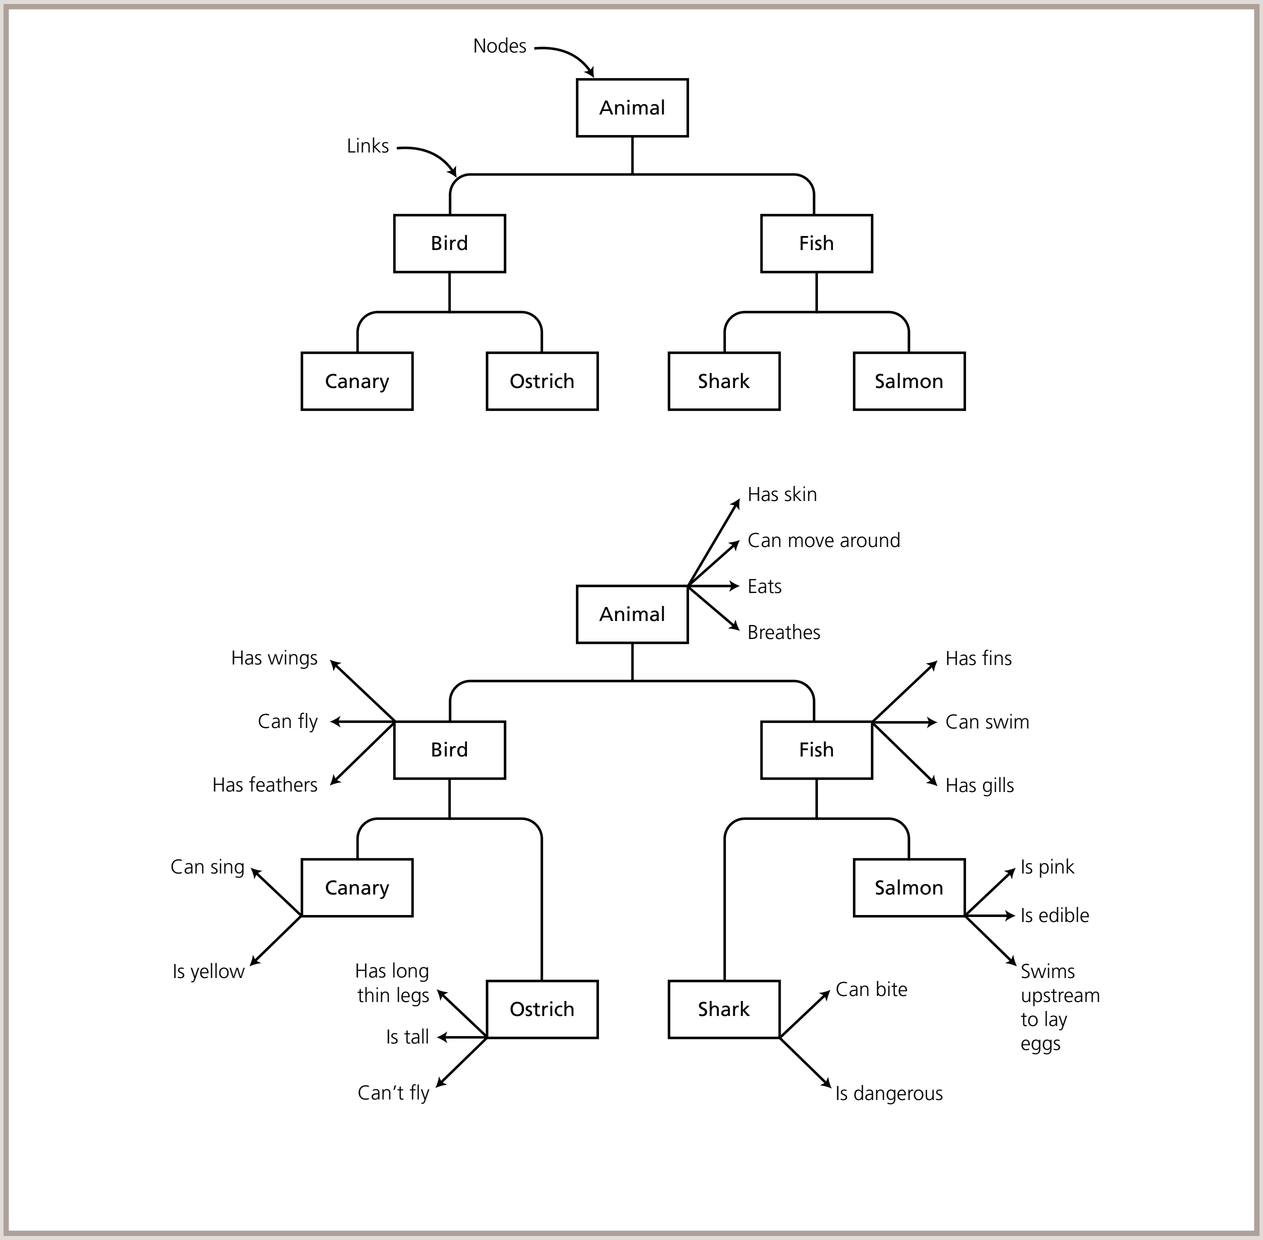

activation should make the search process shorter and the decision quicker for the first sentence. Similarly, a proposition such as, "A canary has skin" should be verified more slowly than "A canary has feathers" since all animals have skin, but only birds have feathers (see Figure 7.3).

FIGURE 7.3

Sample Results from Collins and Quillian's (1969) Experiment
People verified whether or not certain simple sentences were true. The data shown are mean response times for verifying sentences about properties (upper line) and sentences about category membership (lower line). In both cases, response times increased with the number of links that hypothetically should be traversed in a semantic memory network before the statement could be verified.
From Goldstein, E. B. (2005). *Cognitive Psychology*. Belmont, CA: Thompson Wadsworth, p. 289.

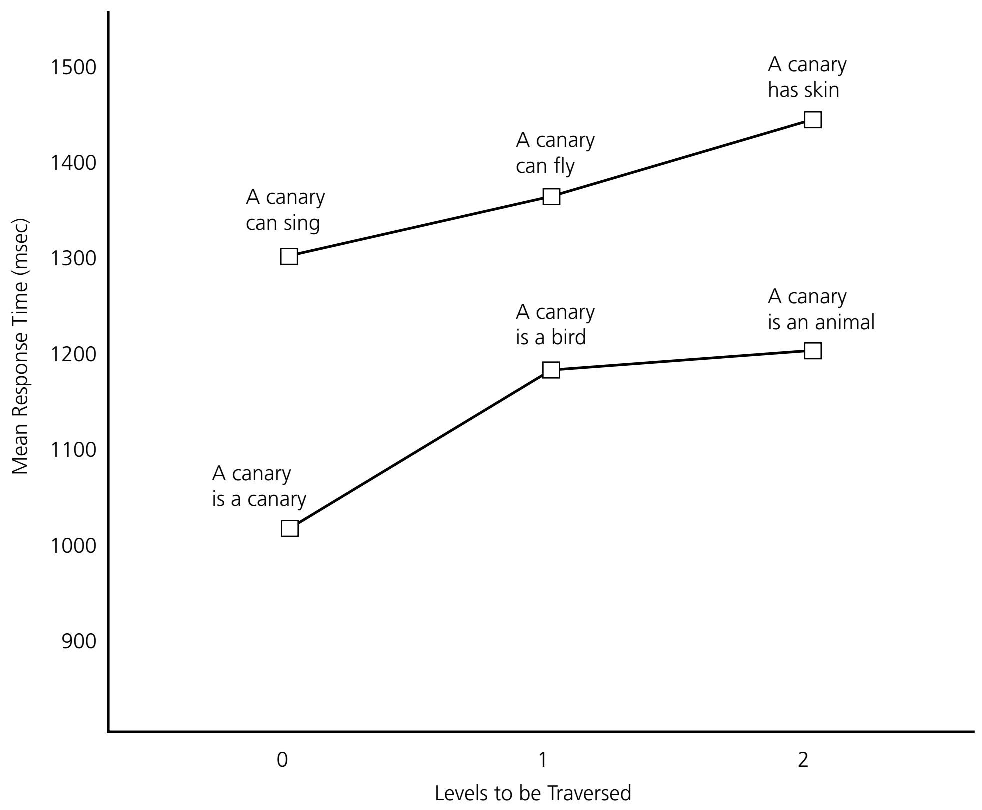

Cognitive economy A principle adopted by Collins and Quillian (1969) in their network model of semantic memory; asserts that any property that is true of more than one concept node in the hierarchy should be stored only at the highest-level (i.e., most general) node in the network over which the property is found

Collins and Quillian (1969) argued in favor of **cognitive economy** by asserting that it would make more sense to store general properties at their highest level of generality rather than at the level of every exemplar of a category that shares the property. That is, the property "skin" should be stored with animal, and the links between canary, bird, and animal would have to be traversed before the fact that canaries have skin could be evaluated as true. The proposition, "A canary has feathers" would involve only the links between canary, bird, and feathers before the proposition could be verified, thus the time should be shorter for determining that a canary has feathers than that it has skin.

These and other predictions were tested in experiments involving human subjects. The results showed consistencies between the structures and processes of the network model and the times that human subjects took to verify sentences. This demonstrated support for the Collins and Quillian theory. That is, by adding a certain amount of time for each stage of processing in the computer network, it was possible to simulate human behavior involved in understanding and verifying simple sentences.

7.3.2 Feature-Based Models of Semantic Memory

Of course, it was not long before problems were identified within the Collins and Quillian (1969) theory. For example Conrad (1972) disputed the generality of the principle of cognitive economy. It seems that unlike in an optimal computer network, properties in human memories are not stored only at the highest level of generality. Rather, they might be stored at multiple levels and with multiple specific concepts, especially if they are salient or prominent features of the concept (e.g., a peacock's feathers or a shark's teeth). One usually does not have to decide that a shark is an animal and that most animals have teeth before the fact that a shark has teeth comes to mind. Conrad's thesis presaged a more general attack on such network models by arguing in favor of distributed information storage as opposed to localized storage within semantic memory. That is, more recent neural network theories of semantic memory replace the notion that a concept node encapsulates all of our knowledge about some object. Rather, it is more likely that a variety of nodes and their interconnections code the meanings of common objects and events in a more realistic portrayal of the neural basis of human memory.

Other problems for Collins and Quillian's theory include the fact that all links in human memories are apparently not equal. Sentences such as, "A robin is a bird" are verified much faster than ones like, "A penguin is a bird," a phenomenon known as the **typicality effect** (e.g., Rosch, 1975). In general, typical members are identified as belonging to a category more quickly than are less typical members. This is true even if care is taken to make sure that the words themselves are matched in frequency of usage in English. The typicality effect has been accounted for in revisions of the theory by making links between typical items and their category nodes shorter or by allowing activation to spread more freely along them than along links from less typical category exemplars (e.g., Collins & Loftus, 1975). Further, people sometimes make errors—for example, in saying "true" to sentences such as, "A bat is a bird" or "A whale is a fish," or saying "false" to sentences such as, "A sponge is an animal." Often people realize such errors and try to correct them, but other times errors like these seem to be based on faulty knowledge represented in memory. These results indicate that there needs to be something more than nodes and links to account for such semantic similarity and typicality effects on sentence comprehension.

Typicality effect The fact that some members of a category are judged to be better exemplars of a category than are other members; also reflected in speed of judgments (e.g., "A fox is a mammal" is judged to be true faster than "A whale is a mammal")

An alternative to simple network models was proposed by Ed Smith and his students (e.g., Smith, Shoben, & Rips, 1974). Their model replaced nodes and links with a more homogeneous, multi-dimensional semantic space. Concepts are represented by specific locations in this hypothetical space, and their meanings by the values that each concept has on the various dimensions. One could imagine that related concepts would be located near one another, and that relative differences between concept meanings could be captured by their distances computed along the different relevant dimensions. To illustrate this notion, an example is shown in Figure 7.4 using familiar kinship terms. In this figure, the terms are differentiated along dimensions of gender, relative age, and relative closeness of relation. Arbitrarily putting the term *son* at the origin of this 3-D space, all eight kinship terms can be identified with different points in the space, with their meanings indicated by their respective locations. Further, distances are short between kinship concepts that are close, such as between son and sister and between father and mother, and greater distances reflect the lesser degree of kinship between, say, son and aunt. Of course, more general realizations of semantic memory would require many more dimensions of meaning and thousands of points within the space.

The feature model of Smith et al. (1974) can handle typicality effects, since robins should be located closer to the central concept of bird than penguins in the feature space. Also, the similarity effect, such as the fact that people are slower or more likely to make errors for sentences like "A bat is a bird" could be explained by the relative closeness of the bat and bird concepts in the feature space. However, it is unclear that distance between two concepts is sufficient to serve as the sole basis for the sentence verification process. To handle this problem, Smith and his colleagues differentiated between so-called characteristic features of concepts that influence their similarities, such as flying for both bats and birds, and defining features, such as laying eggs for birds versus giving birth for bats, which are important for the final decision about categorization processes.

Although the network representation in Figure 7.2 and the multidimensional feature space represented in Figure 7.4 are oversimplifications of parts of semantic

FIGURE 7.4

An Example of a Feature-Based Model of Part of Semantic Memory

A 3-D representation of English kinship terms, with dimensions of age, gender, and immediacy of family relations

The terms are defined relative to the "son."

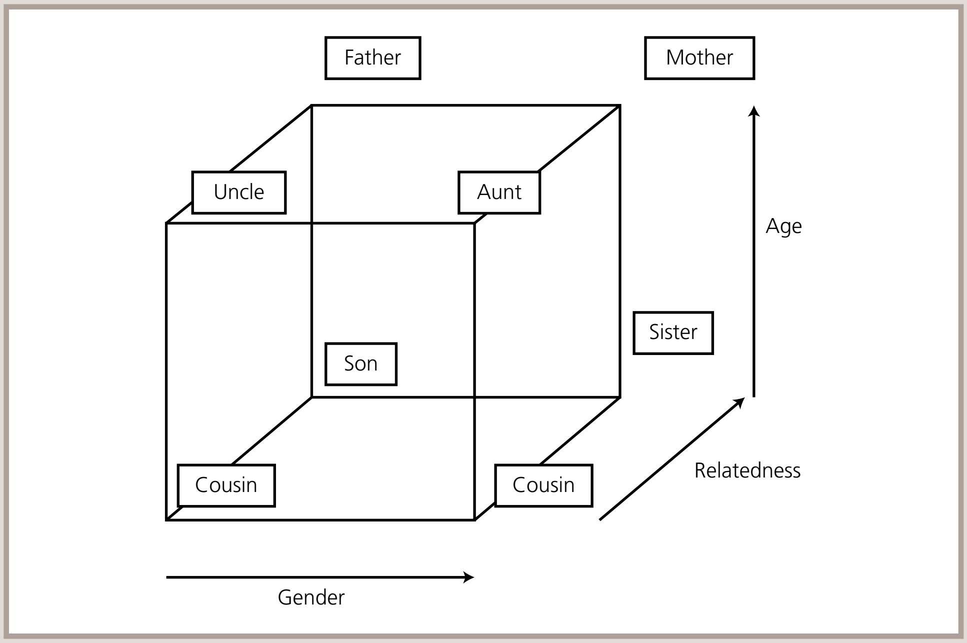

memory, the theories can be extended to arbitrary complexity in principle. Networks can theoretically include nodes for thousands of different concepts with many different kinds of links among them. Similarly, a feature space could include dozens of dimensions of meaning with a range of values represented along each of them. Either model could represent millions of nuances of meaning. However, both models share a property that has come under increasing scrutiny as an unlikely candidate for a structural description of human memory, namely local storage of meaning. If the meaning of a concept were defined as a single node or the relations it has with other nodes, then a literal interpretation of the model would locate conceptual meaning within a few neurons and their links to others. Similarly, the **multidimensional feature model** would appear to locate meaning within a structure as a single point in space, again apparently equating it with a single neuron or small set of neurons in the brain. Although such interpretations may not have been intended by the theories' authors, they seem to follow naturally from the structural descriptions. Such localist ideas have been rejected on numerous grounds, including the facts that a concept's meaning can be accessed in many different ways (from any sense modality, for instance). Further, forgetting and cognitive aging effects show partial or intermittent recall of concepts and their meanings, along with graceful degradation of information storage in aging or injured brains.

Multidimensional feature model A model proposed by Smith et al. (1974) as an alternative to network; representing concepts, such as locations in a space, in which the dimensions are different aspects of meaning (semantic primitives), and features are the specific values that any concept has on the various dimensions

7.3.3 Distributed Model of Semantic Memory

The major alternative to such localist conceptions of semantic memory is a distributed model with connectionist architecture (see Box 3.2). Such neural network theories specify different layers of nodes—typically an input layer, an output layer, and one or more intermediate, or hidden, layers. Each node is connected to other nodes by links that have weights between –1 and +1. Negative weights are inhibitory, and positive weights are excitatory. The network "learns" by adjusting weights until the input results in an output that is acceptably close to the correct response after a large series of training trials with feedback. Memory is represented by the set of weights in the network when the system has reached an equilibrium state, thus the memory for a particular item, such as a concept, is distributed over a network rather than being identified with a small unit in the architecture.

Arguments supporting connectionist theories include their superficial appearance: the networks look like and behave somewhat like part of a human brain. Experiments with the computer networks (i.e., Monte Carlo simulations) are almost as useful as experiments with human subjects in testing theories about perception, learning, and memory (Quinlan, 1991). On the other hand, the networks are complex, and how they work often becomes as difficult to understand as the human mental processes that they are intended to model. Further, it is not certain that they produce unique solutions, or that the solutions are psychologically meaningful. Theorists, such as Pinker (2002), have argued that although "… connectionist networks can manage impressive displays of memory and generalization … they are simply too underpowered to duplicate more realistic feats of human intelligence like understanding a sentence or reasoning about living things" (p. 79). Also, researchers often find it easier to talk about important psychological phenomena, such as priming, discussed next, in the context of a simpler network of the Collins and Quillian (1969) type, than in the more abstract language of connectionism.

7.4 Priming in Semantic Memory

Natural events, including physical, social, and linguistic interactions, include structures and sequences of related items. That is, we tend to see tables and chairs together, people congregate in theatres and restaurants and behave in predictable ways, and adjectives tend to precede nouns in English. It is extremely likely from any reasonable theory of human memory that our internal representations of objects, people, events, and language should reflect these natural relations in the external world. Therefore, it should not be surprising that in semantic memory models, concepts in memory are linked to related concepts, and they can activate each other when any one of them is activated. The natural situation of perceiving related words or objects is mimicked in the laboratory by presenting pairs of stimuli that are either related in some way or unrelated. Then, the speed or accuracy of the response to the second member of the pair is measured. If responses are faster and more accurate for trials with related pairs than for unrelated pairs, the difference is a measure of the amount of **priming** produced by the related stimulus.

Priming The effect that any stimulus (prime) has on the speed or accuracy of responding to another stimulus (target) that is presented at the same time as the prime or later; effects on the target can either be beneficial (positive priming) or detrimental (negative priming) (Priming can also be perceptual —prime and target related in form—or conceptual—prime and target related in meaning or associated to each other.)

In a typical laboratory task involving words, the prime stimulus is a word presented before or at the same time as a target stimulus, and the target word requires a response of some kind. Tasks that have been used are categorization (is the second word an animal name?), naming (say aloud the second word), and making lexical decisions (is the second item a real word or not?), with response time (RT) being the primary measure. The usual result found is that related primes produce faster responses to the target (facilitation or benefits), whereas unrelated primes can result in a small amount of inhibition (costs) when compared with a no-prime control condition. In the context of network models, priming is explained by automatic spreading activation produced by the first item of the pair, and this activation makes related concepts more available in memory. If these activated concepts match the target, it is processed more rapidly, leading to higher levels of performance than in no-prime control trials. If the prime is unrelated to the target, costs can occur because of the diversion of memory activation to parts of semantic memory that are irrelevant to the target, resulting in dilution of resources needed to identify it.

Among the first studies of priming effects on access time for words in semantic memory was that of Meyer and Schvaneveldt (1971). They used a lexical decision task in which participants had to decide whether or not two letter strings were both words. That is, they were to say "yes" when both strings were words, and to say "no" if either or both of them were nonwords. They found that the time to say "yes" on trials when both strings were words was almost 100 ms shorter when the two words were related (e.g., nurse, doctor) than when they were unrelated (e.g., bread, doctor; see Figure 7.5).

FIGURE 7.5

Sample Stimuli and Mean Response Times (RT, in milliseconds) for the Types of Stimulus Pairs Used by Meyer and Schvaneveldt (1972).

The task is to say "yes" only if both letter strings are words and to say "no" otherwise. Responses were fastest if the stimuli were two words related in meaning.

"YES"	RT	"NO"	RT
nurse	855	nurse	1087
doctor		doctum	
bread	940	doctum	904
doctor		bread	
		doctum	884
		corday	

A similar result was reported by Neely (1976), who also used the lexical decision task to measure the speed of word recognition. In Neely's experiment, the prime was presented before the target by a variable amount of time, the so-called stimulus onset asynchrony, or SOA. That is, the onset of the prime occurred between 200 and 2000 ms before the onset of the letter string that was to be judged a word or nonword. Similar to Meyer and Schvaneveldt's results, related priming words produced facilitation, and unrelated priming words produced inhibition relative to a neutral prime condition ("XXXX" presented before the target). Priming was maximal at a prime-target SOA of about 600 ms. This result indicates that it takes on average just over half a second for spreading activation to maximally excite the concept represented by a related target word, and that activation spontaneously dissipates afterwards.

In another experiment, Neely varied the relation between the prime word and the target word or nonword letter string that followed it. He told his participants that, in the normal condition, (1) a prime word would be likely to be followed by a word from the primed category (e.g., "BIRD" followed by a bird name, or a nonword). In a somewhat perverse condition (2) a prime word would be likely followed by a word from some other specific category (e.g., "BIRD" followed by a name of a body part or a nonword).

The interesting condition is the one in which a category prime was usually followed by a word from another category. That is, "BIRD" was usually followed by the name of a body part, like "ARM," but sometimes it was followed by a word from the same category as the prime, like "CANARY," and of course, many trials included nonwords, like "BORT." In this case, Neely's subjects knew what category to expect, presumably by thinking something like, "Bird—oh—that means body part" after the cue. Although when "BIRD" was followed by a real word, it nearly always was the name of some body part, but on a few rare trials, it was actually followed by the name of a bird. Neely found that naturally related words that occurred only occasionally ("BIRD" followed by "CANARY," for example) showed priming at short SOAs (200 to 400 ms), and "reversed" priming effects ("BIRD" followed by a more commonly-occurring and expected example, "ARM") showed up only at longer SOAs (700 ms or more).

The importance of these results is that intentional swapping of a priming cue's meaning for another (BIRD to BODY PART) takes time, but eventually, the reinterpreted cue produces priming much like a naturally-associated category label and an exemplar. However, the translation process cannot prevent automatic priming from initially facilitating the recognition of a related word, even if it is unexpected. In summary, related words produce fast automatic priming, and unrelated but rule-based primes result in slow, voluntary priming. In general, priming helps us to deal with sequences of events that occur in physical actions as well as in normal sentences. That is because sequences in both contexts usually contain related items. These relations can be used to prepare (prime) us for upcoming words or events, such that they are interpreted and understood more rapidly than if no priming had occurred.

7.5 The Neuroscience of Semantic Memory

Since semantic memory lies at the heart of the meaning of things and events, it is not surprising that it is intimately linked with the semantic aspects of language. In fact, one of the most common ways in which semantic memory is activated is through internal, lexical representations of written and spoken words. Warrington (1975) was among the first to provide clinical evidence supporting Tulving's distinction between semantic and episodic memory systems. She reported data from three patients with progressive dementia due to neurological disease who showed specific semantic memory impairments. Whereas these patients suffered little loss in intellectual ability, perception, and expressive language, they demonstrated severe losses in understanding the meaning of objects represented by words or pictures. Many of these results can be summarized by describing the losses in terms of accessing stored semantic memories, not in losses of the actual representations themselves. Further, certain types of semantic information were more vulnerable than others [i.e., property relations ("a canary has feathers") were more impaired than were subset/superset relations ("a canary is a bird")]. Subsequent research with patients suffering from strokes or diseases, such as herpes encephalitis, have also found semantic dementias characterized by difficulty in producing names of objects represented by pictures, by verbal descriptions, or by categories from which patients were to cite examples. All of the patients were subsequently shown to have damage in the left temporal lobe of the brain (Hodges, Patterson, Oxbury, & Furnell, 1992; Martin, 2001).

These conclusions have been supported by neuroimaging techniques using non-patient subjects. A critical manipulation has been to present pictures of familiar versus meaningless, nonsense objects while the amounts of metabolic activity in different brain areas are measured. The idea is that both types of stimuli should activate brain regions associated with processing visual input, but only those areas involved with processing conceptual and verbal information about objects should be active for the real items. The research showed two major brain areas that are more active for the familiar objects, namely the left inferior frontal cortex (Broca's area) and the posterior temporal lobe. Temporal lobe activity was bilateral in some studies (e.g., Martin, Wiggs, Lalonde, & Mack, 1996) and more heavily left-lateralized in others (e.g., Zelkowitz, Herbster, Nebes, Mintun, & Becker, 1998). Similar studies showed greater temporal lobe activity for silently reading real words as opposed to consonant letter strings (Price, Wise, & Frackowiak, 1996). In a review of these results, Martin (2001) concluded that primary storage of semantic information about words and objects is in the left temporal lobe for most people, whereas the left frontal cortex is responsible for directing retrieval of information from semantic memory (see also Moscovitch et al., 2007; Wagner, Bunge, & Badre, 2004).

Further studies of brain-damaged individuals have shown that semantic difficulties can sometimes be very specific. For example, some patients show selective deficits for naming animals, while retaining the ability to name tools, whereas others show a reverse effect: greater difficulty for manufactured artifacts and less for living things (Forde & Humphreys, 1999). Other areas of the ventral part of the temporal lobe seem to be particularly involved in face recognition (e.g., Kanwisher, McDermott, and Chun, 1997), and damage there can result in prosopagnosia—the inability to recognize faces, even those of one's own family. These differences might be due to the fact that different parts of the temporal lobes are associated with categorical information about visual features (relevant for animals) and with information about functional use (relevant for tools). This conclusion is supported by evidence that patients who show difficulty in identifying animals also have difficulty in identifying other objects that are defined by small visual differences, such as precious stones (Haxby, 2004; Martin, 2001).

In a recent review of the neuroscience of semantic memory, Yee, Chrysikou, and Thompson-Schill (2013) have argued that there is evidence for at least two theories of semantic memory formation and use. One is domain-specific, in the sense that semantic categories of the Collins and Quillian (1969) type exist in a generalized, modality-free structure that has abstracted the general properties of living things, artifacts, tools, etc. and their individual exemplars. The other type of theories emphasize the link between sensory-motor based interactions with environmental objects and their resulting general properties as food items, tools, clothing, predators, etc. For example, different parts of the temporal lobe are activated more for animal terms (an area associated with visual information processing) and for tools (am area related to action information). Studies of brain-damaged individuals have supported the claim that sensory-motor areas in the brain are associated with at least some components of semantic memory, whereas other areas, particularly the anterior temporal lobe, are associated with more abstract categorical relations. The picture that emerges is that semantic memory is distributed over a variety of different brain areas that are differentially activated depending on the task environment (e.g., naming, describing a function, use in a sentence, etc.) in which a common object or its name are processed.

In summary, deficits associated with specific types of brain damage, and neuroimaging studies of non-patients have lead to similar conclusions. Many of the brain areas associated with production and comprehension of language are the same as, or closely linked with, areas involved in semantic memory. This relationship is consistent with the main function of language: the exchange of meaningful ideas among members of a linguistic community.

7.6 Interactions Between Semantic and Episodic Memory

Semantic and episodic memories interact in comprehending and remembering everyday events. Research has shown that semantic memory performance is better in tests of items that also have some associated episodic memories (e.g., Westmacott, Black, Freedman, & Moscovitch, 2004). Further interactions have been observed in abstract memories for structures larger than mere concepts and relations (i.e., from repeated experiences with complex events or event sequences). That is, it is believed that not only do the specific elements (e.g., concepts and relations) of semantic memory derive from experience with common objects and events, but larger memory structures are also derived from repeated experiences with common, everyday occurrences. In theory, common experiences can lead to higher-order structures called **schemata** or **scripts,** which serve both to predict and understand events in real time as well as to organize, store, and retrieve information in long-term memory. The term *schemata* (or schemas) usually refers to familiar structures, such as an internal map of your home town, whereas a script is used to refer to a familiar sequence of actions, such as what you might do to get ready for school or work in the morning. However, these structures can incur some costs, as when we "remember" things that simply did not happen, because they plausibly fit with familiar experience.

The idea of memory schemata derives from the work of Bartlett (1932/1967). He reported the first detailed studies of memory for folktales, stories, and pictures. The procedure that he used was to provide a single study episode followed by repeated memory tests hours, days, and years later. Perhaps the most famous of the materials

Schemata Theoretical structures in semantic memory that have developed to represent commonly-occurring objects and events, such as stories, maps or routes for familiar journeys, and layouts of familiar structures, such as faces, kitchens, buildings, and gardens

Scripts Theoretical structures in semantic memory that have developed to represent commonly-occurring sequences of events, such as dining out at a restaurant or attending a theater, lecture, or party he used was the Native American story, *The War of the Ghosts,* although he chose different kinds of materials that varied in their familiarity to the Cambridge students who participated in his research over 80 years ago (see Box 7.1).

BOX 7.1

Bartlett's *The War of the Ghosts*

F. C. Bartlett published an influential book called *Remembering: A Study in Experimental and Social Psychology* (1932/1967). In this book, he described a variety of ideas and experiments that he undertook to study some of the vagaries of memory evidenced by his colleagues and students at Cambridge University in England. One of the most famous examples concerns his use of a native North American folk-tale. In his studies he used the method of an initial reading (actually two readings at the subject's own pace) followed by repeated recollections at various intervals. The original version of the story, as translated by Franz Boas, reads as follows:

> *The War of the Ghosts*
>
> One night two young men from Egulac went down to the river to hunt seals, and while they were there it became foggy and calm. Then they heard war-cries, and they thought: "Maybe this is a war-party." They escaped to the shore, and hid behind a log. Now canoes came up, and they heard the noise of paddles, and saw one canoe coming up to them. There were five men in the canoe, and they said:
>
> "What do you think? We wish to take you along. We are going up the river to make war on the people."
>
> One of the young men said: "I have no arrows."
>
> "Arrows are in the canoe," they said.
>
> "I will not go along. I might be killed. My relatives do not know where I have gone. But you," he said, turning to the other, "may go with them."
>
> So one of the young men went, but the other returned home.
>
> And the warriors went on up the river to a town on the other side of Kalama. The people came down to the water, and they began to fight, and many were killed. But presently the young man heard one of the warriors say: "Quick, let us go home: that Indian has been hit." Now he thought: "Oh, they are ghosts." He did not feel sick, but they said he had been shot.
>
> So the canoes went back to Egulac, and the young man went ashore to his house, and made a fire. And he told everybody and said: "Behold I accompanied the ghosts, and we went to fight. Many of our fellows were killed, and many of those who attacked us were killed. They said I was hit, and I did not feel sick."
>
> He told it all, and then he became quiet. When the sun rose he fell down. Something black came out of his mouth. His face became contorted. The people jumped up and cried.
>
> He was dead.

Bartlett used stories like *The War of the Ghosts* in his research because they provided a number of useful features of theoretical interest for his research. First, the stories represented a fairly alien culture and social environment with respect to his Cambridge students. He was interested in what might happen to a story as it is transmitted from one culture to another. Second, he was interested in how apparently unconnected events might be interpreted and recalled by his subjects. Third, he thought that the supernatural occurrences and vivid imagery that were evoked by some descriptions in the stories might influence memory for them.

Bartlett used the method of repeated reproduction, in which the first attempt at a complete recall of the story was made within about 15 minutes after the two initial readings. Further recalls were attempted at irregular intervals, but typically included two weeks, six weeks, and further months and years after the initial readings with no additional study of the original material. Bartlett noted several changes that occurred in the initial recalls of the story, and these generally became more pronounced as the recall interval increased.

(1) Although the initial recalls tended to be fairly accurate, the stories were often shortened, mainly by omissions, and the vocabulary and grammar tended to be recalled in more modern forms.

(2) The title and proper names all dropped out either with the initial recall or soon afterward.

(3) There was a strong tendency to rationalize unusual or supernatural events, so that they were either omitted or explained in a more natural, satisfying way. Such inferences and explanations are generally added to make the stories more coherent.

(4) More rarely, elaborations and other constructions were offered, usually to give a more dramatic sense to the story and add to its interest and coherence.

In general, the successive reproductions from different readers tended to converge on acceptable, understandable, comfortable, and straight-forward narratives with their puzzling elements removed. Literal accuracy of recall was clearly a rare exception, with both language and content shaped toward the more familiar and conventional. Elaborations and constructive inferences were more present in later recall, indicating that they were added as part of the recall process itself, and not as part of the original attempt to understand the story. Eventually, the story was actually replaced with a constructed revision of it, and this construction tells us about as much about the reader's mental schemata, beliefs and biases as it does about the story on which it is based.

Bartlett discovered many basic facts about memories for episodes that are active research topics today. One is that immediate or short-term recall tends to be largely reproductive, with most errors being omissions. Also, early memory tests can show **hypermnesia**, in which repeated tests without additional study can sometimes produce better recall in the short term. Delayed recall tends to be more reconstructive, with major errors introduced by combining story information with related information in long-term memory. Further, reconstructive errors tended to **normalize** story recall (with regard to cultural norms), making the stories more modern and conventional, by adding distortions and including information that was not in the original story. Normalization processes are evident in that different people's delayed recalls tend to become more alike within a given community or culture than they were in their initial recall. That is, they tend to rationalize or explain odd and supernatural occurrences in the story according to their common internal components of semantic memory.

Hypermnesia The tendency in repeated recollections of a certain event or story to introduce things present in the original that were not recalled in earlier retellings

Normalize The tendency to change characters, settings, and events in recall of a story to make them more consistent with one's own knowledge, beliefs, and cultural norms

Bartlett introduced the term memory schemata to describe sources of changes and additions observed in people's memories for stories. A schema is an active organization of related past experiences, and as such, it is a higher-order structure in semantic memory than simple concepts, associations, and propositions. Any new experience is compared with and sometimes integrated into an existing schema. Often, the new information is modified in some way to explain away discrepancies between the input episode and its interpretation. Similar processes are known to happen in pattern and object recognition, as when we recognize a car or person based on an incomplete or distorted image (e.g., if the object of interest is partially occluded by other objects or viewed briefly as it passes by a small window). Similarly, Gestalt principles of organization tend to improve an internal image of a perceived object to make it appear smoother, more continuous, and more symmetric than it actually might be. Both perception and memory tend to regularize our experiences to make them more like the average of our past encounters with similar objects and events.

Walter Kintsch (1977), discussed memory processes that can cause differences between the actual stories, movies, novels, and real-world events experienced by people and their later recall. He discriminated among (1) *reproductive memory*—which is more-or-less verbatim, accurate recall for some aspects of the original, (2) *constructive memory*—which includes additional *inferences made during the original event* that are incorporated into our memories, and (3) *reconstructive memory*—which includes additional knowledge incorporated into the original after the fact and added as *elaborations during recall* of the original.

Inferences and elaborations presumably arise from relevant knowledge structures in semantic memory that are activated both when the original material is experienced, and later, when it is recalled. Bartlett's schemata are theoretical structures that can affect memory for stories primarily through their influence on story recall. Other structures could influence both initial encoding as well as later recall, as Kintsch suggested. Schank and Abelson (1977) used the term scripts to describe certain sequences of events that occur in common situations (a birthday party, dining out, going to the movies, attending a lecture, and many other familiar events). In their theory, scripts have *headers* that alert us to particular major components of the script. These include the initial situation or plan of action that prepares us for what we should encounter as we enter a restaurant, attend a social event, or enter a sports arena. They can also be the initial acts of a speaker or a movie that leads us to generate certain expectations about a theme, genre, or story line. A script has *frames* or *slots* that are subcomponents executed in a stereotyped order. Thus, when we enter a restaurant, we expect to be greeted by someone, then shown to a table, handed a menu, etc., all in a fixed order that we have learned to anticipate. The various events are slotted into the frames of the overlying script structure as we proceed through the components of the event. If certain expectations are unconfirmed or omitted, *default values* are assigned to occupy the frames unless they are later replaced by actual objects or events in the episode (see Box 7.2).

The benefits or higher-order structures in semantic memory, like schemata or scripts, are that they facilitate prediction and comprehension of frequently occurring events (and almost all events that we experience share components with past experiences). Scripts also provide a retrieval plan so that we can recall events as they have occurred. The costs of such structures include the occasional activation of an incorrect script, perhaps by failing to recognize an important header, or the choice of a default value when a particular event is omitted or the event includes an unusual entry. Bower, Black, and Turner (1979) presented their participants with short stories that included familiar scenarios, such as a visit to a doctor's office. They were later tested with sentences that were or were not included in the original story. When the data for new sentences were examined, they found that people were much more

BOX 7.2

An Example of Following a Script

This is an example of the use of a script to encode a common experience that resulted in an embarrassing mistake. I (JJ) was driving my car in Lawrence, Kansas, and needed to stop for gasoline. I pulled into a gas station that I had used many times, and stopped next to the pump, opened the lid to the gas tank, and removed the pump handle in order to fill up the tank. Unfortunately, the pump had not been cleared from the previous purchase, and I could get no gas. After a few failed attempts to get the attendant's attention, I went inside the shop and asked the attendant to clear the pump, which he did. I then filled the tank of my car and drove off.

My next stop was across the street at a grocery store, and I picked up a few items and returned to my car. I was somewhat surprised to see a police car parked next to mine, with a female officer looking in the windows. "Did you forget something?" she asked me. I thought for a moment before the horrible realization struck that I just might have driven off from the gas station without paying. I admitted it, and she said that she didn't think I was actually trying to get away with theft, since I had immediately parked just across the street. To be sure, however, she accompanied me back to the gas station.

That gave me the opportunity to offer a logical explanation for my absent-mindedness. "You see," I told her, "I am a psychology professor, and we believe that we have these structures called scripts in our heads that prepare us to do a certain number of things in a certain order when doing some routine task. One of the things you do at a gas station is to go inside to pay the attendant for the gas you pumped. Well, I had already gone inside once to ask the attendant to clear and turn on the pump, and when I had finished, I had already filled the slot in the script for going into the store to take care of the bill. Therefore, I drove off thinking that all the required acts had been taken care of." She was so impressed with my story that she stayed to make sure I paid the bill, in case I absent-mindedly followed some other script!

likely to say that such a sentence had actually been presented if it could logically be included in the script than if it were a new, non-script based sentence. That is, people were much more likely to say that they remembered a sentence like, "The patient removed his clothes," than one like, "There was a large plant in the waiting room," even though neither sentence or anything like it had actually been included in the story about a doctor's visit.

Results supporting the use of higher-order memory structures indicate that these structures have been incorporated into our memories because they generally serve useful purposes. That is, we live in an orderly world with numerous scenarios that follow a regular script. Just think about the routines you go through in getting up and getting ready for school or work in the morning, driving to work or school, parking and entering your building, encountering friends and co-workers, executing familiar tasks, carrying on a light-hearted, perhaps cliché-riddled, conversation, and all the other routine tasks that fill our days. The use of scripts and schemata can help us to deal with these situations by providing maps and action sequences that we can choose to follow. These guidelines enable us to interact with the environment in a more efficient way, as most details do not need to be processed to high levels if they can be correctly anticipated. More processing resources then should remain available for processing the unexpected anomaly or less typical events that occur and demand our attention, such as dangers on the road, or being asked a question. Unfortunately, as Kintsch (1977) and Bower et al. (1979) have shown, higher-order knowledge can also sometimes lead us astray. In some situations we can actually produce false memories that can be "remembered" with as great or even greater certainty than actual events if they are consistent with the relevant script. The idea of false memories, as well as forgetting of actual memories, is much more the province of episodic than of semantic memory, and they will be covered in the next chapter.

Another characteristic of semantic memory that distinguishes it from episodic memory is that it appears to be more permanent, and resistant to change. The most common changes in memory are learning and forgetting. These processes are slow developing in semantic memory, but the addition to, and apparent loss of information from, episodic memory is a very common aspect of our lives.

7.7 Non-Declarative Memories

Although much of what we mean when we talk about memory includes things like encyclopedic knowledge and recollection of events, it is clear that we learn and remember many other things that cannot be so easily be described. Squire (1992; see also Schacter & Tulving, 1994) and others have distinguished between **declarative** and **non-declarative** memories (see Figure 7.1). Declarative memories are things that we have learned and can retrieve with awareness. As such, we can describe them verbally, such as the meaning of a word or where we ate dinner last night. Both semantic and episodic memories are included within the set of declarative memories. Non-declarative memories, on the other hand, are often acquired implicitly, without necessarily an intention or even the awareness of learning. Such memories include sensory motor skills, such as riding a bicycle or catching a ball, and broad generalizations about specific experiences, such as biases, preferences, and judgments. Priming effects, such as repetition (or perceptual) priming (e.g., Biederman & Cooper, 1991) and semantic (or conceptual) priming (e.g., Meyer & Schvaneveldt, 1971), are also examples of changes that can occur in memory without our intention or awareness. Another difference between declarative and non-declarative memories is that the types of amnesia resulting from medial temporal lobe damage are largely restricted to declarative memories, whereas the retention and acquisition of procedural memories remain largely intact (Moscovitch, et al., (2007).

Declarative memories Things that are learned and retrieved with awareness and intention, and their retrieval promotes verbal descriptions of their contents; both semantic and episodic memories are declarative memories

Non-declarative memories Things that are learned and retrieved without awareness and intention, such as perceptual-motor skills, biases, habitual responses and value judgments

Chapter Summary

Human long-term memory is a rich repository of an amazing variety of information that seems to have no limits on storage capacity nor a proven means of information loss. In order to understand its complexity, psychologists have used the strategy of *divide and conquer* in order to limit investigations to manageable components. Tulving argued that a major subdivision of long-term memory should be semantic versus episodic memories. Semantic memory is the structure that enables us to summarize past experiences and interpret new ones based on abstracted commonalities of experience. It is the basis of understanding, symbolic thought, and communication, and of necessity, it is similar across individuals. On the other hand, episodic memory is the record of the individual experiences themselves, and is therefore unique to each of us.

Semantic memory has been modeled as a neural network of interconnected concepts and relations among them, as a multidimensional feature space with concepts defined as locations in the space, or as a parallel-distributed network of nodes and associated weights. Whatever the theoretical representation, it is clear that semantic memory is highly associated with language and is represented in the left temporal lobe of the brain for most of us. In any case, the theories need to account for basic facts, such as how we decide whether propositions expressed in simple sentences are true or false, how we understand and remember stories, and how we understand and interact with other people and objects in natural environments. Theorists have found the need to hypothesize the existence of high-order structures, such as schemata or scripts, to explain how easily we interact with common situations and remember them later. Such structures aid our understanding and memory because they are usually adequate architectures to fit with experience, but when the fit is not perfect, these structures can lead to misunderstanding of situations or errors in memory of them. Although most memory research has concentrated on declarative memories, such as episodic and semantic memory, researchers are becoming increasingly aware of the vast amount of information that can apparently be learned implicitly, without awareness and without verbalization during recall, such as the learning and retrieval of perceptual-motor skills.

Review Questions

- **Is human long-term memory a unitary structure, or are there multiple types of memories?**

  Although parsimony dictates that we should not adopt a theory more complicated than necessary to explain the existing data, it is clear that the data from research on human memory points to the existence of different kinds of memory systems. In the Atkinson-Shiffrin theory, a distinction had been made between sensory, short- and long-term memories. Within long-term memory, Paivio's dual-coding theory presumed the existence of different kinds of memories for verbal and visuo-spatial information. Tulving further argued that semantic and episodic memories demonstrate different structural and processing components. Finally, there is evidence that declarative knowledge which we deliberately learn and can describe verbally, is different from procedural knowledge, which we often learn implicitly and usually cannot describe verbally.

- **How is meaning represented in memory?**

  There is general consensus that meaning is represented in a type of semantic memory that is identified with left-temporal lobe processes in most people. Semantic memory theories include primary forms of representation, which entail the conceptual base of knowledge. This conceptual base has been represented alternatively as a simple network of concept nodes and relations among them, a multidimensional feature space in which concepts are locations in the space defined by the values that they have on the various dimensions, or a complex neural network in which concepts are represented by sets of weights distributed over a population of nodes existing at different levels. In any case, these primitive concepts must be combined in some way into propositions and higher-order structures called schemata or scripts. Propositions code simple ideas for expression or comprehension, whereas schemata and scripts code abstract generalizations of large structures or event sequences that have occurred with some regularity in our environments.

- **How are simple sentences represented and understood?**

  Simple sentences are formed by one or more propositions linked into a meaningful set of phrases. Let us consider a simple sentence of the form "A weasel is an animal." In order to understand this sentence, it is necessary to have some internal representation of the meanings of "weasel," "animal," and the relation ("is") between them. We can, for the moment, ignore the articles "a" and "an" as pertaining to individual noun categories. If the meanings can be thought of as locations in a semantic memory space (or activations in a web of interconnected neurons), there must be some way for determining the relation between them and whether this relation is consistent with the verb in the sentence. In a simple network model, the relations among concepts in the memory model (a weasel is a mammal, and a mammal is an animal) can be compared with the relation expressed in the sentence. If they are analogous, the sentence is understood by being consistent with existing knowledge. If there is an inconsistency, the sentence, as it is understood, is deemed to be false.

- **Are there representations in memory for large structures like maps or stories?**

  Researchers have provided evidence for the existence of large-scale structures, such as schemas and scripts in human semantic memory. The evidence has come from studies of story comprehension and recall, which indicate that people make predictable errors in their recalls after various retention intervals. The errors tend to be consistent with interpretations made at the time of reading as well as reconstructions made at the time of recall. Both are influenced by normalizing memory structures that shape perception and memory toward an ideal form based on sets of related experiences. Thus, in reading a story, we expect to be introduced to a setting and characters, the development of themes and problems that need resolution, and an eventual resolution and conclusion. That is, we expect stories to fit into a general scheme, or predictable structure. The same is true for maps, blueprints, diagrams, automobile dashboard controls, and apartment and home layouts. In addition, there are certain sequences of events that we have experienced many times that result in scripts being entered into semantic memory. We have gone to a restaurant many times and expect the sequence to include being greeted by a host or hostess, being shown to a table, being offered a menu and a drink, placing an order, receiving some food, eating it, paying for the meal, and leaving. Repeated sequences like this presumably lead to script-like structures that serve the dual purpose of helping us to understand each episode as it unfolds in time while we are experiencing it as well as to recall the events at a later time should we choose to do so. Thus, as Bartlett argued almost 80 years ago, such higher-order memory structures have been acquired because they help us to understand and remember experiences in our daily lives.

- **What brain processes are involved in storing and retrieving meaningful information?**

  Much research has shown that many parts of the brain are involved in encoding, storing and retrieving memories. Results from brain-damaged individuals and imaging studies of normal, control subjects have converged on painting the same complicated picture. Memory begins with sensory coding areas in the temporal lobes for hearing and the occipital lobes for vision. These sensory events result in perceptual experiences that are stored in many parts of the brain. From these, semantic memory systems seem to be localized in the left temporal lobe for most people, and spatial relations seem to be bilaterally represented in the parietal lobes. Much of the actual work in storing and retrieving information is performed in the frontal lobes, with some evidence to support the notion that the left temporal lobe is more involved with verbal materials, and the right temporal lobe is more involved with spatial materials.

Key Terms