CHAPTER 2

Computer Models of Long-Term Memory

Production system and connectionist models of long-term memory (LTM) are described below, including both spreading-activation and featural approaches. Chapters 5, 6, 7 (particularly Sections 1, 2, and 5), and 8 in the text describe material related to this content and should be read as preparation for the following pages. If the reader is not familiar with the basic principles of production systems or artificial neural nets, then the chapter on computer functioning (on the net) should also be covered.

Schema models, which emphasize higher-order processing structures, were described in the text in Chapter 10.

1. Anderson's ACT: A Spreading-Activation Model

Anderson and Bower (1973) designed a network, spreading-activation model of reasoning and memory called *human associative memory* (HAM). The theory was later developed across a series of different forms of ACT (Anderson, 1976), called ACT* (Anderson, 1983) and ACT-R (Anderson, 1993; Anderson & Lebiere, 1998; Anderson, Budieu, & Reder, 2001; Budiu & Anderson, 2004). ACT-R stands for adaptive control of thought, rational.

According to ACT theory, higher-order cognitive processes are controlled by a single set of principles. The rules that direct mathematical operations also direct verbal reasoning, or the coding of material into memory—and so for all higher-level cognitive functions. The exception to this general rule involves perceptual processes such as vision and hearing. Anderson thus opposes the view that different kinds of memory depend on specialized, underlying structures. Instead, it is posited that the same general-purpose system operates, using different modes of processing to produce different codes.

1.1. Encoded and Retrieved Information in ACT

Anderson believes that the best candidate for a high-level, unitary-function model of cognition and memory is a production system (PS). A production system is a symbolic artificial intelligence (AI) model that emphasizes the role of working memory (WM).

ACT includes declarative-memory, procedural-memory, and working-memory constituents. In this model, declarative memory consists of information stored in LTM. It is latent information, which requires active processing to “unpack” it. This active processing is achieved through procedural memory.

Critical to the ACT hypothesis is the view that when information is activated in memory, this will lead to the selective activation of other, related memory content. This is the “spreading-activation” assumption. It is a means by which the system can move appropriately through the memory network, as against randomly searching the net as a whole. That is, if the system has activated a certain body of information, then it is likely that the activation of other, related information will be more useful than the activation of unrelated content.

1.2. Representations in ACT

In older versions of ACT, representation was viewed as involving propositional codes only. But in ACT* and the current ACT-R, several coding formats are assumed. Order information is handled by *temporal strings*. An example of such a string would be the ability to remember that a series of events occurred in a specific order. When string codes are involved, material can be inserted into the string without change in the previously existing information. Thus, in the string A, C, F, it would be possible to insert B, producing A, B, C, F. The insertion of B would leave A, C, and F unchanged,

The second code in ACT involves spatial imagery. The spatial image code possesses a complex hierarchical structure. For instance, when we see a stimulus, its location may be coded at a higher level in the hierarchy than its shape. In the case of weak memories (due, for instance, to brief exposure to the stimulus), we might therefore be able to recall that an object was seen to our right, but not recall the shape of the object.

The third type of information posited in ACT involves general abstract information. (The Atlantic is a cold ocean. Or, I plan to go on holiday in July.) This takes complex propositional form.

1.3. Processing and Organization in ACT

Different processes operate on the three representational types described above. However, all can be found in the same memory. A particular incident may be recalled in the form of abstract general content (you met Joe yesterday after a period of years), image information (you see Joe's face in memory) and order information (first you saw Joe, then you introduced yourself, then you went for a cup of coffee).

Storage and retrieval are achieved in an identical way for the three types. There are constituents of each type that are called *cognitive units*. There are limits on the amount of information that can be encoded into a cognitive unit. Thus, complex information is coded hierarchically. For a given body of information there will be a higher-order node, associated with lower nodes, which will in turn be associated with yet lower. In this structure, access can occur either from the high nodes down or the bottom nodes up. In some cases, access is blocked. If this occurs, the blocked associative paths will not lead to retrieval.

1.4. Retrieval in ACT

In ACT, content in WM is highly activated LTM content. In other words, if any LTM information becomes activated above a certain level, it will enter WM.

Working memory includes material of which the individual is conscious, as well as material on the fringes of consciousness.

Retrieval occurs as follows. Suppose you are trying to remember the content of a film, *The Man Between*, seen many years ago. A goal requirement would be established in WM. The goal is of course retrieval of *The Man Between*. A goal requirement entails a high level of activation in WM (of the goal specification).

ACT is a production system. A production system operates on an IF/THEN basis. Here, IF certain conditions specified in WM are matched in LTM, THEN the LTM content will be contacted and strongly activated, and probably retrieved. Thus, if a goal involves remembering a *film* called *The Man Between*, and IF this specification in WM is matched with a LTM representation corresponding to Film, *The Man Between*, the title will be retrieved from LTM, accompanied by some associated information such as, for instance, the film's major theme.

Once the title and the theme have been recalled, this will lead to the selective activation in LTM of more content associated with the film. This occurs on the basis of similarity relations between the first-activated content, now in WM, and other content in LTM. For instance, content involving a betrayal theme may lead to content involving a character involved in betrayal, and that may lead to further information about the character, and so on. Some or all of this additional material may be retrieved into awareness.

The matching function in ACT is critical to the success or failure of the attempt to recall the target information. At the beginning of the process, various tests will be run to determine whether the goal specification matches with some body of LTM content. These matching tests can only be performed on information that is activated to some extent in LTM. The higher the level of activation of the LTM content, the faster the matching process can be achieved. There is a limit on the amount of time that the system will attempt to achieve a match. If no match is found within that window, the target content will not be retrieved. The testing for a match occurs at all stages throughout the period of retrieval.

Thus, successful recall depends on the level of activation of the LTM content. If the content is highly activated, the match will be successful. If activation is weak, the matching function attempt may end before retrieval is achieved. And there will be no attempt to match inactive LTM content. The complexity of the match slows down the process, however, such that matching is a function of the degree of complexity (causing slower processing) and the level of activation of the LTM content. In ACT, constituents in LTM can weaken over time, although they are never wholly lost from the system.

1.5. The Provision of Activation

Activation in ACT theory can be provided in three ways. First, any percept will activate the corresponding LTM representation. This activation will be maintained for as long as the relevant stimulus is physically present. Thus, if you are looking at a table, the table concept will be activated, as will information relevant to that particular table.

A second source of activation is provided by the similarity factor described above. Once a body of content, perhaps reflecting Node X, becomes active in LTM (above baseline level), it will lead to the selective activation of related information, also in LTM. The strength of this activation depends on the strength of Node X and the strength of the nodes that show a pattern of connection with X. The node's strength will be manifested in the following way. The greater the strength of the node, the greater the activation of the links leading from it. The overall level of activation will be divided among the links. Thus, the more concepts associated with X (and so the more links leading from X), the lower the level of activation on each link.

The concept of activation through similarity is extremely important. What is being suggested is that if some body of content, perhaps involving information about the country of Brazil, has been activated, then the system will begin to activate other information related to Brazil. This provision of the model means that such activation will not be random. Material related to the current focus of thought or memory will come to the fore. Obviously, a mechanism of this kind would obviate the need for extended searches that might reach only irrelevant content. Also, once similar content has been activated, the retrieval information in WM may find a match with any part of that content—and so bring such content into awareness. Note that this assumption does not involve spatial distance. If X has been activated, it will not activate a representation that is "close" to it in some spatial fashion in LTM (such a relation probably does not even exist). X will selectively activate related, similar information. Also note that the spread of activation occurs at an unconscious level within LTM and is not controlled by the cues in WM. It is controlled by the relations that hold among the various bodies of content in LTM itself.

The third source of activation in ACT centers on the structures in WM established by goals. When an individual wishes to recall a certain memory, the act of recollection of that information content becomes a goal. Goal specifications in WM remain constantly activated, unless a decision is made to change the goal. They can then match with LTM content corresponding to their specifications, at any time.

Attention also plays a role in the theory. When attention is directed toward a certain body of content in WM, this will provide more activation than will obtain for unattended content.

Suppose I am again trying to recall the content of *The Man Between* and I am attending to my effort to remember that content. Given the focus of my attention, I may remember something about *The Man Between*, while not "seeing" a table that is directly in front of me.

Activation once established in the network spreads very rapidly and also decays rapidly. Content within the sphere of spreading activation is generally capable of being recalled. This is the case because a high level of activation makes matching possible.

Anderson has developed a series of equations to express the factors that come into play to produce increased or decreased activation in memory, as the system operates. The model has been able to show results that parallel human performance across a range of general and specific areas, such as priming and fan effects, among others.

1.6. The Fan Effect

According to the ACT* and ACT-R models, nodes possess a certain strength or level of activation. When a given node is accessed, activation will spread into LTM along whatever links extend from that node. The level of activation is fixed, at any given moment in time, and is limited. This activation quantity is divided among the links leading from the relevant node. Thus, if three links lead from that node, the activation level along each will be only one third of the amount that would obtain if only one link extended from the node.

Anderson (1974) reported that the more random facts participants learn about a target individual, the slower they are to identify these facts in a recognition test. The phenomenon was identified as a *fan effect*.

The Anderson study involved statements involving a person in a location, presented in sets of the following kind:

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

The numbers in brackets indicate the number of times either the person was present in the set, and the number of times the location was present. Participants were tested with a mixture of original sentences and foils. The foils re-paired the people with new locations. For example, "The lawyer is in the bank." Network representation of the four original statements is shown in Figure W2.1.

Anderson (1976) showed that fan effects can be obtained not only with newly learned statements such as those described above, but also with familiar material. In

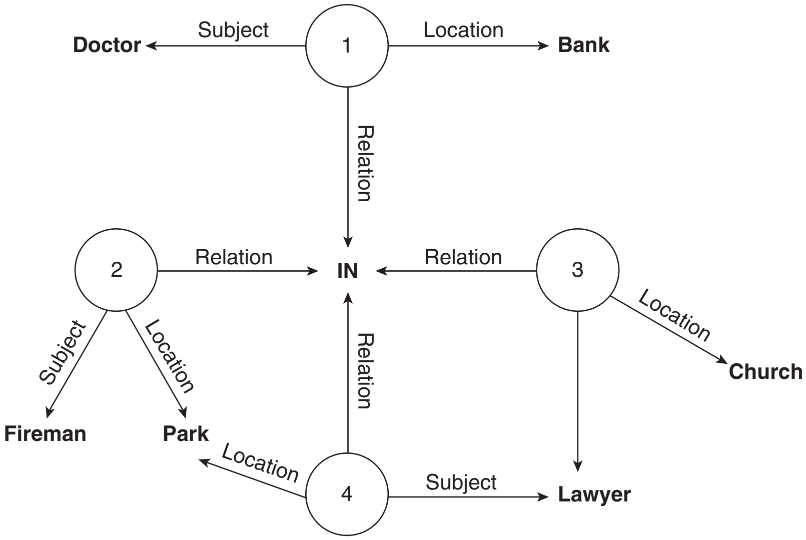

Figure W2.1 A Propositional Diagram of the Four Statements in the Anderson (1974) Study

NOTE: When an individual is characterized as being in a location, this is seen as a description of the individual, and a "subject" link is used to connect the individual to the proposition node.

this study, fantasy facts, such as "Napoleon Bonaparte was from India" were presented, as were true facts, "Napoleon Bonaparte was an emperor." From 1 to 4 different facts were presented for both the fantasy and real materials, across different conditions. The recognition test involved the presented true facts, the presented fantasy facts, or foils. There were fan effects for both the real and the fantasy material.

Reaction time was significantly faster, however, for the real facts as compared to fantasy facts. And, critically, the more fantasy facts participants had learned, the slower they were to identify the real statements.

The data can be explained as follows. Familiar concepts have a greater strength within the system. The corresponding nodes thus generate a higher level of activation, and responses are therefore faster to the realistic information than to the weaker fantasy information. However, in both cases the larger number of facts to be recalled (having been presented within the experimental session), the more the available activation was divided among them. Even when fantasy and true facts were involved together for the same person, some of the activation was channeled into the links associated with the imaginary facts—and so recognition of even the strongly coded "real" information was slowed.

The theory behind fan effects provides an explanation of similarity-based interference. Where similarity or identity is involved, correct order and correct pairings become more difficult to retain. This is because the spread of activation required to express the order (or the pairings) is reduced when the same concept is linked to more than one other concept. (Note that this approach provides a different interpretation of interference from the response competition assumptions described in Chapter 2.)

1.7. Fan Effects or Situation Model Effects?

In the studies described above, an equal latency was found for fan effects associated with a person or a place. Latency is a measure of how long a participant takes to respond to a test question. Here, it made no difference whether the question focused on a person or place; only the number of links leading to the target made a difference.

Radvansky, Spieler, and Zacks (1993), however, using a somewhat different methodology, reported slower reaction times for objects than for location and slower reaction time for location than for animate entities.

Also in the Radvansky et al. study, when an object was described in three different locations response time was slowed significantly more than when three different objects were described as being in a single location.

Zwaan and Radvansky (1998) offered an interpretation of these findings in terms of a situation model effect (see Chapter 10). Participants could establish, for instance, a location as a specific context and then imagine three objects in that location. In contrast, if they had to imagine three separate places—three separate models of context—this would provide a greater drain on memory. The situation argument is further strengthened when a processing structure component is assumed. We routinely construct information about places in which many objects are present. If a location structure exists to achieve this, it is likely that it would be specialized to retain many objects easily.

At first the Radvansky data appeared to weigh against the fan (divided activation) assumption, since the fan hypothesis would not predict weaker fan effects for objects than for locations. Anderson and Reder (1999) reported, however, that if it was assumed that more attention is given to the entities that are presented repeatedly (that is, more to the three objects in a single location and less to the location), and if this attentional factor is entered into equations predicting the final outcome, then the fan hypothesis predicts the same outcome as found in the Radvansky studies.

A number of alternative explanations for the fan effect have been offered (Anderson & Spellman, 1995; Conway & Engle, 1994). The question of whether such effects are due to the division of available activation or to other causes is currently an area of ongoing research.

2. Connectionist Models of Memory

In connectionist models, input stimuli involve sets of features. For instance, in a net designed to identify birds, the input features "small, brown, medium-thick beak, country dwelling, town dwelling" might lead to the identification of the stimulus as a sparrow. Identification can also occur when only some of the relevant features are present. Equally, if an entity such as "sparrow" is provided as input, the net can generate a description of a typical sparrow.

Information in a neural net is coded through a very large number of simple units that are massively interconnected. A memory is a pattern of facilitated activation among a set of such units (as is believed to be the case for memories coded by neurons). According to McClelland (1995), memories of this kind have two critical properties. If an episode, say Episode X, has been coded, then it will be represented in the pattern of activation among the units established by that episode (the input provided by that event). But in addition, the units involved in coding for this episode will also have interconnections with other units not directly corresponding to Episode X. Cues involve the input of some component of the memory. For instance, if I try to recall, "What did I eat for breakfast this morning?" these cues will match with a memory including the information "breakfast this morning," and will activate the rest of the pattern. This pattern will tend to correspond to the one created when the memory was first formed. Perhaps I drank coffee and ate cereal, and this information was coded at the time of the event and activated again when I recalled the event. However, since there will be associations between the units involved in the breakfast memory and other information in LTM, the playing through will not always be exactly the same. If the connections establishing "ate cereal" are weak and I generally have an egg for breakfast, the information "ate egg" may become more strongly activated than "ate cereal," and the result could be a reconstructed memory that differs from the original. The details of how such an event could occur are explained below.

2.1. The Jets and the Sharks

McClelland (1981) developed a model to illustrate the property of generalization in human recall, and also the capacity to provide new or inferred content, both true and false. The model centered on a group of small-time criminals, members of either the Jets or the Sharks gang. Each individual has certain properties, or features, such as being a burglar or having a high school education. Figure W2.2 shows a simplified diagram of the connections between these features, arranged into "property units" and central "instance units." An instance unit links the different properties of a given individual together, via associative connections.

If an instance unit were activated, activation would spread to the properties associated with it. For instance, the unit on the top right-hand side of the pool provides associative relations with Rick, 30s, a burglar, divorced, a Shark, and someone possessing a high school education.

Equally, if an individual is encountered who is in his twenties, single, a Jet, and possesses a junior high school education, these input features will strongly activate the instance unit that is associated with Lance. Other instance units will not be as strongly activated, since few of the input properties correspond to them. Also, there is lateral inhibition within these units and within the property units. What this means is that as one unit becomes more and more activated, it will inhibit the activation of other units within the same domain. In the present instance, as the Lance unit becomes highly activated, this will weaken the activation units of all the other individuals in the two gangs.

The system can identify people on the basis of yet smaller amounts of information. If the input, "Who is a burglar, in his thirties?" the 30s feature will activate units associated with Ralph and Rick, but the burglar feature will only activate the

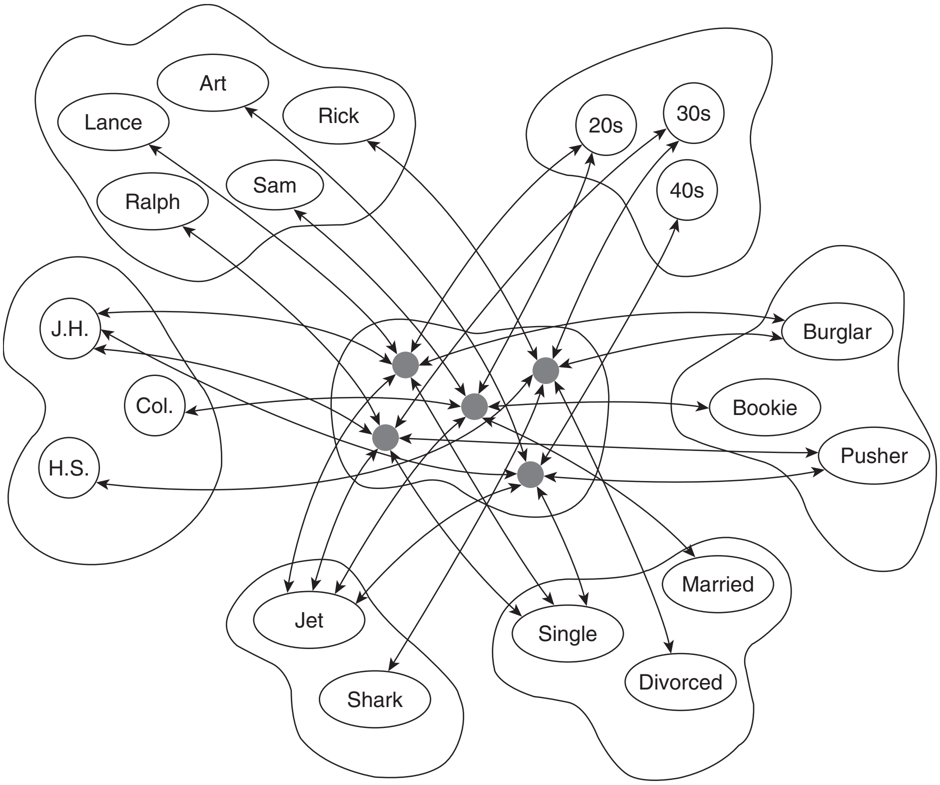

Figure W2.2 Network for the Jets and the Sharks

SOURCE: Reprinted with permission of J. L. McClelland.

NOTE: The black circles are the instance units. An instance unit connects all the information for a given Jet or Shark together. The property units are shown in the outer areas. Activation along the links travels in both directions.

feature associated with Rick. That unit will therefore become the most activated, and the individual will be identified as Rick.

Information is inferred from the net on the basis of similarity. An event of this kind is most likely to occur when information has been "forgotten" or not coded in the first place. Suppose the link between Lance's instance unit and "burglar" was not present in the system. When Lance is activated, the activation will travel to "Jet, 20s, high school education, married." This information is associated with other individuals in the net. These individuals are burglars. Thus, activation will move to "burglar" and be sufficiently strong for an output. The system will thus output the information that Lance is a burglar. This has been achieved on the basis of generalizing from the fact that other individuals similar to Lance are in fact burglars.

Here the inferred information happens to be correct. But suppose Lance was really a bookie, but this information has been lost from the system (forgotten). The same events would play through, and the system would again determine that Lance was a burglar. But here the inferred information would be incorrect.

2.2. Representations Formed Through Gradual Learning

Hinton (1981, chap. 6) developed a network capable of accepting input in the form of propositional statements and slowly forming a complex representation of meaning of that input. The program might be told, "Fish are animals," "Fish can swim," and "Fish have gills." Each of the relevant concepts could also participate in other complexes of information. For instance, the network might code, "Animals can breathe," "Animals have nervous systems," or "Gills enable breathing," and so on.

As work in this area progressed, programs were developed that could provide a more extended body of appropriate information than could be generated by a person trying to "think through" all the relevant properties. After much programming, the meaning of a given entity would be expressed as a particular pattern in the network, say Pattern X. Pattern X would be similar to the patterns established for similar concepts, and it would be different from the patterns established for dissimilar concepts (Rumelhart, 1990; Rumelhart & Todd, 1993). The program generalized, or inferred, information, as part of its natural mode of functioning. For instance, Rumelhart found that when the net had been trained on a large set of propositions concerning canaries and robins, he had only to input the statement, "A sparrow is a bird," for the net to be able to identify the properties of sparrows: that they can fly, have wings, have feathers, lay eggs, and so on.

An example of patterns formed by a network when it is first exposed to a description of some entity (for example, "An oak is a living thing") and when it has been exposed to a large amount of information concerning that entity, is shown in Figure W2.3. Note that at first the patterns for dissimilar concepts are much the same, but as the net acquires more information, the patterns shift and become distinctive.

A neural net can thus be seen as a powerful model of human conceptual representation. But it lacks one critical property. The net can only establish bodies of information of the kind described above on a slow basis: With each new input, all the weights among the units must find the correct adjustment to continue to respond to all possible inputs appropriately. Again, this can be achieved, but only through gradual accommodation of the net as a whole. An approach that works within this context is called *interleaved learning*. Here new learning occurs very gradually, with the system being exposed to new material, alternating with reexposure to old material (interleaving), so that appropriate weights can be found for both.

The problem is that if many new inputs are provided one after the other, the net cannot adjust quickly enough. Content similar to the new input will be disrupted. For instance, if you input a body of information about robins, involving perhaps color and diet, the net is at risk of losing information about the colors and diets of other birds. This outcome has been labeled *catastrophic interference* (McCloskey & Cohen, 1989). Yet humans do not suffer catastrophic interference; we can learn new things about robins, and this will not disrupt well-established information about other birds.

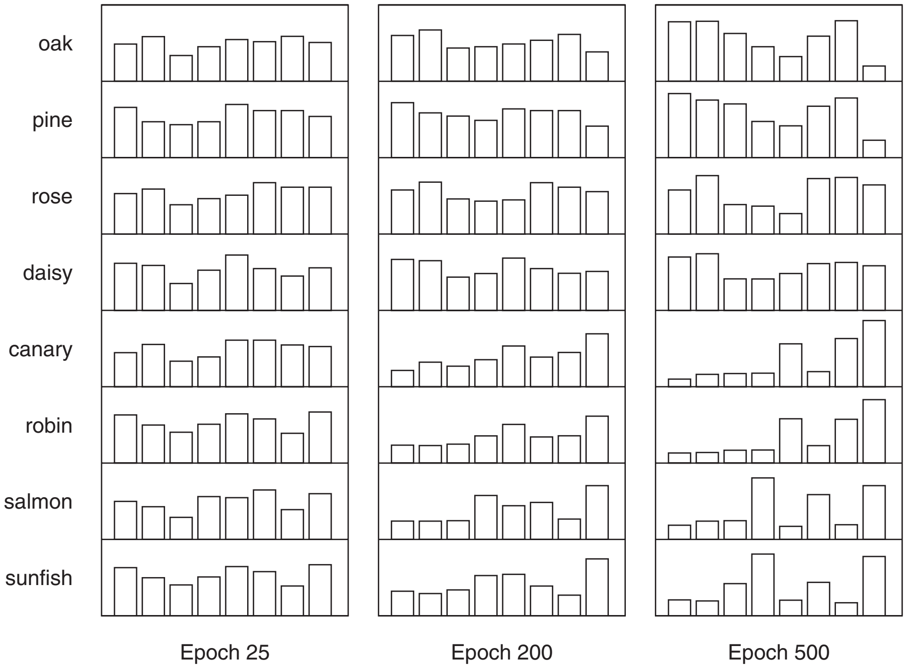

Figure W2.3 Responses of a Network Coding for Semantic Concepts

SOURCE: From McClelland, J. L., McNaughton, B. L. & O'Reilly, R. C., "Why There Are Complementary Learning Systems in the Hippocampus and Neocortex: Insights From the Successes and Failures Of Connectionist Models of Learning and Memory," Technical Report PDP.CNS.94.1, March 1994, *in Psychological Review, 102,* copyright © 1995, American Psychological Association. Reprinted with permission.

NOTE: The network accepts propositions about each concept. When the first propositions are presented, the net shows little differentiation between "oak," "canary," and "salmon." As more propositions are presented, the net increasingly settles into patterns that are different across concepts (different for oak, canary, and salmon) but similar for exemplars of the same superordinate concept: that is, similar for oak, pine, and rose, since they are all plants.

2.3. A Model of Intermediate and Long-Term Memory

In response to the issues described above, McClelland, McNaughton, and O'Reilly (1994) and O'Reilly and McClelland (1996) have proposed that human memory in fact involves two systems. It is well known that damage to certain structures in the brain produces a condition in which the individual can form no new long-term memories, although recollection of information learned before the injury remains intact. The individual may thus not recall a conversation held 10 minutes earlier, although she can remember going on holiday the summer before she suffered the trauma to her brain.

In some cases, remote memory (memory for the holiday) is also impaired, but the critical point is that new learning can be wholly disabled, while past learning, from before the time of the injury, is unimpaired. These data clearly suggest two nonidentical forms of memory, mediated by separate structures within the cortex.

Damage to the medial temporal lobes of the brain produces the outcome described above. McClelland et al. therefore posited a memory system, centered on the medial temporal area, that codes for new memory content. To be precise, individuals suffering from amnesia can retain information for a few minutes, but not longer than a few minutes. The hypothesized system would involve the function that codes new information over periods extending beyond a few minutes. This function is known as intermediate memory.

Under the present hypothesis, the medial temporal lobe system operates to integrate the constituents of new, entering content together. Models that can perform such integration have also been developed (Nystrom, Leigh, & McClelland, 1992). The new information can be coded quickly. It is subject to a degree of similarity-based interference, but there is no complete breakdown of earlier knowledge, since earlier knowledge is handled by a different system.

The new content, coded in the medial temporal lobes, is gradually integrated into this second system, based in other areas of the neocortex. This is the function that provides very long-term (remote) memories and semantic knowledge. But here a much slower process is involved, such that the brain can accommodate new content little by little (in an interleaved fashion). Each reactivation of the material provides the opportunity to integrate it with other background knowledge. This system operates like the Hinton (1981) and Rumelhart (1990) models described above, gradually accumulating an extended body of knowledge that is not subject to similarity-based interference, once it has been strongly coded, but is capable of providing a powerful inferential capacity. The knowledge coded into the neocortical areas will continue to be developed throughout life.

A dual system of this kind could clearly explain the data involving amnesia, in which new memories cannot be formed, although remote memories remain intact.

A major goal of the two-memory model described above was to avoid the problem of catastrophic interference. Other researchers, such as Grossberg and Stone (1986) and Lewandowksi (1992), have developed models that also escape this difficulty.

2.4. Conclusion

The connectionist models introduced here posit distributed representation: that is, meaning content coded by many processing units, with dissociative connections between them. By “dissociative connections” is meant associative relations that may form, and also be lost, such that new relations may later be established. The approach is also fundamentally constructivist in nature. An experienced memory is here the result of reconstructive activities occurring at the time of retrieval. New information may thus be established relevant to a given episode, such that the original memory is changed.

3. Kintsch’s Construction Integration Theory * * * * *

Walter Kintsch has developed a model of prose recall and recognition called construction integration (CI) theory. Like the ACT hypothesis, the approach has evolved through several forms (Kintsch, 1988, 1992a, 1992b, 1998; Kintsch and

Van Dijk, 1978; Kintsch & Welsch, 1991; Van Dijk & Kintsch, 1983). According to CI theory, text comprehension is achieved first on the basis of an automatic, bottom-up mode of processing, in which activation spreads from the input representations, via links in a memory network, to other representations. This stage is described as context insensitive: That is, the higher or general context of the input plays no role here. The input concepts, expressed in words, have preestablished connections with other concepts in LTM, and these relations simply play through.

In addition to the activation of words associated with input words, the system will generate representations of the meaning of the input text reflecting the content of entire phrases and sentences.

The first stage is followed by a second, "constraint-satisfaction" stage, in which higher-order context comes into play. Now activated units in LTM that are irrelevant to context will be overridden, while relevant units will be further activated. In addition to information derived from the original text, information will be provided from LTM as it is needed to generate a fully coherent body of memory content (the situation model).

3.1. Codes in the CI Model

Kintsch assumes that there are many coding formats in human representation, from direct perceptual content to the level of formal abstraction. The one at which full semantic expression is first involved, however, is propositional in structure. In the case of memory for text, both the representation of the original text and the added LTM content will be expressed as propositions.

3.2. Knowledge Nets

Memory within the present model consists of a network of associated nodes. A node may reflect any of several forms of content, but all have propositional structure.

The links in the network are unlabeled. Any given node is connected to others on the basis of *retrieval structures.* That is, the activation of any node will provide a retrieval structure that will selectively activate other nodes; this process is described as a link. The links are powerful and stable, although they may vary in strength. The model is thus an activation model.

Kintsch supports the variable-meaning hypothesis for conceptual representation. This is the view that a concept consists of an extended body of information, but that, in any given memory, only a subset of that information will be represented. Thus, a concept node in the CI model will have associated with it a large set of propositions. In any given situation, a subset of these propositions will be activated as the meaning of the concept. In most cases, a yet larger number of the set will not be activated. The particular propositions activated depend on the general context, resulting in the operation of retrieval cues appropriate to that context.

As the same concept is repeatedly recalled, the meaning will change to some extent. This is due to the fact that the retrieval situation will sample different propositions within the conceptual set, across different retrieval attempts. A probabilistic component is also involved here (reflecting which propositions happen to be sampled at any given moment), and so also the relationship between activated cues and content present in LTM. However, much of the meaning will remain constant, given that the cues will largely remain constant and the information is being drawn from the same propositional set.

Information in human memory shows a high level of coherence and organization. For instance, incoming material is smoothly integrated with background knowledge, and the result is a meaningful body of content. Within the Bartlettian tradition, it is believed that this structured content is achieved by means of schemas, which direct the process of comprehension.

Kintsch's CI model reflects the assumption that traditional schema theory is incorrect on this point. Instead, the processes that operate in the forming of a memory are posited to be associative and bottom-up. Nodes that are linked with other nodes will spread activation on that basis alone. The result is a loose organization, from which schema-like properties typically emerge. But there are no schemas directing the ongoing processing activities, top-down. The organization is instead a function of preexisting associative relations among the nodes. It can be seen as an emergent. At the least, the claim is made here that schematic control may not be as tight or definite as classic schema theory has suggested. Kintsch makes the argument that human memory is too flexible and too context sensitive to be ruled by fixed control structures (Kintsch, 1998, p. 94).

Material is established in LTM on the basis of weak production system rules. They are also "dumb" production system rules because they know nothing of context. Spreading activation throughout the story memory content (a process also reflecting production system rules) will lead, however, to the dominance of context-relevant information over irrelevant information. (How this occurs is explained below.) Spreading activation here occurs as in connectionist networks.

3.3. Latent Semantic Analysis

The representations from which propositions are formed all have meaning content in human cognition. If I read, "The dog jumped the fence," a semantic code for DOG will be activated, and so on. Researchers working with computer simulations thus face the task of programming the meaning of any word that might occur in a text. To develop a memory of this scope within the program would be extremely difficult. Kintsch selected a different approach, first developed by Landauer and Dumais (1997), known as latent semantic analysis (LSA). Here a body of text, such as an encyclopedia, is used. For each segment of the text, a count is made of all the words that co-occur. Then a multidimensional semantic space is developed, in which co-occurring words are situated close to one another. The meaning of any given word is then expressed in terms of its closeness to, or distance from, other words. This is the approach to the representation of word meanings used in the CI model.

3.4. Encoding and Retrieval

In the CI model of text recall, an input text is transformed into propositions. Propositions are generated on the basis of a set of rules. A generalized version of the text (the macro level) is also developed. For instance, if the text involved various details in which an individual went to a supermarket, had trouble with a cart, and bought milk, coffee, vegetables, and meat, the macro level would code, "X went shopping."

A further set of rules generates relevant inferences. Such inferences, derived from LTM knowledge, can be of many kinds (causal, spatial, etc.). This inferred content, added to the facts provided explicitly by the text, constitutes the situation model. A macro level of discourse is also provided to express the situation model. All these constituents (input statements, inferred content, and macro descriptions) are associatively linked together.

Inferred content may draw on scripts. For instance, if an individual was described as being in a restaurant, the associated knowledge present in the restaurant script would be activated and would come into play as a constituent of the total activation spreading in the network reflecting the encoding of the input text. Imagery may also be generated in the construction of a text. Further rules operate for integrating the various propositions with one another. These link individual propositions. Such links also function within a complex proposition, in which the constituent simple propositions are also bound together.

Activation spreads from this body of material on the basis of the many associations that it possesses with other material in LTM. In the present model, the meaning of each concept is expressed by the network of associations (with other concepts) provided by the LSA formulation. Activation will spread among these constituents. At first the spreading activation is quite chaotic; all associations that exist will play through. These include associations reflecting the different meanings of ambiguous words. As associated bodies of knowledge relevant to context come into play, however, they will begin to provide greater activation to context-relevant material than to context-irrelevant material. Only the higher activation content is likely to be retrieved. Thus, context will begin to organize the material into a coherent body of information. This is the Integration stage of the model.

Material will later be retrieved from LTM as a function of the cues deployed in working memory. For any given representation, such cues can involve the representation node itself. Context also functions as a cue or cues. Context can reflect the general nature of the retrieval task (for instance, the representation of the identity of a target story) and the story theme. In addition, goals can serve as retrieval cues (for instance, the goal of recalling the target story), as can the individual's relevant past experiences.

3.5. Working Memory in the CI Model

Kintsch's model of WM was described in Chapter 5. Working memory is understood here as being of strictly limited capacity. The general meaning or theme of the text is entered into WM. As propositions from the text are developed, they are added to the theme and integrated with one another. Typically, WM is engaged in the active processing of one sentence and of any information needed to understand it. The integrated content is then entered into LTM.

When reading a natural story, the individual may feel that the entire general content of the story is held active in her WM. According to the present model, however, this is not the case: The capacity of WM is instead strictly limited. The process of retrieving information into WM from LTM, however, is continuous: And this can give the impression that a large body of content is being actively maintained in awareness.

3.6. The CI Model and Human Data

Kintsch has suggested that if a top-down schema approach directed processing, then the schema would block context-irrelevant meanings. If this was the case, the "wrong" meaning would not come into play—and could not prime related words. In contrast, under CI theory, both meanings of an ambiguous word (both sets of associations with other concepts) would be activated. It has been shown that when experimental participants read sentences such as, "The earthquake destroyed all buildings in town except the mint," items associated with mint as candy (such as "candy") are in fact primed. Further evidence that bottom-up associative processes continue to work (even after context has been established) has been reported by Rayner, Pacht, and Duffy (1994).

Garrod, Freudenthal, and Boyle (1994) conducted a study to examine how the elements referenced by pronouns are identified in text material. In many cases, there is no means of distinguishing between referents on a strictly syntactic basis. Instead, the reader must draw on world knowledge. The following passage provides an example.

> Flying to America, Jane wasn't enjoying the flight at all. The dry air in the plane made her really thirsty. She noticed the stewardess coming down the aisle with the drinks trolley.

Continuation A: Right away she ordered a large glass of Coke. Continuation B: Right away she poured a large glass of Coke.

Human experimental participants identify the "she" in Continuation B as meaning the stewardess. This is clearly based on background knowledge concerning the roles of stewardesses and passengers on planes.

The same input was simulated in the CI model. Here "stewardess" would activate long-term memory knowledge concerning the nature and roles of stewardesses. The system would also infer that Jane was a passenger and activate the nature and roles of passengers. When Continuation B was input, the Jane node originally possessed more activation, this being the referent in the discourse focus. As the processing cycles continued, however, activation from Jane became less in comparison to the activation of stewardess, as information concerning the association of stewardesses and drink pouring extended into the memory net. The pattern of activation finally settled on the proposition “stewardess poured glass of Coke.”

A network representing the coding of the passage in LTM is shown in Figure W2.4. The activation levels for the two continuation sentences are shown in Figure W2.5.

The data described above thus support the claim that patterns of text comprehension reflecting general knowledge, in this case of scripts, can be simulated by an associative, bottom-up process.

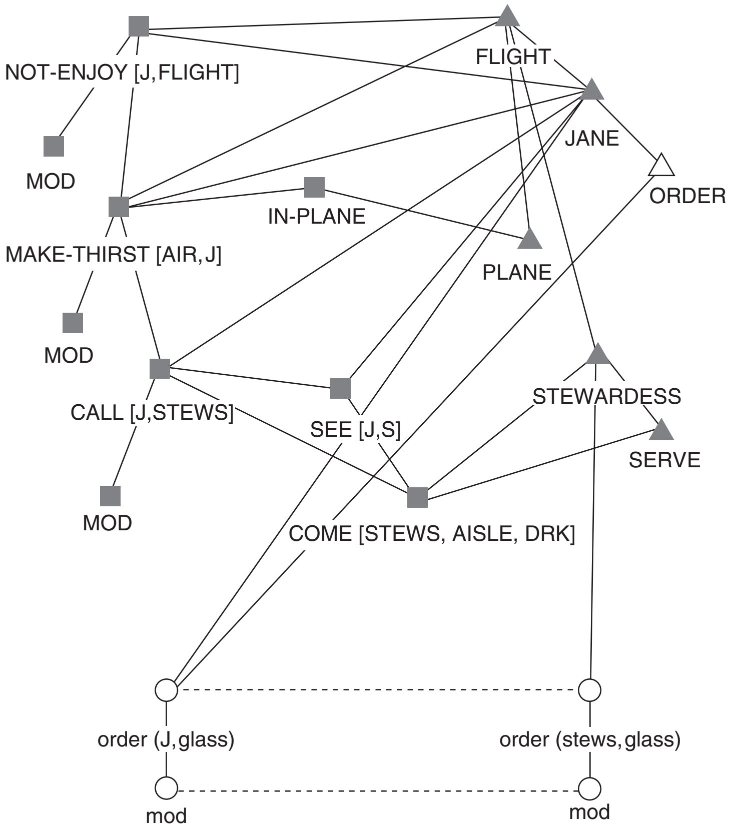

Figure W2.4 Kintsch’s Network Coding for Memory Content Concerning an Airplane Journey

SOURCE: From Kintsch, W., *Comprehension: A Paradigm for Cognition.* Reprinted with permission of Cambridge University Press.

NOTE: The network embodies the information that Jane did not enjoy the flight, that the dry air on the plane made her thirsty, and so on. The final pronoun “she” in “she ordered” remains ambiguous at this point. It could further activate Jane or the stewardess. Activation coming in from the network providing background knowledge about who performs what activities on a plane will result in Jane being the most strongly activated referent for this proposition.

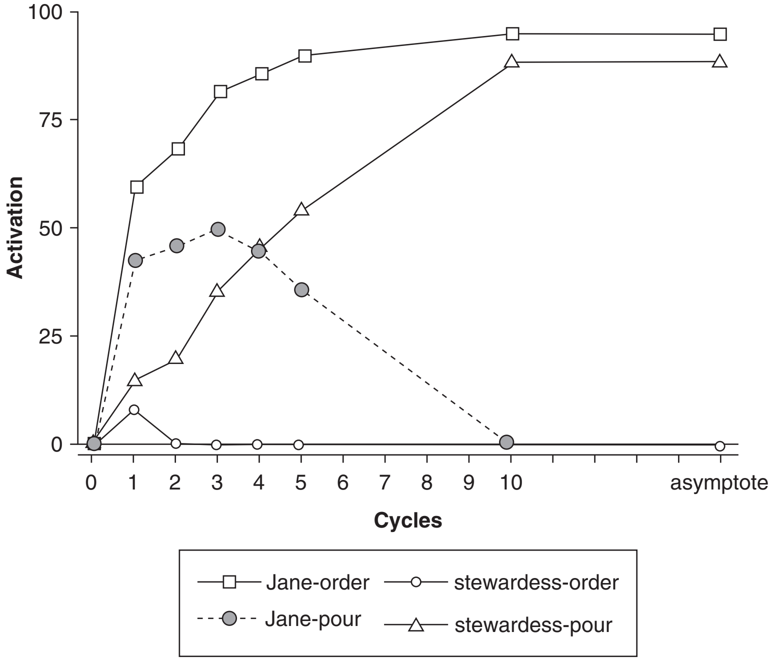

Figure W2.5. The Activation Levels in Kintsch's Simulation for the Two Possible Interpretations of the Statement "She Poured"

SOURCE: From Kintsch, W., *Comprehension: A Paradigm for Cognition.* Reprinted with permission of Cambridge University Press.

NOTE: Across cycles, "Jane pour" receives less and less activation. "Stewardess pour" receives more and more activation.

4. Search of Associative Memory

Gillund and Shiffrin (1981) and Raaijmakers and Shiffrin (1981, 1992) developed a computer simulation of recall called *search of associative memory* (SAM). In 1984, Gillund and Shiffrin published a new version of SAM, intended to model both recall and recognition and to show that these two functions are related: That is, they operate to a considerable extent in the same way.

Probably the most common model of recall reflects the two-stage, or generate-recognize, assumption (Anderson & Bower, 1972, 1974; Kintsch, 1970, 1974). The first stage is a search process, in which material is activated in LTM. This occurs as information spreads through a memory net. In the case of word list learning, search processes are believed to be slow. The second stage involves a "decision" process, in which the system determines whether some activated content does or does not correspond to a memory target.

Here recognition involves only the second stage. The system does not need to search for a target item, since that item has been presented to the individual, and direct access of its representation in LTM should follow.

Gillund and Shiffrin (1984) noted, however, that recognition is known to exhibit properties that are generally associated with a search process. For instance, list length, list organization, and depth of processing all influence recognition as well as recall. Yet if recognition involves the simple, direct access of the target items, none of these factors should play a role. Suppose the item BIRD is a test item for recognition. The associative relation between the test cue (BIRD) and the word item BIRD stored in long-term store should determine whether the item is recognized, and it should not make any difference whether a lot of other items (HOUSE, ROSE, CAT, BOY, etc.) are also on the list or whether only a few items are on the list.

A simple response to the present issue might be to posit that a search process is used in recognition. But there is a major difficulty with this assumption. Subjects can often respond very quickly to a foil (a nontarget item), identifying it as an item they had not encountered in the learning session. Search processes require time. The system would have to search extensively before it could determine that an item had *not* been presented before, and this should eliminate the possibility of a very speedy response.

Given these data, some researchers have proposed that in a recognition test, a swift, direct-access approach is employed. This is based on familiarity. A highly familiar item will be recognized, an item with a very low familiarity count will be rejected, and items with an intermediate familiarity level will give rise to a search process (Atkinson & Juola, 1973, 1974; Juola, Fischler, & Wood, 1971; Mandler, 1972, 1980). This would explain why variables that are believed to influence a search process (such as list length, etc.) are found in tests of recognition. Early models tended to posit that a familiarity test was followed, under some conditions, by a search process. More recently the view has been widely supported that the two processes (which are of different kinds) operate in parallel.

Gillund and Shiffrin (1984) reported a study in which speed of response had been varied. When subjects were required to make a speedy response in a recognition test, it could be assumed that a direct-access, familiarity-based process was involved. Under slow response conditions, a search process could come into play, if search processes do in fact operate in recognition. If search processes do not operate, then the longer period given for a response would simply permit direct access to operate across that longer period. The researchers manipulated a number of additional variables that could be expected to influence recognition. For instance, the foil items were varied between items that were not similar to the target items, and items that were synonyms of the test items or similar to them either graphemically or phonemically.

The longer periods given for response produced superior recognition performance. Critically, however, the manipulated variables described above all influenced recognition in the same way under either speedy-response or slow-response conditions. The authors reasoned that if two qualitatively different processes were at work (direct access and search), then it could be predicted that the manipulated variables should affect them in different ways. By the same logic, if the same process was at work across the quick and slow response conditions, then the manipulated variables should affect response in the same way.

The authors noted that it was possible that different processes were at work across the two response conditions, and that it simply happened that these processes responded to the manipulated variables (such as phonemic similarity or semantic identity, etc.), in the same way. But the data on the whole supported the view that the same process (a direct-access process) was involved.

Gillund and Shiffrin (1984) next posited that there may exist a form of direct access that can be described as complex, rather than simple. This complex process could result in data similar to that found when search processes operate. For instance, it could affect recognition depending on amount of rehearsal, list length, list organization, and so on. On this basis, it might be possible to explain the fact that both recall and recognition often show similar properties and sometimes show different properties. The complex direct-access assumption was then incorporated into SAM, as described below.

4.1. Encoding in SAM

When SAM has learned a list of word items, the items are represented in LTM as sets of features, called *images*. Each image comprises one item on the target list. The sets reflecting a single list are closely interconnected, and each functions as a unit. The sets once formed are permanent. Images contain three kinds of information. Type *a* involves context information: that is, the association between the item and the context in which it was learned. Type *b* involves the associations between a given item and other items on the list. Type *c* involves the association that obtains between an item and itself. (With regard to type *c*, for instance, if an item such as HOUSE is presented for recognition, there will be an associative relation between that test item and the HOUSE image stored in LTM.)

Context can involve setting and temporal factors related to the learning of the list. The identity of the list as a whole is often expressed on the basis of context factors. For instance, it might be conceptualized as "the list I learned this morning," or "the list I learned in the psychology lab."

During learning, items are maintained in a limited-capacity buffer (corresponding to short-term store or more recent versions of WM). About four items can be held in the buffer. During rehearsal of the items, associations are formed between the items and the context, between the items that are rehearsed in the buffer together, and between each item and itself. Rehearsal results in the corresponding codes being entered into LTM. If two items are not rehearsed together, they still possess some small degree of association, on a preexperimental basis.

Another type of cue is a *category cue*. Thus, if an item of a list had been SPARROW and the prompt BIRD (the category cue) was given, the prompt would show a strong association with the target image. The relationship is of course also preexperimental. (This property of SAM means that when items are established for rehearsal in the buffer, features reflecting their meaning are present in the rehearsed image.)

SAM incorporates a *retrieval structure*. This consists of the associative relations between all possible cues and the list items. Thus, the relations between context and each item, and between any given item, Item X, and other items on the list, and the self-association of each item with itself, is represented in the retrieval structure.

Also, the low associations that exist among items that were not rehearsed together, but that exist because of preexperimental associations, are present.

An example of a retrieval structure is shown in Table W2.1a. An example of the recovery probabilities is shown in Table W2.1b. A learning task with four images is shown. Context is identified as C, and each image as I. D represents items not present on the to-be-learned list. The left-hand matrix shows the associative strengths of context with each image and of each image with each of the other images. The next-to-left matrix shows the strength of probe sets in which context plus image cues are used together, as a probe set, to contact the set of images in long-term store. For instance, C + Ii reflects the product of the context association to Image 1 (0.5) and the C + Ii shows the association between the context plus Image 1, working together (0.3, the Image 1-Image 1 association, and 0.5, the Image 1-context association, giving 0.3 × 0.5 = 0.15), and so on.

4.2. Recognition in SAM

In SAM, when items are presented for a recognition test, two cues come into play. The first is the relevant context cue, and the second is the representation of the test item. The context and target item cues together constitute the probe set. (In verbal terms, the subject is asking something like, "Was the word BOAT present on the list that I learned this morning?")

The complex aspect of the direct-access process here reflects the following. The operating cues in the probe set will not simply make contact, providing a certain level of activation, with a possible target image stored in long-term store. They will instead make contact with all the images in long-term store. For instance, whatever associative connections exist between the probe set cues and other images stored in long-term store will also come into play, providing activation. The resulting familiarity value, then, on which a response will be based, depends on the activation of all the images in memory. (Of course this level of activation will depend on the associative strengths between the probe cues and the images. If the test item was in fact learned earlier, then the greatest contribution to the activation will be provided by the associative strength between the operating cues and the target image in long-term store, with a smaller contribution being provided by the associative strength between the test item and other, noncorresponding images in LTS.)

When a model incorporates the assumption that cues contact all content in LTM and that the resulting total activation (of all cues with all the content) is used by the system to determine the response of the system, the approach is known as a *global memory model*. This contrasts with the view that response will be based only on the activation provided by the cues and the target information. The global model assumption provides a form of complex direct access, as against a simple form.

If the activation in the recognition test in SAM is above a certain level, a positive recognition response will be given. The subject determines the cutoff, or "criterion," level. For instance, if it is important not to make a mistake, only a high familiarity level will be accepted.

CUES	RETRIEVAL STRUCTURE Memory Images I_1*	I_2*	I_3*	I_4*	PROBE SETS	STRENGTH TO PROBE SET I_1*	I_2*	I_3*	I_4*	SUM OF STRENGTHS Familiarity of Probe Set	PROBABILITY OF SAMPLING AN IMAGE WITH THE PROBE SET I_1*	I_2*	I_3*	I_4*
C	.5	.3	.8	.4	C	.5	.3	.8	.4	(2.0)	.25	.15	.40	.20
I_1	.3	.3	.4	.1	C+I_1	.15	.09	.32	.04	.60	.25	.15	.53	.07
I_2	.3	.4	.1	.1	C+I_2	.15	.12	.08	.04	.39	.39	.31	.20	.10
I_3	.4	.2	.7	.3	C+I_3	.20	.06	.56	.12	.94	.21	.06	.60	.13
I_4	.1	.1	.2	.4	C+I_4	.05	.03	.16	.16	.40	.13	.08	.40	.40
D_1	.1	.05	.1	.1	C+D_1	.05	.015	.08	.04	.185	.27	.08	.43	.22
D_2	.2	.1	.3	.1	C+D_2	.10	.03	.24	.04	.41	.24	.07	.59	.10

Table W2.1a Context c and Images 1 to 4, and the Association Value Between Each Image and Every Other in the Retrieval Structure

SOURCE: From Gillund, G. & Shiffrin, R., "A retrieval model for both recognition and recall," *Psychological Review, 91,* copyright © 1984 American Psychological Association. Reprinted with permission.

NOTE: These values are established during learning.

	Sum of Strengths				Recovery Probabilities			
	I_1*	I_2*	I_3*	I_4*	I_1*	I_2*	I_3*	I_4*
C	.5	.3	.8	.4	.39	.26	.55	.33
C+I_1	.8	.6	1.2	.5	.55	.45	.70	.39
C+I_2	.8	.7	.9	.5	.55	.50	.59	.39
C+I_3	.9	.5	1.5	.7	.59	.39	.78	.50
C+I_4	.6	.4	1.0	.8	.45	.33	.63	.55
C+D_1	.6	.35	.9	.5	.45	.30	.59	.39

Table W2.1b The Associative Strengths Involved in Recovery and the Recovery Probabilities

SOURCE: From Gillund, G. & Shiffrin, R., "A retrieval model for both recognition and recall," *Psychological Review, 91,* copyright © 1984 American Psychological Association. Reprinted with permission.

The global memory assumption is able to explain why searchlike processes appear to operate in recognition tasks. This occurs because images of items other than the test item are included in the activation that determines the system's response. As a result, items other than the target items will influence the success or failure of the recognition process. With regard to list-length effects, for instance, in the case of a long list of items, the probe set will be contacting many items with which the set has very low associative connections, as well as the target image where the associative connection is strong. In the case of a shorter list, fewer very low-association strength items will be contacted.

Figure W2.6 shows a flowchart of the recognition process in SAM. F is the activation (familiarity value) established on the basis of the context cue and the item cue contacting all images in long-term store.

4.3. Recall in SAM

When recall is attempted in SAM, the long-term store is entered originally on the basis of context cues only. Context operates as a cue because it was associated with the list items during learning, as described above. The strength of the associative relation determines how much potential any cue has to activate an image in LTM. Thus, a context cue with a 0.6 association with a target item will make a stronger contact than a cue with only a 0.2 association with the target image. This first step in recall can be understood as reflecting what occurs when an experimental participant directs herself to recall "the list I learned this morning."

Recall begins with a search process. The first step in this process is described as *sampling*. Here the context cues make contact with the images in long-term store. For any given image, its likelihood of being sampled depends on a ratio rule. The ratio involves, as numerator, the activation of that image alone, in long-term store, in response to all the operating cues. The denominator involves the activation of all the images in long-term store in response to all operating cues. Note that the denominator involves global access. The ratio rule can thus be shown as follows:

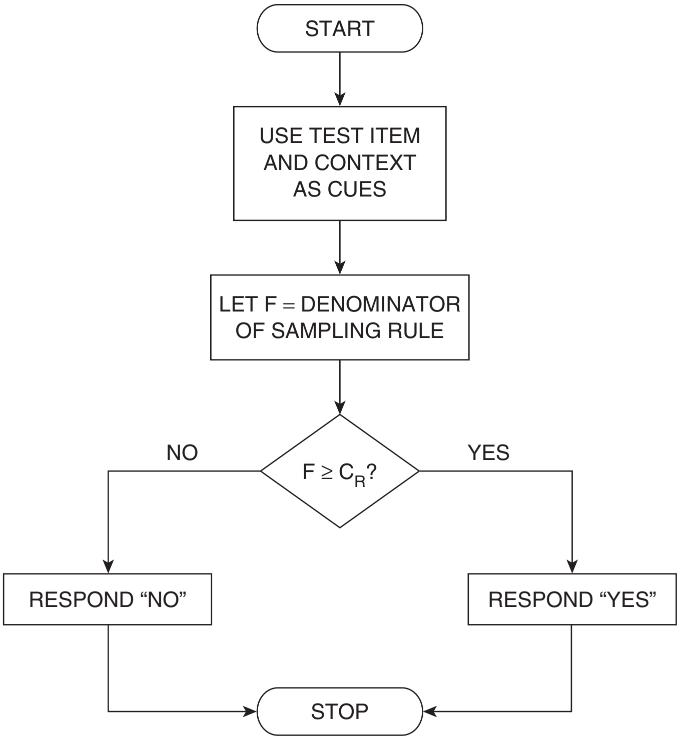

Figure W2.6 A Flowchart of the Recognition Process in SAM

SOURCE: From Gillund, G. & Shiffrin, R., "A retrieval model for both recognition and recall," *Psychological Review, 91,* copyright © 1984 American Psychological Association. Reprinted with permission.

$$\frac{\text{activation of all the images in long-term store in response to all operating cues}}{\text{activation of a given image in response to all operating cues}}$$

The image that will be sampled first is the image with the strongest activation: that is, the image that is most strongly related to the cue or set of cues currently in operation.

When an image is sampled, the system will attempt to recover relevant information from the image. For instance, an attempt would be made to recover the name of the word item represented by the image. A decision process is built into this second stage of the recall function. The system must determine whether the recovered information is relevant to the task requirements. If the information is relevant, the item is retrieved. In the present case, the name of the item would be relevant.

Once an item has been recalled, both the original context cues and that item will be used in a new probe set to contact long-term store again. These new cues will provide the sampling of another image, based on the ratio rule shown above. A recovery and possible retrieval process will then again follow. The system will continue the process of sampling until it repeatedly fails to retrieve any item. At that point, the attempt at recall would normally end.

A flow chart of the events involved when SAM models recall is shown in Figure W2.7.

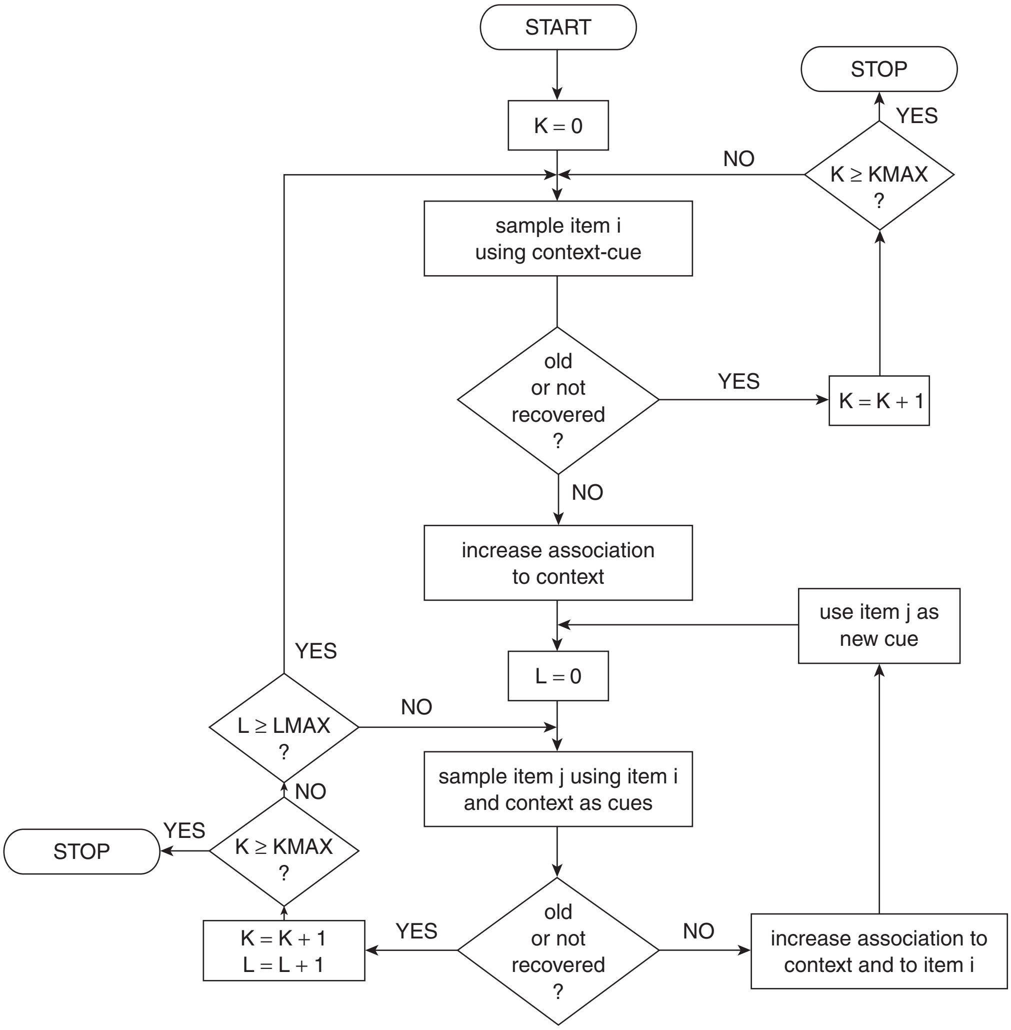

Figure W2.7 A Flowchart of Retrieval in SAM Under Free Recall

SOURCE: From Raaijmakers, J., & Shiffrin, R., "In search of associate memory," *Psychological Review, 88,* copyright © 1981 American Psychological Association. Reprinted with permission.

4.4. SAM as a Cyclical Retrieval Model

SAM is similar to the earlier generate-recognize models in that recall again involves more than one stage, including a decision/evaluation process. It differs from those models in that sampling and recognition occur on a global basis, as against a situation in which retrieval depends solely on the relation between the operating cues and the target, when access first occurs. The global access involved in SAM can explain the searchlike properties found in many recognition tasks. In fact no search is involved, but because the cues contact all the items in long-term store, nontarget items can influence the retrieval outcome.

A second major difference between SAM and many two-stage theories is that SAM does not involve spreading activation in LTM. Although the model posits a limited associative connection between images in LTM, this connection plays little role in retrieval. The critical factor for retrieval is the relation between the operating cues and memory content. There is a cyclical retrieval process in which retrieved images operate as secondary/compound cues to repeatedly enter long-term store and retrieve new material.

4.5. Simulation Data

As noted earlier, there are variables that affect recall and recognition in the same way, and other variables that affect these functions in different ways. SAM has been extremely successful in explaining these phenomena. For instance, in the case of list length, both recall and recognition are impaired in the case of longer as against shorter lists. SAM posits that this identical effect is in fact due to entirely different processes. In the case of recall, the effect is due to the sampling rules. As a list increases in length, the probability of sampling any given item decreases. As a result, fewer items are retrieved.

In the case of recognition, the mean difference between target items and distractor/foil items, in terms of the two distributions of familiarity (for the OLD, target items, and for the NEW, foil items) remains the same as list length increases. But the additional words increase the number of items that are only weakly associated with the test item cue. The extra items thus increase the variability of the two distributions. With this increased variance, the distributions overlap to a greater extent than occurs in the case of short lists. The area of overlap is the area where false recognitions are likely to occur, and so accurate recognition is impaired.

SAM has been able to simulate the human data when a wide range of other variables are manipulated, including presentation time, variability in encoding, context shifts, and testing delays. As with all current computer simulations of learning, however, SAM has not been able to duplicate the full range of available human data.

A major contribution of the model remains the hypothesis of global contact between cues and all material in long-term store, such that nontarget content may influence the probability of recall.

5. MINERVA

Hintzman (1984, 1986, 1988) introduced a model of retrieval, MINERVA 2, in which it was assumed that all stored memory content consists of episodic traces only. An episodic trace involves a trace that represents a particular event. According to this view a separate semantic memory store does not exist.

Under the present model, when rehearsal occurs during the learning of a word list, the rehearsal of a given item does not strengthen or change a unified memory trace of that word in secondary memory. Instead, each rehearsal creates a separate trace (of the same item). At the time of recall, cues contact all the stored traces simultaneously. Each trace is activated depending on its similarity to the operating cues. The traces respond in parallel, with the information that is retrieved reflecting their summed output.

Cues contact all content stored in secondary memory. However, they will generally show a high level of similarity only with the target traces, such that the targets alone are likely to be activated and retrieved.

In MINERVA, when traces in secondary memory are activated, they pass the activation to their constituent features. Also, traces are connected to other traces, by either positive or negative links. Activation can travel among them on this basis.

The information retrieved into primary memory is called the *echo*. The echo has two properties: intensity and content. Intensity reflects the summed activation of all the traces that have been contacted in secondary memory. It corresponds to the notions of familiarity or frequency, as described within the context of signal detection. A well-learned item will have a high level of intensity/familiarity. The second property of the echo is content. Content involves the information present in the echo. Information in MINERVA is coded in the form of feature vectors. Each feature may be represented as +1, −1 or 0. 0 indicates that the relevant feature was never stored. Each individual feature is coded with a certain probability (L, the learning vector), and each feature is learned individually. That is, some features in a given trace may be learned, although others are not.

Retrieval is never spontaneous in MINERVA, but is always produced by a cue or cues. The cue set is called a probe.

The retrieved echo can be used as a second probe, to enter the LTM store again. The echo cue thus contacts the stored information again, on the basis of similarity between the echo and the stored information. The match is generally better than during the first contact, such that the echo then retrieved corresponds more closely to the original input. This retrieved second echo can then be used to enter the long-term store yet again. The result is often the final retrieval of content that is almost identical to the originally stored content.

Rehearsal is of course known to strengthen the memory function. If a given, invariant trace for a word item does not exist and so cannot be enhanced by rehearsal, how does the model explain the improvements provided by rehearsal? The effect in fact occurs because rehearsal creates multiple traces, and the cues contact all traces. The activation of the traces is summed. Thus, many contacted traces will provide a higher intensity to the echo than only a few traces.

MINERVA 2 posits that all representations are episodic, individual traces. The retrieval of what might appear to be unified semantic or generic information is achieved through a construction involving the summed response of the traces that are most strongly activated by the cues. For instance, the model has successfully simulated a phenomenon known as the *schema abstraction task* (Posner & Keele, 1968). In the Posner study, subjects were shown patterns of dots that had been created by taking a specific shape, called a prototype, and moving some of the dots out of line. The prototypes were familiar figures such as a triangle or a square. When they were transformed, the resulting pattern was either close to the original prototype (a low level of change), or fairly different from the original (a high level of change). Subjects were trained to categorize the transformed shapes into one of three categories. The categories had been defined by the prototypes. In short, subjects were being asked to identify that all forms of transformed square belonged in the same category, and so on. Critically, however, they were never shown the actual prototypes during training. They were then tested with items identical to the training items (OLDs), or new transformed items, (NEWs), either with a low or a high level of transformation, or on the original prototypes. Again, they had not seen the prototypes earlier in the study.

Under both immediate and delayed testing, subjects performed better with OLDs than with NEWs, and better with low-level change figures than with high-level change figures. The critical finding, however, involved the response to prototypes. Under immediate testing, OLDs were handled better than prototypes, but under delayed testing, prototypes were handled better than OLDs. In short, the subjects had identified the underlying single pattern on which the various transformations had been based and could recognize it more easily (in delayed testing) even than figures that they had previously studied. This was seen by the authors as an example of a single pattern, abstracted out from the various transformed versions and stored in LTM.

MINERVA 2 simulated this study and obtained data identical to the human data. Yet the program had not stored an abstract pattern. The result was achieved by summing the intensity of response of the various features present in the stored picture items, to the operating cues (the test picture). Hintzman's conclusion was that what may appear to be a single, abstracted representation is in fact a product of the relationship between cues and individually stored, separate traces.

The data could be seen as supporting Hintzman's general position—that a semantic memory, involving unified or abstract representations of meaning, does not exist. It may be that what appears to be unified representation (such as a representation of the concept TABLE) is in fact a phenomenon achieved during recall, based on the concerted activation of multiple separate traces.

MINERVA simulates recognition and cue-based recall. Like the models described earlier, its performance corresponds to some, but not all, of the human data. Its success in the area of visual patterns, described above, is striking. It cannot be assumed, however, that a finding relevant to a visual pattern can automatically be generalized to semantic concepts. This question remains to be answered.

Summary

In a production system, functions operate on an IF, THEN basis. In ACT* and ACT-R, cues in working memory may find a match with corresponding content in LTM. IF this occurs, and IF the content is at a sufficient level of activation, THEN the content will be recalled. If the content is at a low level of activation, the cues cannot match it. Cues increase the activation level of material that they do contact. When several links lead from the same node in LTM, the available activation is divided among them. This is known as a fan effect. Fan effects provide a possible explanation of similarity-based interference.

In connectionist models, the net consists of units with links connecting them. The links have weights and the units have thresholds. If the activation (the sum of the operating link weights) exceeds the threshold of a unit, activation will be passed on. Representation is distributed: Many units provide a given representation, and these can also be used as constituents in other representations. Activity will spread through the net based on the pattern of linked units. There is no boundary to a given memory. If a name unit, X, possesses properties in common with other name units, A, B and C, and A, B and C share yet another property, ZZ, the activation in the net is likely to spread from X to ZZ. The ZZ property is thus inferred for X.

Some recent neural net models posit that two major memory stores or functions are involved in long-term store. One reflects new learning and the other information that was learned in the past. This approach avoids the net's breaking down when exposed to a high level of new information in a short period of time.

Kintsch's CI model is a model of understanding prose and of memory for prose. Input content, information generalized from that content, and relevant background knowledge are all coded. Meaning content takes propositional form. Spread of activation within the system is based on associations among its elements: That is, it is "bottom-up." At first the associations show no sign of context effects. However, as elements increasingly activate context information, through simple associative processes, context comes into play. There are no preexisting, higher-order structures in the model. Simulations of prose understanding show that the model can explain how we use background knowledge to identify the referents of pronouns.

SAM is a model of the recall and recognition that occurs following the learning of a list of random words. It is a multistage model. The first stage in recall involves cues sampling (contacting and activating) material in LTS. The items stored in LTS are called images. Sampling is achieved on the basis of a ratio rule, in which the denominator involves the activation of all material in long-term store that responds to the cues. This means that nontarget as well as target content plays a role in determining retrieval. Models that posit this effect are known as *global memory models.* Once an item has been sampled, information is recovered from it, and an evaluation process then determines whether the image will be retrieved.

Recognition involves a direct-access process in which the relevant context cue and the test item operate together to provide activation of content in LTS. This is achieved on a global basis. If the resulting activation operates at or above a criterion level, the target item is recognized. Recognition appears to show searchlike properties because the decision stage is based on global contact. This means that nontarget items can influence the decision process, an outcome normally associated with searches.

MINERVA posits that all memory content consists of episodic traces. According to this view, distinct episodic and semantic stores do not exist. Nor do unified semantic structures. For instance, if a word item is rehearsed, a "conceptual" structure of this kind is not strengthened. Instead, multiple separate traces representing the item are established. This has the effect of strengthening recall and recognition, because cues contact the total set of traces, each of which adds to the activation present in the retrieved "echo." MINERVA has also shown that effects believed to involve the abstraction of a single visual shape can be simulated on the basis of cues contacting multiple, similar shapes in secondary memory.