The Evolution of Semantic Memory and Spreading Activation

J. S. Markovitch

11-23-1981

Abstract

The purpose of this paper is to demonstrate that it is possible to deduce the structure of human semantic memory by mathematically analyzing the environment which through evolution has shaped it. The theory arrived at is similar to the spreading-activation theories of Quillian, and Collins and Loftus, but it contrasts with the above in that it involves a rigidly restricted activation that employs two distinct types of linking and three distinct types of intersection search. These three types of intersection are then used to explain the facilitation of lexical decisions, the nature of polysemy, riddles, several production experiments by Loftus, as well as the effect of word order on meaning and paired-associate learning.

The Evolution of Semantic Memory and Spreading Activation

The purpose of this paper is to demonstrate that it is possible to deduce the structure of memory by mathematically analyzing the environment in which it evolved.

The theory arrived at employs three very restricted forms of intersection search, which in turn are used to explain various facilitative effects (e.g. Brown & Block, 1980; Meyer & Schvaneveldt, 1971; Tulving & Gold, 1963), some production experiments (Freedman & Loftus, 1971; Loftus 1973; Loftus & Loftus, 1974), the nature of polsemy (Anderson & Ortony, 1975), riddles, and the effects of word order on meaning and paired-associate learning. The theory is similar in outlook to the spreading-activation and intersection theories of Collins and Loftus (1975), and Quillian (1967, 1968, 1969), but it contrasts with the above in that the intersections employed are of three distinct types (deriving from the use of just two types of links), none of which are ever-widening searches, and all of which operate under restrictions imposed by evolutionary considerations.

For an interesting review of the subject in general from the point of view of two authors who argue persuasively that there is a need to integrate the information-processing approach with restrictions derived from evolutionary principles, and who present some valuable preliminary ideas as to what these principles should be, consult Lachman and Lachman (1979).

Methodological Outlook

The information that humans must perceive, store, and retrieve does not occur randomly. Instead it is patterned, and it is these patterns that offer memory the opportunity to make "sense" of the world. As a concrete example, consider the two dimensional picture that the external world imprints on the retina. Taken as raw data it is no more than a collection of textures, lines, and areas of shading; but add to this two dimensional picture a set of assumptions about how the world is constructed (e.g. Waltz, 1975; Horn, 1975) and that two dimensional picture can with reasonable accuracy and consistency be transformed into a three dimensional view of the world, replete with the objects, edges, shadows, and cracks that constitute perceptual understanding.

Strictly speaking, deducing what these assumptions are is not psychology--it is the study of those properties and constraints existing in the physical world that are available for exploitation by a perceptual machine. But since the human mind is a perceptual machine designed by evolution, such research must inevitably be relevant to the study of perception.

This paper will take the method of Waltz and Horn one step further, by applying it to the evolution of cognitive organization in general. If, as their evidence suggests, the environment has shaped perception, then why not cognitive organization and memory as well? If human memory is a cognitive machine designed by evolution to survive in a particular kind of physical environment, then a study of the properties and constraints available for exploitation in that environment should reveal something about the information-processing strategies used by memory.

The possibility of deducing the strategies employed by memory without experimentation was stressed by Simon (1975):

> Discovery of what subjects learn can be approached experimentally, but important preliminary insights can be gained by analyzing the structure of the task itself to determine the possible alternative ways of performing it....Different subjects may in fact learn different things in the same task environment, and a formal analysis of the environment and help define the range of possibilities. (p. 268)

Although Simon's use of the word environment is not specifically a reference to an evolutionary environment, there is no reason why it cannot be taken as such, as the environment in which man evolved is the ultimate "task environment." In what follows it will be demonstrated that some of the most fundamental information-processing techniques employed by memory can be deduced using the above method.

The Statistical Information Available to Memory

Two Incomplete Definitions

It is not enough that memory links together what it perceives; it must also organize its perceptions so as to reflect the physical nature of what it is seeing--the particular objects and events--or otherwise the information it gathers will be useless.

The raw data it gathers are perceptions of such things as roundness and squareness, redness and greenness, smoothness and unevenness, darkness and brightness, which in themselves mean nothing, and which are only of value because they occur in patterns that reflect the objects and events of the physical world. Accordingly, before memory can begin the task of anticipation, it must disambiguate the stimulus it is perceiving; it must, in other words, make a guess at the physical event that it is being confronted with (disambiguation), and only then form a guess as to which physical event is going to occur next (anticipation).

This process necessarily requires that memory take the colors, shapes, and textures it perceives and combine them into single wholes that have a one-to-one correspondence with the physical objects and events of the world. These wholes must in turn be linked with each other according to the dictates of experience, thereby building in memory an internal representation of the external world that allows it to make accurate guesses about what object or event is most likely to occur next.

All of this defines an important rule regarding the organization of memory: Memory must have the ability to link together color, shape, and texture in such a way as to form a new whole capable of forming links distinct from the links of its parts.

The simplest way to see what this means in practice is to imagine a world in which there are insects that are either red or green, and either long or short. Now if it conveniently turned out that all those insects which were red were edible, and all those green were not, then an animal trying to decide which to eat would be able to make up its mind knowing color alone. Shape would be irrelevant.

But what if it turned out that those which were red and long were edible, while those which were red and short were not, and in addition that those which were green and short were edible, while those which were green and long were not? Obviously, consulting shape or color alone would tell memory nothing about edibility. Instead memory would have to make use of the procedure outlined above. Memory would have to combine color with shape in order to establish new entities capable of forming links distinct from the links of their parts. Thus, four new entities or wholes would be established in memory: "red-long" and "green-short" (each linked to the concept edible), and "red-short" and "green-long" (each linked to the concept inedible).

This procedure is, of course, exactly how memory operates. Thus the strongest associations that emanate from the concept "cross" are religious associations. Once it is joined with the concept "red" forming "red cross," however, new associations emerge, in this case medical associations, while at the same time the former religious associations disappear. In other words, when "red" and "cross" are combined into the single entity "red cross," a new whole emerges, and the links that this new entity forms over time can be expected to be distinct from the links of its parts.

In what follows the type of link that "red" has with "red cross" (that is to say the concept "red" with the image "red cross") will be referred to as a part-whole link, since "red" tends to yield not "cross," but rather "red cross," an entity memory treats as a single whole capable of forming links distinct from the links of its parts.

In contrast, the link that exists between the concept "red cross" and a concept such as, say, "disaster" will be referred to as a whole-to-whole link, since "red cross" and "disaster," though linked, are not merged by memory into a single whole having links distinct from the links of its parts, but rather remain separate.

The following are some examples of part-whole links: "red" with "fire engine"; "tube" with "straw"; "car" with "Cadillac"; "tree" with "oak"; "powder" with "cocaine"; "sharp" with "spear"; and "liquid" with "water," "alcohol," "mercury," etc.

The following are some examples of whole-to-whole links: "Africa" with "zebra"; "Russia" with "vodka"; "Italy" with "pizza"; and "car" with "oil," "gasoline," "grease," etc.

An examination of the above examples clearly reveals how part-whole links may in practice be distinguished from whole-to-whole links. Thus "red" is contained within "fire engine" (same color); "tube" is contained within "straw" (same shape; i.e. "straw" is merely a particularized version of "tube"); cocaine is clearly powdery (same texture), etc.

In contrast the concept "zebra" does not contain with it "Africa," nor vice versa (i.e. they are not the same shape, color, or texture); nor "vodka" within it "Russia" (again no sameness of shape, color or texture); nor "pizza" within it "Italy"; nor "oil" within it "car," etc.

One of the best ways to see the difference between part-whole and whole-to-whole linking is to consider the words horse, shoe, and horseshoe. If horse yields horseshoe it must do so by a part-whole link; if horse yields shoe is must do so by a whole-to-whole link. Accordingly, if a subject responds shoe to the stimulus horse, there is no way to tell if he did so by a part-whole or whole-to-whole link without first knowing whether horse has brought about in this subject's memory the word shoe or horseshoe. If the subject were to follow up with the additional response blacksmith (or laces) this uncertainty would be dispelled.

For more on how to tell a part-whole link from a whole-to-whole link see the section on link identification.

Note that in many instances one cannot identify the type of a link by simple inspection. (Eg. if "deception" yields the concept "placebo," does it do so by a part-whole or whole-to-whole link?) This important issue of link identification will receive further treatment in the course of this paper.

Note also that the concepts of part-whole and whole-to-whole linking are similar to the concepts of chunking (Miller, 1956), vertical versus horizontal linking (Anderson & Bower, 1973; Estes, 1972; Johnson, 1972; Wickelgren, 1976a, 1976b, 1977, pp. 18-27, 243-247, 251), configuring (Razran, 1971), and unitization (Hayes-Roth, 1977; Raaijmakers & Shiffrin, 1981). Actually the concepts of part-whole and whole-to-whole linking will not be the basic building blocks of the theory offered here. After all, if a concept A yields B, what is really accomplished by simply attributing that yielding to one type of link or another? Instead this paper will go a step further by analyzing how these links would interact with one another if human semantic memory evolved so as to maximize its ability to anticipate and disambiguate (this type of memory will henceforth be referred to as an optimized memory). In the course of doing this the properties of partwhole and whole-to-whole links will be further defined. The above definitions are, therefore, incomplete.

Using Two Sources of Statistical Information

Once one accepts the idea that memory makes use of more than one type of link, than the question inevitably arises as to what happens when more than one set of links is activated at a time.

Since two types of links are used here, there are three types of situations that may arise. Two sets of part-whole links may be activated at the same time (this will be termed a Type I situation); or a set of whole-to-whole links may be accompanied by a set of part-whole links (this will be termed a Type II situation); or two sets of whole-to-whole links may be activated at the same time (this will be termed a Type III situation).

An example of a Type I situation would be if memory ambiguously perceived something that was bright and yellow. Obviously this would be a case where one set of part-whole links emanated toward those concepts containing "brightness" (e.g. "light bulb," "flashlight," "sun," etc.), while another set of part-whole links emanated toward every concept containing "yellowness" (e.g. "butter," "banana," "sun," etc.), where an optimized memory could be expected to disambiguate "bright yellow" as some concept common to both lists--thus "sun" might be yielded. To be more specific, all other things being equal, memory would disambiguate "bright yellow" as that concept containing "brightness" and "yellowness" that it has seen most often.

An example of a Type II situation would be if memory perceived something ambiguous, after having earlier acquired some pertinent whole-to-whole information. For instance, if memory perceived a hospital and later ambiguously perceived a knife, it could use the whole-to-whole links of "hospital" (which tend to yield "doctor," "nurse," "scalpel," etc.) in combination with the part-whole links of "knife" (which tend to yield "pocketknife," "kitchen knife," "scalpel," etc.) to arrive at "scalpel" as a reasonable disambiguation of "knife."

In what follows it will be argued that in such a Type II situation memory's best strategy is to disambiguate "knife" as that form of "knife" most strongly linked to "hospital."

An example of a Type III situation would be if memory unambiguously perceived two concepts, A and B, where emanating from each was a set of whole-to-whole links. For instance, if "Florida" (whole-to-whole links with "sunshine," "beaches," "orange juice," etc.) occurred successively with "eggs" (whole-to-whole links with "toast," "bacon," "orange juice," etc.), memory could use the information offered by each profile to form a single composite profile of what is likely. Such a profile might (or might not) indicate "orange juice" as the most probable. "Orange juice," as the only concept receiving statistical support from both profiles, is by definition the <u>intersected</u> concept, and whether or not it would be evaluated and yielded as most likely would depend on the particular probabilities involved.

In the above Type I situation "sun" is the intersected concept; in the above Type II situation it is "scalpel."

There are a few technical points that must be noted about the above situations.

First, there are two ways memory may treat a Type I situation such as "bright yellow." It could, on the one hand, simply temporarily establish in memory "brightness" and "yellowness" without yielding anything in particular; or it could, on the other hand, yield by intersection the particular concept "sun," a concept available for retrieval from long-term memory. Similarly, given a Type I situation such as "black man," memory might simply establish in memory an image of a black man; or it could instead retrieve by intersection an image of a particular black man, say Muhammad Ali. These two processes are distinct in that the first involves merely putting in memory the information described (this will be termed construction), whereas the second goes a step further by retrieving from memory a particularized version of this description (this is intersection). In this paper all Type I situations will be assumed to involve an available intersected concept, as will all Type II and Type III situations.

Note that the various types of intersections may in many instances be recognized by simple inspection. Thus the intersection of "sun" by "bright yellow" is clearly a Type I intersection as "bright" and "yellow" are clearly contained within the concept "sun"; and in the same way "black man" clearly yields the concept "Muhammad Ali" by a Type I intersection, inasmuch as "Muhammad Ali" is merely a particularized version of "black man." Some other clearcut Type I intersections are "white powder" ("cocaine"), "green gas" ("chlorine"), and "red bird" ("cardinal").

Type II intersections such as "hospital blade" ("scalpel"), "Africa striped" ("zebra"), "hospital vehicle" ("ambulance"), "car liquid" ("gasoline" or "oil"), and "French drink" ("wine") are identifiable by the fact that the two parts of the intersection provide different types of information to memory. Thus it is the job of the concept "blade" to restrict what may be yielded to a concept from the class of blades; "hospital" then selects from this class a particular form of "blade," perhaps "scalpel." Thus, unlike in a Type I intersection, the parts of a Type II intersection perform in different roles.

The Type III intersection, in contrast to all of the above, involves two parts neither of which restricts. In other words, in yielding "orange juice," neither "Florida" nor "eggs" is contained within "orange juice." The "Florida eggs" intersection of "eggs" is interesting in that it involves two concepts intersecting a third concept. Good examples of this type of intersection are difficult to find as two concepts usually either single out a large number of intersected concepts, or no intersected concepts. Much easier to find are Type III intersections in which two concepts intersect an ambiguous word, as, for instance, the concepts "computer" and "horse" intersect the word bit. In the same way, "smoke" and "water" intersect pipe; "plaster" and "actors" intersect cast; "eye" and "student" intersect pupil; "comma" and "stomach" intersect colon; "month" and "army" intersect March; and "gas" and "army" intersect tank.

Note that there are three ways to identify the type of an intersection. The first is by simple inspection, as described above. The second is by examining what in memory has occurred contiguous with what, and then deducing from this what types of intersections are possible (this technique is treated in the section on the requirement of contiguous occurrence). And the third is by backward inference (this is where one observes how successful memory is at various intersection tasks, and then infers from the results what types of links must have been involved; this technique is treated in the section Link Identification). Each of these techniques complements the others.

Intersection in Theory

Inasmuch as the various types of intersections have been introduced, it is now possible to proceed with an analysis of how an optimized memory can best treat each one. It will in particular be demonstrated that a memory optimized by evolution for accurate anticipation and disambiguation can be expected to treat Type III intersections very differently from Type I and Type II intersections. Attention will in particular be focused on the mathematics of the Type II and Type III intersections, inasmuch as the best strategy for memory to use in treating the Type I intersection is fairly obvious (i.e. memory should always yield that concept that contains the specified features, and that been observed most often--hence "bright yellow" is disambiguated as "sun").

Inasmuch as both Type II and Type III intersections use whole-to-whole links, it is necessary to begin by examining how a set of whole-to-whole links may be used to store information on what is likely.

Storing Statistical Information

If one assumes that memory evolved so as to maximize its ability to anticipate, then a concept such as "eggs" should tend to yield that concept which has in the past occurred most often with it, all other factors being equal. Accordingly, if the concept "eggs" were to occur, it might yield "toast," or perhaps "bacon" or "orange juice," as being most likely to occur next.

But the whole-to-whole links emanating from "eggs" should be able to contain more information than merely what, based on past experience, is most likely to follow "eggs"; by varying their strengths, they should also be able to reflect relative likelihoods. Thus if "eggs" has a probability of .0001 (i.e. "eggs" occurs once in every 10,000 percepts), then during the course of a lifetime it should occur several thousands of times, with each of these occurrences constituting a sampling of what is likely given the presence of "eggs"; that is to say, "eggs" might develop strong links with "toast," bacon," and "orange juice," while developing no links at all with the vast majority of concepts, and in this way the link strengths emanating from "eggs" could come to constitute a statistical profile on what is likely to follow the perception of "eggs." (Technically speaking, such profiles are samples taken from the environment under biased conditions, where the "bias" takes the form of the conditional occurrence of the percept.)

Note that in what follows it will be assumed that the sum of all of the whole-to-whole link strengths emanating from each concept equals one. The purpose of this rather artificial assumption is to simplify matters by seeing to it that a particular link strength corresponds to a particular probability regardless of which profile it is taken from.

Using the above statistical profiles as a base, it is now possible to analyze Type II and Type III intersections.

An Analysis of a Type III Intersection

Assume that the concept "eggs" has a probability of .0001 and that given its occurrence the probability of "toast" being seen next is .5, "bacon" .25, and "orange juice" .04. These probabilities may be conveniently expressed as follows:

"eggs" ⟶ "toast" (.5)
⟶ "bacon" (.25)
⟶ "orange juice" (.04)

And assume as well that the concept "Florida" occurs with a probability of .0001, and that its profile is:

"Florida" ⟶ "sunshine" (.2)
⟶ "beaches" (.04)
⟶ "orange juice" (.01)

Now it is clear from the "eggs" profile that if an optimized memory perceived only "eggs," "toast" would be yielded as being most likely to occur next, as its probability (.5) is substantially higher than any of the other probabilities.

But if the concept "eggs" occurred, having been immediately preceded by the concept "Florida," then memory should not merely give great weight to the statistical information offered by "eggs," but should also give some weight to the information offered by "Florida," on grounds that as a concept that occurred almost as recently as "eggs," the information it provides should be of considerable value;

having two sources of information on what is likely, an optimized memory is obligated to use both in order to form the best possible view of what might occur next.

If one assumes that memory treats each set of probabilities as of equal value, then to form a single profile telling what is likely given the occurrence of both "Florida" and "eggs" requires the summing of corresponding probabilities and dividing by two. The only important difference between such a composite profile and each profile taken separately is that, relative to the other concepts, "orange juice" moves up somewhat. In the "eggs" profile "toast" is 12.5 times as likely as "orange juice" (.5 versus .04), while in the composite profile "toast" is only 10 times as likely (.25 versus .025). This upward shift occurs because "orange juice" is the only concept on both lists.

It is important to notice how small this upward shift is, in particular that it is nowhere near enough to allow "orange juice" to go to the head of the list. The reason for this can be seen clearly by noting that the strongest link of all those emanating from "Florida" and "eggs" is the one running from "eggs" to "toast." Accordingly, "orange juice's" two sources of support in combination would have to be greater in value than "toast's" one source of support if "orange juice" is to be yielded as being most likely.

"Florida" ——————→ "toast" (zero probability)

——————→ "orange juice" (.01)

"eggs" ——————→ "toast" (.5)

——————→ "orange juice" (.04)

A look at the above probabilities yields the following simple calculation expressing the ratio of the likelihood of "toast" to the likelihood of "orange juice" if both profiles are treated as being equally important:

$$\frac{P(\text{toast}|\text{Florida}) + P(\text{toast}|\text{eggs})}{P(\text{orange juice}|\text{Florida}) + P(\text{orange juice}|\text{eggs})} = \frac{.0 + .5}{.01 + .04} = \frac{10}{1}.$$

This calculation shows clearly and simply that "orange juice," though receiving support from both profiles, is still only one-tenth as likely as "toast" and therefore would not be evaluated by memory as being most likely.

This point is important because "orange juice," as the only concept common to each list, is the intersection of all those concepts having whole-to-whole links with "Florida," with all those concepts similarly linked to "eggs." The calculations shown here demonstrate that, given the above probabilities, an optimized memory which treated "Florida" and "eggs" as two sources of statistical information would find that the phrase Florida eggs suggested to it "toast," rather than the intersected concept "orange juice." In other words, if such a memory were asked the cognitively simple question What is associated with both Florida and eggs?, it would find "Florida" and "eggs" suggesting to it "toast" rather than "orange juice," which is to say that the first concept that it thought of would not be the answer, and memory would therefore be forced to employ some sort of trial-and-error search strategy to arrive at it. This key point will be explored later in detail.

Note that if memory treated the "eggs" profile as more important than the "Florida" profile, which is less recent, the above difficulties would increase rather than decrease.

<u>An analysis of a Type II Intersection</u>

Now it is possible to demonstrate that plugging the same values used above into the corresponding Type II intersection completely alters the expectations of an optimized memory, and turns the task of intersection into one easy for memory to carry out.

The corresponding Type II intersection involving the concepts "hospital" and "blade" requires that the concept "hospital" occur with a probability of .0001 (i.e. once in every 10,000 percepts as was the case with "Florida"), and that "blade" in one form or another (e.g. "pocketknife," "breadknife," or "scalpel," etc.) likewise occur with a probability of .0001 (the same as "eggs"). It is also necessary to assume that over the course of a lifetime the following profile is established for "hospital":

"hospital" ⟶ "breadknife" (zero probability)

⟶ "scalpel" (.01)

And that the following profile emanates from the concept "blade":

"blade" ⟶ "breadknife" (.5)

⟶ "scalpel" (.04)

Note that these two profiles incorporate probabilities that are identical with those used with "Florida" and "eggs." The important difference introduced there is that the part-whole links of "blade" have replaced the whole-to-whole links of "eggs," a small modification that makes necessary a new mathematical analysis.

It is clear from the above "blade" profile that if memory perceived only the presence of some sort of "blade," but did not know what kind it was, "breadknife" would, by a factor of 12.5, be a better guess than "scalpel." The key question is, however, what memory would do in disambiguating "blade" if the concept "hospital," with all of the statistical information it contains, were present as well?

To answer this, consider that "blade" indicates that "breadknife" (probability .5) is 12.5 times as likely as "scalpel" (probability .04). Accordingly, if the addition of "hospital" increased the probability of "scalpel" by a factor greater than 12.5, then "scalpel" would become more likely than "breadknife."

The probability of "scalpel" given the presence of neither "hospital" nor "blade" is P(blade) x P(scalpel|blade), which is (.0001)(.04), or .000004. According to the "hospital" profile the probability of "scalpel" jumps to .01 with the occurrence of "hospital," a 2500-fold increase. It follows that the 12.5:1 edge that "breadknife" has over "scalpel" in the absence of "hospital" becomes a 1:200 edge in <u>favor</u> of "scalpel" when "hospital" occurs.

All of this can be explained in a way that is perhaps somewhat easier to understand. "Breadknife" occurs with a probability of P(blade) x P(breadknife|blade), which is (.0001)(.5), or .00005; "scalpel" occurs with a probability of P(blade) x P(scalpel|blade), which is (.0001)(.04), or .000004. This means that given the conditional occurrence of neither "hospital" nor "blade," "breadknife" has a .00005 to .000004 edge over "scalpel," which is 12.5:1.

Now if "hospital" occurs the probability of "scalpel" increases from .000004 to .01, as is indicated by the "hospital" profile, leading to the new ratio of .00005 to .01, which is a 1:200 edge in favor of "scalpel." This means that if "blade" were ambiguously perceived, "scalpel" would be more likely than "breadknife" by this 200-fold factor. Actually the edge would be even greater since the zero probability that "hospital" ascribes to "breadknife" indicates that its probability is probably below .00005. In other words the addition of "hospital" would dramatically alter the probabilities so as to make the intersected concept, "scalpel," the preferred disambiguation of memory.

Now note that this is in sharp contrast to what happened earlier when these same values were plugged into the equivalent Type III intersection. The most striking fact to emerge from that analysis was that the presence or absence of "Florida" as a concept preceding "eggs" had only the slightest effect on the probability of the intersected concept "orange juice"; the probability of "orange juice" rose relative to the other concepts, but only by the slightest amount.

From all of this a key principle emerges: A memory optimized for prediction should find Type III intersection tasks (e.g. name something associated with both "Florida" and "eggs") more difficult to carry out than Type II intersection tasks (e.g. name a form of "blade" associated with "hospital"), even though the only difference between the two tasks is that Type III intersection involves intersecting two sets of whole-to-whole links, whereas Type II intersection involves intersecting a set of whole-to-whole links with a set of part-whole links. More will be said on this issue later.

Memory's Need for Two Types of Links

It has up to now been assumed that evolution would force memory to use two types of links in order to maximize its ability to anticipate and disambiguate. What follows is a mathematical analysis of what would happen to a memory that used only one type of link and was therefore forced to treat Type II intersections and Type III intersections in the same way.

Assume that the concept X occurs with an unknown probability $P(X)$, that Y occurs with a probability $P(Y)$, and that the profiles emanating from these concepts are in part as follows:

$$
\begin{array}{ll}
X & \longrightarrow N \\
 & \longrightarrow I \\
Y & \longrightarrow N \\
 & \longrightarrow I
\end{array}
$$

Assume also that $P(I|X)$ and $P(I|Y)$ are non-zero probabilities (thereby making I, by definition, an *intersected* concept), that $P(N|X)$ equals zero (thereby assuring that N, by definition, is a non-intersected concept), that P(N|Y) is the highest probability in the Y profile (thereby assuring that N is the concept most strongly linked to the concept Y), and that P(X) is sufficiently high to allow the X profile to contain accurate probabilities (i.e. X has occurred sufficiently often for its whole-to-whole links to constitute an accurate profile of what is likely given the occurrence of X).

Now if X is the concept "Florida," Y the concept "eggs," N the concept "toast," and I the concept "orange juice," then calculating how likely toast would be in comparison to "orange juice" given the occurrence of both "Florida" and "eggs" would involve the equation:

$$\frac{P(\text{toast}|\text{Florida eggs})}{P(\text{orange juice}|\text{Florida eggs})} = \frac{P(N|X) + P(N|Y)}{P(I|X) + P(I|Y)} = \frac{P(N|Y)}{P(I|X) + P(I|Y)},$$

which leads to the definition

$$\underline{\text{Type III Ratio}} = \frac{P(\text{non-intersected concept})}{P(\text{intersected concept})} = \frac{P(N|Y)}{P(I|X) + P(I|Y)}.$$

And, as regards a Type II situation, if X is the concept "hospital," Y the concept "blade," N the concept "breadknife," and I the concept "scalpel," then calculating how likely "breadknife" would be in comparison to "scalpel" given the occurrence of both "breadknife" and "blade" would involve the equation:

$$\frac{P(\text{breadknife}|\text{hospital blade})}{P(\text{scalpel}|\text{hospital blade})} = \frac{P(Y)P(N|Y)}{P(I|X)},$$

which leads to the definition

$$\underline{\text{Type II Ratio}} = \frac{P(\text{non-intersected concept})}{P(\text{intersected concept})} = \frac{P(Y)P(N|Y)}{P(I|X)}.$$

Now earlier it was noted that shifting from a Type II intersection to a Type III intersection, or vice versa, altered memory's expectations by a large factor. This factor will be termed the shift ratio (S.R.) and will be defined as the Type III ratio divided by the Type II ratio:

$$\underline{\text{S.R.}} = \frac{P(N|Y)/[P(I|X) + P(I|Y)]}{P(Y)P(N|Y)/P(I|X)} = \frac{P(I|X)}{P(Y)[P(I|X) + P(I|Y)]}.$$

If one then assumes that $P(I|X)$ and $P(I|Y)$ are approximately equal, which roughly they would tend to be, then the above equation yields the formula:

$$\underline{\text{S.R.}} = \frac{1}{2P(Y)}.$$

Since human experience encompasses tens of thousands of recognizable percepts, it would be common for $\underline{P(Y)}$ to be less than .0001, which means that shift ratios of 5000 or more would be common. This means that a memory that treated Type II and Type III intersections alike would frequently incorrectly estimate probabilities by a factor as high as 5000, and that still larger errors would be inevitable as well.

Needless to say, such a poorly designed memory would have a hard time competing for survival with well-designed memories that took advantage of part-whole and whole-to-whole linking in order to treat Type II and Type III intersections properly. It follows that one is justified in expecting that it is likely memory treats Type II and Type III intersections differently, inasmuch as such an expectation is supported by a mathematical analysis of the physical world in which memory evolved, an analysis that also supports the more particular conclusion that memory stores its information in part-whole and whole-to-whole links.

Invalidation and Residual Activation

Actually there is another reason to expect that memory stores its information in part-whole and whole-to-whole links. Consider again the concepts "hospital" and "blade." It is a fact that the whole-to-whole links emanating from "hospital" provide information that retains its usefulness for some time--that is, once one has perceived "hospital," those concepts associated with it ("doctor," "nurse," "scalpel," "beds," etc.) have a higher probability for some time afterward, probably for many minutes. But the same cannot be said for the information offered by the part-whole links of "blade." Thus if "blade" is perceived and successfully disambiguated as, say, "scalpel," that successful disambiguation would invalidate "kitchen Knife," "pocketknife," "butter knife," etc. as concepts to be regarded as more likely than usual. In contrast, none of the whole-to-whole links emanating from "hospital" would in any way be invalidated by the occurrence of, say, "doctor"; in fact, the concepts "nurse," "scalpel," and "beds" would, if anything, be more likely by virtue of the occurrence of "doctor" supporting memory's original expectations.

Put another way, if one saw a hospital and expected to see doctors, nurses, scalpels, and beds, etc., one would regard this class of expectations as even more valid if one of them were to occur.

In contrast, if one saw a blade and expected it to be either a scalpel, kitchen knife, or pocketknife, etc., one would regard this class of expectations as invalidated if one of them were confirmed as the blade seen. If a blade were ultimately correctly disambiguated as a scalpel, one would not then go on attributing higher probabilities to kitchen knives and pocketknives--they would no longer have higher than normal probabilities.

Accordingly, if a memory optimized for anticipation and disambiguation perceived "hospital" and "blade" in immediate succession, ideally it would use the whole-to-whole links emanating from "hospital" to disambiguate "blade" and would then for some time afterward attribute elevated probabilities to all those concepts having whole-to-whole links with "hospital"; but such a memory ideally would not go on to attribute higher than normal probabilities to the various forms of "blade"--instead the part-whole links emanating from "blade" should be made to cease to exert influence as soon as "blade" itself had ceased to be immediately present in memory.

As it turns out, what is true for the part-whole links of "blade" and the whole-to-whole links of "hospital" is also true for part-whole links and whole-to-whole links in general. Part-whole links, such as those that colors, shapes, and textures have with the concepts they are a part of, fairly consistently supply information that is invalidated by accurate disambiguation. Whole-to-whole links, in contrast, in general supply information that is not invalidated by accurate disambiguation or anticipation. Accordingly, in an optimized memory a set of whole-to-whole links, once activated, should remain activated for some time afterward, with a resultant heightened expectation being established in memory for all those concepts singled out by those whole-to-whole links. Part-whole links, in contrast, should not give rise to any form of persisting activation, but rather should instead merely act by restricting what may be yielded in memory. Hence if an ambiguously seen red shape (part-whole links with "fire engine," "apple," etc.) were perceived by memory, memory's disambiguation of that shape should be restricted by "redness," in the sense that memory should limit its range of possible disambiguations to the class of concepts containing "redness," but no persisting activations should be given rise to.

The principle that whole-to-whole links should give rise to persisting activations, but not part-whole links, will be termed the principle of residual activation.

It is worth noting that in the logogen model of Morton (1969) a stimulus does not exert any influence after it has ceased to be immediately present in memory. Inasmuch as a stimulus inevitably has part-whole links with the various percepts it tends to give rise to, Morton's restriction is a specialized version of the more general restriction introduced here, namely that no part-whole link exerts a persisting influence. Of course in Morton's theory there is a "Context System" that does exert a persisting influence.

Type IV Interactions

It should be noted that memory has the ability to comfortably handle familiar phrases as if they were single words. Thus if horseshoe is rewritten horse shoe, it harkly matters to memory--the two words may still be treated as a single word. And in the same way, a phrase such as Military Police need not yield its concept by an intersection of the concepts "military" and "police," inasmuch as Military Police may be treated as one word. And lastly, once someone is told that Florida eggs intersects "orange juice," that person will then have a link running from the phrase Florida eggs taken as a single whole to "orange juice." Accordingly, for that person the phrase Florida eggs would readily yield "orange juice," but not by Type III intersection; instead it would do it by the whole-to-whole link that runs directly between the two.

Whenever two or more words act as a single whole in yielding a concept a Type IV interaction will be said to be at work. Naturally any familiar phrase sets up a potential for Type IV processing, as any group of words, if it is short enough and familiar enough, may be treated as a single word.

The Requirement of Contiguous Occurrence

It is interesting to note that for A and B to yield a concept Z by either a Type I or Type II intersection, or by a Type IV interaction, it is necessary for A, B, and Z on some prior occasion to have occurred all three together in memory.

In other words, if in a particular memory A, B, and Z have not occurred together, then there is no way there can exist in memory an entity Z that contains--i.e. has part-whole links with--both A and B, and so a Type I intersection of Z is an impossibility (to see this example clearly, imagine A as "yellowness," B as "brightness," and Z as "sun"; for "yellowness" (A) and "brightness" (B) to be established in memory as part of the particular image "sun" (Z) requires that at some time all three occur together).

And similarly, if A, B, and Z have not occurred together, then there is no way A can have formed a whole-to-whole link with BZ, nor B a whole-to-whole link with AZ, and so a Type II intersection of Z is an impossibility (to see this example clearly, imagine A as "hospital," B as "blade," and BZ as the particular form of "blade" "scalpel"; for a whole-to-whole link to be established between "hospital" (A) and "scalpel" (B and Z) requires that at some time all three occur together.)

And lastly, if A, B, and Z have not occurred together, then the whole AB cannot have formed a whole-to-whole link with Z, which means that Z may not be yielded by a Type IV interaction (to see this example clearly imagine A as the word horse, B as the word shoe, and Z as the concept "horseshoe"; for horseshoe (A and B) to form a whole-to-whole link with concept "horseshoe" (Z) requires that all three at some time occur together.)

It follows from all of the above that two concepts A and B cannot be expected to cooperate in yielding a third concept Z by a Type I or Type II intersection, or a Type IV interaction, unless at some time in the experience of the memory in question A, B, and Z have all occurred together. This requirement, which must be met if Type I, Type II, or Type IV processing is to be a possibility, will be termed the requirement of contiguous occurrence.

The obvious advantage of the above principle is that it is an aid in distinguishing Type III intersections from Type I and Type II intersections. Thus if it is known that in a particular memory "Florida" tends to yield "orange juice," and "eggs" tends to yield "orange juice" (i.e. "Florida" and "eggs" are each known to be linked either by part-whole or whole-to-whole link to "orange juice"), and it is known that "Florida," "eggs," and "orange juice" have never occurred all three together in memory, then it follows that "Florida" and "eggs" cannot yield "orange juice" by either Type I or Type II intersection. It follows that "Florida" and "eggs" must single out "orange juice" by a Type III intersection, which means that memory will quite possibly find yielding "orange juice" a difficult task, especially as a Type IV interaction is also ruled out as a possibility.

Intersection in Practice

This section will concern itself with the practical application of the theoretical ideas just explained. It is interesting to see how the concepts of part-whole and whole-to-whole linking, which were introduced for the purpose of explaining how one would expect the environment to have shaped memory through evolution, help make clear as well why psychology almost alone among the experimental sciences has been unable to advance beyond first principles. Psychology, quite simply, has been monitoring the wrong variables.

Comparing Activation Theories

It is important to realize that the principle of residual activation is not a spreading activation theory of the type employed by Collins and

Loftus (1975) or Quillian (1967, 1968, 1969), inasmuch as they require that once a concept is activated all of its surrounding concepts must become activated, and these concepts must in turn activate to some degree all of the concepts that surround them, and so on.

In contrast to the above theories, the principle of residual activation states that once a whole is activated, only the whole-to-whole links emanating from that entity residually activate other wholes, and no further activations take place until a new whole is yielded.

Thus, according to the theory of spreading activation (see in particular Collins & Loftus, 1975), if "Florida" occurs, it would follow that "beaches," "Miami," and "orange juice" would be temporarily activated in memory, *and these concepts would further activate to some* degree still more concepts: that is, the spreading activation, having reached "orange juice," would then continue to spread to "eggs," "toast," "salt," and so on.

In contrast, according to the principle of residual activation, if "Florida" occurs, it would follow that "beaches," "Miami," and "orange juice" would be residually activated, but no further activations would take place until a new concept was yielded. Accordingly, "eggs," "toast," and "salt" would join the above in residual activation only when (and if) "Florida" yielded "orange juice."

Another important difference between the theories of spreading activation and residual activation concerns the treatment of part-whole links.

According to the theory of spreading activation, if a concept such as "red" occurs, it should activate "apple" and "fire engine," etc., and these concepts should in turn activate such concepts as "supermarket" and "firehose," inasmuch as there is no rule forbidding part-whole links from initiating a spreading activation.

In contrast, residual activation requires that if a concept such as "red" occurs, it would not residually activate such concepts as "apple," and "fire engine," etc., inasmuch as part-whole links are forbidden to give rise to persisting activations of any kind.

It should be clear from the above comparisons that there is an overall similarity between the principle of residual activation and the concept of spreading activation as employed by Collins and Loftus, as well as certain clearcut differences. This overall similarity will be discussed in the next section, and attention will then be focused on the differences between the two theories.

Do Secondary Activations Occur?

It would be a tedious, and somewhat unproductive task to analyze all of the data that have accumulated supporting the conclusion that secondary activations of some kind occur. In summary it may be said that both residual activation theory and spreading activation theory argue that when a concept is subjected to a secondary activation, memory's ability to retrieve that concept should be facilitated (i.e. primed).

Such facilitative effects appear to have been observed for lexical decisions (Meyer & Schvaneveldt, 1971; Meyer, Schvaneveldt, & Ruddy, 1975; Schvaneveldt & Meyer, 1973), lexical decisions primed by ambiguities (Holley-Wilcox & Blank, 1980; Onifer & Swinney, 1981; Simpson, 1981), naming (Brown & Block, 1980; Warren, 1977), the tachistoscopic identification of words (Tulving & Gold, 1963), intersection (Loftus, 1973; Loftus & Loftus, 1974), and interference in a Stroop (1938) task (Warren, 1972, 1974). There is also an apparent facilitative effect at work that causes perceptual slips (see the examples of Celce-Murcia, 1980; Garnes & Bond, 1980).

To take just one of the above paradigms, consider the general issue of lexical decisions. According to both residual and spreading activation theories, making a lexical decision on a letter string such as NURSE (i.e. judging whether NURSE is a word) should facilitate making a subsequent decision on DOCTOR, inasmuch as the word nurse should give rise to a secondary activation of the word doctor, which should facilitate its retrieval. This effect has been clearly observed in the studies mentioned above.

Accordingly, in what follows it will be assumed that secondary activations of some kind occur, and attention will instead be focused on evaluating how they occur. In particular, attention will be paid to the issue of whether there really are two types of links, one of which residuallly activates and one of which does not, as this is a key point on which spreading activation theory and residual activation theory differ.

The Effect of Order on Type II Intersection

A production experiment was conducted by Freedman and Loftus (1971) in which subjects were presented noun categories (e.g. animal, fruit, president, etc.) followed .5 to 5 sec later (or preceded .5 to 5 sec earlier) by restricting letter, and were asked to carry out an intersection. For example, animal followed by Z (or Z followed by animal) was supposed to provoke a response such as Zebra. It was clearly demonstrated that a quicker response was possible when the noun category was presented first and the letter second (reaction times were measured from the presentation of the second item).

Freedman and Loftus interpret this result by arguing that if the noun animal occurs first, a certain portion of the intersection process may be carried out even before Z occurs, whereas if the letter occurs first, there is no corresponding activity that may be profitably undertaken by memory, and therefore memory must bide its time while waiting for the noun animal to occur. This portion of wasted time is, according to Freedman and Loftus, the reason for the difference in reaction times.

The theory of part-whole and whole-to-whole linking can account for the above asymmetrical reaction times in two ways. First, it is possible to argue that whole-to-whole links (such as the one that runs between the concept "animal" and the word Zebra) take longer ot activate than part-whole links, perhaps because unlike part-whole links they may undergo residual activation. It follows that in intersecting a noun category with a letter, memory saves more time if the noun is presented ahead of time than if the letter is, because activating the whole-to-whole links of the noun is a more time-consuming procedure than activating the part-whole links of the letter. In other words, memory saves more time by completing the activation of the whole-to-whole links ahead of time than the activation of the part-whole links, because activating the whole-to-whole links is the more time-consuming procedure.

The assumption that whole-to-whole links take longer to activate will be termed activation asymmetry.

In addition, the theory offered here argues that part-whole links (such as the one that runs between the letter Z and the word Zebra) cannot under residual activation, and that if they are activated first, memory will inevitably find that trying to keep them activated disrupts the process of acquiring new information (such as the whole-to-whole links of the relevant noun category), and consequently memory will find intersection a slower and more difficult process. This disruption will be termed activation interference.

Of course, either activation asymmetry or activation interference taken separately is sufficient to account for the results of Freedman and Loftus (1971). Because the yielding of the word Zebra by animal-Z is a Type II intersection, one can simply argue that the whole-to-whole links of the noun category should precede the part-whole links of the letter so as to avoid activation interference, or to take advantage of activation asymmetry, or both. But both activation asymmetry and activation interference are useful, and perhaps necessary, if one wishes to account for the full range of effects order may have on Type II intersection.

To see why, consider Loftus and Loftus (1974). In this experiment subjects were able to respond to simultaneously presented category-letter pairs .13 sec faster than letter-category pairs. Inasmuch as simultaneous presentation was involved, the difference in reaction times does not lend itself to explanation through activation asymmetry--it can, however, be readily attributed to activation interference. Thus one may argue that the stimulus animal-Z more quickly obtains the response Zebra than does Z-animal, because memory finds it difficult to keep the part-whole links of the letter Z activated while going on to read the word animal.

It is interesting that when the same experiment was repeated with a time interval inserted between the category and the letter, subjects intersected category-letter pairs .27 sec faster than letter-category pairs. In other words, the effect that order had on reaction time--what will be termed the order effect--approximately doubled when a time interval was inserted between the category and the letter. This increase is readily accounted for by arguing that the addition of the time interval allows activation asymmetry to have an effect along with activation interference. Accordingly, both activation interference and activation asymmetry are useful in explaining Loftus and Loftus (1974).

It should be noted, as Loftus and Loftus point out, that the relatively small effect (.13 sec) of varying order given simultaneous presentation may actually be the consequence of simultaneously presented items being read in the correct order even when they are presented in the wrong order (this would be an effective way of evading activation interference). This raises the possibility that the doubling of the order effect that occurs when an interval is inserted may be caused by this trick becoming no longer available to memory, which in turn raises doubt about the existence of activation asymmetry. In any case, however, from the point of view of psycho-evolutionary analysis the issue is somewhat academic, inasmuch as psycho-evolutionary analysis, though clearly calling for the existence of activation interference (whose existence the above experiment appears to confirm), is more or less neutral on the separate question of whether activation asymmetry must exist as well.

Before beginning the next section it is necessary to dispose of a possible alternative explanation for the results of Freedman and Loftus (1971). One could argue that a noun category followed by a restricting letter is the preferred order for yielding an intersected word because a noun must be recognized and its concept yielded, whereas a mere letter need only be recognized (i.e. the letter, unlike the noun, does not yield a concept). Thus the simple kind of activation asymmetry could be assumed to be the cause of the difference in observed reaction times.

Fortunately, however, Freedman and Loftus also tested their subject's ability to carry out intersections using phrases such as food white, seasoning white, and fuel liquid (i.e. a noun category followed by a restricting adjective), and obtained essentially the same order effect they observed earlier for a noun category with a restricting letter, although admittedly the effect was somewhat smaller.

That this order effect should manifest itself again makes sense, as the restricting adjectives used (see Freedman & Loftus, Table 1) could in general be expected to operate by part-whole links (e.g. white, yellow, liquid, green, enormous, small, long, etc.), and the noun categories used in general could be expected to operate by whole-to-whole links (e.g. the concept "food" probably has whole-to-whole links with the various types of food; and the concept "president" probably has whole-to-whole links with all of those diverse people who have been president, and so forth).

The order effect probably would have been larger if care had been taken to choose word pairs that must be treated as Type II intersections (as, for example, Africa striped intersects the concept "Zebra"). Thus some of their word pairs appear to allow Type I intersection. For instance, it is conceivable that memory treats bird followed by yellow as a Type II intersection, and yellow followed by bird as a Type I intersection, thereby evading an order effect. This is a possibility because the concept "bird" probably has whole-to-whole links with the different forms of "bird" (e.g. "sparrow," "ostrich," "hummingbird," "vulture," etc.), as well as part-whole links with all those forms of "bird" that are prototypically "bird-like" (e.g. "sparrow," "robin," and "canary," which have in common pretty much the same shape).

There is every reason to expect that memory makes liberal use of such redundancies in storing information, as there is no rule forbidding a complex concept such as "bird" from having whole-to-whole links with the different forms of "bird," as well as an additional set of part-whole links with those forms of "bird" involving the same shape and design. It would be interesting to see if the above experiment yielded a larger order effect when restaged without such ambiguous intersections.

The Optimal Word Order for a Language

From all of the above it follows that if memory wishes to carry out the Type II intersection of, say, Africa and striped, the most effective order (i.e. the order that will most easily yield an intersected concept such as "zebra") would be Africa striped, as it is natural for memory to carry out Type II intersections with whole-to-whole links activated first, and part-whole links second. In carrying out this Type II intersection memory should find the word order Africa striped easier, not because of any rule of the English language (actually the phrase Africa striped is not grammatical and can hardly be regarded as a standard way of denoting "zebra"), but because by placing the part-whole links of the concept "striped" second, activation interference is avoided.

It is important to note that not all phrases are as simple for memory to deal with as is Africa striped. With such a phrase ther is no doubt that for a Type II intersection to take place the whole-to-whole links must be provided by the concept "Africa" (or the word Africa), and the part-whole links by the concept "striped." Whatever the word order this must be so, as the concept "striped" can hardly be expected to provide the whole-to-whole links.

But some Type II phrases (i.e. phrases that give rise to Type II intersections) are congnitively ambiguous in the sense that memory has a choice not merely of word order, but also a choice as to which word will provide the whole-to-whole links and which the part-whole links.

As an example of such a phrase consider company plane. If "company" provides the whole-to-whole links, and "plane" the part-whole links, then a concept such as "Learjet" might be yielded; in contrast, if "plane" supplies the whole-to-whole links, and "company" the part-whole links, then the intersected concept would tend to be a type of company that manufactured planes (e.g. "Lockheed"). Thus the meaning of the above phrase (i.e. the concept it tends to yield) is profoundly affected by whether memory chooses to have the first word supply the whole-to-whole links and the second the part-whole links, or vice versa.

Now it is important to realize that in making the above decision memory is, in effect, deciding a point of grammar relating to word order. In English, of course, such phrases are treated according to a rule which requires that company plane suggest "Learjet," and plane company "Lockheed," which is another way of saying that the first word in such phrases is to supply the whole-to-whole links, and the second the part-whole links. In choosing this word order English is, it is very important to note, avoiding the activation interference that would arise if it decided instead to have the first word supply the part-whole links. It follows that it is not unreasonable to argue that the rule of word order adopted by English for phrases such as company plane may be taken as further evidence in favor of the existence of activation interference.

Of course, if the above reasoning is really valid, then a study of the world's diverse languages should reveal a distinct preference for this arrangement, especially among those languages that have little or no inflection, and which therefore are especially dependent on word order to convey meaning. Thus it is important that the word order used in Chinese, the world's most important uninflected language, is as predicted.

Also note that languages which, like English, employ the optimal word order should form compound words more freely and confidently than languages, such as French, which do not. This is because the non-optimal word order, which is normally a nuisance to memory, becomes a double nuisance when incorporated in a compound word, especially when it is written. Thus while pomme de terre (roughly "apple of the earth" or potato) is merely a psychological inconvenience, pommeterre (appleearth) the equivalent compound word, is a real psychological calamity for a memory attempting a Type II intersection of its component parts to obtain "potato." Memory first has to recognize that it is a compound, then locate where the second word begins, read it, then read the first word, and then carry out the intersection. Or memory can read the first word, recognize that it is a compound, and then try to keep its part-whole links activated while the second word is read and its whole-to-whole links are established, a process that inevitably involves activation interference. With the correct order, of course, a compound may simply read straight off without any special treatment (according, earthapple is an entirely convenient term for "potato"; whereas appleearth is not).

Admittedly some languages appear to form compounds quite comfortably in the reverse order, for instance the American Indian language Yana

(Sapir & Swadesh, 1960); but psycho-evolutionary analysis predicts that such languages should be exceptions to the rule. Actually, in languages such as Yana, compounds almost certainly are processed through Type IV interactions (i.e. they are treated as single entities).

Type II and Type III Intersection

The psycho-evolutionary calculations offered earlier suggest that Type II intersections should be easy for an optimized memory, but not necessarily Type III intersections.

That Type II intersections actually are easy is shown by Freedman and Loftus (1971), Loftus (1973), and Loftus and Loftus (1974), each of which found that subjects can intersect a noun category with a letter in about 2 sec. Freedman and Loftus also found that Type II intersections involving noun categories and adjectives (e.g. food white, intersecting, say, "cauliflower") also posed subjects no particular problems (intersections were carried out even faster than their noun category-letter counterparts). Thus the results of experimental research are in line with what psycho-evolutionary reasoning predicts for Type II intersections.

Unfortunately similar data are not available for Type III intersections, but it is readily apparent to those who try that solving a Type III intersection task usually requires a process of trial-and-error that is not necessary when solving a Type II intersection task. This is because in a Type II intersection task the first concept suggested to memory is usually an intersected concept (i.e. Florida drink yields "orange juice" before it yields anything else); whereas in a Type III intersection such as Florida eggs, "toast" or "sunshine" or any number of concepts might be yielded before "orange juice" is yielded (if it is yielded at all).

Thus the Type III question What is associated with Florida and eggs? (answer: "orange juice") should be answerable, but should frequently require a process of trial-and-error that would be unnecessary if the question were instead What kind of drink is associated with Florida?

The prediction that Type II intersections should prove easy for memory, and Type III intersections difficult, is a key point of psycho-evolutionary theory.

Objections to Spreading-Activation

It is interesting that while Freedman and Loftus (1971) offer no clear explanation of their finding that noun categories behave differently in memory from adjectives and letters, Collins and Loftus (1975) do attempt such an explanation. They point out that the various concepts denoted by a noun category usually are "interlinked" in a way that the various concepts or words denoted by an adjective or letter are not. Thus the different kinds of fruit, for instance, usually occur together and so become interlinked. Collins and Loftus then argue that this makes a spreading-activation more effective among the members of the "fruit" category than it is among the members of the class of "red" concepts, and that it is this time-consuming spreading-activation, which may begin only after the concept "fruit" is activated in memory, that is responsible for the difference in reaction times. And of course the same lack of interlinking that prevents spreading-activation from working effectively among those concepts containing "redness" can also be expected to hamper spreading-activation among the class of all words beginning with the same letter (see Collins & Loftus, p. 416).

The above explanation is clearly a possible way of explaining why a noun category followed by a restricting letter or adjective is a more convenient order for memory than the same items in reverse (basically they are postulating a type of activation asymmetry). It is less effective in explaining why the order effect should also be present for simultaneously presented items (Loftus & Loftus, 1974), although admittedly they could argue that even simultaneously presented items are read successively, and that therefore it is still somewhat better to have the noun category first so as to allow its time-consuming spreading-activation a small head start it would lack if it were placed second. Whether this very small head start can be regarded as sufficient explanation of the .13 sec order effect observed for simultaneously presented items is debatable however.

A more direct and compelling objection to the Collins and Loftus outlook is that by not using two types of links they cannot explain how the "Florida eggs" intersection of "orange juice" differs structurally from the "Florida drink" intersection of "orange juice." To do so they would have to draw a distinction between the relationship "eggs" has with "orange juice" and the relationship that "drink" has with "orange juice." Thus they might say that "eggs" operates by spreading activation whereas "drink" tends to yield "orange juice" by a kind of restriction; but if one is going to do this then why not go all the way and postulate the existence of two types of links? After all, to introduce a type of restriction for use with spreading-activation is only an implicit way of saying what should be said explicitly, namely that the most parsimonious explanation of the data involves using two types of links. This would allow the abandonment of their "interlinking" explanation of the order effect, an explanation which required that there be spreading-activations emanating from concepts such as "red," but that these spreading-activations be inefficient because of a lack of interlinking, an argument that was somewhat clumsy to begin with.

It is a clumsy argument in that it requires that the concept "striped" in the intersection "African striped" activate not merely "striped" concepts such as "zebra" and "barber pole," etc., but also indirectly such concepts as "haircut" and "scissors" (both linked to "barber pole"). To assume such a spreading-activation is not only counter-productive in that it implants in memory expectations that are gratuitous and misleading; it is also unnecessary. A simple restriction of some kind would do as well without creating all of the unjustified expectations.

It would be interesting to see if a letter string such as STRIPED can facilitate lexical decisions on letter strings such as HAIRCUT and SCISSORS (Meyer & Schvaneveldt, 1971). If an extended spreading-activation emanates from "striped," it should. Psycho-evolutionary reasoning, of course, predicts no such facilitation.

It should be noted that if Collins and Loftus (1975) is interpreted as employing a spreading-activation in combination with a type of restriction, then their theory becomes very similar to the two-link theory offered here, with intersection through two acts of restriction corresponding to Type I intersection, intersection through an act of spreading-activation and an act of restriction corresponding to Type II intersection (where the more convenient order would be spreading-activation first, restriction second), and intersection through two acts of spreading-activation corresponding to Type III intersection.

Paired-Associate Learning

It should be noted that some types of experiments not usually thought of as involving intersection--paired-associate learning for instance--may be redesigned so as to explicitly become intersection experiments. Thus if visual concepts having a great many whole-to-whole links but few part-whole links (e.g. "army," "White House," "France," "ship," "farm," etc.) and auditory concepts having few whole-to-whole links by many part-whole links (e.g. "squeak," "hum," "bang," "squeal," "crackle," etc.) were employed in a paired-associate learning experiment, the subjects should find the "correct" word order Army squeak and White House hum more suggestive of imagery and therefore easier to memorize than the same phrases in the "incorrect" order, squeak Army and hum White House, inasmuch as the nature of the chosen concepts forces memory to use Type II intersection to arrive at mediating imagery. This is because Type I intersection, which works well independent of order, has been rendered almost impossible by matching visual concepts with auditory concepts.

Note that of the abundance of paired-associate learning tasks described in places such as Paivio (1971), not one concerns itself purely with Type II intersection as does the above experiment. Clearly by careful regulation of part-whole and whole-to-whole linking, new twists can be given to a large number of traditional cognitive tasks.

Polysemy

It was demonstrated by Anderson and Ortony (1975) that a phrase may give rise to a mental representation that is more detailed than the individual words of the phrase would seem to justify, and that this may be viewed as essentially a problem of polysemy. For instance the word container in the phrases cola container and apple container might in the first instance be particularized in memory as "bottle," and in the second as "basket," even though these concepts are not actually mentioned.

Similarly, there exists a more "subtle form of polysemy" (Anderson & Ortony, 1975, p. 177) that apparently allows a word such as kick to be represented in memory in as many different ways "as there are animals that kick and objects that can be kicked," where the particular image yielded would depend on the contest in which kick was embedded.

The problem with kick, of course, is that it is extravagant to assume that there are as many definitions of kick as there are particularized versions of kicking--just as for instance it would be extravagant to assume that there are as many definitions of red as there are red objects. In each case the remedy is to assume that the word, whether kick or red, has as its definition a general meaning, and that it is this general meaning which yields the various particularized versions of "kicking" and "redness." This same solution would, of course, work for container, since one may assume that the word container yields concept "container," and that it is this general concept which tends to yield "bottle," "basket," and the other forms of "container."

The advantage of this simple outlook is that it accords with common sense and is theoretically economical. It is also flexible enough to handle words such as ball, which must have at least two general meanings ("round object" and "dance"), inasmuch as the word refers to a set of objects and a set of events each of which has virtually nothing in common with the other.

The above position is with some modification that of Anderson and Ortony (see pp. 177-178). Anderson and Ortony go on to assert that an intersection search (concerning which they cite Quillian, 1968, 1969) might be the means by which memory arrives at a particularized image. This would mean that cola container yields "bottle" by an intersection search emanating from the concept (or perhaps the word) "cola," and from the general concept "container." And similarly, the words kick and ball would be disambiguated by an intersection search that exploited the overall context in which they were embedded.

As it turns out, their position, whose only clearcut weakness is that it does not clearly define the intersection search undertaken by memory, may be restated in the conceptual framework of part-whole and whole-to-whole linking, with the result that a much needed set of restrictions is imposed on the intersection search employed by memory.

Recall that for intersection to reliably to take place it must involve either two sets of currently activated part-whole links (e.g. "bright yellow" yields "sun"), or a set of residually activated whole-to-whole links and a set of currently activated part-whole links (e.g. "hospital" followed some time later by "blade" yields "scalpel").

It follows that interpreting Anderson and Ortony's examples in terms of part-whole and whole-to-whole linking forces the conclusion that the word kick yields by whole-to-whole link the concept "kick," which in turn yields by part-whole links all those particularized images that arise from "donkey kick," "human kick," etc. And in the same way, container yields by whole-to-whole link the concept "container," which in turn yields by part-whole links "bottle," "basket," etc.

As regards ball, it may refer to either "dance" or "round object." The simplest and most obvious way to account for this would be to assume that ball has a whole-to-whole link with each of these concepts. Unfortunately, it cannot be that simple as a sentence such as Cinderella went to the ball would then set up in memory a Type III situation, with the concept "Cinderella" and the word ball each tending to yield by whole-to-whole link the concept "dance." That would mean that the meaning "dance" might very well not be yielded, since Type III situations, unlike Type II situations, cannot be relied upon to yield an intersected concept. And since clearly the above sentence causes no such problem, it follows that a Type II intersection must be taking place. Accordingly, ball must have a part-whole link with ball in the sense of "dance" and another with ball in the sense of "round object," each of which must in turn be linked by whole-to-whole link to its general concept. Thus when ball is perceived, it is first disambiguated as either the word ball(dance) or ball(round object), then a whole-to-whole link from one of these yields the corresponding general concept, and that concept is disambiguated.

It follows that Type II intersections may cascade (cf. Anderson & Ortony, 1975, p. 178) in the sense that a Type II intersection might first be used to disambiguate a word and then be used to disambiguate the general concept denoted by that word. Thus, upon hearing or seeing the phrase Yankee pitcher, memory could use the whole-to-whole links of "Yankee" to disambiguate pitcher as pitcher(baseball)--as opposed to pitcher(water)--and memory could then use the same whole-to-whole links to disambiguate the concept "pitcher" as a specific Yankee pitcher.

Intersection Tasks versus Memorization

It should be noted that the above examples (e.g. apple container, donkey kick, etc.) are less than ideal because they leave some doubt as to whether a Type I or the more interesting Type II intersection is at work. This is one of the reasons deliberately artificial phrases such as hospital blade (or better yet Africa striped intersecting "zebra") are more interesting intersection tasks than casually ambiguous phrases such as apple container. If "hospital blade" yields "scalpel" there is little doubt that "hospital" does so because of the whole-to-whole link it has with "scalpel." With "apple container," however, there is doubt as to whether "apple" operates by a part-whole link with the image of a basket containing apples, or if "apple" operates by a whole-to-whole link with an image of basket alone--or if perhaps both links exist. This kind of flexibility, which undoubtedly memory possesses, makes the job of the experimenter that much more difficult.

Even so, it is almost certain that the data from intersection tasks tell the theorist far more about cognitive organization than do the data from memorization and learning tasks. This is because an intersection task probes cognitive organization by concentrating on only one aspect of memory, information retrieval, whereas a memorization task simultaneously involves information processing, storage, and retrieval, three sets of variables instead of one.

The point is that a memorization task involving phrases such as apple container and donkey kick might allow the subject to use any of several images to store the same information, and these images might be arrived at by Type I, Type II, or Type IV processing, and might be retrieved by Type I, Type II, or Type IV retrieval. Thus by its very nature the memorization task must gratuitously introduce added levels of complexity that prevent the theorist from developing a clear picture of what is really going on in memory.

In contrast a carefully constructed intersection task involves information retrieval only, with the memory under examination forced to act along predetermined channels. In this way the job of the theorist, normally an impossible one, becomes just manageable.

Link Identification

It follows from all that has been said so far that one of the major goals of psychology should be to identify by means of intersection experiments (e.g. Freedman & Loftus, 1971) what in memory is linked to what, and then to identify which of these links are whole-to-whole links and which part-whole links. Thus if A and B yield by intersection Z, A and B must have some kind of link with Z. Identifying which type they are requires the use of the following three principles (note that these principles are merely convenient reformulations of principles deduced earlier).

First principle: In any consistently or unexpectedly easy intersection of two concepts, there must be at least one set of part-whole links involved (i.e. a Type I or Type II intersection is at work).

Second principle: Type II intersections are more easily carried out with the whole-to-whole links first and the part-whole links second.

Third principle: If memory is consistently unsuccessful or experiences unexpected difficulties in yielding an intersected concept, there must be two sets of whole-to-whole links (i.e. a Type III situation) involved.

Of these three principles the least serviceable is the third, as it involves a search for consistent failure to yield an intersected concept.

Fortunately for the experimenter the first and second principles work very well together. Thus for anyone wanting to know what part-whole links (if any) emanate from a concept such as "deception," the first principle suggests the idea that a concept such as "hospital," which has many whole-to-whole links, should be made to precede (this is required by the second principle) the concept "deception," in a presentation to a subject who is then to attempt to intersect them. If the subject reports that hospital deception yields, say, "placebo," then the first principle would indicate that "deception" must have operated by a part-whole link.

Accordingly the experimenter would then be justified in concluding that "deception" has a part-whole link with "placebo," or at least with a concept closely related to "placebo." Then by varying the associative context from "hospital" to "military," to "White House," to "communist," etc., the experimenter could arrive at a collection of intersection responses from his subjects, a collection that profiles the various part-whole links of "deception."

The same procedure could then be applied to other concepts suspected of having part-whole links, with the end product being a list telling what is linked to what and how. This network could then be used to make accurate predictions of how memory processes, stores, and retrieves information.

Conclusion

The psychologist's normal method of working requires that he gather data about memory and then deduce from these data how memory is organized. The method employed within these pages breaks with that tradition by seeking the organization of memory not within memory itself, but in the forces that helped mold and shape it. This line of reasoning inevitably leads to the conclusion that the evolutionary forces of the environment have fashioned memory with an infrastructure employing two types of links, and that even though this infrastructure plays a direct and rather straightforward role in many cognitive actions, it is in effect made invisible by virtue of its efficiency and adaptability.

Accordingly, the theorist is confronted with a formidable problem precisely because memory itself is so flexible, and any number of cognitive experiments, however well executed, will not allow him to map out its cognitive structure, unless of course they are specifically designed to exploit those points of inflexibility predicted by psycho-evolutionary theory. It is with intersection tasks that memory finds itself most restricted, and they must inevitably play an expanding role in the attempt to illuminate the organization of memory. Such tasks must be used to discover what part-whole links and whole-to-whole links are at work in memory, for these are the links that steer cognitive flow. Any theory that ignores these links cannot hope to account for the actions of memory at its deepest level.

References

Anderson, J. R., & Bower, G. H. Human associative memory. Washington, D.C.: V. H. Winston, 1973.

Anderson, R. C., & Ortony, A. On putting apples into bottles--A problem of polysemy. Cognitive Psychology, 1975, 7, 167-180.

Brown, S. W., & Block, R. A. Context facilitation and disruption in word identification, Bulletin of the Psychonomic Society, 1980, 15, 242-244.

Celce-Murcia, M. On Meringer's corpus of "slips of the ear." In V. A. Fromkin (Ed.), Errors in linguistic performance: Slips of the tongue, ear, pen, and hand. New York: Academic Press, 1980.

Collins, A. M., & Loftus, E. F. A spreading-activation theory of semantic processing. Psychological Review, 1975, 82, 407-428.

Estes, W. K. An associative basis for coding and organization in memory. In A. W. Melton & E. Martin (Eds.), Coding processes in human memory. New York: Wiley, 1972.

Freedman, J. L., & Loftus, E. F. Retrieval of words from long-term memory. Journal of Verbal Learning and Verbal Behavior, 1971, 10, 107-115.

Garnes, S., & Bond, Z. S. A slip of the ear: A snip of the ear? A slip of the year? In V. A. Fromkin (Ed.), Errors in linguistic performance: Slips of the tongue, ear, pen, and hand. New York: Academic Press, 1980.

Hayes-Roth, B. Evolution of cognitive structures and processes. Psychological Review, 1977, 84, 260-278.

Holley-Wilcox, P., & Blank, M.A. Evidence for multiple access in the processing of isolated words. Journal of Experimental Psychology: Human Perception and Performance, 1980, 6, 75-84.

Horn, B. K. P. Obtaining shape from shading information. In P. H. Winston (Ed.), The psychology of computer vision. New York: McGraw-Hill, 1975.

Johnson, N. F. Organization and the concept of memory code. In A. W. Melton & E. Martin (Eds.), Coding processes in human memory. New York: Wiley, 1972.

Lachman, J. L., & Lachman, R. Theories of memory organization and human evolution. In C. R. Puff (Ed.), Memory organization and structure. New York: Academic Press, 1979.

Loftus, E. F. Activation of semantic memory. American Journal of Psychology, 1973, 86, 331-337.

Loftus, G. R., & Loftus, E. F. The influence of one memory retrieval on a subsequent memory retrieval. Memory and Cognition, 1974, 2 467-471.

Meyer, D. E., & Schvaneveldt, R. W. Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. Journal of Experimental Psychology. 1971, 90, 227-234.

Meyer, D. E., Schvaneveldt, R. W., & Ruddy, M. G. Loci of contextual effects on visual word recognition. In P. M. A. Rabbitt & S. Dornic (Eds.), Attention and performance V. New York: Academic Press, 1975.

Miller, G. A. The magical number seven, plus or minus two: Some limits on our capacity for processing information. Psychological Review, 1956, 63, 81-97.

Morton, J. Interaction of information in word recognition. Psychological Review, 1969, 76, 165-178.

Onifer, W., & Swinney, D. A. Accessing lexical ambiguities during sentence comprehension: Effects of frequency of meaning and contextual bias. Memory and Cognition, 1981, 9, 225-236.

Paivio, A. Imagery and verbal processes. New York: Holt, Rinehart & Winston, 1971.

Quillian, M. R. Word concepts: A theory and simulation of some basic semantic capabilities. Behavioral Science, 1967, 12, 410-430.

Quillian, M. R. Semantic memory. In M. Minsky (Ed.), Semantic information processing. Cambridge, Mass.: M.I.T. Press, 1968.

Quillian, M. R. The Teachable Language Comprehender: A simulation program and theory of language. Communications of the Association for Computing Machinery, 1969, 12, 459-476.

Raaijmakers, J. G. W., & Shiffrin, R. M. Search of associative memory. Psychological Review, 1981, 88, 93-134.

Razran, G. Mind in evolution. Boston: Houghton Mifflin, 1971.

Sapir, E., & Swadesh, M. Yana dictionary (M. R. Haas, Ed.). Los Angeles: University of California Press, 1960.

Schvaneveldt, R. W., & Meyer, D. E. Retrieval and comparison processes in semantic memory. In S. Kornblum (Ed.), Attention and performance IV. New York: Academic Press, 1973.

Simon, H. A. The functional equivalence of problem solving skills. Cognitive Psychology, 1975, 7, 268-288.

Simpson, G. Meaning dominance and semantic context in the processing of lexical ambiguity. Journal of Verbal Learning and Verbal Behavior, 1981, 20, 120-136.

Stroop, J. R. Factors affecting speed in serial verbal reactions. Psychological Monographs, 1938, 50 (5, Whole No. 225), 38-48.

Tulving, E., & Gold, C. Stimulus information and contextual information as determinants of tachistoscopic recognition of words. Journal of Experimental Psychology, 1963, 66, 319-327.

Waltz, D. Understanding line drawings of scenes with shadows. In P. H. Winston (Ed.), The psychology of computer vision. New York: McGraw-Hill, 1975.

Warren, R. E. Stimulus encoding and memory. Journal of Experimental Psychology, 1972, 94, 90-100.

Warren, R. E. Association, directionality, and stimulus encoding. Journal of Experimental Psychology, 1974, 102, 151-158.

Warren, R. E. Time and the spread of activation in memory. Journal of Experimental Psychology: Human Learning and Memory, 1977, 3, 458-466.

Wickelgren, W. A. Network strength theory of storage and retrieval dynamics. Psychological Review, 1976, 83, 466-478. (a)

Wickelgren, W. A. Subproblems of semantic memory: A review of "Human Associative Memory" by J. R. Anderson & G. H. Bower. Journal of Mathematical Psychology, 1976, 13, 243-268. (b)

Wickelgren, W. A. Learning and memory. Englewood Cliffs, N.J.: Prentice-Hall, 1977.