Chapter 6
Building a Semantic Network

A semantic network is a propositional knowledge structure consisting of a set of nodes that are selectively connected to each other by links labeled by the relationship between each pair of connected nodes (Fahlman 1979, 1981; Quillian 1968; Stillings et al. 1987, chap. 4). Links in such networks may assert category membership or a part-to-whole property. These can be described as *is-(a)* and *has-(a)* relationships such as "A trout *is a* fish" and "A trout *has* gills" (figure 6.1). Many other kinds of relationships among nodes may also be represented by labeled links. For example, a link could represent the concept *is made of*, as in "a crowbar *is made of* steel," or a relationship of relative weight, as in "an elephant *weighs more than* a mouse."

When a synaptic matrix serves as a biological instantiation of a semantic network, the word-tokened inputs to its mosaic cells and the tokened class cells of its output correspond to nodes. Selectively augmented synaptic transfer weights (ϕ) on the appropriate filter cells correspond to the links among the nodes. The relationships among connected nodes might be labeled in at least two ways. One way would be to have a separate synaptic matrix devoted to each kind of relationship between nodes—for example, a matrix that processes only *is-(a)* relationships and one that processes only *has-(a)* relationships. The other approach, the one adopted here, is to have the internodal relationship represented by the predicate node. Thus, instead of "A trout (node 1) ‹*has* (labeled link)› gills (node 2)," the synaptic matrix representation would be "A trout (node 1) ‹ϕ link› *has* gills (node 2)." In this scheme, an object and an implication about the object taken together constitute a predicate category that is signaled by the discharge of a single physically indexed cell.

The Synaptic Matrix as a Semantic Network

Consider the following sequence of sentences that could provide the kind of information captured in the network illustrated in figure 6.1:

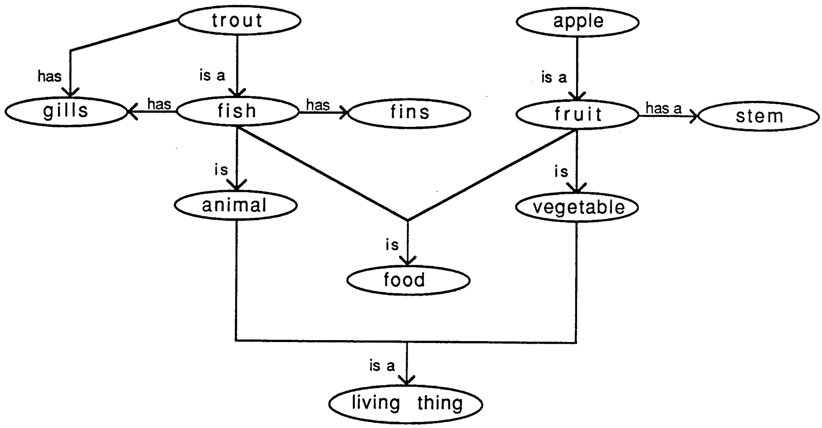

Figure 6.1
A semantic network composed of abstract nodes designating entities and abstract links designating relationships between connected nodes. Notice that if there were no direct *has* link between *trout* and *gills,* traversal of the *is a* link from *trout* to *fish* and the *has* link from *fish* to *gills* allows the proposition *trout has gills.*

(1) A trout is a fish.

(2) A fish has gills.

(3) A fish has fins.

(4) Fish is food.

(5) Fish is animal.

(6) An apple is a fruit.

(7) Fruit has a stem.

(8) Fruit is food.

(9) Fruit is vegetable.

(10) An animal is a living thing.

(11) A vegetable is a living thing.

Each string that composes one of the sentences consists of a sequence of visual or auditory patterns that must be registered, parsed, and subjected to lexical classification by means of neuronal activity. Some kind of analytic brain mechanism, shaped by learning within a particular cultural context, must be able to perform a syntactic decomposition of the string so that subject and predicate can be designated and processed appropriately.

Let us consider how a synaptic matrix might process sentences that are presented as visual stimuli. It is assumed that the pattern of marks that constitute the words of a printed sentence can, like any other visual pattern, be mapped into sets of indexed cells that would be part of a lexical module. When we considered the synaptic matrix in the context of visual processes, retinal stimuli, after translation to the normal foveal axis, were input to an array of cells in the matrix that were designated as mosaic cells (M). The circuitry for semantic processing is somewhat different. It requires a set of autaptic cells that capture the neuronal tokens of words in short-term memory (working memory); it is the activity of these cells, which we shall call word cells (W), that provides the direct afferent excitation to the mosaic cell array (M). Thus, W_1 and M_1 might represent *apple,* and W_2 and M_2 might represent *is (a) fruit*. The designations for filter cell (f) and class cell (Ω) remain unchanged, as do the biological properties of the detection matrix and the imaging matrix.

The objective is to establish within a synaptic matrix a preferred excitation path from W_1 (*apple*) to Ω_2 (*is (a) fruit*) so that upon any subsequent presentation of the word *apple,* the semantically appropriate class cell (Ω_2) will be discharged. This can be done in a simple learning procedure if the filter cell coupled to the class cell token of the predicate *is (a) fruit* were to discharge concurrently with the W-token input that represents the subject *apple.* Because of the sequential nature of sentence strings, however, the required temporal overlap of subject and predicate does not occur naturally; some means must be provided to compensate for this if the proper association is to be learned.

The desired cooccurrence of the tokens for subject and predicate can be ensured if discharge of the word cell (W_i) representing the subject is sustained beyond the duration of its evoking stimulus and terminated after the predicate cell has discharged. This is accomplished by having each word cell in the input array to the synaptic matrix be an autaptic cell and priming the array so that each word token (W_i) continues to discharge (as a latched cell) until the stimulus sentence is completed and the predicate token (Ω_i) has fired.

Illustrated in figure 6.2 is a modified synaptic matrix with the intrinsic capability to organize a semantic network if it is given a series of simple sentences. Initial input to the synaptic matrix is through a set of autaptic word cells (W_i), each of which (if active) stimulates its coordinate mosaic cell and, by means of an axon collateral, a paired

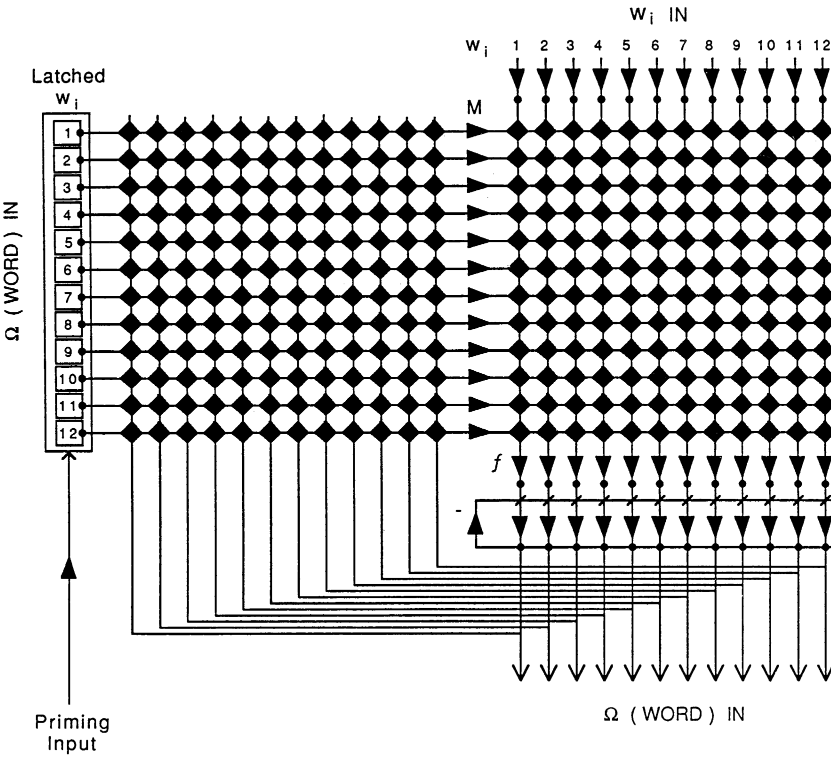

Figure 6.2
Schematic of a semantic network. Open squares (designated W_i) within priming field are autaptic cells that act as lexical tokens. Each autaptic cell connects to a paired mosaic cell (M) and, through an interneuron (not shown), to a paired filter cell (f). Dots represent fixed excitatory synapses; short, oblique slashes represent fixed inhibitory synapses; filled lozenges represent adaptive excitatory synapses. Reset neuron ($-$) generates an inhibitory postsynaptic potential to reset all class cells when discharged.

filter cell (f_i) in the detection matrix. At the start of each sentence, all W_i are primed, enabling the word tokens in the sentence string to be latched. At the completion of each sentence, the priming input is interrupted and all W_i stop firing. The repeated sequence of initiation and interruption of excitatory priming of word cells occurs in phase with the input and completion of successive sentences over any extended lexical message.

Suppose, in the printed sentence *A tiger is a cat,* articles are ignored so that the first recognized character string is *tiger,* which fires W_1, and the second recognized character string is *is cat,* firing W_2. In this case, the mosaic cell M_1 and the filter cell f_1 will discharge concurrently, causing an increase in the efficacy of the synaptic junction between M_1 and f_1 in the detection matrix in accordance with the learning principle. Then mosaic cell M_2 and filter cell f_2 will fire together, causing an increase in the transfer weight of their synaptic junction. At the same time, the synapse connecting M_1 and f_2 will also be modified because M_1 was continuing to discharge (stimulated by the latched autaptic cell W_1) when f_2 fired. In the imaging matrix, the concurrent discharge of class cell Ω_2 and mosaic cell M_1 will raise the efficacy of their synaptic link. After this bit of learning, if *tiger* alone is presented as a stimulus, the response will be *tiger* (Ω_1) *is cat* (Ω_2). If *is cat* is presented alone, it will evoke Ω_2 in the detection matrix, which, in turn, will evoke M_1 in the imaging matrix, resulting in the Ω_1 output *tiger.*

When the semantic network is queried, the discharge of any particular mosaic cell (M_i) is the selected token of a word or phrase taken as the subject of a sentence. The discharge of any particular class cell (Ω_i) is the selected token of a word or phrase taken as the predicate of a sentence—for example, "A *trout* (M_1) is a *fish* (Ω_3)." If we were to ask "What is a trout?" "a fish" would be one of a number of proper replies. If the same question were put to a semantic network previously presented with "A trout is a fish," we expect a similar response. That is, if M_1 (*trout*) is fired, then Ω_3 (*fish*) should fire. Suppose we were to ask, "What is an example of a fish?" The response "a trout" would be appropriate. In this case, the subject *trout* is selected when the predicate *fish* is queried. In terms of the neuronal structure of the semantic network, if Ω_3 (fish) is discharged, then M_1 (trout) should fire.

But we would expect more of a neuronal mechanism that is to serve as a semantic processor in the human brain. In particular, we would expect such a processor to make sense out of combinations of sentences. For example, suppose you were provided the information contained in the following sentences:

(1) A trout is a fish.

(2) A cod is a fish.

(3) A trout lives in freshwater.

(4) A cod lives in saltwater.

If you were asked to give an example of a fish, the responses *trout* and/or *cod* would be appropriate. These could be computed from sentences 1 and 2, where each sentence is treated independent of the other. However, if you were asked to give an example of a fish that lives in saltwater, the information contained in at least two of the four sentences (2 and 4) would have to be combined and used in such a way that the response *cod* is given. From a strictly logical point of view, the fact that a trout lives in freshwater does not preclude the case that it also lives in saltwater. But given the information contained in these sentences and the form of the question asked, the natural inference is that *cod* is the right answer and *trout* is wrong.

If something like this is accounted for by what happens physically in the brain, then the semantic processor must somehow combine and properly relate the information contained in all four sentences. Something roughly corresponding to the following simple bit of reasoning must be taking place: *A trout and a cod are both examples of a fish, but I know only that a cod lives in saltwater, so cod is the right answer.* A semantic network organized by an augmented synaptic matrix can generate "reasonable" inferences of this kind.

An Example of Performance

The following simulation provides an initial example of how the synaptic matrix can perform as a self-organizing semantic network. The neuronal mechanism illustrated in figure 6.2 constituted the processing model that was simulated.

Synaptic changes were based on the same learning principles as presented in chapters 2 and 3. The saturation limit (*Lim*) for the transfer weights (ϕ) was arbitrarily set at a value of 2. All ϕ changes reached the saturation limit during learning because in sentence processing the number of coactive lexically tokened cells is relatively small. Parsing of the words and phrases that composed the sentences typed into the computer was accomplished by a program subroutine that was not part of the simulation of the neuronal model. Thus, input to the word cells (W) attached to the synaptic matrix was mediated by an ordinary program for parsing character strings and linking them to W cells rather than by a simulated neuronal mechanism for recognizing the actual visual patterns of the character strings. Similarly, initiation and interruption of *W* cell priming were controlled by a standard computer subroutine that responded to the initiation of a sentence and the punctuation at the end of it. However, the essential semantic mechanism for systematically organizing the relationships among biological tokens of words and phrases was a simulation of a dynamic neuronal structure (the augmented synaptic matrix).

For the purpose of simplifying the simulation program, the model ignored the article strings *a, an,* and *the,* as well as the verb *is* when it learned its lexicon from the sentences presented. An *is (a)* relationship was taken as the default characterization of a predicate link to a subject whenever an assertion of relationship was required but was not explicitly specified in the output of the semantic network; for example, the output string *trout fish* was transformed by adding the appropriate articles and the *is (a)* link to make the well-formed expression, "A trout is a fish." Other predicate relationships were expressed directly in the output of the semantic network as predicate phrases (hyphenated in both input and output to facilitate phrase parsing); for example, the output string "fish has-gills" required only the addition of the article to make the well-formed expression, "A fish has-gills."

The following sentences were presented to the neuronal model for semantic processing (a space was left between the last word in each sentence and the period to make it easier for the parsing program to sense the end of the sentence and signal the interruption of *W* cell priming):

(1) A TIGER IS A CAT .

(2) A TIGER HAS-STRIPES .

(3) A TIGER IS WILD .

(4) A TABBIE IS A CAT .

(5) A TABBIE HAS-STRIPES .

(6) A TABBIE IS A PET .

(7) A FLAG HAS-STRIPES .

(8) A FLAG IS AN ARTIFACT .

(9) A CAT IS A MAMMAL .

(10) A MAMMAL IS AN ANIMAL .

(11) A DOG IS A MAMMAL .

(12) A DOG IS A PET .

(13) A DOG BARKS .

From this list of 13 sentences, the model automatically built a lexicon consisting of the following 12 items:

1. TIGER
2. CAT
3. HAS-STRIPES
4. WILD
5. TABBIE
6. PET
7. FLAG
8. ARTIFACT
9. MAMMAL
10. ANIMAL
11. DOG
12. BARKS

The fact that the number of items in the lexicon is many fewer than the total number of words in the 13 stimulus sentences follows from the fact that once a particular word or phrase has been learned, it is no longer novel. The synaptic matrix normally learns a current input (selectively tunes a new filter cell) only when it detects a novel stimulus. Also, recall that the articles and the verb *is* were, for convenience, arbitrarily ignored by the parser.

During the presentation of the stimulus sentences, synaptic weights (ϕ) were changing automatically in accordance with the learning principle. After all 13 sentences were typed in, the ϕ distributions in the detection matrix and the imaging matrix were as represented in figure 6.3.

The semantic network was tested in two different response modes, each appropriate to a different kind of query. In one test mode, the network was asked to define a stimulus word that was the subject of a sentence previously presented to it. In the second test mode, the network was asked to infer the subject(s) of one or more stimulus words that were the predicates of sentences previously presented. When the network was required to define a subject, priming of W cells was sustained, and each activated response token (Ω_i) was assumed to reafferent its word token as the next input to the network—for example, $\Omega_9 \rightarrow W_9 \rightarrow M_9$. When the network was required to infer a subject, priming of W cells was interrupted, and the only continuing activation of the network was mediated by Ω cell

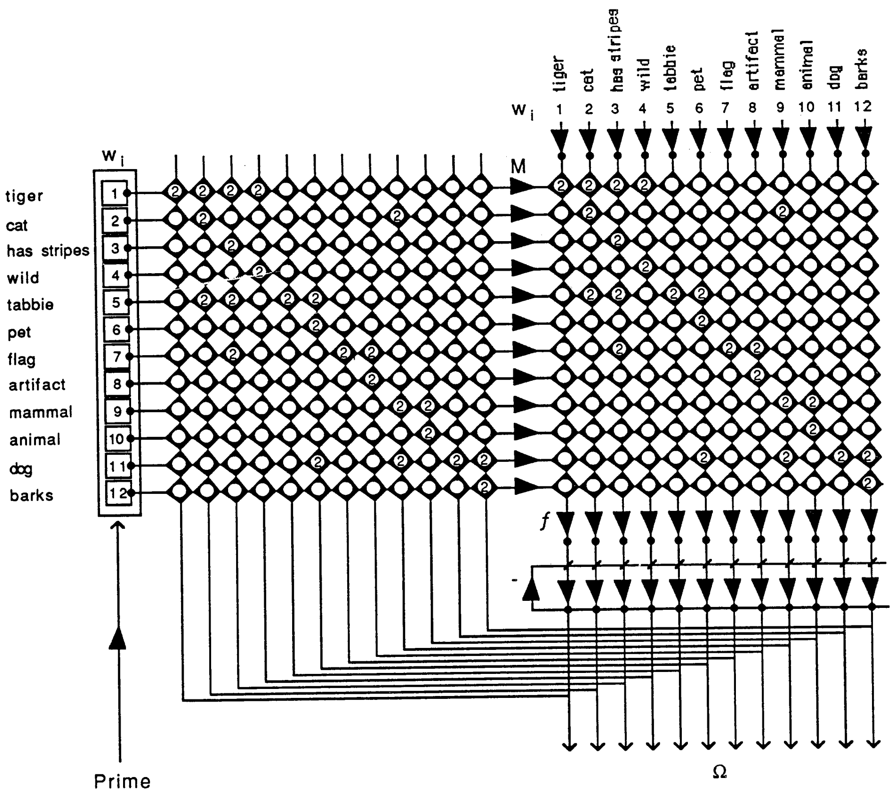

Figure 6.3
State of semantic network immediately after thirteen related sentences were presented to it. Saturation level for synaptic weights (ϕ) = 2. Synaptic junctions without printed weights are those that remained unmodified by learning and are at a uniform low basal value.

feedback through previously augmented synaptic junctions on its mosaic cells—for example, $\Omega_5 \rightarrow M_2$.

In the course of semantic processing, the subject of a sentence is represented by the activity of a mosaic cell (M_i), whereas a predicate is represented by the activity of a class cell (Ω_i). When the semantic network is queried, parallel activation of multiple predicates in response to a given subject can occur only when the synaptic weights (ϕ) that control the respective rates of EPSP integration in Ω cells have the same value (the saturation limit *Lim*). Under such circumstances, it is possible for several class cells to fire together within each Ω reset cycle.

Defining a Subject

After the 13 sentences were presented, the semantic network was asked to define the words in its lexicon that corresponded to the subjects in the sentences. The responses of the model are given in the following listing of questions and answers (phrases in parentheses were not part of the model's lexicon and are added for clarification):

Question: (WHAT IS A) TIGER ?
Response: (A) TIGER (IS A) TIGER
(A) TIGER (IS A) CAT
(A) TIGER HAS-STRIPES
(A) TIGER (IS) WILD
(A) TIGER (IS A) MAMMAL
(A) TIGER (IS AN) ANIMAL

Question: (WHAT IS A) TABBIE ?
Response: (A) TABBIE (IS A) TABBIE
(A) TABBIE (IS A) PET
(A) TABBIE (IS A) CAT
(A) TABBIE HAS-STRIPES
(A) TABBIE (IS A) MAMMAL
(A) TABBIE (IS AN) ANIMAL

Question: (WHAT IS A) FLAG ?
Response: (A) FLAG (IS A) FLAG
(A) FLAG (IS AN) ARTIFACT
(A) FLAG HAS-STRIPES

Question: (WHAT IS A) CAT ?
Response: (A) CAT (IS A) CAT
(A) CAT (IS A) MAMMAL
(A) CAT (IS AN) ANIMAL

Question: (WHAT IS A) MAMMAL ?
Response: (A) MAMMAL (IS A) MAMMAL
(A) MAMMAL (IS AN) ANIMAL

Question: (WHAT IS A) DOG ?
Response: (A) DOG (IS A) DOG
(A) DOG BARKS
(A) DOG (IS A) PET
(A) DOG (IS A) MAMMAL
(A) DOG (IS AN) ANIMAL

The semantic network's responses to the queries are appropriate given the original information (sentences) that it had learned; among its responses to each query was the identity relation ("A TIGER IS A TIGER"); and the model was able to make appropriate inferences from the information provided. Despite the fact that it was not told that a tiger is a mammal or that a tiger is an animal, it was able to infer that if a tiger is a cat and a cat is a mammal and a mammal is an animal, then a tiger is both a mammal and an animal. The same kind of inferences were made about a tabbie. And the model inferred from the fact that a dog is a mammal that it must also be an animal.

The reason that these inferences were made follows from the synaptic structure endogenously organized in the course of learning from the original sentences. When the word cell W_1 corresponding to the character string "TIGER" was discharging, it evoked (through its learning-enhanced synaptic junctions in the detection matrix) the discharge of the following class cell tokens: Ω_1 (TIGER), Ω_2 (CAT), Ω_3 (HAS-STRIPES), and Ω_4 (WILD). These, in turn, were assumed to evoke their word tokens as reafferent stimuli to the synaptic matrix, providing new inputs to the semantic network. In particular, for example, indirect feedback from the active Ω_2 token (CAT) fired W_2 (CAT), which fired M_2 (CAT), and this cell then fired Ω_9 (MAMMAL) (via f_9 in the detection matrix). The response *MAMMAL* was given as part of the definition of *TIGER* though the system was not told that a tiger is a mammal. Moreover, once the response *MAMMAL* was made, the same kind of recursive process induced the discharge of the Ω_{10} token (ANIMAL). So *ANIMAL* was also given as a proper response to the question, "WHAT IS A TIGER?"

A particularly important kind of synaptic change automatically occurs in the semantic network when these recursive neuronal events take place. Since the original subject of the query is latched in an autaptic token in the array of word cells, its coupled mosaic cell remains active and forms an enhanced synapse with each of the filter cells discharged in the inference process. A predicate response that was originally produced by a chain of inference is now directly mediated by a new synaptic link between neuronal tokens of subject and predicate. The self-restructuring of the synaptic matrix that occurred during the course of querying it is revealed in figure 6.4. The synapses marked by asterisks in the figure are the newly formed token-to-token gateways in the semantic network.

Inferring a Subject

In the previous test, the semantic network was asked to define or describe a subject: *WHAT IS A TIGER? WHAT IS A FLAG?* The following test required the semantic network to infer an appropriate subject from a given predicate or a set of predicates. The state of the synaptic matrix corresponded to figure 6.4. Instead of initiating each query by discharging a selected word cell as in the previous simulation, class cells that represented selected predicate tokens were fired either singly or in combinations. Thus, when the model was asked "(WHAT) ANIMAL (IS) WILD? Ω_{10} (ANIMAL) and Ω_4 (WILD) were automatically discharged together. Since the sentences that provided the original information for the naive semantic network contained no plural words, the following queries could be put only in singular form (phrases and letters in parentheses were not part of the model's lexicon and are added for clarification):

Question: (WHAT ARE SOME) ANIMAL(S)?
Response: (A) TIGER
(A) CAT
(A) TABBIE
(A) MAMMAL
(AN) ANIMAL
(A) DOG

Question: (WHAT) HAS-STRIPES?
Response: (A) TIGER
HAS-STRIPES
(A) TABBIE
(A) FLAG

Question: (WHAT) PET HAS-STRIPES?
Response: (A) TABBIE

Question: (WHAT ARE SOME) PET(S)?
Response: (A) TABBIE
(A) PET
(A) DOG

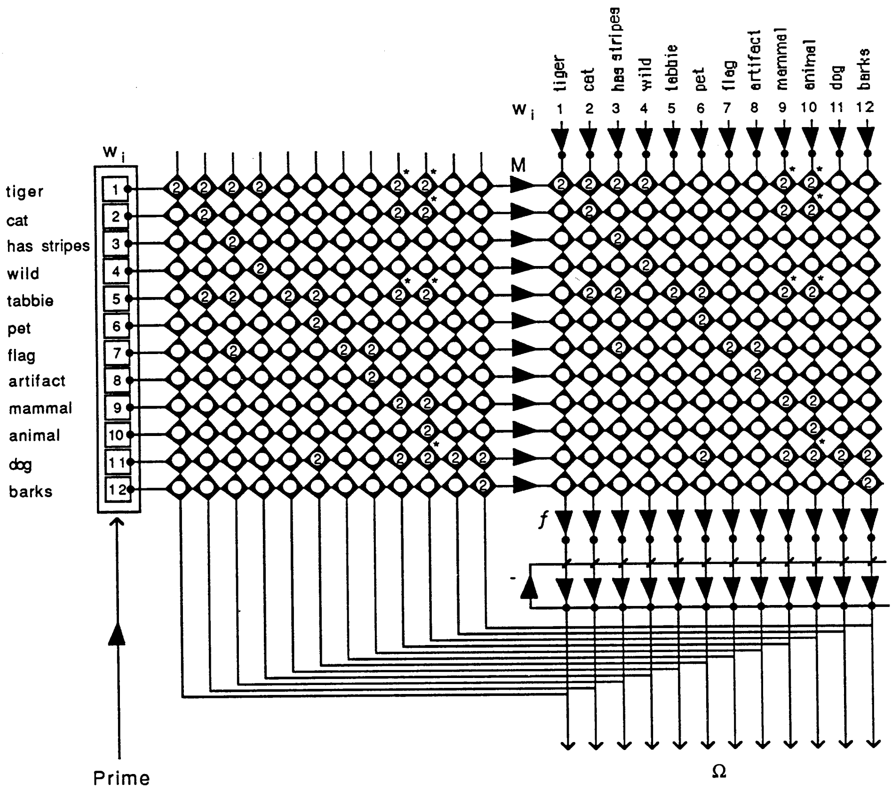

Figure 6.4
State of semantic network after it had been "asked" to define the subjects of the presented sentences. New synaptic changes were automatically induced in the course of responding to queries. The newly strengthened synaptic junctions are marked by an asterisk.

Question: (WHAT) HAS-STRIPES (AND IS AN) ARTIFACT?
Response: (A) FLAG

Question: (WHAT) ANIMAL (IS) WILD?
Response: (A) TIGER

Question: (WHAT) ANIMAL (IS A) PET (AND) BARKS?
Response: (A) DOG

The semantic network that was endogenously organized within the synaptic matrix was able to infer the appropriate subjects on the basis of either simple or complex predicates: *an animal* versus *an animal that is a pet and barks.* And this was accomplished by plausible neuronal mechanisms and architecture.

The reason that the network behaves as it does can be seen in the synaptic structure of figure 6.4. The greater the number of predicate qualifiers (Ω_i) that are "true" of any given subject (W_i), the more learning-enhanced synapses will activate its paired mosaic cell (M_i) when those predicate tokens are fired during a query. It follows that the most appropriate subject tokens (M_i) will exhibit response latencies shorter than those of less appropriate tokens. When a query permits more than one subject as a correct response, all proper M cells are discharged in parallel because each has an equal number of active synapses (see the response to "WHAT ARE SOME ANIMALS?").

Linking the Semantic Network to the Real World

The semantic network model seems to provide a neuronal mechanism that can adaptively organize a lexicon and instantiate logical relationships among lexical items. When it is questioned, it can respond appropriately on the basis of inferences that it makes from previously presented propositional information. But when it asserts that "A TIGER HAS-STRIPES," the word *TIGER* and the phrase *HAS-STRIPES* have no referents other than the character strings that evoke their neuronal tokens. Taken as an isolated module, the semantic network can exhibit logical processing of its physical symbols (tokens), but it can provide no information about their meaning with respect to events in the real world (see Johnson-Laird, Herrmann, and Chaffin 1984).

An appropriate reciprocal mapping is required between the synaptic matrix for pattern recognition at the sensory input level and the lexical tokens within the semantic network. This mapping should provide a neuronal structure wherein a sensory pattern evokes its proper lexical token and the discharge of that token can evoke an afferent image of that sensory pattern. Moreover, the synaptic gateways that instantiate such a mapping between separate processing modules must be able to evolve adaptively in the context of novel sensory environments and new lexical items. A neuronal system that can accomplish the desired mapping for visual-semantic processing is shown schematically in figure 6.5. This architecture is similar to the connectivity scheme illustrated in figure 3.5. It can be adaptively constructed by the backward-chaining mechanism (described in figure 3.6) in the following way.

Mosaic cell analogs (images) of objects at the level of the sensory matrix are mapped to particular object tokens (Ω_i), which provide input to the next processing stage, the synaptic matrix for lexical assignment (designating objects and events by the words that are conventionally used to refer to them). Since, during the learning of a lexical assignment, an object token must be active concurrently with the discharge of its proper lexical token, the axon collateral of the discharging lexical token (Ω'_i) that feeds back to the array of object tokens (Ω_i) will selectively strengthen its synapse with just that object token that was coactive with it (coactivity of Ω_i and Ω'_i). This provides a proper backward mapping from lexical tokens at the stage of lexical assignment to object tokens at the earlier stage of pattern recognition. The same process of backward selection of the proper synaptic couplings occurs between the matrix for lexical assignment and the semantic network because of the coactivity of input tokens from the former and feedback from the output tokens of the latter during the course of building the lexicon. In this fashion, the neuronal architecture shown in figure 6.5 is adaptively constructed as the lexicon is learned and lexical assignments are made. When the semantic network asserts that "A TIGER HAS-STRIPES," it not only states a proposition; it simultaneously evokes internal afferent patterns (images) that are analogs of the objects to which the proposition refers.

When synaptic matrices are organized in this way, the perception of either a specific object in the environment or the character string that names the object will evoke the same lexical tokens in the brain. Conversely, discharge of a lexical token can evoke an image of both its associated object and the object's name. And when lexical tokens are evoked by sentential stimuli, they are automatically related to each other in a logical fashion by selective synaptic coupling within the semantic network. Notice, however, that at the level of the semantic network, some of the components of sentential stimuli could be pictures as well as words, since each (if represented at all) is represented at this stage by its appropriate lexical token. These to-

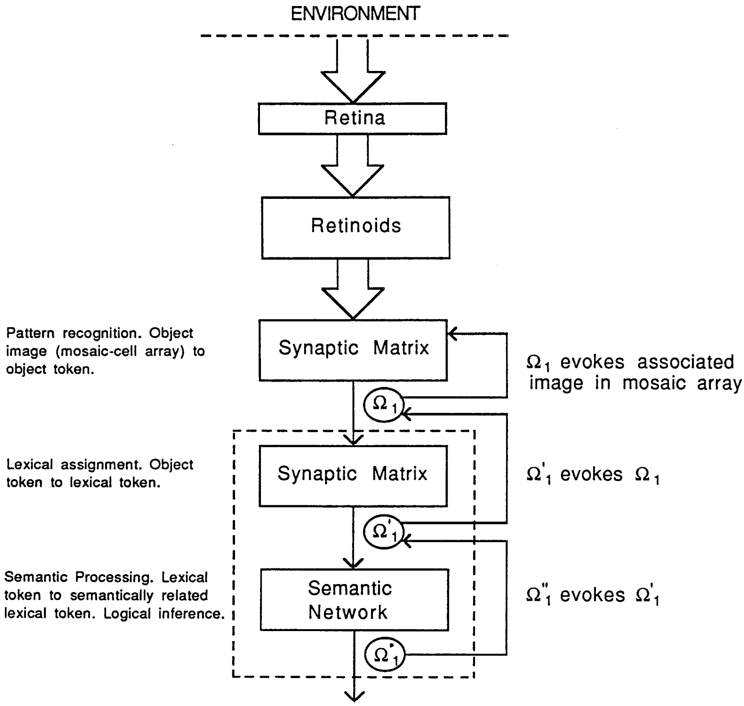

Figure 6.5
Flow diagram depicting forward and backward chaining of Ω tokens. Adaptive mapping binds lexical tokens to appropriate sensory images (including lexical strings).

kens constitute a common conceptual representation for both images and words (Banks and Flora 1977, Guenther and Klatsky 1977, Kroll and Potter 1984, Potter and Faulconer 1975). Thus, the proposed visual-lexical-semantic mechanism behaves in a manner that is consistent with the experimental finding that sentence processing can work directly with internal representations provided by pictured objects as well as with words (Potter et al. 1986).

The Structure of Representation

Fodor and Pylyshyn (1988) have argued persuasively that "mental representations need an articulated internal structure." The properties of the neuronal system suggested here satisfy the requirement for articulated representational structures in a number of ways. First, at the stage where exteroceptive stimulus patterns are learned, filter cells are fully tuned in one exposure to a stimulus. This means that the strength of association between any arbitrary current stimulus and any given class cell token depends on only the correlation between the constituent structures of the current stimulus and the exemplar to which the mediating filter cell was tuned; it is not confounded by the frequency of stimulus occurrence. Second, while individual tokens do not possess constituent structure, they systematically evoke images of their referents (sensory analogs or subsets of other tokens), which can be structurally complex and are decomposable into constituent parts having their own semantic properties. Finally, individual tokens can be activated in combinatorial fashion so that cognitive processing is influenced by the joint effect of separate tokens within a structured representation (as in the query, "What has stripes and is an artifact?").

Moreover, given the capability of retinoid mechanisms to assemble complex, organized representational structures from successive inputs of simpler parts, a visual-lexical-semantic system that incorporates such mechanisms can, in principle, meet the challenge of productivity put by Fodor and Pylyshyn (1988). A system of this kind can not only represent veridical events, but it can inventively synthesize novel images (Trehub 1977, 1987) and sentences. Novel productions can then be projected to synaptic matrices to be learned and stored as newly tuned filter cells in an expanding long-term memory.

In the simulation presented in this chapter, a semantic network was queried by an exogenous source (the tester). Such queries can also be generated endogenously by the appropriate discharge of lexical tokens (self-query). In this way, the neuronal network can monitor the semantic implications of lexical communications in an on-line fashion without mediation or prompting from an outside source.

The communications subject to semantic processing can be self-generated sentences, as well as the lexical productions of others. Because lexical tokens are systematically bound to events in the real world, inferences drawn by the semantic network can provide information on which to base reasonable plans of action. These plans, in turn, can be expressed and learned as new lexical productions—tokened representations of action schemes that are subject to elaboration and interpretation within the semantic network.