HUMAN MEMORY AND COGNITION

Mark H. Ashcraft
Cleveland State University

SCOTT, FORESMAN AND COMPANY
Glenview, Illinois   Boston   London

To the three who matter most:

Mary Helen Homerin Ashcraft,
Jordan Mark Ashcraft, and
Laura Catherine Ashcraft.

Credit lines for copyrighted materials appearing in this work are placed in the "Credit" section at the end of this book.

Printed in the United States of America.
23456 KPF 9392919089

Library of Congress Cataloging-in-Publication Data

Ashcraft, Mark H.
Human memory and cognition.

Includes bibliographies and index.
1. Memory. 2. Cognition. I. Title.
BF371.A68 1988 153 88-15838
ISBN 0-673-18862-0

Chapter 6

SEMANTIC LONG-TERM MEMORY

- **Semantic Memory**
  - The Collins and Quillian (and Loftus) Model
  - Smith's Feature Overlap Model
  - Empirical Tests of Semantic Memory Models
- **Categorization, Concepts, and Prototypes**
  - Concept Formation
  - Natural Language Concepts
  - Concept Formation Reconsidered
- **Priming Semantic Memory**
  - Priming Tasks
  - Semantic Priming Studies
  - Extensions of the Lexical Decision Task
- **The Generality of Semantic Networks**
  - Networks Versus Feature Lists
  - Not for Nouns Only

> *It follows that what we began by calling the 'image,' or 'copy,' of the fact in the mind, is really not there at all in that simple shape, as a separate 'idea.' Or at least, if it be there as a separate idea, no memory will go with it. What memory goes with is, on the contrary, a very complex representation, that of the fact to be recalled* plus *its associates, the whole forming one 'object' . . . , known in one integral pulse of consciousness. (James, 1890, pp. 650–651)*

> *Semantic memory is the memory necessary for the use of language. It is a mental thesaurus, organized knowledge a person possesses about words and other verbal symbols, their meaning and referents, about relations among them, and about rules, formulas, and algorithms for the manipulation of these symbols, concepts, and relations. (Tulving, 1972, p. 386)*

> *Human concepts are probably . . . like hooks or nodes in a network from which many different properties hang. The properties hanging from a node are not likely to be all equally accessible; some properties are more important than others, and so may be reached more easily or quickly. In such a representation, going from one concept to another does not involve scanning a list, but rather activating a path via some property from one to the other. . . . Thus, a concept would be a set of interrelationships among other concepts . . . everything is defined in terms of everything else. . . . In this respect,* [semantic memory] *is like a dictionary. (Collins & Quillian, 1972, pp. 313–314)*

For me, news of semantic memory arrived in the fall of 1972, in the form of a book we used in Graduate Seminar. The book was *Organization of Memory*, edited by Tulving and Donaldson (1972). I had been raised on a diet of conditioning and learning experiments, so in many ways I wasn't prepared for the new possibilities represented by semantic memory research. I was critical of several of the semantic memory chapters in the book (hypercritical, in retrospect), complaining that this idea was vague or speculative, that there were no data to support that conjecture, and so on. My advisor always responded, "Yes, but look at the *scope* of these ideas. These people are taking a *big bite* into things psychology has been ignoring for a long time."

Slowly, it dawned on those of us in the class that this new stuff was interesting, *much* more interesting than the ideas we'd learned were the appropriate ones to think about. It began to sink in that "Hey, we can study what people *know*, and how they know it." Of course, it took us a while to figure out how to start researching "what people know." Furthermore, this new stuff was demanding—all of a sudden, we needed to know linguistics and computer programming, not just psychology. But it was exciting—there was a sense of "being there" as important things were happening, and being a part of those things. Whether or not cognitive psychology was a "revolution" in Kuhn's (1962) sense of that word, the new area of semantic memory certainly felt revolutionary to me. The nonsense syllable study I had done in college instantly became ancient history. And in the space of one year, students in Graduate Seminar switched from studying verbal learning and retention to studying human knowledge.

This chapter continues the description of human long-term memory that began in the last chapter. The present chapter is concerned with a rather different kind of long-term memory than that discussed in chapter 5, however. This chapter is about *semantic memory*, literally "memory for meaning," or to put it simply, knowledge. Semantic memory is our *permanent memory store of general world knowledge*, variously described as a thesaurus, a dictionary, or an encyclopedia. Semantic memory is where your knowledge of language and other conceptual information is stored. It is the permanent repository of information you use to comprehend and produce language, to reason, to solve problems, and to make decisions. Whereas episodic memory is a personal, autobiographical store, semantic memory is a generic storehouse of knowledge.[1] While your episodic memory differs substantially from mine, our semantic memories are thought to be largely similar—maybe not overwhelmingly similar in exact content, depending on our cultural backgrounds, but certainly similar in terms of structure and processes. Thus, to reuse an example from before, your episodic memory record of your mother's maiden name is different from mine, although we all share a highly similar concept in semantic memory, the concept of a maiden name. Likewise, my specific memories of college are quite different from yours, but we all share a general kind of knowledge about being a college student.

As Tulving (1972) noted, the first usage of the term *semantic memory* appears to have been in M. Ross Quillian's doctoral dissertation in 1966. Quillian set himself the task of programming a computer to understand language, that is, to answer a variety of questions in a reasonably human-like fashion, and to be able to paraphrase English text. The inspiration for this work came not from psychology, but instead from computer science and artificial intelligence (AI). Machine translation, as it was known, had been a long-standing goal in computer science, yet progress toward this goal had been surprisingly slow. The overly confident predictions of the 50s had failed to take into account a subtle yet important fact—even the simplest acts of human comprehension require vast amounts of knowledge, much more than the words to be comprehended would suggest. Thus, for computers to understand, answer questions, or paraphrase, they needed to have this kind of knowledge base. It wasn't enough to have dictionary definitions of words—computers needed to have extensive

---

[1] Hintzman's (1978) term *generic memory* is a better one than *semantic memory* for at least two reasons. First, it implies the notion of "general world knowledge" more accurately, just as "generic green beans" implies regular, ordinary green beans. Second, the term *semantic memory* might be understood incorrectly to refer exclusively to language-based memories, and thereby could exclude general world knowledge that is not word-based, for example, images (whether visual, kinesthetic, or otherwise), knowledge of numbers and mathematics, and so on. Nonetheless, the term *semantic memory* was used first, it's the one we have gotten in the habit of using, and it's probably useless to try to change now. Besides, much of the semantic memory research *has* been on truly language-based semantic factors. Other forms of long-term generic memories are being attributed to yet other divisions (e.g., Anderson's "procedural knowledge," 1982).

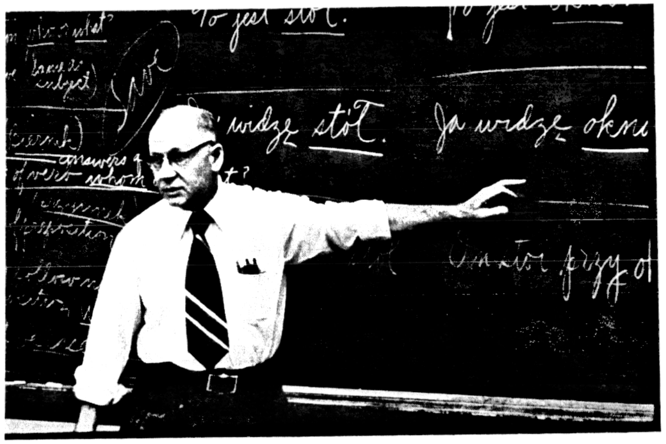

Semantic memory contains our long-term memory knowledge of the world, including our knowledge of words, concepts, and language.

computer

knowledge of the world in order to understand even simple sentences. This is often referred to as *tacit knowledge,* the implied (but not stated) knowledge that is necessary to understand what *is* stated or mentioned.

A clearcut illustration of the importance of tacit knowledge, with its implications both for machine translation and human comprehension, was contained in the Collins and Quillian (1972) chapter in *Organization of Memory.* Although lengthy, it expresses the scope of the issue quite well, and demonstrates the need for a memory containing semantic information:

> *At one time, I was trying to get a computer to be able to read sentences from pre-school children's books. My aim was to have the computer relate these sentences correctly to some body of information it had stored, its memory or "knowledge of the world." One such book, which described crossing a street, contained the sentences* [sic], *"The policeman held up his hand and the cars stopped." Now, suppose one asks what is the minimum amount of information a mechanism must have stored to relate this sentence to, if it is to comprehend it in a reasonably human-like way? In particular, consider whether the machine must have stored the fact that moving cars usually have drivers? One's first thought might well be no, since drivers aren't mentioned or directly involved in the sentence. But, suppose the sentences preceding this one in the book had said that there had just been an earthquake, and that two cars, parked on a hill, had started to roll down it. Then comes the sentence above, "The policeman held up his hand and the cars stopped." Virtually every adult reader of this will wonder: just how did the policeman manage that? In other words, in understanding the initial sentence, it seems that there indeed was some tacit use of the knowledge that cars ordinarily have drivers. If there were not, how can it be that, once a reader is led to believe that a moving car lacks a driver, he will then recognize that something is strange about a policeman being able to stop it just by holding up his hand? (pp. 327–328).*

Quillian's explicit point in this passage was that a computer must have a great deal of knowledge available in memory about the real world, even if it is only to comprehend a seemingly simple sentence. The implicit point in the quotation, and the point in the Collins and Quillian chapter as well, was that an adequate understanding of how *humans* comprehend language must take into account the same vast storehouse of knowledge that a computer would have to possess. The study of that vast storehouse is the study of semantic memory.

This is a big topic, so it will take us two chapters to cover all of the fascinating research that's been done on human semantic memory. In this chapter we'll cover the basics, the fundamental structures and processes investigated in semantic memory research. We'll ask questions like "How is the meaning of a word represented in memory?" and "How are word meanings retrieved from memory?" We'll pick up on some ideas that were introduced earlier in the book, especially the ideas of priming and automaticity. We'll also encounter the first wide-scale use of time as a measure of mental processes, *reaction time (RT),* to respond to simple sentences such as "A robin has wings." Finally, we'll consider two specific psychological models of semantic memory, theories that were advanced to explain what it is people know about words and word meanings, concepts, and their interrelationships. This will set the stage for the following chapter, where episodic and semantic long-term memory are studied together. There, we'll delve into the question of *interactions* between the episodic and semantic systems—for instance, how does our general world knowledge influence our memory for specific events? Throughout both chapters, the theme we will be most concerned with is the representation of knowledge, or as Kintsch (1974) put it in the title of his book, *The Representation of Meaning in Memory.*

▼ Semantic Memory

A study on "leading questions," also discussed in the next chapter, provides a convenient entry into the topic of semantic memory. Loftus and Palmer (1974) showed their subjects several short traffic safety films that involved car accidents. The subjects were asked to describe each accident after seeing the film, and then were asked to answer a series of questions. Unknown to them was that one question in the list, which asked for an estimate of the car's speed, was the critical item. As the authors pointed out, people are notoriously poor at estimating such factors, indicating that there might be some room for leading questions to have an effect. One group of subjects was asked "About how fast were the cars going when they hit each other?" The other four groups were asked virtually the same question, except that the verb "hit" was replaced with either "smashed," "collided," "bumped," or "contacted." As you might expect, subjects who got the stronger verbs like "smashed" in their questions gave higher estimates of speed—the question lead them to a biased answer.

Hold it. *Why* would we expect this effect? Why are we not surprised that people estimated higher speeds for the "smashed" question than the "bumped" or "hit" questions? Our intuitive answer here is something like "Well, 'smashed' implies a more severe accident than 'bumped.'" But, consider this intuitive answer again. *How* did Loftus and Palmer's subjects know that "smashed" implies a more severe accident? It's not enough merely to say that "smashed" *means* more severe; that's as unsatisfactory an answer as "I just know it" for the question "Does a robin have wings?" We're asking a more basic question than that. We want to know what is stored in memory that tells you what "smash" and "bump" mean. How is the difference between those two concepts *represented* in memory, and how do you *retrieve* those concepts when you encounter those words? What codes the fact that "smashed" implies a severe accident, that robins have wings, that moving cars have drivers, or that bananas, canaries, and daisies are all yellow? In short, what is the structure and content of semantic memory *per se,* and how do we access the knowledge stored in it?

As indicated at the beginning of this chapter, one of the earliest systematic attempts to answer such questions (aside from philosophical and strictly linguistic analyses) was Quillian's work in artificial intelligence (e.g., 1969). His model of semantic memory, TLC (for Teachable Language Comprehender), was not a genuine psychological model, but rather a computer program for understanding language.[2] Very shortly, however, Quillian began a collaboration with Allan Collins, and the psychological model they based on TLC became the first serious attempt in cognitive psychology to explain the structure and processes of semantic memory.

The Collins and Quillian (and Loftus) Model

The Collins and Quillian model of semantic memory (1972; also 1969, 1970, and Collins & Loftus, 1975[3]) was an extensive theory of semantic memory, comprehension, and meaning. At the heart of the model were

---

[2] Originally, there was a distinction between Artificial Intelligence and Computer Simulation; AI modeled human processes, but was relatively unconcerned that the modeling stick to known facts about human mental processes, whereas Simulation modeled human processes to discover more about how humans did various mental operations. Quillian's reseach fell into the category of Artificial Intelligence. For a variety of reasons, detailed in chapter 13, the original distinction has been largely abandoned; putting it simply, the AI approach doesn't work very well without a thorough knowledge of human mental processing.

[3] There were differences, of course, in the specifics of the models in these several papers, but for present purposes all are treated as roughly the same. Two issues about which the original model was vague did attract a great deal of experimental attention—the issue of cognitive economy and the representation of what was later referred to as typicality. Debates over these issues, and whether or not the original model did or didn't account for the facts (e.g., Collins & Loftus, 1975; Smith, 1978), tended to degenerate into hair-splitting, in my opinion, so are not treated in any depth here.

two fundamentally important assumptions, one about the *structure* of semantic memory, and one about the *process* of retrieving information from that structure. Because these two assumptions have been typical of all semantic network models since the early Collins and Quillian work (e.g., Glass & Holyoak, 1975), we'll take our time in explaining them, so you'll have a firm grasp of the ideas.

Nodes in a Network As you read in the opening quotations, Collins and Quillian viewed the entries in semantic memory (i.e., concepts) as being nodes in a network. In other words, the structure of semantic memory was said to be a **network,** *an interrelated set of concepts, or interrelated body of knowledge.* Each concept in the network is represented as a **node,** *a point or location* in the semantic system. Furthermore, concept nodes are linked together by **pathways,** *labeled, directional associations between concepts.* This entire collection—nodes connected to other nodes by pathways—is the network. Notice that every concept is related to every other concept in such a structure in the sense that some set of pathways, however indirect and long, can eventually be traced between any two nodes. (At one point, Collins and Quillian used the analogy of a large fishnet, where the knots correspond to nodes, and the strings that go from one knot to the next correspond to pathways.)

Spreading Activation The major process that operates on this structure is **spreading activation,** the *mental activity of accessing and retrieving information from this network.* Concepts are usually in a relatively quiet, unactivated state—they are accessible, but not currently being attended. For example, at this very instant in time, as you're reading this sentence, one of the many concepts in your semantic memory that is probably not activated is "machine." When you read that word, however, its mental representation receives a boost in activiation—"machine" is no longer quiet and unactivated, it's active or primed, "awakened" so to speak. This activation, for Collins and Quillian, *was* the process of retrieval, the process of accessing the meaning of a concept. An important feature of spreading activation is that once a concept becomes activated, the concept then begins to spread that activation to all the other concepts it is linked to. Thus, activation begins at a concept node, and then starts spreading throughout the network along the connecting pathways. In Collins and Quillian's particularly apt description, this spread of activation corresponds to a memory search, in which the "search continually widens like a harmless spreading plague" (1972, p. 326).

Look at panel A in Figure 6–1, a simple diagram of a few concepts and the interconnecting pathways among them. Even such a simple network codes or represents a great deal of information. For instance, one fact or bit of information coded in this network is that ROBIN is a member of the category BIRD; another is that ROBIN has a RED BREAST; another is that CANARY is YELLOW, and so forth. Each of these simple con-

FIGURE 6-1

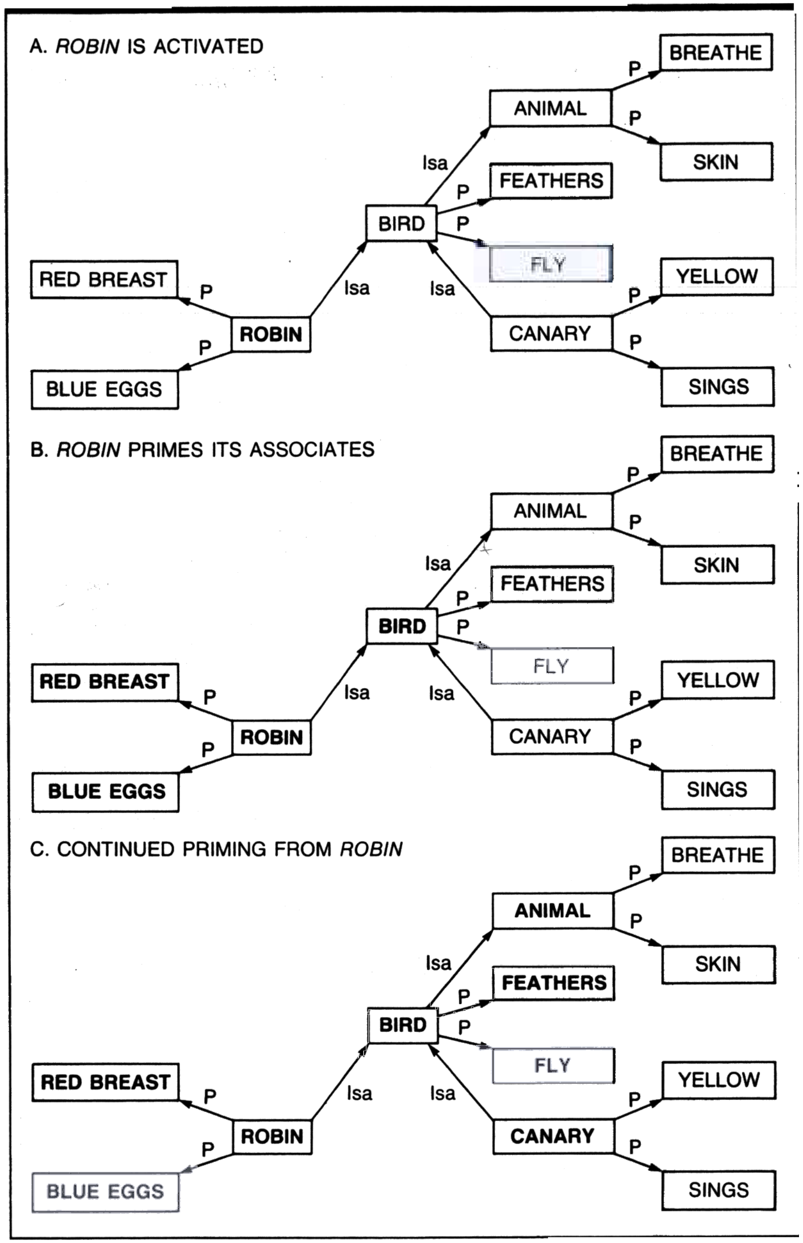

A portion of the semantic network is illustrated. In panel A, the concept ROBIN has been activated, and is shown in boldface. In panel B, the spreading activation from ROBIN has now activated concepts linked to ROBIN, e.g., the boldface BIRD, RED BREAST, and BLUE EGGS. In panel C, the continued spread of activation that originated from ROBIN is depicted.

nections records an elementary fact or **proposition**, *a relationship between two concepts*. Notice further that each of the pathways joining two concept nodes is a *labeled and directed* pathway—each pathway specifies a certain relationship and the direction of that relationship. Thus ROBIN *isa* BIRD—*is a member of the category* BIRD—and BIRD has the *property* FEATHERS.[4] Further, the illustrations in Figure 6–1 show what would happen to this portion of your semantic network when the word "Robin" is presented. First, the concept corresponding to "Robin" becomes activated, illustrated by **boldface** in panel A. After this, the concept ROBIN will begin to spread activation to those concepts it is linked to, the **boldface** BIRD, RED BREAST, and BLUE EGGS concepts in panel B. *Those* concepts continue the spread of activation to *their* associated nodes, as depicted in panel C.

Collins and Quillian proposed that such a spread of activation is triggered each time a concept is activated in semantic memory. Thus, when *two* concepts become activated, then there are *two* simultaneous spreads of activation, one from each concept node. As an excercise, study the top portion of Figure 6–1 and mark which concept nodes will become activated by a sentence like "A robin can breathe." To keep track of the original source of activation, write a 1 next to pathways and nodes that will be activated by ROBIN, and a 2 next to those activated by BREATHE. Take this demonstration through at least two cycles; first, the original node activates its connected nodes, then those nodes activate *their* connected nodes. (In a sense, you are "hand simulating" the computerized memory search in TLC. This is a laborious enough process that you no doubt realize one of the attractive features of computer simulation and artificial intelligence programming. Let the computer do all of the painstaking work of spreading the activation and keeping track of the sources, while you merely wait at your terminal for the program to tell you the final outcome of the search.)

Intersection What did you discover in the previous excercise? If you did it correctly, you should have discovered two important properties of spreading activation search. First, you found that the activation originating with ROBIN eventually primed a node that was *also* primed or activated by BREATHE. This is the *exact* process proposed by Collins and Quillian, and originally programmed into Quillian's TLC model, to explain how information is retrieved from semantic memory. The "harmless spreading plague" eventually encounters another "harmless spreading plague" that came from a different source. When that happens, then a connecting route or set of pathways has been retrieved from semantic

---

[4] The *isa* relationship, reflecting category membership, was a bit of cognitive jargon contributed by Rumelhart, Lindsay, & Norman (1972), in the same book as the 1972 paper by Collins & Quillian—it means "is a," as in "is a member of the category." The importance of the direction of the relationship is well illustrated by *isa*, since the reversed direction for *isa* is not true, i.e., *All BIRDS are ROBINS. (Note: By convention, sentences that are intentionally wrong are prefaced by an asterisk.)

memory. In the terminology of the model, *when the two spreads of activation encounter one another*, an **intersection** has been found between the two concepts, ROBIN and BREATHE. Once an intersection has been found, then a decision stage must operate to make sure that the retrieved pathway is valid, i.e., that it represents the same relationship as specified in the sentence. In other words, a similar pathway to the one between ROBIN and BREATHE would be found between ANIMAL and RED BREAST, but the decision stage would decide that it's not true that "all animals have red breasts." (Incidentally, this should sound somewhat familiar, a search stage followed by a decision stage; see chapter 2 on A Process Model, or chapter 4 on Sternberg's approach if you need a quick review.)

Related Concepts The second property of this kind of search that you should have discovered is that *other* concepts also become activated or primed during the search. That is, while the intersection pathway involved ROBIN *isa* BIRD *isa* ANIMAL *property* BREATHE, many other concepts were also primed during the search; there should also be a 1 next to RED BREAST and BLUE EGGS from the first cycle, a 1 next to FLY, FEATHERS, and CANARY after the second cycle, and so forth. Thus, a spreading activation search not only retrieves the relevant pathway between two concepts, it also activates *related* concepts. These related concepts will not remain activated forever, of course, since activation is always presumed to decay after some amount of time. Nonetheless, for a short period, these related concepts have received a boost in their activation levels, making them temporarily more accessible. This *priming* of related concepts is absolutely key to an understanding of semantic processing—we'll return to it repeatedly throughout the chapter, and indeed throughout the entire book.

Smith's Feature Overlap Model

Given the excitment of studying meaning and how it is represented in memory, it is not surprising that other approaches to semantic memory soon appeared. We'll focus here on only one of those other approaches, the Smith Feature Overlap Model, since it offered a rather clear contrast to the Collins and Quillian model in some basic assumptions, and since it was the most successful challenger to that model.

Feature Lists Smith's model (e.g., Smith, Rips, & Shoben, 1974) was considerably simpler than the Collins and Quillian network model in its assumptions about the structure of semantic memory, but as a consequence was somewhat more elaborate in its assumptions about the process of retrieval. Its most basic structural element was the **feature list**. Smith et al. suggested that, rather than postulating extensive networks of concepts and linking pathways, we might consider semantic memory to be a collection of lists. Each concept in semantic memory was represented as a list of **semantic features** in the model, *simple, one-element characteristics or properties of the concept*. Thus, the concept ROBIN would be represented as a list of ROBIN's features, like animate, red breasted, smallish, winged, feathered, and so forth (see Figure 6–2).

Smith et al. suggested that these feature lists were ordered in terms of a factor they called *definingness*. That is, Smith et al. said that the feature lists stored in memory were ordered in a kind of priority ranking, with the most defining features for a concept toward the top of the list, and the least defining features toward the bottom. Thus, an absolutely *essential feature*, like *animate* for BIRD, was called a **defining feature**, and would be stored near the top of the feature list. Conversely, features that are not particularly defining for the concept, say that a ROBIN *perches in trees*, would be placed toward the bottom of the list. In fact, Smith et al. proposed that these lower features were more appropriately called **characteristic features** of the concept, *features that are only common or frequent, but not essential to the meaning of the concept.* Thus, characteristic features do not *define* what it is to be a ROBIN or BIRD, they are only commonly observed features of the concepts. Note that, according to Smith et al., characteristic features are not particularly important to the central meaning of the concept, whereas defining features are—a robin may or may not eat worms, but to be a robin it absolutely has to be animate.

Feature Comparison The major process of information retrieval hypothesized by Smith et al. was a feature comparison process. Say you were given the sentence "A robin is a bird," and had to make a true/false judgment. According to the model, you would access the two concepts

FIGURE 6–2

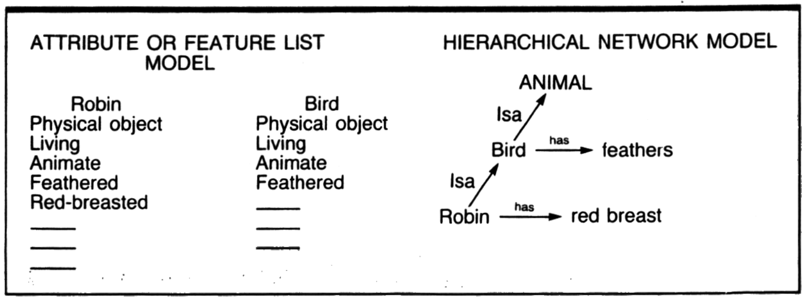

Information in semantic memory is represented differently in feature list models and in hierarchical network models. In feature list models, a concept is represented as a list of simple semantic features; in hierarchical network models, concepts are represented as nodes that connect to other nodes via pathways. The Smith et al. (1974) model is a feature list model, and the Collins and Quillian (1972) model is a hierarchical network model. (Figure adapted from Smith, 1978.)

ROBIN and BIRD in semantic memory, and then would proceed to compare the features on those two lists. This feature comparison process involved a global comparison of the features—some randomly selected subset of features on each of the two lists would be compared in order to "compute" the similarity between the two concepts. This comparison process yielded a *feature overlap score,* simply an index or measurement of the similarity of the two concepts. Of course, for the concepts in "A robin is a bird," the feature lists should overlap a great deal—there are hardly any ROBIN features that aren't also BIRD features. The outcome of this feature comparison process then would be a very high overlap score, so high in fact that you could confidently repond "yes" right after this global comparison of features. Conversely, with a sentence like "A robin is a bulldozer," there should be so little feature overlap that you could respond "no" immediately without any further processing. These "fast yes" and "fast no" responses were called *Stage I responses* by Smith et al.—when overlap scores are either very high or very low, there is no need to continue the search and a response can be made immediately.

Consider two other kinds of sentences, where the stated relationship isn't quite so obvious. First, consider "A chicken is a bird." As before, the process of retrieving information is the feature comparison process, now being performed on the CHICKEN and BIRD features. Most people's intuition is that chickens are a somewhat less "representative" example of the bird category. Compared to your "average" bird, chickens seem rather unusual—they don't perch or make nests in trees, they don't eat worms, they're larger, and so on. Isn't it clear, then, that the Stage I comparison process should find only an intermediate degree of overlap between CHICKEN and BIRD? Smith et al. claimed that when the overlap scores indicated only moderate similarity, a second comparison was necessary. The second comparison was referred to as a *Stage II comparison.*

Unlike the fast, global Stage I comparison, the Stage II comparison was a careful and rather slow one, and it used *only* the defining features to compute its evidence. Thus, for the CHICKEN-BIRD sentence, only the defining features of the two concepts would be compared in Stage II. Since it is in fact true that chickens are birds, presumably there would be a match on all the tested features in this stage, yielding a "slow yes" response, "slow" because it involved Stage II comparison, and "yes" because all of the defining features from CHICKEN would match those from BIRD during Stage II comparison. Finally, consider the sentence "A bat is a bird." This sentence should also yield only a moderate overlap score during Stage I, thus necessitating a Stage II comparison. During Stage II, however, there will be several important *mismatches.* It would seem that only the characteristic features of bats make them similar to birds. Their defining features (mammal, furry, teeth, etc.) give rather convincing evidence, however, that the sentence is false. This would be a "slow no" response, "slow" because it required Stage II processing, and "no" because of mismatches on the defining features.

Empirical Tests of Semantic Memory Models

Most of the early tests of semantic memory models adopted what was known as the **sentence verification task,** in which *simple sentences are presented for the subjects' yes/no decisions* (e.g., "A robin is a bird" or "A canary is green"). Accuracy scores for subjects' decisions about these simple sentences would be rather uninformative, of course—people seldom are in doubt or make mistakes about such simple facts. Thus, reaction time (RT) tests were the usual method of testing semantic memory models; present simple sentences, both true and false ones, and *time* the subjects as they make their yes/no decisions (the speed emphasis is clearly implied in the Smith et al. "fast" and "slow" responses). As is generally the case in the information processing framework, the assumption is that time provides a window into the unseen mental processes, here the processes of semantic memory search and decision.

Notice, however, a distinguishing feature of semantic memory tasks, as opposed to virtually all of the short-term and episodic long-term memory tasks we've discussed. In semantic memory tasks we are testing people on the *knowledge they already possess,* their conceptual, general world knowledge about robins, machines, and so forth. The tasks for both short-term and episodic long-term memory provide a set of stimuli for the subject to master, to hold in short-term memory or to commit to long-term memory. Semantic memory tests, conversely, rely on subjects' existing conceptual knowledge—when subjects walk into the lab, they are not asked to learn a list, but instead are asked to demonstrate what they already know by saying "yes" or "no" to the stimulus sentences. Clearly, the tests rely on the assumption that people have the relevant knowledge in memory, and further, that such semantic knowledge is largely similar among individuals (or at least is similar among the individuals within a language culture).

Collins and Quillian's (1969) earliest report tested an obvious prediction from their model, that two concepts that are closer together in the network should require less time for verification than two that are farther apart. Refer again to Figure 6–1, and to the "hand simulation" you performed. If we assume that your conceptual knowledge about this portion of semantic memory is accurately represented in the figure, then several predictions can be made. For example, which sentence should be faster to verify, "A robin is a bird" or "A robin is an animal?" Right—"A robin is a bird," because it should take less time for the spread of activation to intersect for this sentence than for the "A robin is an animal" sentence. Likewise, if the figure is accurate, it should take less time to verify that canaries are yellow than that canaries can fly or that they breathe, again for the same reason.

Figure 6-3 shows the results of Collins and Quillian's (1969) test of these predictions. In the figure, Collins and Quillian used the symbol S to indicate a superordinate statement, what we've been calling *isa* sentences, and a P to indicate property statements. They tagged on a digit from 0 to 2 to the S or P, to indicate how many levels up in the hierarchy the search had to proceed in order to find the stated concept. As you can see, reaction time increased as the semantic distance between the two concepts increased. For the superordinate sentences, searching for an S2 relationship (e.g., "A canary is an animal") required about 75 msec longer than searching for an S1 relationship ("A canary is a bird"). The same increase in time was also found for property sentences going up 1 versus 2 levels (notice that a P0 sentence meant "up 0 levels," i.e., that the second term in the sentence, the predicate, was stored at the same level in the hierarchy as the subject of the sentence; e.g., "A canary can sing"). It takes longer to retrieve a relationship between two concepts when those concepts are stored farther apart in the semantic structure. Moreover, it appeared from these results that the *isa* pathway was stronger than the *property* pathway, since superordinate sentences were faster overall than property sentences (see Hampton, 1984, for confirmation).[5]

FIGURE 6–3

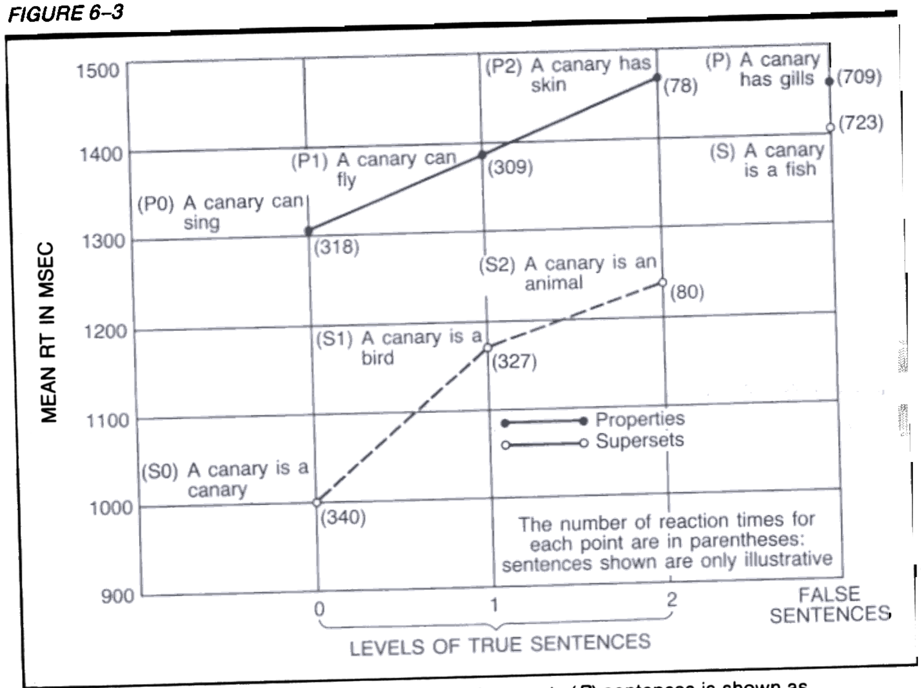

Reaction time to superordinate (*S*) sentences and property (*P*) sentences is shown as function of *levels within the hierarchy*. An S2 sentence involves a superordinate connection *2* levels up the hierarchy; S1 means 1 level up in the hierarchy; a 0 level sentence had the predicate stored at the same hierarchical level. (From Collins and Quillian, 1969.)

While this result sounds like strong evidence for the network model, it turns out that it was by no means a definitive test—it did not distinguish Collins and Quillian's model from the Smith et al. approach, for instance (although it might be argued that Smith et al. designed their model specifically to account for this kind of result). Consider the feature overlap process in the Smith et al. model, and how it would operate on these sentences. It should be clear that, on the average, there should be less feature overlap between any two concepts when those concepts are farther removed from one another in the hierarchy depicted in Figure 6–1 (notice, however, that Smith et al.'s model did *not* structure concepts into a hierarchy, and that the rigid hierarchical scheme was later abandoned by Collins and Loftus, 1975). Sentences like "A robin is an animal," which Collins and Quillian called S2 sentences, would yield lower overlap scores than sentences like "A robin is a bird" (an S1 sentence) during a feature comparison process. Thus, on the average, more S2 sentences would need the extra comparison step of Smith et al.'s Stage II. This would of course tend to slow the comparison process down.

Let's discuss three specific issues that were debated in the literature as evidence for or against these two models: cognitive economy, property statements, and typicality. You'll get a flavor for how cognitive psychology tests models and theories, and in the process you'll encounter the evidence for the strongest generalization about semantic memory, a generalization known as the semantic distance or semantic relatedness effect. Although it was not necessarily clear at the time, all three of these issues are different expressions of the same underlying generalization—semantic memory's structure is based on semantic relatedness among concepts.

Cognitive Economy A clear commitment of Quillian's (1969) TLC model was that redundant information is *not* stored in memory: "The sheer quantity of information involved . . . argues strongly that both the human subject's memory and our model thereof contain as little redundancy as possible and that it [should] contain stored facts only when these cannot otherwise be generated or inferred" (p. 228). This position has been given the name **cognitive economy.** To economize in the number of concepts that must be stored, *only nonredundant facts will be stored in memory.* A rather straightforward example of this principle is the

---

[5] Collins and Quillian, of course, tested many more concepts than just canaries and robins, although the tradition in this area of research is to illustrate the models using these words. In a depressing example of literal mindedness, a student of mine once answered an essay question on semantic memory by saying "Collins and Quillian devised a psychological model to explain what people know about birds."

appearance of FLY only with BIRD in Figure 6–1, instead of also with ROBIN and CANARY. It seems very likely that Quillian adopted this principle because of his work with computer modeling. That is, storage space is at a premium in computer models, so the principle of cognitive economy is a convenient principle to adopt for such work. Furthermore, the simple inference process of retrieving a property from a superordinate could be invoked over and over, and would yield the same outcome—knowledge that a robin does fly, for instance. Although the original statement of the model in a psychological journal (Collins & Quillian, 1969) indicated that an extreme stance on this principle was not being proposed, most early investigators realized that some kind of cognitive economy was intended.

There is clearly a grain or two of truth to the cognitive economy idea. After all, it strains the imagination to suppose that we would fill our memories with such facts as "the philosopher Aristotle had two hands," that we would waste mental effort and space in such a colossal fashion. On the other hand, there are several difficulties with the principle of cognitive economy for more ordinary concepts. One involves the issue of forgetting, or at least a version of it. To illustrate, say that when you were three, you saw a bird in the backyard, and your father said to you, "That's a robin. See it *fly* away?" You might store a connecting pathway between ROBIN and FLY in this situation, and might continue to do so for other birds as well, until you later found out that almost *all* birds fly. The principle of cognitive economy would seem to imply that when you learned the more general fact, and stored a pathway between BIRD and FLY, the separate pathways to FLY from ROBIN and other birds would have to be *erased*. This is a rather difficult position to swallow, given what we know about forgetting from long-term memory. To their credit, Collins and Quillian (1969) tested sentences in their experiment that *seemed* likely to require inferences—for instance, "A canary has skin." Nonetheless, the possibility existed that a different structure, one that was not cognitively economical, might be correct instead.

The best known of the studies that challenged the cognitive economy principle was done by Conrad (1972). In her study, Conrad first collected normative data from a sample of college students, asking them to write down properties of a variety of different words (like "robin," "banjo," "onion," etc.). She tabulated the frequency with which different properties occurred in these written listings, then used the words in sentences for her verification task. Basically, Conrad found that there was little evidence for the economical scheme implied by cognitive economy—properties at various levels in the hierarchy seemed to be stored *repeatedly*, not in the overly tidy, nonredundant fashion implied by Quillian (see also Ashcraft, 1978b). Furthermore, the frequency with which properties were produced in the normative data was a much more powerful predictor of RT performance than the hierarchical scheme used by Collins and Quillian. High-frequency properties, like WINGS for the concept ROBIN,

were verified very quickly, whereas low-frequency properties, like FEET for ROBIN, required more time in Conrad's results (see also Ashcraft, 1976; Glass, Holyoak, & O'Dell, 1974). It seemed very possible, in other words, that Collins and Quillian had obtained their result not because the properties or superordinates were more distant in a hierarchical structure, but instead, because their "more distant" concepts were actually associated more weakly with the concepts. The cognitive economy effect, it seemed, was due to Collins and Quillian's failure to control for the effect of frequency or strength of association.

Property Statements. Another feature of Conrad's (1972) study took on added significance as more semantic memory experiments were conducted. Despite the early testing of both superordinate and property statements by Collins and Quillian (1969), most of the subsequent research focused exclusively on category statement verification, testing sentences that asserted "S is a P," i.e., "S belongs to the superordinate category P." (These important studies are discussed shortly, under "Typicality effects.") When experiments began to probe more thoroughly into the richness of semantic knowledge, **property statements** became an important testing ground. Essentially, a property statement is one that *asserts that some concept, S, has a certain property or characteristic, P;* for instance, "A robin has wings," "A canary is yellow," and so forth. As experiments using property statements began to appear, a difficulty arose in the use of the Smith feature comparison model.

Consider "A robin has wings." The normal Smith et al. comparison process, you'll recall, was to access the feature lists for both concepts, then do a global Stage I comparison on these lists; sometimes, of course, this Stage I comparison was followed by a Stage II evaluation of just the defining features for the two concepts. Whereas this process seemed quite sensible for category statements, it seemed rather implausible when property statements were investigated. That is, the Smith et al. approach suggested that for "A robin has wings," you must access the feature lists for both concepts, your concept of ROBIN and your concept of WINGED THINGS. You would then conduct the regular feature-overlap process on these two feature lists.

An initial peculiarity involves categories like WINGED THINGS, THINGS WITH FEET, and to use a Smith et al. example, BROWN THINGS (as in "An ostrich is brown"). It seems a bit farfetched that we actually have concepts or categories in semantic memory corresponding to WINGED THINGS or THINGS WITH FEET, each with its own feature list. Assuming for a moment that we do have such concepts, however, a second puzzle arises. What features might be on your list for the concept WINGED THINGS? Aside from *has wings*, and possibly *can fly*, it's difficult to imagine what other features would compose this list. Note further that each feature list was assumed to have both defining and characteristic features on it. If we can only think of one or two features for

WINGED THINGS, then the further division into characteristic and defining features seems quite implausible. Yet this was the process advanced by Smith et al. (1974) to explain property statements. (A related peculiarity was that WINGS would probably be a feature stored on your feature list for both ROBIN and BIRD. Why, then, would you need to access a WINGED THINGS concept when WINGS is already a feature stored on your ROBIN list?) The Smith et al. model was extended to three different kinds of property statements; adjective (e.g., brown), relative adjective (e.g., small), and *has* properties (e.g., wings). As Ashcraft (1978a) noted, however, each type required a distinctly different kind of comparison process to render it workable within the feature comparison approach. Furthermore, there was little or no evidence that subjects treated different kinds of properties differently, either in normative or RT studies (e.g., Ashcraft, 1976, 1978b; Glass et al., 1974).

As Collins and Loftus (1975) pointed out, the network approach contains both property and superordinate pathways, and thus has no difficulty in explaining how people verify property statements—the same intersection search process applies to both. Furthermore, the troublesome distinction between defining and characteristic properties was not in the Collins and Quillian model, whereas Smith et al. had some difficulties in convincing researchers that such a distinction even existed, much less that it was as central to semantic memory processes as they had proposed. Finally, Collins and Loftus made one additional telling criticism of the feature comparison model. It seems highly arbitrary that the Smith et al. model stores any and all features that might be learned about a concept *except* for the important superordinate information about that concept, e.g., that a robin is a bird: "While most people may not have learned some superordinate relations (e.g., that a beaver is a mamm~~m~~al, or a sled is a vehicle), there are many they have learned (e.g., that a wren is a bird, and a beaver is an animal). Why would they not use such information if it is stored? How in fact can they avoid using it? It is an unlikely model which postulates that people use information that is less relevant to make a decision, instead of information that is more relevant" (Collins & Loftus, 1975, pp. 425–426).

Typicality Effects The category statement studies we've been discussing usually used the sentence frame "An S is a P" for the stimulus sentences. For such frames, the S term, for subject, and the P term, for predicate, were category members and superordinate category terms, respectively (e.g., A robin is a bird, and a robin is an animal). As many researchers found, more distant superordinate terms led to slower performance.

A variation on this theme quickly came to the forefront in semantic memory research, in which the S term, rather than the P term, was of interest. The extensive set of norms by Battig and Montague (1969), originally collected as a source of stimuli for free recall and clustering studies, showed clearly that there was a stable ranking or ordering of the members of a category. To use our familiar example, "robin" was the most frequently occurring member of the bird category—it was listed by 377 of the 442 subjects; "chicken," on the other hand, was listed by only 40 subjects. As you might suspect, such a range of occurrence says something important about people's representation of the semantic category BIRD (and other categories as well; see footnote 5), and, as a consequence, implies something important about their RT performance.

The important result in studies that used S terms from the whole range of frequencies was that frequently-listed category members were verified more rapidly than others—you can decide "yes" to "Is robin a bird?" more rapidly than to "Is chicken a bird?" Under the standard assumption that time differences can reveal mental processes, this result suggested a new dimension to semantic memory. Unlike the equal-length pathways depicted in Figure 6–1, such a result suggested that pathways to less-frequent category members were longer, making those members "farther" away in the semantic network. Under the Smith et al. scheme, less frequent category members were said to have lower feature overlap with their superordinate than more frequent members. In fact, it was exactly this kind of result (Rips, Shoben, & Smith, 1973; Rosch, 1973) that led Smith et al. (1974) to their feature comparison model. So-called "central" members of the category (Rosch, 1973) were verified more quickly than "peripheral" members, as illustrated in Figure 6–4. This important effect is now called the **typicality effect**—*typical members of a category can be judged more rapidly than atypical members.* (We will return to typicality later, since it is more important than this short discussion implies.)

Semantic Relatedness Figure 6–5 illustrates a modified network representation of part of the bird category, one that incorporates the three issues we've been discussing. Notice first that there is no rigid cognitive economy in the illustration—properties listed for a concept are linked directly to that concept, rather than indirectly via pathways through a superordinate node. Second, notice that these linking pathways are of different lengths, reflecting the results that sentences with high-frequent properties are verified faster than those with low-frequent properties (this result was clearly anticipated by Collins and Quillian; see the introductory quotation). Finally, notice how typicality effects are represented in the figure—typical or central members of the category are connected to the superordinate node by *shorter* pathways, whereas atypical or peripheral members are linked by *longer* pathways. In other words, the network structure of a category is illustrated in the figure, with the length of connecting pathways coding the degree to which two concepts are related. Shorter pathways denote concepts that are closer in semantic space, or to use a different terminology, denote concepts that are more highly related to one another. Longer pathways denote lesser degress of semantic relatedness.

FIGURE 6-

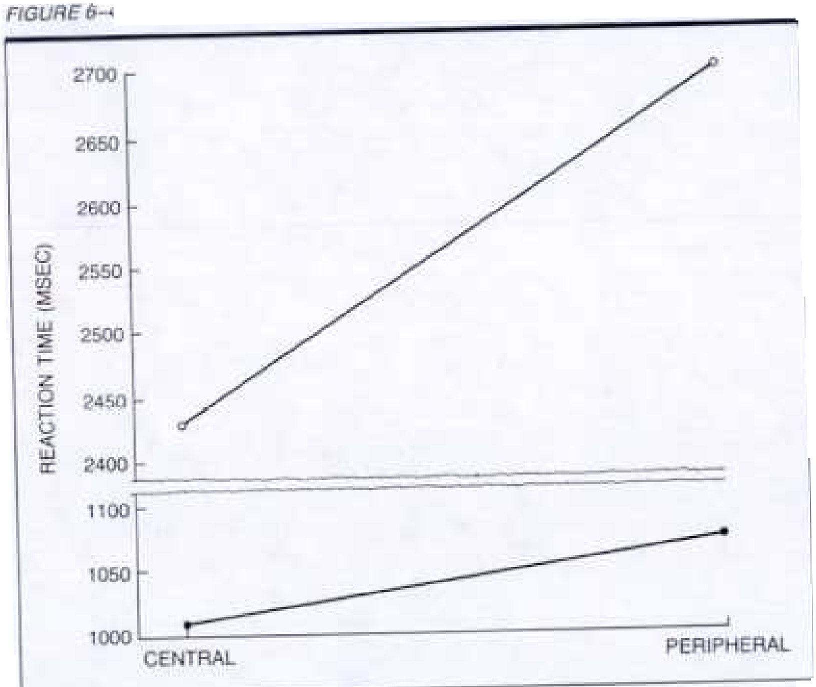

Reaction time to true sentences is shown separately for adults (bottom curve) and children (top curve). Sentences had either a central or a peripheral member of a category as the subject of the sentence; i.e., "A *robin* is a bird" versus "A *chicken* is bird." The term *central* means *typical*, and *peripheral* means *atypical*. (From Rosch, 1973.)

Unfortunately, it is difficult in a two-dimensional figure to illustrate some other features of networks. For example, most researchers would agree that the *strict* hierarchical approach is incorrect. This conclusion is based on evidence that statements like "A beaver is an animal" tend to be verified more quickly than "A beaver is a mammal." Of course, in a strict hierarchy, the mammal sentence should be faster to verify or judge, since mammal is a subset of animal. Nonetheless, RTs to such sentences show a different effect than strict logic or hierarchies would imply (e.g., Rips et al., 1973). In the figure, this relaxation of the strictly hierarchical scheme would mean that there should be a connecting pathway from CHICKEN to ANIMAL, for instance, and that this pathway would be shorter than the CHICKEN *isa* BIRD pathway. Conceiving of such networks in three-dimensional space makes this easier to imagine, but, of course, three-dimensional networks cannot be illustrated easily in a drawing.

FIGURE 6-5

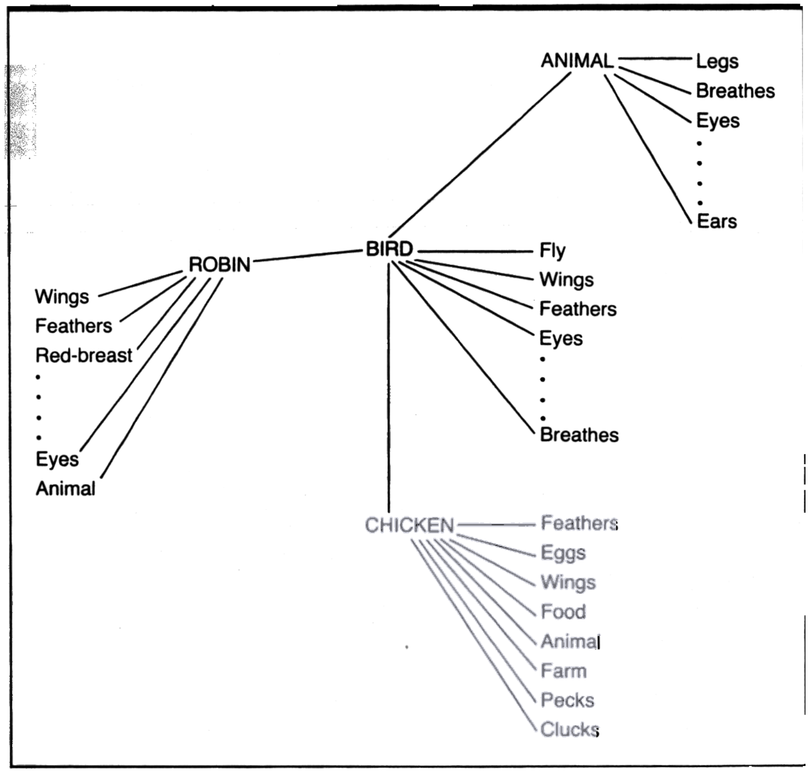

A portion of the semantic network is illustrated, taking into account three empirical effects: (1) there is no strict cognitive economy in the hierarchy, so redundant information is stored at several different concepts; (2) typical members of the category are stored more closely to the category name or prototype member; and (3) properties that are more important are stored more closely to the concept than those of lesser importance.

Regardless of how we diagram the illustrations, the empirical prediction from such a network is that performance, particularly RT performance, will vary directly as a function of the length of the connecting pathway. To state it slightly differently, the higher the semantic relatedness between concepts, the faster you are able to retrieve the connection between them. This is in fact the **semantic relatedness effect;** *concepts that are more highly interrelated can be judged "true" more rapidly than those with a lower degree of relatedness.* The important new ingredient here is that this semantic relatedness principle applies both to category statements (a robin versus a chicken is a bird) *and* to property statements (a robin has wings versus has a beak). To make a long story short, the prediction turns out to be true for both kinds of statements (e.g., Ashcraft, 1978a).

The opposite side of the semantic relatedness effect has also been confirmed, that higher degrees of relatedness will *slow down* performance when the sentence is actually false (Kintsch, 1974). In other words, the large number of connecting pathways (network), or the nontrivial number of overlapping features (feature comparison) between BAT and BIRD, or between WHALE and FISH, make those sentences relatively slow to judge as "false." In confirmation of this generalization, Smith et al. (1974) reported that subjects answered false-but-related sentences such as "Lions have stripes" about 200 msec *more slowly* than unrelated false sentences like "Lions have dials." The only clear exception to this generalization involves direct contradictions, which are usually rejected rather quickly despite strong relatedness between the concepts (e.g., "All dogs are cats"; see Glass & Holyoak, 1975).[6]

Results such as these suggested a great regularity in semantic memory structure—concepts farther removed from one another in the hierarchy required more time for retrieval, and concepts closer together required less time, and were more central to the meaning of the concepts or categories. **Also called the semantic distance effect,** the semantic relatedness effect is an important generalization: The closer two concepts are in semantic memory (distance), or the more related they are, the faster is the mental search process that retrieves information about the concepts. This idea is particularly easy to visualize for network models, where more activation accumulates at the closer concepts because a shorter distance has to be traveled (or alternately, that activation accumulates more rapidly at closer concepts due to shorter distances). A grander implication also follows from this effect as well, that semantic memory—all of our general world knowledge—is structured according to semantic relatedness. In short, semantic memory stores concepts in terms of their relatedness to one another, and this is the basic dimension along which semantic memory is organized.

▼ Categorization, Concepts, and Prototypes

A final series of important studies will nail down this principle of semantic relatedness and the structure of categories in semantic memory. This discussion will also set the stage for several important processes to be considered in later chapters, among them the comprehension of conversation and the acquisition of language. As preparation for this section, you might look up the words "bird" or "flower" in your dictionary, and note any illustration that accompanies the definition.

Concept Formation

Let's begin this section with a quick look at traditional laboratory studies of *concept formation*, since this body of work is often contrasted with semantic memory concepts and categories. In the typical laboratory study of concept formation or identification (e.g., Bourne & Bunderson, 1963; Haygood & Bourne, 1965), subjects were presented with a long series of stimuli such as those illustrated in Figure 6–6. After each stimulus was presented, subjects had to indicate whether the figure was or was not an instance of the concept being tested, and were then given feedback as to the correctness of their decision.

If you're paying attention here, you should be saying to yourself "Wait a minute—*what* concept?" In other words, how can subjects judge whether or not a stimulus is an instance of a concept when the concept itself isn't known? This, in fact, was the main point of this kind of study—by receiving feedback on their choices, subjects would eventually "acquire" or figure out what the concept was. At that point, the concept was said to be formed or identified; another way of phrasing this is to say that the subject had learned to classify the stimulus correctly. Study Figure 6–6 again, paying attention to the answers at the right of each stimulus. Once you've decided what the concept is, check the end of the next paragraph to see if you were right.

An important factor in these studies was the complexity of the concept that was to be identified. For example, a simple concept might be defined as having only one critical feature; e.g., shape might be the valid dimension, and "roundness" the critical value on that dimension. A complex concept might join two or more critical dimensions with a specific combination rule; for instance, the relevant concept might be defined as "round, or containing a plus, but not both features together." As you can infer from the figure, the concepts that people were supposed to identify were combinations of simple features, features that could be manipulated independently as the experimenter invented the stimuli. For our purposes, the important results of such studies can be summarized briefly (e.g., Kintsch, 1970): once the relevant dimension is selected (for instance, ignore shape and size, pay attention only to color), then the concept has essentially been identified; the more independent dimensions along which the stimuli vary (e.g., shape, size, color, location), the more trials it takes to identify the concept; the more redundant, concept-relevant dimensions, the faster the identification (in other words, if either roundness or presence of a plus indicates a positive instance, and roundness and presence of a plus *always* occur together in the stimuli, then you can identify

---

[6] Glass and Holyoak (1975) also reported another kind of statement in which a related, but false pairing was judged quite rapidly, exemplified by "All animals are birds." They argued that the availability of a counterexample (e.g., reptiles are animals, but aren't birds) made such judgments faster. Since their study examined other kinds of quantifiers as well (e.g., "Some animals are birds"), it may be that this exception is best understood in terms of reasoning and decision processes, rather than as a straightforward result pertaining to the structure and ordinary processes of semantic memory.

FIGURE 6–6

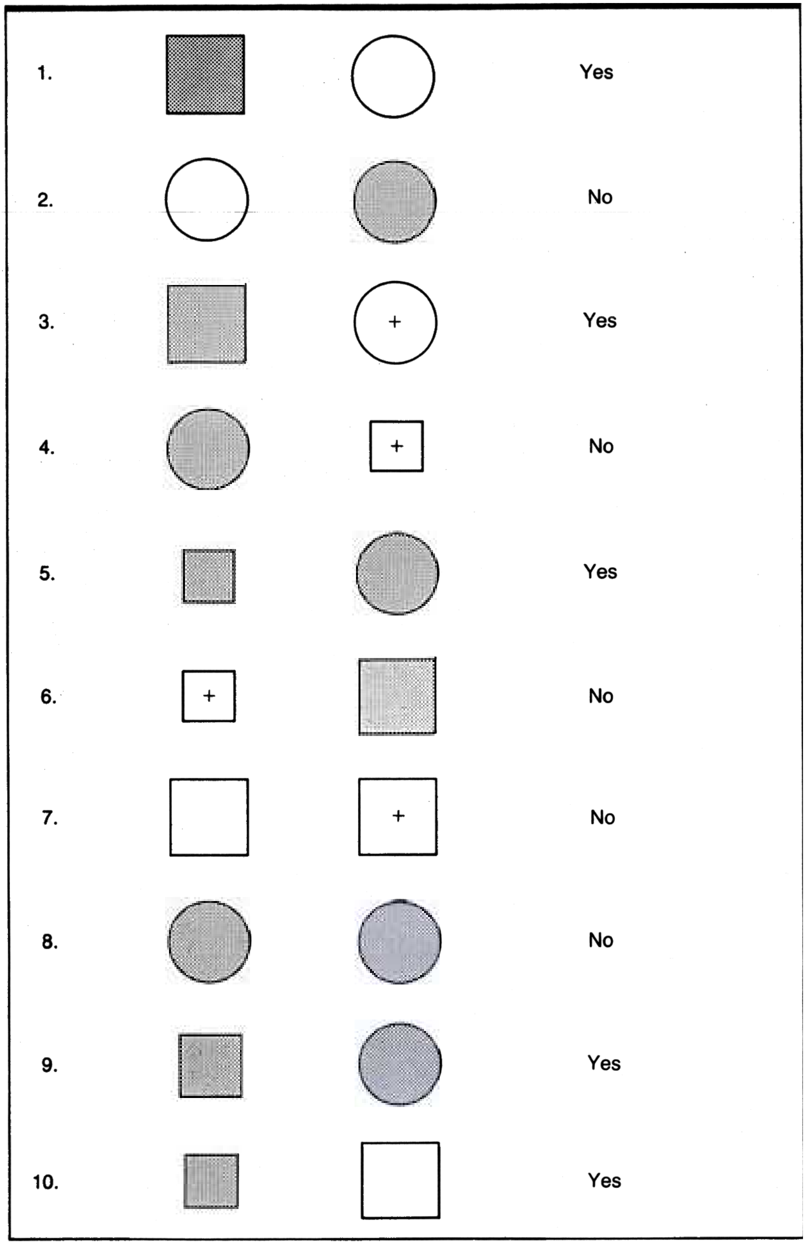

A set of ten concept formation patterns and feedback on which patterns are positive instances of the concept, and which are negative. See text for the rule that expresses the concept.

the concept more quickly). (The concept in Figure 6–6 is "shaded square on the left," a combination of the three dimensions shape, shading, and position.)

Natural Language Concepts

Beginning in the early '70s, a very different approach to the topic of concept formation and classification appeared, largely due to Eleanor Rosch's important research. A central theme in Rosch's reports (e.g., 1973) was that traditional concept identification research "missed the boat" in one especially critical way, having to do with the concepts that people were to acquire. Look at the positive instances of the concept in Figure 6–6. Because the concept defined by these instances is composed of separate, independent features, a question like "which of the positive instances is a *better* member of the category or concept" seems very strange—they are all equal, since they are all members. For example, patterns 3 and 5 are both positive instances, despite the fact that one is large and one is small. Rosch contended that such artificial categories bear little relationship to the natural categories we use in language and thought. **Natural categories,** *concepts and categories that occur in the real world of our experience,* have a complex internal structure, she argued. Questions like "which is a better member of the dog category, collie or Pekingese?" make perfectly good sense for natural categories, suggesting that the category "dog" has an internal structure in which some members are "better" or more representative than others. Given this obvious difference between natural categories and experimenter-defined concepts, Rosch argued that the concept formation literature may have given us "a quite distorted view of how *real* categories are learned and how they function in cognitive processes" (Rosch, 1973, p. 112, emphasis added).

In terms of examples such as those in Figure 6–6, Rosch was essentially saying that real-world categories are quite different from those studied with the laboratory task of concept formation. Real-world category members are not bundles of independent features, classified into categories because of the presence or absence of those features ("is it square, is it shaded, is it on the left?"). Instead, real-world categories show the characteristic of "fuzzy boundaries," ill-defined or uncertain membership for a variety of category instances. For example, is a sled a toy or a vehicle? Isn't a chicken a better member of the farm-animal category than of the bird category? Furthermore, as Collins & Loftus (1975) pointed out, no single feature seems absolutely necessary as a criterion of membership or classification. For example, having a skin or peel is certainly an important feature of the fruit category, yet a peeled orange is still a fruit. Contrast this with pattern 7 in Figure 6–6, in which violation of merely one feature makes the pattern a negative instance.

Rosch claimed, in short, that membership in categories is not an all-or-none affair, where an instance either is or isn't a member. Instead, membership in categories is a matter of degree. Some members are highly representative or "good" instances, close in some sense to an "ideal" member. Some are "poor" instances, on the periphery, in that "grey" area where members seem to blend into another category. This is, of course, the ill-defined or fuzzy boundary idea.

Perceptual Categories To document her point, Rosch presented a fascinating study conducted with members of the Dani tribe in New Guinea (Rosch-Heider, 1972). Both short- and long-term memory tasks were administered, using chips of different colors as the stimuli. She found that the Dani learned and remembered more accurately when the chips were "focal" colors rather than "nonfocal" colors, or, putting it loosely, "a really red red" as opposed to "a sort of red red." The *central*, perceptually salient "good" red, in other words, was a better aid to accuracy than the nonfocal "off-red." One compelling aspect of this study was that the language of the Dani contains only *two* color terms, one for "dark" and one for "light." Nothing in their language expresses meanings like "true red" or "off-red," and yet their performance was influenced by the centrality of focal versus nonfocal colors.[7]

Given the absence of separate color terms in their vocabulary, Rosch reasoned, it must be that there are structured, mental categories of colors in the subjects' semantic memories, categories that do not rely at all on the spoken language. Each category has a central tendency, a focal member that represents the "true" or "good" color, and also has noncentral, nonfocal members that are less representative. Thus, the centrality or "goodness" of a color in its own category made a difference in the memory tasks, even without distinct words to name the different color shades. As you recall from the previous section on typicality, synonymous terms for *focal* are *central* or *typical*, and *peripheral* or *atypical* are synonymous terms for *nonfocal*.

Semantic Categories Once she had demonstrated the structure of such perceptual categories, Rosch went on to demonstrate an analogous structure in natural, semantic categories. For example, in one set of studies (Rosch, 1973, 1975), subjects were asked to rate a list of category members on their "representativeness" or "typicality." She then found that categorization results—for instance, RT to verify that "An *S* is a *P*"—depended on the rated typicality of the instance (Rosch, 1973). This is a result you've already encountered in Figure 6–4. Her later work extended this typicality effect even further. For example, she found that using the category name "prepared" people better for a judgment about a typical

---

[7] Another feature of Rosch's study is interesting as well—she used a traditional paired-associate learning task with the Dani subjects, partly because its standard methodology would yield convincing results in this otherwise unusual and original experiment. Occasionally, it pays to use a thoroughly understood "shovel" when you're digging for something new.

member than for a judgment about an atypical member (1975; more about this priming effect later). Members that are judged typical of a category, furthermore, tend to share more common features than those that are judged atypical. Typical members also tend to share *fewer* features with members of *other* categories (e.g., Rosch & Mervis, 1975; see Rosch, 1978).

Far from finding evidence that category members are "equal" in the concept formation sense, Rosch's extensive program of research revealed repeatedly that natural concepts and categories have an internal structure. Real-world features do *not* occur independently of one another, although that is the way that shape, color, etc., are manipulated in the concept formation task. Instead, features or properties of real objects come in correlated bundles. As a simple example, while shape and color may be independent in concept formation tasks, "wings" and "beak" are anything *but* independent in the real world—the things in the real world that have wings *often* have beaks too. We structure our mental representations of such categories in terms of these correlated features, with *typical* instances of the category stored centrally, at the "core" of the concept's meaning, and with *atypical* instances stored more peripherally

Prototypes The term that Rosch proposed for *the central, core instance of a category* is **prototype**—a "really red red" is a *prototypical* red, a "doggy dog" (1975, p. 198) is a prototypical dog. Rosch advanced the argument, then, that our mental categories are represented in terms of a prototype, with typical members stored close to the prototype, and peripheral members stored farther away. When asked to think of or imagine a dog, for example, you generally think of your prototype—few of us would immediately think of a Chihuahua, but instead would think of a German shepherd, a terrier, or some other more "doggy dog." Think back now to our standard semantic memory category. Did your dictionary have a picture of a rather ordinary looking, typical, yet nondescript bird next to the definition, and the same sort of generic picture of a flower? Such neutral, nondistinctive pictures are probably quite close to the mental representation of a prototype, the central, organizing representation in natural categories.

What other kind of evidence is there that typicality and prototypes are important issues in the representation of word and category meanings? One such demonstration involves *hedges*, statements that limit or restrict a sentence in some fashion. Several investigators have noted how our language permits us to hedge a statement, to qualify it in a sense (e.g., Lakoff, 1972). Thus, one of Smith et al.'s (1974) early points on typicality was that atypical category members are often hedged with the word "technically," as in "Technically, a chicken is a bird." The hedge, of course, indicates that there's something about the statement that might be odd without the hedge. Conversely, statements that are genuinely false, like "A bat is a bird," can be hedged with the phrase "loosely speaking"; the semantically related concept is close in meaning, but mismatches on some important features. Hedging, of course, implies a typicality effect, as well as something about our sensitivity to the typicality of concepts. Notice how wrong the sentences are when the hedges are reversed, for instance *"Technically speaking, a bat is a bird" and *"Loosely speaking, a chicken is a bird." And hedging a typical member strikes everyone as unacceptable, too—no one would say *"Technically, a robin is a bird."

For sheer cleverness, however, none of the typicality investigations matches a study by Rips (1975). Rips used a dramatically different kind of task to obtain his evidence. Subjects read a story about an island inhabited by only eight species of animals: sparrows, robins, eagles, hawks, ducks, geese, ostriches, and bats. One group of subjects then read that a highly contagious disease had been discovered among all of the sparrows; another group read that the robins had the disease, another that eagles were affected, and so forth. Subjects were then asked to estimate the percentages of the *other* animals that would also contract the disease. The estimates yielded overwhelming evidence of a typicality or prototype effect. Species that are rated as quite typical or close to the prototype, such as sparrows, were judged very likely to infect virtually all the other species. Atypical instances, say geese, were judged likely to infect only other atypical members, ducks for instance, and this to a much lesser degree. The underlying issue, of course, was typicality, the representativeness of species within the overall category. Subjects assumed that if a typical instance had an important property, then that was sufficient to predict that *all* instances, typical and atypical, would share the property as well. If the property was only true of an atypical instance, however, subjects tended to doubt that the property would be shared throughout the category.

Internal Structure and Categorization The generalization you should remember from all of this research is that categories and concepts in semantic memory are *internally structured* by semantic relatedness. Categories provide us with a way of classifying objects in our environment, and thereby predict what the objects do, what properties they possess, and how they may be equated under certain circumstances but not under others. To illustrate, tame and wild turkeys are roughly equivalent if you're classifying on the basis of "things to have for dinner," but are not equivalent if the classification is based on "legal game in hunting season." Likewise, while a hammer and a pair of pliers belong to the same superordinate category of tools, they are not equivalent under most circumstances (but according to a principle in my Murphy's law calendar, "in a pinch, the nearest tool becomes a hammer"). We have central, core meanings for each category and concept, and these can be represented as prototypes. Other instances of the concept are arrayed around the prototypes, in a semantic distance sense, and can be thought of in terms of their similarities and dissimilarities to the prototype. Bruner et al. (1956), writing at the very beginning of modern cognitive psychology, claimed that "categorization is the *means by which the objects of the world about us are identified*" (p. 12). To this, we would add that the process of categorization, relying on existing general world knowledge, is guided or even determined by semantic relatedness and category structure.

Concept Formation Reconsidered

As a postscript to this discussion of category structure, classification, and prototypes, reconsider the concept formation stimuli in Figure 6–6. Recall that Rosch (1973) viewed such arbitrary patterns, and the concepts specified by them, as highly artificial and quite dissimilar to natural concepts and categories. Interestingly, however, it seems that people do abstract a prototype-like representation of such stimuli as they "acquire" the tested concept.

A standard demonstration of this point is a study by Posner and Keele (1968). They constructed prototypical patterns, such as a triangle, along with various "distortions" of the patterns, all illustrated in Figure 6–7. Subjects were then shown the distorted patterns and had to classify them into their appropriate category (four separate categories, and four "families" of distortions, were used). After performing in this training portion of the study, subjects were tested in a transfer portion, to see if their training on the categories would transfer to new patterns. Would subjects be able to classify the new patterns accurately? Naturally, the greater the distortion in the stimuli, the slower the original learning. Furthermore, subjects were better at classifying distortions they had already seen than those they had not. The critical result, however, was that subjects classified the prototypical patterns just as easily as they classified the familiar distortions, even though they had *never* seen the prototype patterns. Thus, subjects seemed to have abstracted the central tendency or core of the category from the "flawed" or distorted patterns they had originally seen. They represented the central tendency of the category as a kind of prototype.

Even more remarkable was a study reported by Bourne (1982), whose earlier concept formation research had been criticized by Rosch (1973) on

FIGURE 6–7

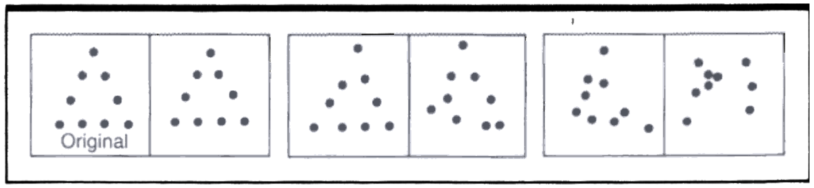

A prototypical triangle pattern is shown on the left along with various distortions of the prototype. The degree of distortion increases going from left to right. In Posner's research, subjects acquired the concept from various distorted patterns; they categorized the prototypical patterns as most representative of the concept, even though they had never seen it. (Adapted from Posner, Goldsmith, & Welton, 1967.)

the grounds of artificiality. Bourne reported a traditional concept formation task, in which two critical features defined the concept. The rule that defined the concept was that the presence of either of the two critical features—call them *x* and *y*— made a stimulus "positive." The interesting manipulation in the experiment was on those stimuli that had *both* of these features present, a condition labeled *xy*. As in all concept formation tasks, subjects' answers were followed by feedback, which told them whether their judgments were correct or not. Bourne varied the feedback on the trials that presented the combination *xy* stimuli; one group was told 100% of the time that these were "positive" instances, one group was given positive feedback on 75% of these trials, one group on 50%, one group on 25%, and one group on 0%. In other words, the 100% group was always told that *xy* stimuli were positive instances, the 0% group was always told that *xy* stimuli were negative instances, and the other groups were given mixed feedback.

Not surprisingly, subjects were rather slow in learning which response to make to *xy* stimuli, since either the *x* or the *y* feature by itself made for a "positive" instance. Early in the tests, virtually all subjects responded "positive" to these stimuli. Slowly, however, the feedback given the mixed and 0% groups began to take effect, so that by the last of the six concept formation tests, subjects in the 0% feedback group called *xy* stimuli "positive" only 1% of the time (i.e., a 1% error rate, given their feedback).

This typical concept formation procedure was followed by several special tests. In one of these, subjects were asked to rank stimulus patterns from the best example of the concept to the worst. Despite the nature of the task and the varying percentages of feedback, subjects rated the *xy* patterns as typical of the concept. In particular, subjects in the 25% condition rated the *xy* patterns as *most typical,* even though these patterns had only been called "positive" on 25% of their trials. Even subjects in the 0% condition, who *never* saw *xy* pairs labeled "positive," rated them as more typical than the other "negative" patterns—by rights, of course, *xy* should have been just as atypical as any other pattern that had always been negative for these subjects. In short, "even with relatively modest feedback to the effect that *xy* instances are *sometimes* positive, subjects, on the average, ranked *xy* instances as the *best* example of the concept that had been learned" (Bourne, 1982, p. 7, emphasis added). The typicality effect applies, apparently, even when supposedly artificial patterns, and the laboratory concept formation task, are used to study classification.

X

▾ Priming Semantic Memory

In chapter 3 you read about automatic processing, in which mental processes occur without any intention or conscious effort. At that point, we introduced the notion of **priming**—a word is presented, for example, and

because of automatic access to its meaning, priming boosts or *causes an activation of both that concept and other concepts that are related to it*. The effect of priming, in that discussion, was to influence subjects' performance in the shadowing task. In particular, it was found that subjects will switch their attention from one auditory channel to another if the unshadowed, unattended channel completes the meaning of the sentence they're attending. Thus, the phrase "While Bill was walking through the forest" was said to prime concepts in memory that are related to forests, e.g., trees and so forth. Because those concepts were primed, the phrase "a tree fell across his path" (see Figure 3–12) intrudes into the subject's shadow. In short, priming resulted in the activation of related concepts, and thus influenced the subjects' selective attention.

In the chapter 3 material we relied on your general intuitions as to what words might be related to a priming word, as in the "forest-tree" example. But now, you've just studied several sections on the topic of semantic relatedness and you have a much better idea of how words and concepts are related to one another. You even read of one specific model that predicts these priming effects, the Collins and Quillian (1972; Collins & Loftus, 1975) network model. Recall from that model that a concept becomes activated and then spreads activation to its neighbors, to those concepts it is semantically related to. You also read that activation will persist for a period of time, but that it eventually will decay.

Priming Tasks

Let's introduce some precise vocabulary for these ideas to facilitate the explanation that follows. *Any stimulus that is presented first*, in hopes that it will influence some later process, is called a **prime**—simple enough, you'll agree. The *stimulus that follows the prime* is the one we expect will be influenced in some way. This later stimulus is called the **target,** since it is the presumed destination of the activation or priming process. Thus, primes precede the targets, and targets are influenced by the primes. Sometimes this influence is beneficial, as when a prime makes the target easier or faster to process. This kind of *positive influence on processing* is referred to as **facilitation,** or sometimes simply **benefits.** Occasionally, the influence is negative, as when a prime is misleading. When the prime slows down performance to the target, the *negative influence on processing* is called **inhibition;** we also say, equivalently, that there were **costs** associated with the prime.

Finally, since priming is a level of mental activation that will eventually dissipate, we often need to keep track of the period of time that intervenes between the prime and the target. In some studies, this period of time is filled with other stimuli or trials. In this case, the **lag** between prime and target, usually the *number of intervening stimuli*, is our index of the separation between prime and target; "lag 2" would simply mean that two trials came between the prime and the target. In other studies, the period of time between the prime and the target is of interest. However nonintuitive, this period of time is often referred to as the **SOA**, the **stimulus onset asynchrony;** if we consider the prime and target to be the two halves of a complete stimulus, then the onset or beginning of the two halves occurs asynchronously at different times. Thus, we might present a prime, and then 500 msec later present the target. This would correspond to a 500 msec SOA interval, where SOA is *the length of time that separates the onset of the prime and the target.*

Priming is generally viewed as the most fundamental process of retrieval from semantic memory. Because of its importance, cognitive psychology has developed a variety of specific methods and tasks to examine priming. To illustrate how these methods were developed, and to explain how these tasks reveal the phenomenon of priming, we'll examine a series of experiments by Posner and his colleagues. Posner began with a simple letter matching task to study how people recognize and categorize these visual patterns. From this, he then developed a priming task to see how advance information—the prime—might influence these judgments. We'll begin with the simple matching task, and then describe the elaboration of that task which incorporates priming. This description should set the stage for an in-depth look at semantic priming, and the many different kinds of experiments that have investigated this effect.

The Posner Matching Task You're in a very easy RT task. You are shown stimuli that contain two letters and your task is to judge whether the two letters are the same or different. If you see a pair of letters like those in panel A of Figure 6–8, the AA pair, your RT to press the button labeled "same" is very fast compared to your RT for any non-matching pair such as AB (Posner, 1969; Posner & Mitchell, 1967; Posner & Snyder, 1975; see Figure 6–8A). This is called the *physical identity* condition, in which you make same/different judgments depending on whether or not the two letters are physically the same.

Now consider a different rule for making the judgments. Judge whether or not the two letters have the same name, the *name identity* condition. Here, of course, we can present either AA or Aa as the stimuli, and both pairs should be judged "same." Because of their physical identity, however, AA will still be processed somewhat faster than Aa, about 70–100 msec faster; see panel B in Figure 6–8. The basic idea here is that you need to process the stimuli beyond a physical point in order to know if they have the same name since nothing about the physical appearance of A and a tell you that they have the same name. Nonetheless, the physical identity of AA speeds your processing of that stimulus pair, even when your judgments are to be made on a name identity basis.

Finally, consider the third condition, a *semantic identity* or *rule identity* condition: say "same" if both stimuli are vowels or both are consonants, otherwise, say "different." Under this rule, pairs like AA, Aa, and Ae should all be judged "same," since they are all vowels. Nonetheless, just as was found for the name identity judgments, a simpler basis for making the decisions is also possible part of the time. As panel C in Figure 6–8 shows, the simpler bases speeded performance considerably. In particular, under rule identity instructions, the pair AA is judged with an RT of 549 msec, Aa takes 623 msec, and Ae requires 699 msec; Bc, a "same" judgment because both are consonants, takes 904 msec, and AB, requiring a "different" decision, takes 801 msec. From these results, Posner (1969) concluded that even for the complex rule-based decision, simpler mental codes—first, a visual code reflecting physical characteristics, then a "name code"—can enter into the subjects' decision making.

FIGURE 6–8

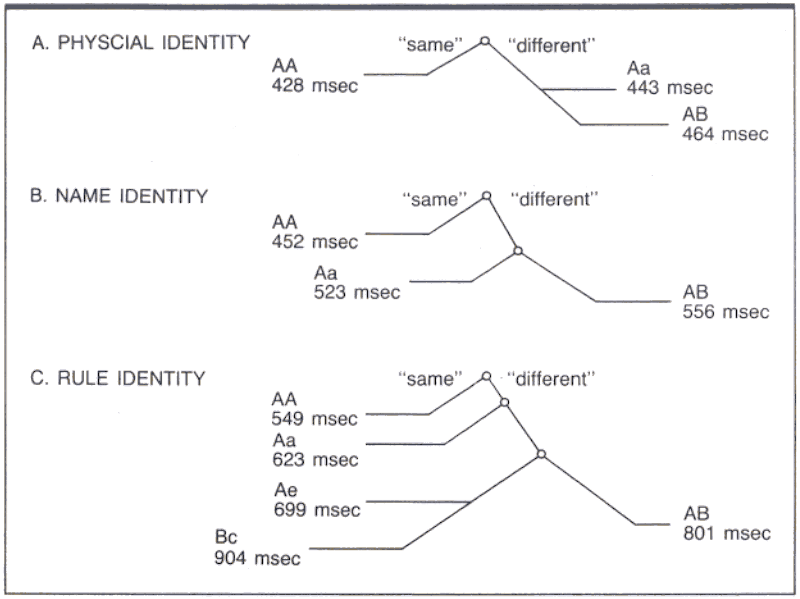

Reaction times to Posner's letter matching task, under three different rules. In panel A, subjects responded "same" if the two letters were physically identical; in panel B, they responded "same" if the two letters had the same name; and in panel C, they responded "same" if both were consonants or both were vowels. (After Posner, 1978.)

Letter Matching with Primes With this as background, let's now consider an elaborated version of this task, where the elaboration involves presenting a prime stimulus followed by some target. Sticking with letters for the moment, Posner and Snyder (1975) presented a series of experiments that examined how priming might affect the matching of letters according to different rules. In the first study, two important variables were investigated, the kind of prime that was presented and the proportion of trials on which the prime was relevant to the upcoming target. On half of the trials in the experiment, the prime stimulus was neutral with respect to the target—a plus sign appeared before the target. On the safe assumption that the plus sign should have no specific effects on letter matching, the plus condition served as a neutral baseline. The other half of the trials presented a prime that was a letter. Part of the time, the prime matched the pair of letters in the target—A was the prime for AA. This condition was labeled the "relevant prime" condition, since the prime was relevant for the upcoming target stimulus. Part of the time, however, the prime did not match the target letters—B as the prime for AA. This is called the "irrelevant prime" condition. As in the simple matching task, subjects were instructed to say "same" if the two letters in the target were physically identical, and "different" otherwise; they were told that the prime could simply be treated as a warning signal.

The other variable of interest in this study was the percentage of the trials on which the letter prime matched the target, i.e., was relevant to the decision the subject would have to make regarding the target. For some subjects, the letter prime matched the target 80% of the time, and didn't match the other 20% of the time; this is the 80–20 condition. For another group of subjects, the prime matched 50% of the time, and didn't match the other 50% of the time—the 50–50 group. For a third group, primes matched only 20% of the time, and mismatched 80% of the time. The purpose of the varying percentages was to lead the subjects to different expectancies; for instance, if only 20% of the primes you see are relevant to the target, you'd eventually come to expect that very few of the primes would be helpful for processing the targets. This expectancy might influence your performance, as would the opposite expectancy, i.e., that most of the primes are relevant and very few irrelevant.

Table 6-1 shows the RT results that were obtained. Look first at the 80–20 condition. With the neutral prime, RT to judge AA the same was 414 msec. With the relevant A prime, however, RT was 329 msec, an 85 msec *speeding* of the AA match due to the useful A prime—we call that an 85 msec facilitation. Conversely, getting B as a prime before AA yielded an RT of 450 msec, a 36 msec *slowing* of the AA match; putting it slightly differently, subjects benefitted 85 msec from the relevant prime, but the irrelevant B prime cost them 36 msec. Facilitation was only a bit smaller—71 msec—in the 50–50 condition, and was still significant, although reduced, even in the 20–80 condition (31 msec). In other words, even when the prime matched the target only 20% of the time, seeing A before AA facilitated the matching process significantly. On the other hand, the 14 msec cost of an irrelevant prime in the 50–50 condition was not significant, and the cost was literally 0 msec in the 20–80 condition.

Table 6–1

Conditions	+ AA	A AA	B AA	Benefits	Costs
80–20	414	329	450	85 (414–329)	−36 (414–450)
50–50	429	358	443	71 (429–358)	−14 (429–443)
20–80	439	408	439	31 (439–408)	0 (439–439)

Benefits are computed as RT to the neutral prime minus RT to the relevant prime; in other words, the amount by which primed RT is faster than neutral RT. Costs are computed as RT to the irrelevant prime minus RT to the neutral prime; in other words, the amount by which irrelevantly primed RT is slower than the neutral prime RT.

Adapted from Posner & Snyder, 1975.

In their second experiment, Posner and Snyder (1975) used only the 80–20 condition, and focused their attention on another variable, the time course of priming. The task was exactly as before, but now the target followed the prime by some interval of time—SOAs ranged from 0 to 500 msec. These results—facilitation and inhibition of letter-matching times across SOAs—are presented in Figure 6–9. Notice first how rapidly the relevant prime began to benefit the subjects' responses. Even by the 75 msec SOA interval, benefits were about 30 msec, and this facilitation remained fairly constant up to the 500 msec SOA interval, a full half-

FIGURE 6-9

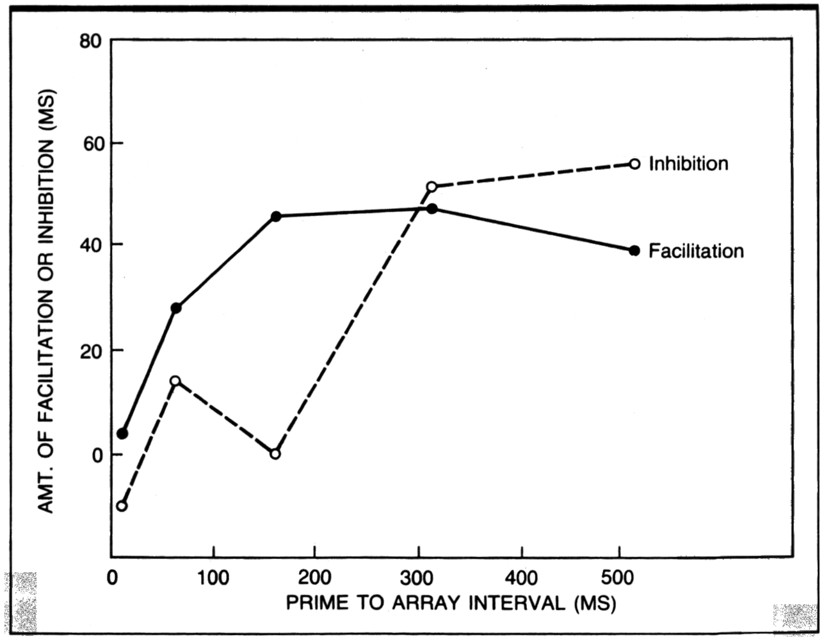

Facilitation and inhibition of matching times due to priming; for the solid curve, a relevant prime *A* speeds processing of the matching pair *A A*, even at the shortest SOA interval. In the dotted curve, an irrelevant prime *B* begins to slow down processing of the pair *A A*, but only after about 150 msec have elapsed. (From Posner & Snyder, 1975.)

second after the onset of the prime (bear in mind that this was a physical identity task). Receiving a misleading or irrelevant prime, however, altered the pattern across time. There was no particular cost to the irrelevant prime until after about 150 msec (the apparent cost at 75 msec was not significant), but thereafter the inhibition was quite strong.

What are the ingredients in this task that enable us to examine priming? First, we have primes and targets being presented with a same/different or yes/no decision required on the target. In the neutral prime condition, RT to the target serves as a baseline—this is how long the mental processing takes under neutral conditions. From this baseline we can now examine the effects of relevant and of irrelevant primes. If a prime is relevant, it activates some information in memory that will be useful for the upcoming target. Thus, RT to the target should be facilitated when the prime is relevant. Facilitation can be examined at very short SOAs, where only automatic processes are thought to be possible, all the way up to rather long SOAs, where conscious processes can be involved. As Posner's results showed, we can expect to observe facilitation under both automatic and conscious processing conditions; that is, automatic priming of related information speeds processing at short SOA intervals, and conscious processes related to the prime have their effect at longer SOAs.

Conversely, if the prime activates information that is of no use during processing of the target, then two effects are obtained. With very short SOAs, no effect of irrelevant primes is expected—the activation of information that is irrelevant to the upcoming target does not influence the processing of the target itself. At longer SOAs, however, we observe an increasing amount of inhibition due to the irrelevant prime—as the conscious expectation takes effect, more and more time is needed to recover from the misleading effects of the irrelevant prime. In short, as the SOA interval grows longer, conscious processing begins and starts to be demonstrated in the RT results.

Manipulating the subjects' expectations about the relevance of a prime consists of a manipulation of the proportion of trials on which the prime is indeed relevant. Importantly, the expectation-based effect is assumed to be under influence of conscious processes. Any effect of this expectancy, then, is interpreted jointly with the SOA interval—if the prime is relevant because of conscious expectation, then this relevance should show up at later SOA intervals, when conscious processes begin to emerge. Conversely, if the prime mismatches the conscious expectation, then this inhibitory effect would be present at the longer SOA intervals.

Semantic Priming Studies

A variety of tasks have been used to examine semantic priming, most of them containing the basic elements displayed in the Posner and Snyder priming study; i.e., primes of varying degrees of relevance to the targets, intervals of time (or intervening trials) between primes and targets, and

sometimes expectancies about the relevance of primes to upcoming targets. We'll examine a direct extension of the Posner matching task first, and then consider several other studies that also assessed priming effects, using tasks such as word naming, sentence verification, and lexical decision.

Priming and Typicality Rosch (1975) adapted the Posner matching task to provide further evidence for her notion of typicality in semantic categories. She used category names as primes, followed after a 2000 msec SOA by a pair of words or a pair of pictures as the target. Subjects were to respond "same" if both target stimuli belonged to the same category, and "different" otherwise (in other experiments in the same report, she examined simultaneous presentation of primes and targets, different SOA intervals, and a variety of other factors; we'll focus on just one of her experiments, however).

Look first at the middle panel of Figure 6–10, when two different members of the same category were presented in the target. For these "same" judgments, typicality had an obvious and strong effect; for all four conditions, pictures vs. words that were primed vs. unprimed, typical members were judged "same" more rapidly than atypical members. Note also that pictures were uniformly processed faster than words, suggesting that the simple line drawings Rosch used were somewhat closer to the subjects' mental representation of a prototype (this was even the case in Experiment 7, where picture and word trials were intermixed).

Now look at the left panel, where the targets were physically identical pictures or words. Just as Posner found, when the two target stimuli are physically identical, responding is even faster, since this simpler basis is available to assist the judgments. Beyond that effect, however, notice that on the physically identical targets, typicality and priming had somewhat different effects than before. For the high-typical pairs, priming facilitated the judgments—primed pairs were processed significantly faster than unprimed pairs. But, for items at the low-typicality level, priming with the category name actually slowed down the processing. For example, even though the target was "peacock peacock," a low typical, physical-identity pair, RT was 100 msec slower when that pair was primed by BIRD than when it was not primed at all. Apparently, the prime word activates the prototype of its category, as well as typical members. This results in a benefit for those items that resemble the prototype; as we said before, the category name "prepared" subjects for typical items. For atypical items, however, it's almost as if the category name is misleading. For items at the medium level of typicality, primes had neither positive nor negative effects.

Word Naming and Sentence Verification An early investigation of word naming and semantic priming was conducted by Freedman and Loftus (1971). These investigators asked subjects to produce a member of

FIGURE 6-10

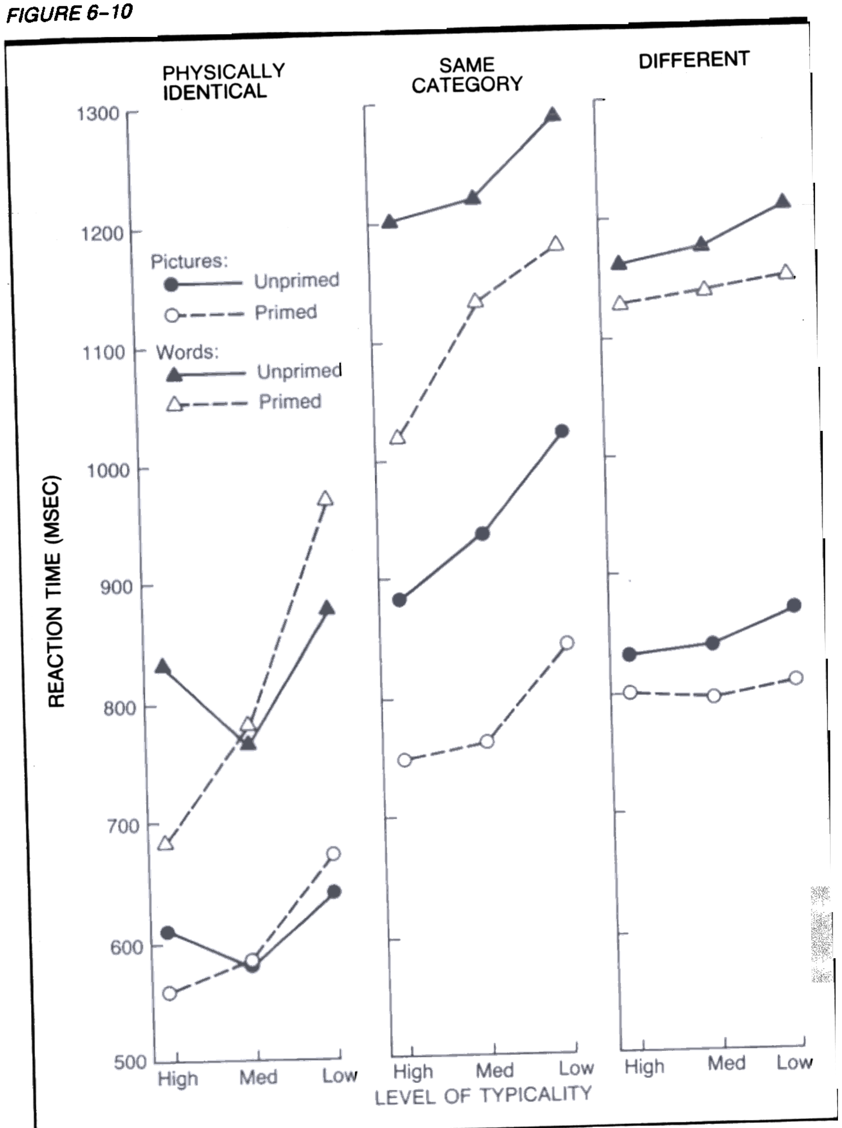

Reaction time in a "same/different" task, in which pairs of pictures or words were shown, and subjects said "same" if both members of the pair were from the same category. Pair members were either physically identical (left panel), from the same category (middle panel), or from different categories (right panel). Pair members were either high, medium, or low in typicality, and were preceded by either a neutral or relevant prime. The SOA was 2 seconds. (From Rosch, 1975.)

---

a category that either began with a certain letter or was described by a certain adjective: e.g., "Name a fruit beginning with P" or "Name a red member of the flower category." Half of the trials showed the letter or adjective first, and then the name of the target category; in other words, the letter or adjective served as the prime on these trials, since it came first. In the other half of the trials, the reverse order was used—the category name was the prime, and the letter or adjective was the target. Freedman and Loftus measured the RT of the subjects' responses beginning with the target, the second bit of information that was presented; see Table 6-2 for a summary of these conditions.

Freedman and Loftus found that performance was significantly faster when the category was used as a prime (fruit-P) than when the letter or adjective was the prime (D-mammal; red-fruit). This suggested that the category name activated its semantic representation and then primed the members of the category. When the letter or adjective targets were then presented, priming from the category name had made it easier to access a member of the category—the members had already been primed. Conversely, receiving a letter or adjective as a prime had very little effect. No relevant activation of potential targets was possible with such primes, so no facilitation of word-naming latency (RT) was found. (If you believe that the subjects *can* rapidly activate words based on a letter prime, try the following demonstration. Check your watch, then give yourself 10 seconds to name as many words as you can that begin with D. You'll be surprised at how few you can name, particularly when you contrast that with naming members of a semantic category, say, fruits.)

Table 6-2 SAMPLE EXPERIMENTAL TRIALS FROM FREEDMAN AND LOFTUS (1971), AND LOFTUS AND LOFTUS (1974)

		Prime	Target	Sample Response
FREEDMAN & LOFTUS				
Letter or adjective first	1	P	FRUIT	Pear
	2	RED	FLOWER	Rose
Category first	3	MAMMAL	D	Dog
	4	WEAPON	G	Gun
LOFTUS & LOFTUS				
Target at lag 0 (#1 to #2)	1.	P	FRUIT	Pear
	2.	RED	FRUIT	Apple
Target at lag 2 (#3 to #6)	3.	B	FRUIT	Banana
	4.	RED	FLOWER	Rose
	5.	D	MAMMAL	Dog
	6.	RED	FRUIT	Apple

In both experiments, the task is to name a word that meets the restrictions stated in the prime and target.

In a similar, but slightly more complex study, Loftus and Loftus (1974) again presented both "category-letter" and "letter-category" trials, and then tested the same category a second time, at lags 0 and 2; these conditions are also summarized in Table 6–2. The results of this study are presented in Figure 6–11. Panel A shows RT to name a target when the category and letter were presented simultaneously (the advantage for the category-letter order here was presumably because the category name was read first, probably because of the physical arrangement of the stimulus terms; this is unspecified in the original report). As the figure shows, when a second retrieval from the category was required, subjects' responses were faster—it took almost 2.2 sec on the initial trial, but just under 1.9 sec on the target trial at a lag of 0. Furthermore, this benefit or facilitation of processing on the target trial was still apparent even when two unrelated trials intervened between the prime and target trials, i.e., at lag 2.

Panel B displays the results from the other condition, when the prime preceded the target by 2.5 sec (a 2.5 sec SOA). While the same pattern of priming was obtained here, notice also that processing was faster overall in this condition. The 2.5-sec interval shaved nearly 400 msec off the RTs, presumably because relevant concepts were activated during the SOA interval. Notice that this study showed both kinds of priming effects: First, the priming of targets within a single trial was examined, looking at performance under a 0 msec vs. a 2.5-sec SOA; second, priming across trials was examined, looking at the speeding of performance at lag 0, then the somewhat lesser degree of facilitation at lag 2.

Similar outcomes have also been reported in the sentence verification task we studied earlier. For example, Ashcraft (1976) found that a sentence such as "A sparrow has feathers" facilitated another sentence about the same category, e.g., "A robin can fly," but only if the target sentence concerned a frequent or important property. Furthermore, this facilitation effect was quite strong at lags 0 and 1, but was weak or nonexistent at lag 4. Unlike the Loftus and Loftus study, where prime and target trials included repetition of the category word, this sentence verification study avoided repeating any words. This eliminated the possibility that the priming effect was due merely to the ease of re-encoding a repeated word.

FIGURE 6–11

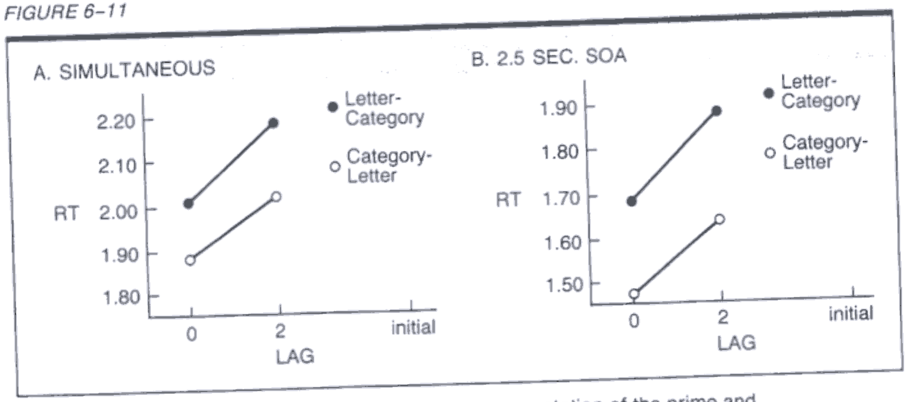

Reaction time in seconds is shown for simultaneous presentation of the prime and target (panel A) and an SOA of 2.5 seconds (panel B). In both panels, the RT to the prime is shown on the right ("initial"); the curves show RT to the targets at lags 0 and 2 (From Loftus & Loftus, 1974.)

Priming Lexical Decisions Another line of research has demonstrated the same kind of results, and has supported the same kind of network interpretations while also showing the intimate and unavoidable relationship between our semantic concepts and the words we use to name them—in other words, between the semantic and the *lexical* entries in memory. You've already encountered this other kind of research back in chapter 2, when we discussed process models. At that point, you read about the *lexical decision task,* in which subjects judge as rapidly as possible whether or not a string of letters forms an English word (retrieval cues: OFFICE, MANTY, MOTOR). This task has been adapted to tests of semantic priming by presenting two letter strings per trial, either of which may be a word or a pseudoword (a pseudoword looks like a real word and can be pronounced). The condition of greatest interest of course is the set of trials when both letter strings are words.

As Meyer and Schvaneveldt (1971; also Meyer et al., 1975) found, two related words such as BREAD BUTTER can be judged more quickly as "words" than two unrelated words, such as NURSE BUTTER. Table 6–3 displays the Meyer and Schvaneveldt (1971) results, and shows this effect quite clearly. Related words were judged in 855 msec, compared to 940 msec for unassociated words.

One particularly interesting aspect of these results, and in fact all results with the lexical decision task, is worth pointing out here. It is not logically necessary for subjects to access the *meanings* of words in the lexical decision task; technically, they need only "look up" the words in some sort of *mental dictionary* or **lexicon**, determining if the word is

Table 6–3 MEAN REACTION TIMES (RTS) OR CORRECT RESPONSES AND MEAN PERCENT ERRORS IN THE YES-NO TASK

Type of Stimulus Pair		Correct Response	Sample Stimuli	Mean RT (msec)	Mean % Errors
TOP STRING	BOTTOM STRING				
word	associated word	yes	nurse-doctor	855	6.3
word	unassociated word	yes	bread-doctor	940	8.7
word	nonword	no	book-marb	1087	27.6
nonword	word	no	valt-butter	904	7.8
nonword	nonword	no	cabe-manty	884	2.6

there or not. Yet, the results repeatedly show the influence of a word's meaning—it is the meaningful connection between BREAD and BUTTER that facilitates this decision, rather than some lexical connection (you might think since both begin with B that there is a lexical basis for the facilitation, but the same benefits are found with word pairs that are quite dissimilar in spelling, for instance NURSE DOCTOR). Apparently, even though the task only requires a judgment of word or nonword, we nonetheless access the words' meanings as part of our normal processing. Therefore, the relatedness of meanings influences performance. This should not be a surprise to you—it is exactly the same effect that Stroop found with his color words (see chapter 1). In fact, the Stroop task is the conceptual "grandfather," so to speak, of all such priming tasks and automaticity effects (though as mentioned in chapter 1, there is no hint in Stroop's 1935 article that it would become so important many years later).

Recently, Marcel (1980, 1983) has reported an impressive extension of such results. Marcel examined word recognition in a priming task from the standpoint of automaticity and conscious awareness of the prime. He presented primes and targets, as did Meyer, but with a twist—the prime was immediately followed by a scrambled visual pattern. The purpose of this scrambled pattern was to mask the prime, that is to present the masking pattern so soon after the prime that subjects were not consciously aware of the prime word at all. You'll recall from the Visual Sensory Memory research in chapter 3 that the circle marker cue in Averbach and Coriell's (1961) research acted as an "eraser" on round letters such as C. This effect was *backwards masking*, when a stimulus prevents further processing of an earlier one by means of interference within the visual system. This phenomenon was exactly the one Marcel used to prevent conscious awareness of the prime. By following the prime so rapidly with a masking pattern, he obtained complete backwards masking. Subjects claimed they had seen no prime whatsoever. Nonetheless, relevant primes such as "child" facilitated lexical decisions about words such as "infant" (see also Carr et al., 1982). Semantic priming, at least in the lexical decision task, appears to occur automatically. Especially dramatic is the fact that priming occurred even though subjects did not realize consciously that they had seen a prime at all.

Extensions of the Lexical Decision Task

Two extensions of the lexical decision task will complete our examination of semantic priming. The first involves a look at automatic and conscious priming processes, where conscious expectations about the target were manipulated. The second concerns ambiguous words—words that have more than one meaning—and how those several meanings are retrieved. The first deals with issues we've been discussing throughout the chapter, whereas the second begins to introduce a new issue, that of context, which we will pursue in several more chapters.

Expectancies and Semantic Priming An impressive set of demonstrations of priming and automatic word retrieval has been reported by Neely (1976, 1977). In one study (1976), Neely selected word pairs so that one of the members was the primary associate to the other at least 40% of the time (based on free association norms). These related pairs were contrasted with unrelated word pairs, and also with a condition in which the neutral letter X was paired with a word. For all trials in the experiment, subjects had to judge whether the target string, the second member of each pair, was an English word—the lexical decision task. As Figure 6–12 shows, the processes involved in making these lexical decisions were greatly facilitated when the prime was a related, associated word. Facilitation grew from 17 msec at the shortest SOA to 56 msec with a 2000 msec (2 sec) SOA. Inhibition was observed for the unrelated word pairs; there was a relatively constant 16 msec cost of receiving an unrelated word as a prime.

In a more thorough examination of priming, Neely (1977) not only tested the effects of semantic relatedness on performance in the lexical decision task, but also the possible effects of expectancy. But rather than manipulate expectancy by altering the proportion of relevant and irrelevant primes, Neely decided merely to tell subjects what kind of targets

FIGURE 6–12

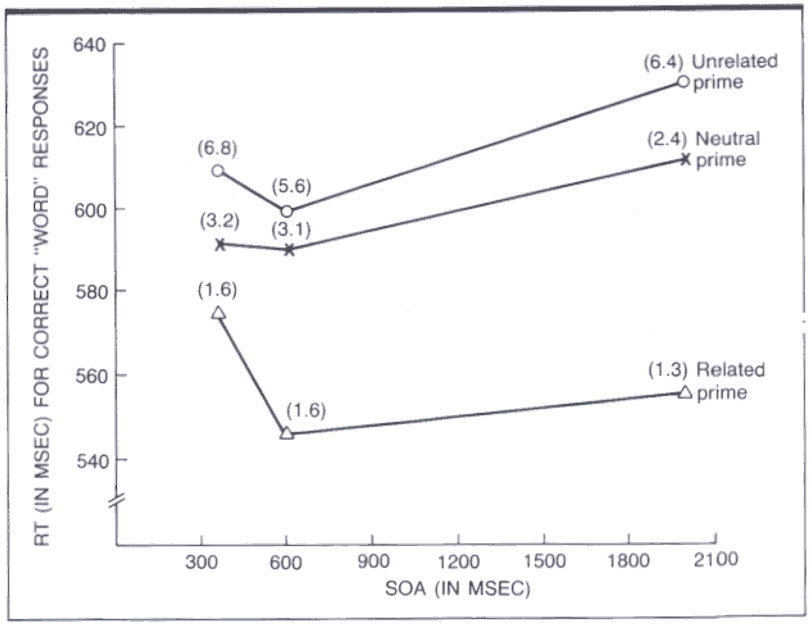

Reaction time to lexical decision targets is shown across SOA intervals for unrelated, neutral, and related prime conditions. The numbers in parentheses are the error rates in each condition. (From Neely, 1976.)

would usually follow which primes. For some of the categories he tested, subjects were told to expect a target from the same category—seeing BIRD was a tip-off that a member of the bird category was very likely to be the target. For other categories, subjects were told to expect a switch—seeing BUILDING tipped the subjects off that a member of the body-part category was likely to appear in the target position, and vice versa.

The clever part of the study (although the subjects probably didn't think so) was that part of the time the tip-offs were correct and part of the time they were misleading. In other words, you were told that when you saw BIRD you'd then see a member of that category as the target. That usually happened. You were also told that when you saw BUILDING, you'd then probably see a body part as the target, and that when you saw BODY, you'd probably get a part of a building as a target. Most of the time this happened as you expected—BUILDING was followed by LEG, BODY was followed by DOOR or WINDOW. But part of the time it didn't happen as you expected—BUILDING would then be followed by WINDOW (or BODY followed by HEART), even though you expected the other category. The cleverness here, of course, is that BUILDING would normally be a relevant prime for WINDOW, but because of the expectancies, subjects assumed they'd see a body part. Table 6–4 shows the various conditions Neely tested, along with examples of the stimuli.

Neely's results were fascinating (see Figure 6–13). For regular priming (BIRD-ROBIN), facilitation of word recognition was observed across the whole range of SOAs, a replication of the Posner and Snyder (1975) results with semantically related words instead of letters (and without the repetition of the prime in the target—BIRD then ROBIN as opposed to A then AA). When BODY was the prime, there was facilitation for the expected BUILDING targets, but *only* at the longer SOAs. In other words, it took a while for the expectation of a category switch to take effect, but after it did there was observable facilitation. The truly impressive part of the results was that BODY yielded significant priming to HEART at the short SOAs, even though the subjects had been told to *expect* BODY to be followed by a part of a building. In other words, without enough time for the expectation of a category shift to affect processing, the prime facilitated its own related category members rather than those in the "expected" category. On the other hand, at long SOAs, when the expectation did have time to affect performance, BODY yielded inhibition on its own category members when the expected shift to parts of a building *didn't* occur (BODY-HEART). This inhibition was similar to the BODY-SPARROW condition, a switch that the subjects *never* expected.

Table 6–4 CONDITIONS AND SAMPLE STIMULI IN NEELY'S (1977) STUDY

Condition	Sample Stimulus
No Category Shift Expected	
No Shift	BIRD-robin
Shift	BIRD-arm
Category Shift Expected from Building to 'Body Part'; from Body to 'Part of a Building'	
No Shift	BODY-heart, BUILDING-window
Shift to Expected Category	BODY-door, BUILDING-leg
Shift to Unexpected Category	BODY-sparrow

FIGURE 6–13

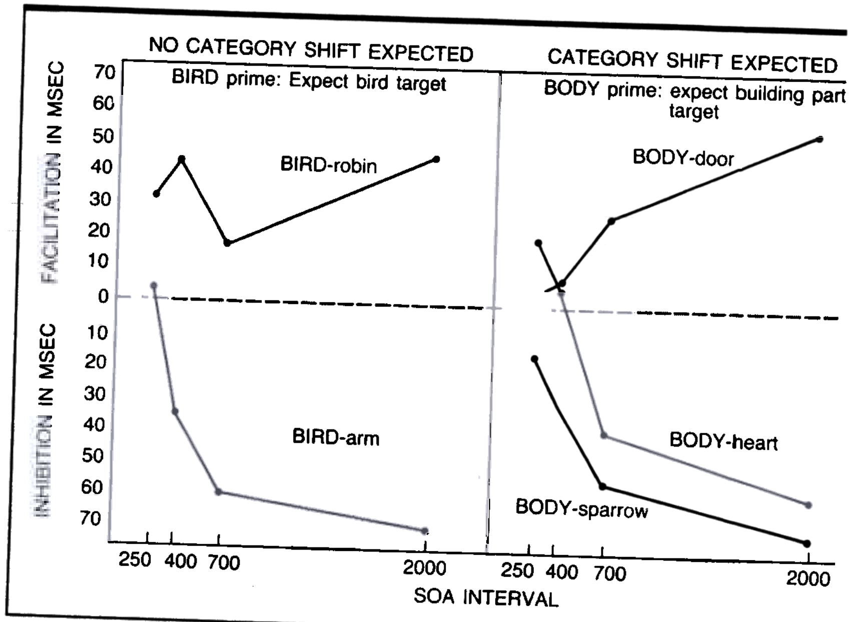

Reaction time to lexical decision targets. In the left half, subjects saw a prime and did not expect a shift in category; sample stimuli are BIRD-ROBIN for a relevant prime, and BIRD-ARM for an irrelevant prime. In the right half, subjects expected the target to come from the part of a building category if they saw BODY as a prime, and from the body part category if they saw BUILDING as a prime. When the shift in category occurred as expected, then RT was facilitated at longer SOAs. When the expected shift did not occur, there was facilitation when the prime was relevant (BODY-HEART). Inhibition occurred when the shift was completely unexpected (BODY-SPARROW). (From Neely, 1977.)

Neely's research provided a dramatic confirmation of the role of priming in semantic retrieval. Even with only the shortest of exposures, a relevant prime such as BIRD facilitates the processing of a related concept, ROBIN. This facilitation remains significant throughout a large range of SOAs. A relevant prime also facilitates the processing of related concepts even when subjects expect a switch in categories, but this facilitation only occurs immediately after the prime, before the conscious expectation of a category shift can have an effect. Finally, priming with an unrelated concept inhibits subjects' responses to the target, unless the unrelated target came from a category that the subject had been told to expect. This latter kind of facilitation only began to show up in the later SOA conditions, so was presumably due largely to conscious processing effects.

Lexical Ambiguity The priming task has also "done duty" in the lexical ambiguity area of research. **Lexical ambiguity** refers to the fact that *many words*, such as bank, pen, will, play, and so forth, *have more than one meaning*. From the standpoint of semantic memory, it is interesting to ask, "How are ambiguous words retrieved from memory?" That is, when we access a concept in semantic memory, we retrieve the meaning of that concept. When the word name for the concept points to two or more distinct meanings, do we access all meanings? Do we access the most frequent meaning first? Or does the context of the word determine which meaning we access? From the standpoint of comprehension of language, ambiguous word studies provide a way of understanding the role of context, for instance, how a strong context can bias the interpretation of words as we read or hear.

The basic task here is to present the ambiguous word in the prime and then present one of several kinds of words as a target. Just as BIRD primes related information, our ambiguous prime should activate the meanings and concepts it is related to. The difference, of course, is that the ambiguous prime is related to two distinctly different meanings—BANK is related both to the concept MONEY and the concept RIVER. If both meanings are activated, even when the surrounding context of the prime is biased toward one meaning only, then this tells us one thing about context effects; if *only* the biased meaning is retrieved, i.e., the meaning that fits the context, then this tells us something quite different about the effects of context. Thus, the logic of the research is to examine which meaning or meanings of an ambiguous word become activated by looking at RTs to target words that are associated with the different meanings.

As a concrete example, consider Simpson's (1981) report in which lexical decisions were made on a variety of target words. In one experiment, ambiguous or nonambiguous primes preceded the targets (e.g., BANK-MONEY vs. BOAT-MONEY); of course, the ambiguous prime BANK could also have been followed by RIVER, related to the other meaning of the word. In the second experiment, the prime was an entire sentence. In one condition, the prime sentences were ambiguous, not particularly related to either meaning. In another condition, prime sentences were biased toward the frequent meaning of the ambiguous word, called the *dominant meaning*, and in a third condition, primes were biased toward the infrequent meaning, the *subordinate meaning*. Furthermore, the prime sentences themselves were either weak or strong in the way they biased the context (see examples in Table 6–5). In both studies, of course,

Table 6–5

Sentence Type	Example	Related Lexical Decision Targets	
		Dominant	Subordinate
(1) Ambiguous	We had trouble keeping track of the *count.*	NUMBER	DUKE
(2) Weak dominant	The musician kept losing track of the *count.*	NUMBER	DUKE
(3) Strong dominant	My dog wasn't included in the final *count.*	NUMBER	DUKE
(4) Weak subordinate	The king kept losing track of the *count.*	NUMBER	DUKE
(5) Strong subordinate	The vampire was disguised as a handsome *count.*	NUMBER	DUKE

half of the targets were semantically related, and half were unrelated to provide a baseline measurement.

The results of this study are shown in Figure 6–14. First, note that there was no effect of any of variables on lexical decision times for unrelated words—all of the baseline conditions had an RT of about 650–660 msec. In other words, when context sentences are unrelated to the target, the RT is simply a function of the characteristics of the target word itself. Now look at panel A of Figure 6–14, where the context sentence was ambiguous. In this completely unbiased condition, the sentence prime included the ambiguous word (COUNT), and this was followed by a word/nonword judgment on either the dominant (NUMBER) or subordinate (DUKE) meaning. In this condition, both meanings of the ambiguous word showed facilitation, i.e., a decline in RT from the baseline (unrelated) condition. But, the more frequent meaning showed greater activation than the subordinate meaning—in both pairs of curves (this was a replication condition), the dominant target was judged more rapidly. This is similar, in principle, to the greater activation that accumulates at a frequent or important property vs. an infrequent property in the sentence verification research you read about earlier (e.g., Glass et al., 1974).

When the priming sentence was biased toward the dominant meaning (panel B, e.g., COUNT-NUMBER), the pattern was largely the same—a bit of facilitation to the less frequent meaning, but much more to the dominant meaning. Finally, when the context sentence biased the *subordinate* meaning (panel C), but only in a weak fashion (e.g., "The king kept losing track of the count"), both the "number" and the "royalty" meanings showed about equal facilitation. Context here seemed to be compensating for the overall dominance effect by activating the subordinate meaning more than it was by the ambiguous prime sentence. This last conclusion was supported by the results from the "strong bias" condition. With strong bias toward the dominant meaning (right half of panel B), no facil-

FIGURE 6-14

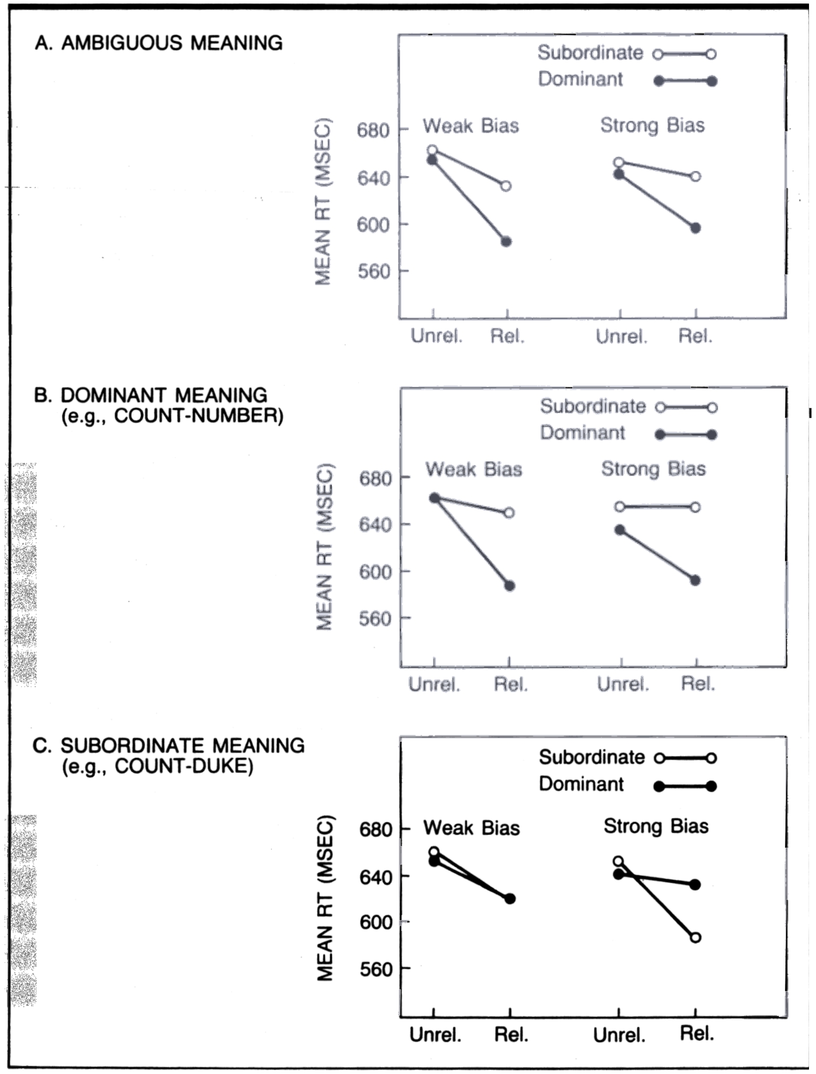

Reaction time to make lexical decisions is shown for prime sentences that were ambiguous (panel A), prime sentences that biased the dominant meaning of the ambiguous word (panel B), and prime sentences that biased the subordinate meaning of the ambiguous word (panel C). Each pair of curves shows that unrelated targets required about 650 msec to judge; this time was unaffected by the other variables. When targets were related, then facilitation was observed to both meanings, although less to the subordinate meaning. Only when the sentence prime was strongly biased toward the subordinate meaning was RT to the subordinate word faster than to the dominant word. (From Simpson, 1981.)

itation was found at all for subordinate meanings. With strong bias toward the subordinate meaning (e.g., "The vampire was disguised as a handsome count"), the subordinate meaning was much faster than the dominant meaning (right half of panel C).

After reviewing many such experiments, Simpson (1984) concluded that activation usually spreads to all of the meanings of an ambiguous word, but that the amounts of activation depend on at least two factors, the dominance of the different meanings and the strength of the surrounding context. With no context, meanings are activated to a level that depends on their normal dominance. But with biased context, the biased meaning receives an extra boost in activation. Thus, findings such as the subordinate priming pattern described earlier clearly suggest an important role for context in the more ordinary processes of accessing word meanings. That is, in routine processing of word meanings, context seems to have a strong influence on activation and retrieval from semantic memory (e.g., Kintsch & Mross, 1985; Schwanenflugel & Shoben, 1985). As Simpson has noted, such a conclusion is strengthened by the "special case" provided by ambiguous words, since context effects cannot be easily separated from the normal process of meaning retrieval when the word being retrieved has only one meaning.

▾ The Generality of Semantic Networks

Let's conclude this chapter with a somewhat broader perspective on semantic memory and semantic networks than is suggested by the research you've been reading about. We claimed at the outset that semantic memory is your permanent knowledge store about concepts, yet all of the examples we've used, and all of the research we've talked about, involved *nouns*. You might mistakenly conclude that networks are not useful for explaining the "smash" versus "bump" verb concepts that we started with. For that matter, you might also mistakenly conclude that semantic memory is only for *words*. The problem is that many of our concepts do not correspond to a simple, convenient one-word label (like that orange-purple-grey color of the sky at sunset). You'll recall that Rosch's subjects didn't have vocabulary for the various colors, yet showed that the differences among shades did have an effect on their memory performance. But we haven't talked about nonverbal concepts at all in this chapter. Furthermore, we left at least one very important issue up in the air: the usefulness of networks vs. feature lists as ways of representing semantic knowledge. We need to resolve the debate between Collins and Loftus (1975) and Smith et al. (1974). And finally, if semantic memory is truly the place where our general world knowledge is stored, then how is more complex knowledge retrieved and used? What does the priming of ROBIN by BIRD tell us about complex concepts and their retrieval, and about our memory for complex events?

Networks Versus Feature Lists

Let's begin with the network issue, first in terms of the specific models we discussed before, then at a more general level. Collins and Quillian (1972; Collins & Loftus, 1975) proposed that semantic memory is an organized network of interrelated concepts, and that activation spreads to concepts when we retrieve information from the network. Smith et al. (1974) proposed that semantic memory is composed of feature lists, one for each concept, and that retrieval consists of comparisons of the features between two lists. Which model is right, or at least closer to the truth?

In an interesting footnote to this debate, Hollan (1975) pointed out that feature-based and network-based models are not necessarily incompatible or contradictory, that in some sense the two approaches are merely superficially different ways of accounting for the same underlying knowledge (see also Collins and Loftus, 1975). While this theoretical point may be true, it is also true that the specific models we discussed differ at more than a superficial level.

In particular, consider how the models differed on the critical effects of priming, a topic you are surely expert on by now. At a general level, all of the priming effects you have studied can be easily accommodated within network models of semantic memory, since every network model proposed in cognitive psychology has involved the principle of spreading activation (e.g., Anderson, 1976, 1983; Anderson & Pirolli, 1984; Glass & Holyoak, 1975). And priming, after all, was specifically predicted by the Collins and Quillian network model. On the other hand, there were no mechanisms within the Smith et al. feature model to account for such priming, especially when the priming operates from one trial to another (e.g., Ashcraft, 1976; Loftus & Loftus, 1974). The Smith et al. model yielded RT predictions for processes occurring within a single trial, but was simply not equipped to handle priming effects from one trial to the next. As Eysenck (1974) put it, the Smith et al. model was "remarkably restricted" to the simple sentence verification task, especially to the verification of category statements.

In a way, it's unfair to criticize a model for failing to predict an effect that it was never intended to predict in the first place. On the other hand, a significant model of semantic memory must make predictions about the important empirical phenomena, or else it is irrelevant. Given that priming became a critically important factor in later research, a phenomenon demonstrated over and over again, it was without a doubt a decided weakness that the Smith et al. model was so silent on this topic.

But are all feature comparison models equally silent on priming, in principle? We don't know. The history of this debate in cognitive psychology went something like this. The original flush of success enjoyed by the Collins and Quillian model was cut short by empirical demonstrations of typicality and semantic relatedness, and by the evidence against strict hierarchical organization. While these factors did not genuinely violate any of Collins and Quillian's stated principles, their model was indeed quite vague on mechanisms to account for those variables. The Smith et al. model portrayed an alternate approach, one that seemed to do a very decent job of accounting for those results with testable, specific mechanisms. But then, more and more research began to look at priming effects, a difficulty for the premier feature model of the day. Furthermore, there was little effort on the part of feature list theorists to modify their models in order to account for such priming effects. From this perspective, the feature list approach to semantic memory died of inattention, rather than of some fatal empirical blow.

Just as important, more investigators in the mid- and late-70s began to look at much more complex semantic relationships and representations than simple noun categories. Researchers began to explore comprehension of paragraphs, connected text, and spoken conversations; people became interested in large-scale semantic representations that involved distinctly episodic as well as semantic knowledge. In all of these cases, serious attempts were made to model the obtained effects in computer simulation or AI models. The (unspoken) consensus in most of this work was that network approaches provided a convenient, flexible, and powerful way of attacking these psychological processes. Networks had, in a sense, proved their usefulness in the basic semantic memory research, and were now being expanded into more complex areas with great success. We'll discuss some of these areas in the next chapter, and you'll see that the general network approach is used by virtually all of the models. From this perspective, networks gave cognitive psychologists an enormously useful format within which significantly more complex questions could be addressed. Feature list approaches, inherently more limited it would seem, didn't keep up the pace in these developments.

Not for Nouns Only

As an illustration, consider semantic concepts that are not nouns, such as the "smash" and "bump" verbs we began with in this chapter. By the time Gentner (1975), for example, began to investigate the semantic structure of verbs, it seemed quite natural to conceive of verb concepts as nodes in a network that demonstrate standard network effects, with concept nodes connected to other nodes by pathways, activated by priming in a spreading activation sense, subject to the regular "laws" of semantic relatedness, and so forth. In other words, the systematic relationships between members of verb "families," for instance, Gentner's verbs of possession (give, take, buy, sell, trade, etc.), seem to lend themselves to a "node and pathway" representation. In contrast, it certainly seemed difficult to dream up possible lists of features that would define the concepts and express their interrelationships as would be necessary for a feature list approach. Whereas feature list approaches seemed cumbersome, network approaches seemed tailor-made for such concepts.

Indeed, this is exactly the kind of representation that Gentner proposed. Her empirical data on accuracy of usage across childhood clearly demonstrated two results that were entirely compatible with network representations. First, the simpler verbs like GIVE and TAKE involve fewer nodes and pathways in their semantic representations, and are mastered much earlier in a child's language use. Second, the more complex verbs such as BUY and SELL overlap considerably with some of the simpler concepts. SELL, for instance, can be decomposed into the simpler concepts GIVE and TAKE, which imply "possession" and "transfer of ownership," along with the extra notion of PAYMENT. Not surprisingly, children showed poor performance on the more complex verbs. At an early age, they treated "sell" as if it meant "give," then slowly added the semantic relationships and extra concepts necessary for a full expression of the meaning of SELL. These complexities were easily expressed within network notation systems; to my knowledge, no one ever tried expressing them in a feature list model that made empirical predictions.

Moreover, semantic networks are not just limited to word concepts, to concepts that can be easily labeled with a convenient, one-word name. Just as the research has done, we have simplified our study of semantic memory in this chapter by discussing only concepts that have simple one-word names. This is not at all necessary, however. As an illustration, consider sensory concepts. A favorite example of mine was provided by Collins and Quillian (1972), the semantic concept of "the sound a rooster makes." To paraphrase their discussion, we have several semantic concepts that relate to this idea—the auditory concept of what a rooster actually sounds like, the verbal concept we call "cock-a-doodle-doo" in English, and, for that matter, the auditory concept of the *sound* that our verbal name has. Each of these is a distinct, but related concept in semantic memory, and each is accessible by means of the normal retrieval process.

Network approaches have also been proposed for describing the process of letter identification (McClelland & Rumelhart, 1981), for explaining the effects of emotional state on learning (Bower, 1981), and for representing people's knowledge of simple arithmetic facts (Ashcraft, 1982; Campbell & Graham, 1985). But, by far, the most impressive applications of network models are those such as Anderson's ACT model (1983), which account for semantic *and* episodic memories (and others as well). As you study these "second generation" theories in the next chapter, you'll discover that cognitive psychology is completely "hooked" on the notion that permanent knowledge is represented as interconnected nodes in a network.

Don't lose track of the basic elements of semantic networks—nodes linked by pathways that specify the connecting relationship. In principle, as several authors have noted, we want to propose a theory of semantic memory that is as flexible as people are in expressing relationships. In other words, to use another memorable Collins and Quillian (1972) example, people are good at expressing the often subtly different relationships among words. For instance, "people like dogs" is different from "dogs like people" is different from "cats like mice." Our models must be just as good as people are at representing that knowledge. Further, our semantic memory structure must be compatible with episodic memory, since it is our basic semantic system that guides our understanding of and memory for daily events. How semantic memory serves as the knowledge base for comprehension of language and of experience, of those daily events, is the topic of the next chapter.

CHAPTER SUMMARY

1. Semantic memory contains our long-term memory knowledge of the world, including our knowledge of words, concepts, and language. Early studies of the structure and processes of semantic memory generated two kinds of models, network approaches and feature list approaches.
2. The Collins and Quillian network model claimed that concepts are represented as nodes in a semantic network, with connecting pathways between concepts. Accessing a concept involved the process of spreading activation—activation spreads from the originating node to all those nodes connected to it by pathways. Several early studies, using the sentence verification task, supported an early version of this model, although the disconfirmation of the cognitive economy hypothesis was viewed as evidence against the model.
3. Smith et al. claimed that semantic concepts are lists of semantic features. Verification in their model consisted of accessing the feature lists and performing a comparison process on the features. While the Smith et al. model initially seemed more able to explain typicality effects, the model did not provide a particularly plausible account of how we verify property statements. More importantly, the phenomenon of semantic priming was not predicted by the model. When this phenomenon became a central focus of semantic memory research, the feature list approach fell into relative neglect.
4. Concept formation research suggested that our concepts consist of the presence or absence of simple, independent features. Rosch challenged this notion and demonstrated that real-world concepts and categories involve "fuzzy boundaries," i.e., that some members of categories are judged as more typical than others. Typical members are said to resemble the prototype of the category. Interestingly, typical members show greater effects of priming than do atypical members. Typicality

effects support the generalization known as the semantic relatedness effect, that "yes" judgments will be speeded up if the concepts being judged are more highly related to each other.

5. A variety of tasks have been developed to examine the role of priming in the retrieval of information from semantic memory. In such research, a prime of some degree of relationship to a target is presented first, and then RT to the target is measured. When the prime is relevant, RT to the target is usually speeded up, even at very short time intervals (SOAs). This is generally taken as evidence that semantic priming is an automatic process. When the prime is irrelevant to the target, RT is generally slowed down at longer SOAs. This is usually interpreted as evidence that irrelevant primes generate a conscious expectation that slows down processing when the expectation is misleading. Among the tasks used to examine priming in semantic memory are word naming, sentence verification, and lexical decision tasks.

6. Because of the importance of priming, and because of the relative neglect of the feature list approach, cognitive psychology has now devoted considerable attention to the network approach to semantic memory. This approach is now found in a wide variety of topics, and has been extended far beyond the simple noun concepts that were so heavily investigated in the 1970s.

SUGGESTED READINGS

I still have not found a more useful introduction to semantic memory than the 1972 Collins and Quillian paper. It's thought provoking, insightful, clever, and delightful to read. I would strongly encourage you to read it, to see how these authors manage to discuss generalization and discrimination, basic memory retrieval, imagery, metaphoric language, language comprehension, and computer simulation while maintaining a lively, often humorous style. It's a classic.

Beyond that paper, several books and articles will fill in some gaps, or extend your understanding significantly. The entire second half of Tulving and Donaldson's 1972 book is devoted to papers on different aspects of semantic memory; this also includes Tulving's article on the episodic-semantic distinction. Somewhat later, Smith (1978) reviewed the field of semantic memory research and theory, and elaborated considerably on models that represent hybrid approaches, as well as on related kinds of research that semantic models should account for. And finally, there has recently been an attempt to integrate network and feature list approaches within one *hybrid model* (Kounios, Osman, & Meyer, 1987). In this model, both a network search and a feature list comparison process occur simultaneously; the one that finishes first is the one that governs RT to the stimulus. This approach suggests that at some deep level, network and feature list approaches are entirely compatible.

Research continues on the structure and processes of semantic memory. Several investigators have found that Rosch-inspired notions of basic level categories and prototypes generalize to quite different domains; for instance, to social categories such as WORKER, EMPLOYER, and POLITICIAN (Dahlgren, 1985); to "event categories" such as school activities and types of shopping (Rifkin, 1985), and to abstract concepts in computer science (Adelson, 1985).

Recently, L. C. Smith (1984; no relation) has reported on the "semantic satiation" effect, that after continued repetition of the same concepts, the spread of activation can actually slow down due to excessive priming. Interestingly, semantic satiation affected word retrieval, but not lexical decision times. And Bowles and Poon (1985) have found both semantic and orthographic priming effects, where the latter is due to similarities in spelling patterns between prime and target. Smith and Osherson's (1984) recent paper discusses the concept of prototypes for "combined" concepts, for instance *pet fish* or *brown apple*. See the Suggested Readings at the end of the next chapter for additional papers that relate to semantic memory.