Barrett (natural science) have attempted to reconstruct the development of Dawin's thinking about evolution. The focus is on the critical period of 1837–1839, shortly after the voyage of the *Beagle* and two decades before the publication of the *Origin of Species*. During these years Darwin was an active and respected scientist, writing up the zoological results from the *Beagle* expedition and presenting his theories of the growth of coral reefs and the ecology of earthworms, but his notebooks reveal that he was simultaneously struggling with larger questions about evolution, variation, man, mind, and materialism. Here is a fascinating chronicle of false starts, errors reluctantly abandoned, transitional ideas, redefinition of the problem, and fear of persecution. Gruber notes that "even as he began to realize he might not want to say certain things aloud, he felt a need to say them as clearly as possible to himself." There seems to have been no sudden insight, even after reading Malthus, but there was a crucial cognitive reorganization when Darwin gave up the effort to explain variation within species and treated such variation as a premise of his argument about natural selection. The case of man and his accomplishments was treated as a unique scientific opportunity by Darwin, and his notebooks record many psychological speculations about the social human animal at least three decades before the publication of *The Descent of Man*.

Gruber presents the relevant facts and contemporary background, and he occasionally debates alternative explanations for the discontinuities in Darwin's progress. More often he presents a generalization about thinking — usually from Wertheimer, Köhler, or Piaget — and applies it to particular events. This idiographic procedure is not designed to yield any new generalizations about creative thinking, of course, nor is there any comparison with the thought processes of ordinary people; the purpose is rather to enable the reader to follow, and perhaps to identify with, the twists and turns of a great man's thoughts at a great moment in the history of science. Chapter 12, apparently intended as an interpretive summary, is, unfortunately, more Gruber than Darwin. But overall, Darwinophiles will enjoy the inside story of their hero's intellectual adventures, and it can be recommended as a corrective for psychology students who think of creativity in terms of writing unusual uses for a penny.

The last part of the volume is for Darwin scholars. It contains some of the notebooks, including one labeled "Old and Useless Notes," and other source materials. In a two-page foreword Jean Piaget shows how the present work can be subsumed under the heading of genetic epistemology.

Donald M. Johnson, *Michigan State University*

The Representation of Meaning in Memory

By Walter Kintsch. New York: Halsted Press, 1974. Pp. 279. $14.95.

Walter Kintsch has for some time been evolving a theory of memory. In 1970 Kintsch proposed a theory principally concerned with the recall and recognition of words (in *Models of Human Memory*, ed. D. Norman, pp. 307–373). However, to account for effects of semantic organization and for logical reasons, he elaborated that theory with some ideas about the semantic structure of the lexion. By 1972 the theory had become transformed into a distinctly propositional representation and had begun to study issues of sentence and paragraph

memory (in *Organization of Memory,* ed. E. Tulving and W. Donaldson, ch. 7). In 1974 his book *The Representation of Meaning in Memory* is a comprehensive report of the current state of his theory and an impressive variety of experiments it has generated.

The book has the sensible organization of three initial chapters devoted to theory and a subsequent seven to reports of various lines of experimental investigations. In addition, there are introductory and concluding chapters.

The theory

The principal object of analysis for Kintsch is a text, a series of one or more sentences. A text is represented as a list of atomic propositions. So, for instance, Kintsch proposes to represent the sentence *Cleopatra's downfall lay in her foolish trust in the fickle political figures of the Roman world* by this list of eight propositions:

1. (Because, α, β)
2. (Fell down, Cleopatra) = α
3. (Trust, Cleopatra, Figures) = β
4. (Foolish, Trust)
5. (Fickle, Figures)
6. (Political, Figures)
7. (Part-of, Figures, World)
8. (Roman, World)

Note that one proposition may be part of another proposition. So propositions 2 and 3 (α and β) are part of proposition 1. In this way Kintsch can create hierarchical nestings of propositions.

Kintsch gives considerable care to explicating the structure of propositions. He uses an n-ary relational structure. This representational system is quite powerful and similar systems are also used by several other researchers. In his system, relational terms (e.g., verbs) and predicates (e.g., adjectives) are written first and are followed by a list of their arguments. For instance, in proposition 3 Trust is the relational term and Cleopatra and Figures are the arguments. The arguments are either nouns or other propositions.

Most aspects of his representation seem adequate, but I find Kintsch's decision to use nouns as arguments unsatisfactory. Consider his representation of *The old man smiled:* (Old, Man) & (Smiled, Man). This identifies only two propositions underlying the sentence. However, the analysis by predicate calculus of this sentence identifies a third proposition: $(\exists x)$ [Old (x) & Man (x) & Smiled (x)]. That is, there existed someone, x, who was old, a man, and smiled. There is a separate proposition that x is a man. Thus, we can negate the original sentence by saying no, you're mistaken, the old person who smiled wasn't a man. Thus, the noun Man is not the argument of the predicates Old and Smiled. Rather, it is also a proposition about an unspecified individual (represented by the x in predicate calculus). I think Kintsch would have been wiser to use nouns as predicates as is done in predicate calculus. Actually, in Kintsch's theory nouns can sometimes serve as predicates. For instance, he represents *A robin is a bird* as (Bird, Robin).

Another reason for using nouns as predicates just as adjectives is that the distinction between nouns and adjectives is not clear-cut. The same term will

be used as noun or adjective in different contexts. For instance, *square* is an adjective when one is referring to picture frames, but it is a noun in a geometry text. Similarly, *red* is usually an adjective, but it can be a noun in Las Vegas. Given that Kintsch (p. 15) is only trying to represent the logical-semantic structure of language, and not the pragmatic considerations that determine a word's usage as a noun or adjective in a particular context, it seems that Kintsch should represent both nouns and adjectives identically.

Note that Kintsch's propositional notation is different from that we used in 1973 (Anderson and Bower, *Human Associative Memory*) or that of Norman, Rumelhart, and the LNR Research Group in 1975 (*Explorations in Cognition*), in that it is linear and not a network structure. He argues that there is a good reason for preferring linear structures, it being much easier to represent long and complex paragraphs with them than in a network notation. Network structures when they become large tend to get very messy, with crossing branches and such. This point is purely notational but still important. Network structures are useful for certain types of analyses, such as of the process of searching a proposition, while linear structures have advantages for analyses of other tasks.

Kintsch wants to establish the expressive power of his formalism. To do this, he reviews such problematical features of natural language as definite versus indefinite description, quantification, modality, implication, presupposition, location, time, and tense. For each such feature he attempts to show that it can be properly expressed in his representation. I think Kintsch is fairly successful with most of these features, but his treatment of quantification is inadequate. He proposes to treat quantifiers as predicates. So he will represent *All men die* as (Die, Man) & (All, Man). Similarly, *Some citizens complained* is represented as (Complain, Citizen) & (Some, Citizen). However, this will not deal with a phenomenon known as scope of quantification. Consider the sentences *All philosophers read some books* and *Some books are read by all philosophers.* The most common reading of the first sentence is that every philosopher reads some books although not necessarily the same ones as other philosophers, while the most common reading of the second sentence is that there is a specific set of books such that each one is read by all philosophers. Kintsch would represent both sentences as (Read, Philosopher, Book) & (Some, Book) & (All, Philosopher). Thus, his representation does not have the ability to represent the potential difference in meaning of these two sentences.

He offers a procedural analysis of many features of language. One example Kintsch gives is the difference between the indefinite and definite article. That is, the first time a particular woman is mentioned in a text we will refer to her as *a woman* but on later appearances as *the woman;* for example, *A woman bought steak yesterday* and then *The woman was shocked by the price.* Thus, the meaning of *a* versus *the* resides in the procedures that use these terms to decide reference in language comprehension or the procedures that choose between *a* and *the* in text generation.

Historical analysis

Kintsch includes in his theoretical section an evaluation of the suitability of associationism as a model for the sorts of cognitive behavior that concern him in his book. He points out that the associationism of Mueller could not account

for directed, purposeful thought. Kintsch reproduces Selz's arguments for labeled links or relations as replacements for associations. He sees his use of propositions as providing the same sort of structure as labeled links.

Kintsch further argues that labeled links are antithetical to associationism and scolds us for stating otherwise: "There have been attempts to update associationism by making relations out of associations while still retaining the term association. Anderson and Bower (1973) have chosen such a terminology, thereby turning the concept of an association into the opposite of what it has stood for in the last 2000 years. The title of their book is 'Human Associative Memory,' but associationists from Aristotle to James would have to disavow their efforts" (p. 42). As we argued in our book, it is very hard to list any strong tenets which were held by all associationists. I do not think that Kintsch can really defend his claim that all associationists would disavow relations. I can find no evidence that Aristotle was of that opinion (or of the other). Hume and other associationists proposed similarity, contiguity, and causality as primitive relations that could connect ideas. Thomas Brown and Herbert Spencer were both very explicit about the need for relations. It is hard to decide what, if anything, is meant by associationism when such varied opinion was held by those who are called associationists. Although we tried to find some meaning for the term in our book, we could be wrong. Perhaps the term associationism has no meaning.

Pattern matching

In chapter 4 Kintsch discusses episodic memory with particular attention to list learning and to the encoding-specificity phenomenon of Tulving and Thompson (*Psychological Review,* 1973, 80:352–373). This chapter is important to the rest of the book because here Kintsch spells out his ideas about pattern matching and completion, concepts reappearing at various points in his later discussions of prose memory. He proposes that each episode in memory is stored as a set of elements. Elements can refer to abstract propositions or to sensory events. The question Kintsch addresses is how input Y will be recognized as old. Kintsch proposes that Y is encoded as a set of elements and that this set is simultaneously matched to all episodes in memory. Let χ be the set of all episodes in memory and let X_j be a particular episode. Then Kintsch proposes that the probability of recognition will be a function of the maximum overlap between Y and the elements in χ. As he writes,

$$P_r\,(\text{match } Y \text{ to } \chi) = \max N(X_j \cap Y)/N(Y),\ X_j \Sigma \chi,$$

where $N(X_j \cap Y)$ is the number of elements in the intersection of X_j and Y, and where $N(Y)$ is the number of elements in Y. Note that Kintsch is measuring overlap solely in terms of set intersection. This equation specifies that the probability of recognition is a function of amount of overlap. If the input Y is recognized, it will be recognized as corresponding to the memory episode X_j with the maximum overlap. Of course, if the maximum overlap is small, Y will not be recognized.

A pattern-matching model is essential to a theory of memory that is going to deal with prose. I think Kintsch has made a good beginning here, but I suspect his formulation will run into difficulties because it does not make reference to the structure of the episodes. That is, Kintsch measures only how many

elements in two sets are the same. He does not test whether they are in the same structural relation.

Kintsch considers propositions as elements to be matched. Presumably, one proposition must perfectly match another. So (Give, John, Mary, Ball, Yesterday) would completely mismatch (Give, John, Mary, Bat, Yesterday). This seems wrong. The pattern-matching process should have access to the internal structure of propositions and be able to detect overlap. Another problem with Kintsch's structureless notion of pattern matching can be seen by considering these sentences along with their propositional representations:

1. *The man who teased the young dog hit the woman.*
   (Tease, Man, Dog) & (Young, Dog) & (Hit, Man, Woman)
2. *The man who liked the beautiful girl hit the woman who teased the young dog.*
   (Like, Man, Girl) & (Beautiful, Girl) & (Hit, Man, Woman) & (Tease, Woman, Dog) & (Young, Dog)
3. *The man who teased the young dog hit the woman who liked the beautiful girl.*
   (Tease, Man, Dog) & (Young, Dog) & (Hit, Man, Woman) & (Like, Woman, Girl) & (Beautiful, Girl)

Consider the third sentence as an input and contrast its match with the other two sentences in memory. It overlaps with each in three propositions and so has the same overlap with both as measured by Kintsch's equation. But surely the second sentence is a poorer match to the third than the first is. The proposition (Beautiful, Girl) occurs as part of the modification of Man in the second sentence but Woman in the third. Similarly, the location of the proposition (Young, Dog) is reversed from modifying Woman to Man. However, the measure of overlap is insensitive to these differences in the location of propositions in the overall structure.

The experiments

Kintsch reports a large number of experiments in an attempt to illustrate the usefulness of the concepts of his theory. These are not really tests of the theory. As he concludes at the end of his experimental section, "The reader has surely realized by now that the experiments are not tests of strict deductions from the theory. The theory is not specified completely enough to permit such strict deductions. It merely sketches in the relevant logical-linguistic background, but lacks a detailed processing component. Linguistic-logical considerations are powerful enough to constrain our ideas about how knowledge can or must be represented in memory, but they fail to provide us with more than the most general notions about psychological processing. The experiments reported here are studies in search of a processing theory rather than tests of one" (p. 243).

One of the features of propositions is that they are abstractions from text. So two different texts can assert the same set of propositions. In one series of experiments Kintsch had subjects study two types of paragraphs, one syntactically simple and one complex but both asserting the same base set of propositions. He reasoned that it should take longer to read and comprehend the complex paragraphs because it is harder to extract the underlying abstract propositions.

However, once comprehended, it should make no difference whether the original paragraphs were simple or complex when the subject was required to make an inference from those paragraphs, since the inference is made from the identical set of abstract propositions in both cases.

Kintsch did find that it took longer to read the complex paragraphs than the simple ones and that once the paragraphs were comprehended there was no difference between the two types in the time it took the subjects to make an inference. However, Kintsch also found across four experiments that subjects made more errors in inferences from complex paragraphs. While none of the individual effects was significant, the pattern is suspicious. Given the dangers of a trade-off between speed and accuracy (see Pachella, *Human Information Processing,* 1974, ed. B. H. Kantowitz, ch. 2), I would say Kintsch is too willing to accept the null hypothesis here. He seems to fall prey to what has been termed representativeness error — to equating no significant difference with no difference.

Kintsch in a later chapter looks at another, similar prediction about inference making. He had subjects study paragraphs such as *A burning cigarette was carelessly discarded. The fire destroyed many acres of virgin forest.* He reasoned that in comprehending such a paragraph a subject must make the inference *A discarded cigarette caused a fire.* This would be stored with the propositions explicitly affirmed in the paragraph. Thus, in a later test of the truth of the inference, subjects should be just as fast as if the inference had been actually stated. This would be another example of the abstract nature of the representation in memory, in that the same representation is set up for a number of different texts.

Kintsch contrasted subjects' verifications of sentences that had been explicitly stated and sentences that would have had to be inferred. He found, contrary to prediction, that subjects were faster to accept sentences that had been actually stated. In a later experiment he found that when he delayed his test 20 min or 48 hr, the difference between the two types of sentences became insignificant. He argues that the advantage of the explicit sentences in the immediate test was because surface representations of those sentences were still available, are more rapidly accessible than propositional representations, and so permitted more rapid verification.

There are problems with Kintsch's data analysis for this experiment. In the crucial experiment he contrasted verification of explicit versus implicit statements at delays of 0 sec, 30 sec, 20 min, and 48 hr. He found significant differences between explicit and implicit statements at delays of 0 or 30 sec but nonsignificant differences at delays of 20 min or 48 hr. However, explicit statements were verified still somewhat faster than implicit ones at delays of 20 min or 48 hr. Once again, Kintsch regards a failure to reject the null hypothesis as evidence of no difference. What he should have done to provide evidence for his theory is show that the difference in verification of explicit and implicit sentences at delays of 20 min or 48 hr was significantly less than that at the delays of 0 or 30 sec. This he did not do. Also, when the test was delayed there were many more errors (19%) on the implicit than the explicit statements (4%). Once again, one should worry about a trade-off between speed and accuracy.

Tests of representation

In addition to showing that the idea of abstract propositions is useful, Kintsch wants to show that the details of his representation are useful in understanding data. He provides a number of demonstrations of this point.

He shows that the time it takes to read a text is a function of the number of propositions that his analysis assigns to the text. That is, when the number of words in two texts is held constant, it takes subjects longer to read the text that contains more propositions. Another predictive variable uncovered by his analysis is the position of the proposition in a hierarchical representation of the text. Subjects are more likely to recall the higher propositions. In another demonstration, Kintsch shows that when the number of words in sentences is controlled, subjects are less likely to recall sentences based on two or three propositions as a unit than they are to recall sentences based on a single proposition. This is again taken as evidence of the psychological reality of the proposition as a unit of analysis.

A constant problem in all his research derives from the fact that Kintsch is contrasting different types of sentences under different conditions. Differences he ascribes to the structural variables in the sentences may really be due to other, uncontrolled differences in his materials. For instance, he contrasts the recall of sentences like *The settler built the cabin by hand* and *The crowded passengers squirmed uncomfortably.* Although both have four content words, by Kintsch's analysis the first has just one underlying proposition and the second three. Therefore, as explained above, subjects should tend to recall the former more as a unit and the second in a more fragmentary fashion. However, that difference between the sentences is confounded with lexical choice. The first sentence consists of three nouns and a verb; the second sentence consists of an adjective, a noun, a verb, an adverb. Adjectives, verbs, and adverbs tend to be less imageable than nouns, and as a consequence, they are probably harder to recall. If they *are* being more poorly recalled, one would expect more fragmentary recall of the sentences in which they participate.

Related to this problem is Kintsch's refusal to apply the statistics advocated by Clark (*Journal of Verbal Learning and Verbal Behavior,* 1973, 12:335–359) for testing whether effects are consistent across materials. As Clark demonstrated, apparent differences between conditions can be due to peculiarities of the particular sentences in the condition and would not replicate with other materials. Clark advocates procedures to test whether the effects will generalize over materials as well as subjects.

The Clarkian statistics tend to reduce the significance of results. As it stands, Kintsch seems to be too often accepting the null hypothesis. If he had used Clark's procedure, most of his remaining effects would also have become insignificant, as he admits. The obvious solution would be for him to have performed more powerful experiments with larger samples of materials and subjects. However, this would have been extremely expensive, working with paragraphs as his material. Kintsch typically gets one observation per paragraph, which is much less efficient than in the experiments on verbal learning where one gets one observation per word. Nonetheless, if one is going to do research on paragraphs, one might as well do it right. In my opinion, Kintsch's experimental section would have been stronger if he had reported fewer but more powerful experiments.

Effects of paragraph length

Kintsch looks at the effects of paragraph length on inferences, truth judgments, and sentence recognition. Analogous to the results of Atkinson and Juola for word lists (*Contemporary Developments in Mathematical Psychology,* 1974, ed. R. C. Atkinson et al.), Kintsch finds that the longer the paragraph, the slower subjects are to make these judgments. He proposes that all the propositions in these paragraphs are examined in parallel and that the subjects respond when they have examined the last proposition. Assuming that different propositions take varying amounts of time to be examined, the time to process the paragraph will depend on how long it takes to examine the most complex proposition and will increase with the length of the paragraph.

Semantic decomposition

A currently appealing idea in both generative semantics and computational linguistics is that complex words are decomposed into semantic primitives. So, for instance, Lakoff (*Synthese,* 1970, 22:151–271) analyzes *persuade* as *cause to come about to intend.* There have been suggestions (e.g., Anderson and Bower, 1973; Norman et al., 1975; Schank, *Cognitive Psychology,* 1972, 3:552–631) that when we hear *John persuaded Bill to hit Mary,* we store in memory something on the order of *John caused Bill to come about to intend to hit Mary,* in which there is no remnant of the term *persuade.* Both Bower and I and Norman et al. were a bit cautious in our advocacy of semantic decomposition, suggesting that it is an option that the subject may use if he wants a deep, detailed representation of the sentence's meaning.

Kintsch acknowledges some of the linguistic arguments for decomposition and admits it might have advantages in a computer implementation, but he argues against semantic decomposition as a psychological model. Intuitively, it seems to him "a poor hypothesis to assume that in memory and comprehension processes all concepts are decomposed into some small set of features, given that language has evolved to where we use complex word concepts" (p. 12). He wonders whether decomposition would ever stop — whether the final atomic primitives would ever be reached.

He puts to test a hypothesis he sees as deriving from the idea of semantic decomposition. The hypothesis is that it should be more difficult to process sentences that contain semantically complex words, that is, words that require considerable decomposition. In various experiments Kintsch looks at sentence-initiation time, sentence-completion times, phoneme monitoring, and sentence memory. In none of these does Kintsch find evidence that semantically complex words like *convince* are more difficult than less complex words like *believe.* In my opinion, these results present something of a challenge to theorists proposing semantic decomposition. It seems incumbent on them to show that there are differences in the processing of complex and simple words.

John R. Anderson, *University of Michigan*

Note: Preparation of this review was supported by Grant GB-40298 from the National Science Foundation.