Chapter 13
Narrative Comprehension and Other Aspects of the Semantic Network

Chapter 6 described the neuronal structure of a self-organizing semantic network based on the operating principles of the synaptic matrix. A simulation of the network showed that exposure to a set of simple sentences resulted in the automatic learning of lexical tokens (M_i, Ω_i), and when the lexical tokens were evoked by sentential stimuli (queries), they were automatically related to each other in a logical fashion through selective synaptic couplings induced by the intrinsic operations of the network. Thus, after learning a body of information conveyed by sentences that characterized the properties of named objects, the model was able to respond correctly to questions about the material learned. The simulations presented in this chapter test the performance of the putative semantic network with somewhat more complex semantic tasks.

When the semantic network is stimulated by proper sentences (character strings), it builds a lexicon and selectively organizes an internal pattern of synaptic enhancement. The distribution of synaptic weights established during learning then determines the input-output routing for neuronal activity later induced by word stimuli that correspond to items in its lexicon. In its basic form, the semantic network simply establishes privileged neuronal links between the subject (initial character string) and predicate (following character string) of each sentence presented to it. If there is a true logical structure in what the network is "told," it can capture that structure and make valid inferences that are reflected in its response to queries.

Unlike the visual model, which can parse objects out of a complex visual field, a neuronal mechanism for parsing lexical strings has not yet been formulated for the semantic network model. In all simulations of the semantic network, parsing is accomplished by a subprogram that detects alphabetic characters, spaces, and sentence terminators. After lexical strings are parsed, however, subsequent processing is carried out by the intrinsic neuronal properties of the simulated network.

Comprehending a Narrative

There are many formidable problems to be solved in the effort to develop even an abstract computational model of natural language processing (Clifton and Ferreira 1987, Miller and Johnson-Laird 1976, Stillings et al. 1987, chap. 11). The difficulty is compounded if we constrain our theorizing within the bounds of biological plausibility. The neuronal (semantic) network detailed in chapter 6 is by no means a broadly competent processor of natural language, but it does display some essential capabilities that make it a strong candidate for one of the modules in the brain's language processing system.

One of the central problems in comprehending a narrative relates to the assignment of pronoun reference. In normal discourse, pronouns are commonly used to achieve economy of expression and to convey a sense of continuity. For example, "Laura went to the library. *She* was looking for a particular book. *She* recalled that *it* had an unusual fifteenth-century map of the world. *This* was *her* special interest." If we are to understand these sentences as more than a sequence of independent assertions, we must be able to assign the pronouns to their appropriate references. We must know that *she* refers to *Laura* and not to *the library; it* refers to *a particular book; this* refers to *fifteenth-century map of the world* and probably to old maps in general; and *her* refers to *Laura*. This means that before the passage can be understood as a narrative, one must have prior knowledge of at least the rough equivalence classes for the pronouns used. However, in the general case, determining the reference of a pronoun can be a difficult problem because, except for number and gender, there are no exact criteria for pronoun substitution (Miller and Johnson-Laird 1976). Context and prior knowledge determine what a given pronoun can stand for in a narrative passage. Suppose we teach the semantic network some general facts about the meaning of words (including pronouns) that it will later encounter. Can we then present it with a simple narrative, ask it questions, and get responses that correspond to what one would expect if the passage were understood? The following simulation test demonstrates that, after it is given relevant general information, the semantic network can be "shown" a narrative passage and respond appropriately when it is later questioned about the passage.

The same neuronal mechanism that was described in chapter 6 was simulated, and the following list of sentences was typed in to be learned for background knowledge:

A MAN IS A HE .

A WOMAN IS A SHE .

BOB IS A MAN .

JENNY IS A WOMAN .

JENNY-AND-BOB IS A THEY .

A NON-HUMAN IS AN IT .

A CAR IS NON-HUMAN .

These sentences provided the network with information about the gender attachment of two proper names (Bob, Jenny) and two pronouns (he, she), the impersonal attachment of one pronoun (it), the plural attachment of another pronoun (they), and the impersonal nature of a common noun (car). Following the presentation of this essential general information, the following brief narrative was typed in:

> BOB DID NOT HAVE TO WORK THAT EVENING . HE PHONED JENNY . HE SUGGESTED A MOVIE . SHE SUGGESTED DINNER AND A MOVIE . SHE OFFERED TO DRIVE . HE INSISTED ON DRIVING . SHE SUGGESTED THAT THEY EAT AT BUB'S RESTAURANT . SHE HAD WINE WITH HER MEAL . HE HAD BEER . THE GLOBE WAS THE THEATRE THEY WENT TO . CABARET WAS THE MOVIE THEY SAW . THE CAR WAS OLD . IT BROKE DOWN ON THE WAY BACK . SHE WENT HOME BY TAXI . HE TOOK A BUS . THEY LAUGHED ABOUT IT LATER .

After the passage was presented to the semantic network, it defined all subjects of the input sentences by discharging each of the word cells (W_i) that constituted autaptic tokens of its learned lexical items. This resulted automatically in additional selective synaptic modification that constrained the network to make proper chains of inference on the basis of the sentences it had learned. Then the model was asked questions about the narrative:

Question: WHO DID NOT HAVE TO WORK?
Response: BOB DID NOT HAVE TO WORK

Question: WHO PHONED JENNY?
Response: BOB PHONED JENNY

Question: WHO SUGGESTED DINNER AND A MOVIE?
Response: JENNY SUGGESTED DINNER AND A MOVIE

Question: WHO JUST SUGGESTED A MOVIE?
Response: BOB SUGGESTED A MOVIE

Question: WHO OFFERED TO DRIVE?
Response: JENNY OFFERED TO DRIVE

Question: WHO INSISTED ON DRIVING?
Response: BOB INSISTED ON DRIVING

Question: WHAT WAS OLD?
Response: (THE) CAR WAS OLD

Question: WHO SUGGESTED THAT THEY EAT AT BUB'S?
Response: JENNY SUGGESTED THAT THEY EAT AT BUB'S

Question: WHO HAD BEER?
Response: BOB HAD BEER

Question: WHO HAD WINE?
Response: JENNY HAD WINE

Question: WHAT WAS THE THEATRE THEY WENT TO?
Response: (THE) GLOBE WAS THE THEATRE THEY WENT TO

Question: WHAT WAS THE MOVIE?
Response: CABARET WAS THE MOVIE

Question: WHAT BROKE DOWN?
Response: (THE) CAR BROKE DOWN

Question: WHO WENT HOME BY TAXI?
Response: JENNY WENT HOME BY TAXI

Question: WHO TOOK A BUS?
Response: BOB

Question: WHO LAUGHED ABOUT IT?
Response: JENNY-AND-BOB LAUGHED ABOUT IT

After learning the relevant lexical concepts, the semantic matrix was able to comprehend a simple narrative, judged by the appropriateness of its answers. Notice, though, that the passage permitted no ambiguity about the referents of *he* or *she*. These pronouns could only refer to Bob or Jenny. Similarly, *they* could refer only to *Bob-and-Jenny*. The pronoun *it* was less constrained but posed no problem in the context of the questions asked. What would happen if different nouns shared common characteristics for pronoun substitution? The following passage was presented to the semantic network, and its responses to two queries revealed its inadequacy in understanding narratives in which pronominal reference is ambiguous:

JIM MET FRED AT THE BALL PARK . THE YANKEES WERE PLAYING THE RED-SOX . THE GAME WENT INTO EXTRA INNINGS . JIM WAS AN AVID FAN . HE STAYED UNTIL THE END OF THE GAME . FRED WOULD RATHER WATCH FOOTBALL . THE PATRIOTS-AND-MIAMI WERE ON TV THAT NIGHT . HE LEFT AFTER THE NINTH INNING . THE GAME ENDED IN THE TWELFTH INNING .

Question: WHO STAYED UNTIL THE END OF THE GAME?
Response: JIM STAYED UNTIL THE END OF THE GAME
FRED STAYED UNTIL THE END OF THE GAME

Question: WHO LEFT AFTER THE NINTH INNING?
Response: JIM LEFT AFTER THE NINTH INNING
FRED LEFT AFTER THE NINTH INNING

A person reading such a passage resolves the ambiguity in pronominal reference by employing the heuristic of searching for and finding the most recent appropriate antecedent of each pronoun encountered. This means that the narrative must be learned and stored (at least in short-term memory) as an episodic sequence to allow a backward search for proper antecedents. A competent brain model must have the same capability.

Such a capability is provided by linking the semantic network to a synaptic matrix for episodic learning and recall by means of an adaptive interfacing matrix. In this case, words that evoke lexical tokens in the semantic network will, at the same time, establish corresponding class cell tokens within the matrix for episodic learning, and the temporal sequence of these episodic tokens will necessarily follow the order of lexical stimulation. In addition, since the discharge of each token in the semantic network will be paired with the concurrent discharge of its corresponding token in the episodic matrix, these tokens will be synaptically coupled through the interfacing adaptive matrix. When the semantic network is later queried about a narrative, its intrinsic logic evokes each of the lexical tokens that could be proper antecedents for pronoun substitution (see the responses *JIM* and *FRED*). The system then initiates a backward search (activation of successively "earlier" autaptic cells in the recall ring) over the episodic matrix starting at the location of the queried (activated) predicate. When it discharges the first token that matches one of the proper antecedents evoked in the semantic network, that nominal token in

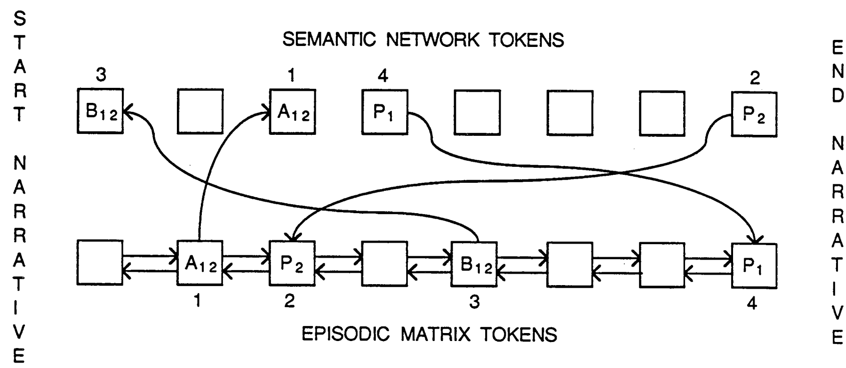

Figure 13.1
Semantic network augmented by episodic matrix. Modules are synaptically linked through an interfacing adaptive matrix by the normal coactivation of corresponding lexical tokens (represented by squares) during the learning of a narrative passage. Lexical tokens in semantic network are logically related but not temporally organized. Lexical tokens in episodic matrix are temporally organized but not logically related. P_1 and P_2 are pronominal predicates. A_{12} and B_{12} are different subjects, but each is a logically proper referent of both P_1 and P_2, and both subject tokens are evoked in the semantic network when either predicate is queried. Numbers above and below the tokens indicate the relative temporal order in which each occurs in a narrative passage. If predicate P_1 is queried, a backward search over the episodic matrix encounters B_{12} first. Discharge of B_{12} in episodic matrix elevates output of B_{12} (the appropriate antecedent) in the semantic network. If predicate P_2 is queried, a backward search encounters A_{12} first. Discharge of A_{12} in episodic matrix elevates output of A_{12} in the semantic network.

the network (the most recent with respect to the predicate that is activated) will receive an increment of stimulation that raises its activation level above its competing token(s), and the final (overt) response will be the correct one. The dynamic coupling between the semantic network and the episodic matrix is schematized in figure 13.1. The constraints on lexical response imposed by this hypothesized neuronal subsystem correspond to the natural heuristic for pronominal reference.

The augmented system was simulated and its comprehension of the narrative passage about Jim and Fred at the ballgame was tested again:

Question: WHO MET FRED AT THE BALL PARK?
Response: JIM MET FRED AT THE BALL PARK

Question: WHO WERE PLAYING THE RED-SOX?
Response: (THE) YANKEES WERE PLAYING THE RED-SOX

Question: WHO STAYED UNTIL THE END OF THE GAME?
Response: JIM STAYED UNTIL THE END OF THE GAME

Question: WHO LEFT AFTER THE NINTH INNING?
Response: FRED LEFT AFTER THE NINTH INNING

Question: WHO WERE ON TV THAT NIGHT?
Response: (THE) PATRIOTS-AND-MIAMI WERE ON TV THAT NIGHT

The responses to the questions "Who stayed until the end of the game?" and "Who left after the ninth inning?" show that the previous failure to disambiguate pronominal reference was corrected when the temporal processing capability of the episodic matrix was utilized with the logic of the semantic network. The competence of the augmented model resulted from its ability to instantiate the heuristic of selecting the most recent one of different logically appropriate pronominal antecedents.

Negation

In all of the sentences considered so far, each predicate or predicate phrase asserted the existence of an attribute, an action, or an event—*is a man, took a bus, were on TV*. The meaning of a sentence can be fundamentally changed by negating a stated predicate—for example, "Fred was not an avid fan." "What has feathers but does not fly?" "He had no tickets to the concert." The words *not* and *no* are general-purpose lexical signs that assert that the next word or phrase is false with respect to its referent. The negational modifier is an important aspect of human communication and reasoning and must be instantiated in the neuronal mechanisms of the brain's semantic network.

Figure 13.2 presents an elaboration of the basic semantic network that provides it with an intrinsic competence for the logical processing of negation. During learning or when a subject (W_i) is to be defined, the autaptic cell S (top left of figure) is automatically activated. This inhibits discharge of the autaptic cell N if the negators "NOT" or "NO" are encountered in a sentence, and these words will have no privileged effect on lexical processing. However, if a predicate token is initially activated in order to infer a subject, the S cell will be bypassed, and, thus, it cannot inhibit the N cell. In this case, recognition of the lexical string "NOT" (or "NO") fires the autaptic cell N, which gives an excitatory bias to all of its target autaptic cells (n), each discretely paired with a cell in the mosaic array (M). A short collateral excitatory output from each M cell synapses with its paired n cell. Each n cell makes an inhibitory synapse on its paired M cell.

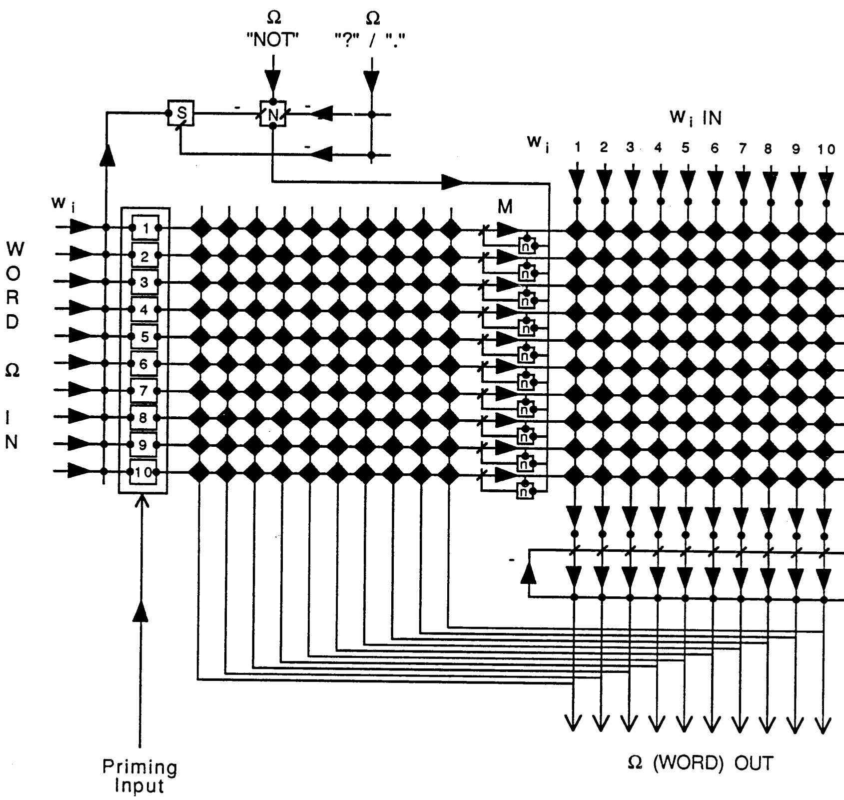

Figure 13.2
Semantic network for the logical processing of predicate negation.

When the network is queried by activating one or more predicate tokens (Ω_i) and before "NOT" is encountered, the predicate collaterals in the imaging matrix will discharge those mosaic cells corresponding to their associated subject tokens (via previously learned synaptic links). The activation of each subject will proceed in normal fashion. As soon as "NOT" is detected, however, all *n* cells are automatically primed. Because of the priming, the discharge of the next mosaic cell (subject) by its proper predicate token will then fire its linked inhibiting *n* cell, and that subject token will immediately be squelched. Notice that negation does not inhibit the negated predicate but does inhibit any subject evoked by that predicate. Thus, if the network has learned no more than "An apple is a fruit" and is then asked "What is a fruit?" it will respond "(An) apple"; but if it is asked "What is not a fruit?" it will give no response because the proper subject, "apple," will have been squelched. Notice at the top of figure 13.2 that detection of a sentence terminator inhibits both the *N* cell and the *S* cell. This resets the negation circuit to its initial condition by removing excitatory bias from all *n* cells and allows subject tokens to discharge normally until the next "NOT" is encountered in the context of a predicate query.

The semantic network, including the neuronal circuitry for processing negation (figure 13.2), was simulated. The following list of sentences was presented to the model:

JUMBO IS AN ELEPHANT .

AN ELEPHANT IS BIG .

AN ELEPHANT IS A MAMMAL .

A MAMMAL IS AN ANIMAL .

AN ANIMAL IS A LIVING-THING .

A MOUSE IS A MAMMAL .

A MOUSE IS SMALL .

AN ELEPHANT IS GRAY .

A MOUSE IS GRAY .

A CUCUMBER IS A VEGETABLE .

A CUCUMBER IS GREEN .

A VEGETABLE IS A PLANT .

A PLANT IS A LIVING-THING .

A POPPY IS A FLOWER .

A FLOWER IS A PLANT .

A POPPY IS RED .

A DANDELION IS A PLANT .

A DANDELION IS YELLOW .

A ROCK IS A MINERAL .

A MINERAL IS A NON-LIVING-THING .

A RUBY IS A GEM .

A GEM IS A MINERAL .

A RUBY IS RED .

A GEM IS VALUABLE .

After these sentences were typed in and the definitions of all subjects had been tested, the network was queried as follows:

Question: WHAT IS RED?
Response: (A) POPPY
(A) RUBY
(A) GEM

Question: WHAT IS RED AND IS A NON-LIVING-THING?
Response: (A) RUBY
(A) GEM

Question: WHAT IS RED AND IS NOT A NON-LIVING-THING?
Response: (A) POPPY

Question: WHAT LIVING-THING IS RED?
Response: (A) POPPY

Question: WHAT LIVING-THING IS NOT RED?
Response: JUMBO
(AN) ELEPHANT
(A) MAMMAL
(AN) ANIMAL
(A) MOUSE
(A) CUCUMBER
(A) VEGETABLE
(A) PLANT
(A) FLOWER
(A) DANDELION

Question: WHAT IS GRAY?
Response: JUMBO
(AN) ELEPHANT
(A) MOUSE

Question: WHAT IS GRAY AND SMALL?
Response: (A) MOUSE

Question: WHAT IS GRAY AND NOT BIG?
Response: (A) MOUSE

Question: WHAT IS GRAY AND NOT SMALL?
Response: JUMBO
(AN) ELEPHANT

Responses were correct not only when the network was required to infer a subject on the basis of negated predicates after exposure to a list of definitions. Negation was also appropriately processed by the network when it was tested for narrative comprehension. For example, the following passage was presented to the model:

> LARRY LOOKED FORWARD TO THE CONCERT . ANN LOOKED FORWARD TO THE CONCERT AS WELL . LARRY ARRANGED TO MEET ANN AT THE AUDITORIUM . THEY HAD BOUGHT THEIR TICKETS IN ADVANCE . ANN ARRIVED ON TIME . LARRY ARRIVED A BIT LATER . ANN HAD A TICKET WHEN THEY MET . LARRY WAS EMBARRASSED . HE HAD FORGOTTEN HIS .

Question WHO LOOKED FORWARD TO THE CONCERT?
Response: LARRY LOOKED FORWARD TO THE CONCERT
ANN LOOKED FORWARD TO THE CONCERT

Question: WHO BOUGHT THEIR TICKETS IN ADVANCE?
Response: LARRY-AND-ANN BOUGHT THEIR TICKETS IN ADVANCE

Question: WHO ARRIVED WITH A TICKET?
Response: ANN

Question: WHO ARRIVED WITH NO TICKET?
Response: LARRY

In answering the last question, the semantic network did not have enough background knowledge to base its response on the meaning of "embarrassed" or "forgotten." Thus, it could not infer from the fact that "Larry was embarrassed" or that "he had forgotten" that it was Larry who did not have a ticket when he arrived for the concert. When the network responded correctly to the question "Who arrived with no ticket?" it reached its decision simply because it had information that Ann arrived and had a ticket, but it had no information that Larry had a ticket when he arrived. Therefore, the intrinsic logic of the semantic network did not suppress the response "Larry" as a possible subject for one who had arrived and had no ticket, whereas the response "Ann" was inhibited because it contradicted the predicate condition of having no ticket.

If the sentences "Larry was embarrassed" and "He had forgotten his" were not included in the narrative passage, the semantic network would still have said "Larry" in response to the question "Who arrived with no ticket?" But in this case, the question posed would have been inappropriate or deceptive.