DIMENSIONS OF INFERENCE

Gail McKoon
Roger Ratcliff

I. Introduction

Most investigations of the inferences that occur in reading have been motivated by a single question: Is an inference constructed during reading or isn't it? In this article and in recent research, we have proposed a new framework with which to view inference processes. According to this framework, inferences can vary on several dimensions, including the time required for them to be formed and the strength and specificity with which they are encoded. Consideration of these dimensions generates research concerned with the theoretical processes by which inferences are constructed, as well as the circumstances under which the results of inference processes can be observed empirically and the methods with which they can be measured.

The primary purpose of this article is to present evidence for each of several dimensions of inference, reviewing data from our laboratory that support each one. A secondary purpose is to specifically acknowledge that some of the notions of dimensions of inference are inherited from current global memory models. While the global models are used only at a metaphorical level, they provide valuable suggestions for understanding the encoding of inferences and offer the possibility of more quantitative development in the future.

THE PSYCHOLOGY OF LEARNING
AND MOTIVATION, VOL. 25

II. The Time Course of Retrieval for Inferences

One dimension of inference that has received considerable attention is the time required for retrieval. When an inference is completely encoded or instantiated during reading, then it should be available at a later test in the same way that explicitly stated information is available. However, even if an inference is not at all encoded during reading, then it still may be generated from cues given at the time of a retrieval test. For example, Corbett and Dosher (1978) showed that *hammer* could be an effective retrieval cue for a sentence about pounding nails, even when it would not have been encoded as an inference during reading (see also Singer, 1978).

There are several lines of evidence that indicate that the generation of inferences at retrieval takes more time than would be required by simply matching a retrieval cue against an inference that was actually encoded in the mental representation of a text. A strongly encoded inference can affect recognition performance within only 650 msec of processing time (see discussion below). The generation of an inference from a retrieval cue is unlikely to be accomplished in this same amount of time. The retrieval cue would first have to access information from the relevant text in memory, and then inference processes would have to work from the text information and the cue to produce the inference. Generation of an inference in this manner is unlikely to happen in as little time as 650 msec, as evidenced by experiments that have investigated the time course of retrieval processing.

In a series of articles, it has been shown that early in the time course of processing, only information about independent items (e.g., single words) is available. Information about the relations between items, information of the kind that would be used in generating inferences, is not available until 600 to 700 msec of processing has elapsed. This pattern of a delay in the availability of relational information has been shown to apply to several kinds of relations, including agent versus object relations in simple sentences (Ratcliff & McKoon, 1989), paired-associate relations for words (Dosher, 1984; Gronlund & Ratcliff, 1988), and set inclusion relations for categories (Ratcliff & McKoon, 1982). If these kinds of information are not available as early in processing as 650 msec, then it seems reasonable to assume that generated inferences would not be available that early either. Thus, inferences generated at retrieval would not be available until later in processing than inferences encoded in the mental representation of a text.

McKoon and Ratcliff (1986) presented indirect evidence about the time course of retrieval processing for inferences. Table I shows an example of the materials used to investigate inferences about predictable events. The predicting text was written to predict that the actress would die, and *dead*

TABLE I

EXAMPLE OF MATERIALS USED TO INVESTIGATE INFERENCES ABOUT PREDICTABLE EVENTS

Predicting:	*The director and the cameraman were ready to shoot closeups when suddenly the actress fell from the fourteenth story.*
Control:	*Suddenly the director fell upon the cameraman, demanding closeups of the actress on the fourteenth story.*
Cued recall with *dead* as a cue:	23% recall with predicting versus 4% recall with control (statistically significant difference)
Speeded recognition of *dead*:	34% errors with predicting versus 23% errors with control (marginally significant difference)

was used as a test cue to represent this inference. The control text was written to include all of the words from the predicting text that individually might be semantically associated to the test word *dead* (e.g., *fell, fourteenth story*). The inferences due to the meaning of the predicting text as a whole could thereby be separated from inferences due to associations of the test word with the meanings of the individual words in the text.

In one experiment with these materials (McKoon & Ratcliff, 1986), the amount of time for retrieval was relatively long because the task was cued recall. Subjects read a list of sentences, and then they were given a list of cues to recall the sentences. A sentence was presented in the list in either its predicting or control form, and the cue for each sentence was the word that represented the inference about the predicted event for its predicting text. Subjects could spend as much time as they liked on each cue. The results are shown in Table I: A cue was much more likely to lead to retrieval of a predicting sentence than a control sentence. The data suggest that, given this amount of time and this task, inferences that related cues to predicting sentences could be generated.

In a second experiment (McKoon & Ratcliff, 1986; see also McKoon, 1988; McKoon & Ratcliff, 1989a), retrieval time was severely limited. The task was recognition of single words. Subjects read a list of four sentences, and then they were given a short list of test words. For each test word, they were asked to decide whether or not it had appeared in one of the studied sentences, responding "yes" if it had and "no" if it had not. A signal was given 350 msec after each test word,and subjects were trained to respond 300 msec after the signal, so that total response time was limited to 650 msec. Each sentence was presented in either its predicting or its control version, and the test word was the predicted event (the correct response was "no," because it did not explicitly appear in any studied sentence). If the predicted event was strongly encoded into the representation of the predicting text in memory, then correct negative responses for the predicted word should be difficult, and so there should be more errors with the predicting text than with the control text. However, as shown in Table I, there was only a small difference in error rates (significant in one experiment but not in another).

Examining these two experiments in combination, it appears that in recall, the predicted event can function as a recall cue for its sentence, but in recognition, the predicted event makes little contact with its sentence in memory. The predicting condition increases recall by over 500% relative to the control condition. In contrast, the predicting condition increases errors in recognition by only about 50% (an especially small amount relative to variance, as evidenced by marginal significance). This difference shows that the inferences have a greater effect on recall than on speeded recognition. One interpretation would be that inferences about predictable events of this kind can be generated with the time available in recall, but they are not strongly encoded enough in the mental representations of their texts to have a large effect on speeded recognition. Thus, the results suggest that different information is available at different points in the time course of retrieval from an inference cue to the mental representation of a text. Not surprisingly, information that can be generated is not available as early as information about what is actually encoded in the mental representation. Another conclusion to draw from this is that studies aimed at measuring the encoding of inferential information must limit the time course of retrieval so that inferences are not generated during testing. When the time course is limited in this way, then other dimensions of inference such as the degree of encoding become apparent.

III. The Degree of Encoding of Inferences

A. GLOBAL MEMORY MODELS: GOODNESS OF MATCH

As mentioned at the beginning of this article, most investigations of inference processes have been designed to address the question of whether or not some specific kind of inference was constructed during reading. Phrasing the research problem in this way has often led to the presupposition that the answer to the question is "yes" or "no." However, our current framework proposes instead that the degree of encoding is variable. This framework is derived from recent models of memory.

The current models of memory provide descriptions of the processes by which information is retrieved from memory and account for a broad range of empirical data (Gillund & Shiffrin, 1984; Hintzman, 1986; Murdock, 1982; Ratcliff, 1978, 1988; Ratcliff & McKoon, 1988b). For the relatively fast processes involved in recognition of single words, all the models assume that a test word presented for recognition is matched against all items in memory in parallel. The result of this matching process is an overall measure of the goodness of match between the test word (in its experimental context) and memory. This measure is referred to as familiarity (Gillund & Shiffrin, 1984), resonance (Ratcliff, 1978), or echo intensity (Hintzman, 1986).

The retrieval models can be applied qualitatively to inference research by considering the matching of a potentially inferred concept to memory. The models suggest that such a concept would not necessarily match in an all-or-none manner, but rather that goodness of match could be a matter of degree. With a very high degree of match, the inference represented by the concept might be said to be instantiated. With a very low degree of match, the inference might not have been encoded at all. And with some degree of match that was between these extremes, the inference would be partially encoded. The important question to address then becomes what are the factors that govern the degree to which an inference is encoded? We have proposed two such factors, the strength with which the inference is encoded and the specificity of the inference.

B. THE STRENGTH AND SPECIFICITY OF INFERENCES

An inference is assumed to be represented in the mental representation of a text as a set of features (or propositions) of meaning. For a partially encoded inference, the features of the set do not completely or explicitly instantiate the inference. The degree of encoding is proposed to be a function of two factors: the specificity of the features making up the inference and the strength with which they are encoded. The degree of specificity, or focus, of an inference refers to the diversity of the different features or propositions that are included in the set making up the inference. Some inferences may be highly focused on one concept or event; others may be compatible with a variety of different possibilities. The strength of an inference and the specificity of an inference are assumed to be independent factors (though obviously there will be covariation in practice), and stronger inferences are assumed to be more accessible in memory than weaker inferences. Thus, an inference may be only partially encoded because its features are focused but weakly encoded, or strongly encoded but diffuse, or both weakly encoded and diffuse.

The effect of increasing the strength and specificity of an inference is shown in Table II (McKoon & Ratcliff, 1989a). The predicting text includes several words (*seamstress, threaded,* and *needle*) that are semantically associated to the predicted event, represented by the word *sew.* Because the text points so specifically to the event of sewing, an inference about sewing should be strongly focused. And because of support from the strong associations between the words of the text and the inference, the inference should be relatively strong. This high degree of encoding, shown by the high error rate with the predicting text compared to the control text, can be contrasted with the results (in Table I) obtained for similarly predictable events that do not have support from semantic associations (see McKoon & Ratcliff, 1989a).

A second example of increasing the degree of encoding of an inference through the use of well-known associations is shown in Table III (McKoon & Ratcliff, 1989c). The most likely exemplar of the category of animals that are milked on farms is *cows,* and the explicit description of this category gives the inference enough support to allow it to be encoded to a high degree, as shown by the large proportion of errors with the predicting text. Another possible exemplar of the category, *goat,* does not appear to match at all the encoded information in memory for the predicting sentence.

TABLE II

INCREASING THE DEGREE OF ENCODING OF AN INFERENCE THROUGH USE OF WELL-KNOWN ASSOCIATIONS: I

Predicting:	*The housewife was learning to be a seamstress and needed practice, so she got out the skirt she was making and threaded her needle.*
Control:	*The housewife was a careless seamstress, and when she dropped an unthreaded needle on the floor, she didn't find it until she stepped on it days later.*
Speeded recognition of *sew:*	59% errors with predicting versus 40% errors with control

TABLE III

INCREASING THE DEGREE OF ENCODING OF AN INFERENCE THROUGH USE OF WELL-KNOWN ASSOCIATIONS: 2

Predicting:	*The old man loved his granddaughter and she liked to help him with his animals; she volunteered to do the milking whenever she visited the farm.*
Control:	*The old man loved his granddaughter and she liked to help him with his animals when she visited the farm; she also liked the milk and cookies her grandmother provided.*
Speeded recognition of *cow:*	66% errors with predicting versus 40% errors with control
Speeded recognition of *goat:*	41% errors with predicting versus 41% errors with control

Although *goat* is compatible with the predicting sentence when it is presented in the immediate context of the predicting sentence (as judged by subjects' ratings; see McKoon & Ratcliff, 1989c), there is no higher error rate for *goat* in the predicting condition than in the control condition. One interpretation of this finding is that the encoded inference about *cows* is not a general description of the animals that can be milked on farms (with *cow* the most typical). Instead, the inference may be an encoding of the specific information that relates the text to *cows,* including, for example, information about the granddaughter helping with cows, the granddaughter visiting the cows, and so on.

These results are consistent with the notion that the degree of encoding of an inference can be manipulated and that one variable that increases degree of encoding is the amount of well-known information that is present in the text to support the inference. Although the inferences about *dead* and *sew* are rated about equally predictable from their predicting texts

(McKoon & Ratcliff, 1989a), only the inference about sewing leads to a significantly higher error rate in the predicting condition relative to the control condition. And even though *goat* is judged to be compatible with its predicting text (McKoon & Ratcliff, 1989c), there is no increase in error rate with the predicting text as there is with the most typical exemplar, *cow*. However, although these results are indicative, they do not allow a complete description of the encoded inferences. A complete description depends on comparisons of the match between a test word and memory in one retrieval context to the match between that same test word and memory in other retrieval contexts.

IV. The Retrieval Context of an Inference

A. CUE-DEPENDENT RETRIEVAL

In the previous section, it was claimed that inferences were not encoded in an all-or-none fashion, but instead could be encoded to varying degrees. This notion of variable encoding implies variable interactions between encoded inferences and retrieval contexts. Such interactions are at the heart of current memory models and so these memory models can be used to conceptualize the process of matching test cues against memory.

In the years since the case was made that memory could not be assessed independently of retrieval context (e.g., Craik & Tulving, 1975; Tulving, 1974), cue-dependent retrieval processes have become an important component of global memory models. Whether the models use distributed (Hintzman, 1986; Murdock, 1982) or single-node-per-concept (Gillund & Shiffrin, 1984) representations, whether they are intended to model the information in memory or the time course of retrieval (Ratcliff, 1978), they have all adopted as a basic mechanism some measure of the goodness of match between a cue (or test item) given at retrieval and information in memory.

Retrieval context includes not only the experimental situation and one test cue to memory, but also "priming" cues. In Gillund and Shiffrin's (1984) model and in Ratcliff and McKoon's (1988b) account of priming phenomena, a target test item and a prime item presented immediately prior to it are assumed to combine in a "compound" (with the prime weighted less than the target). The response to the target is determined by the goodness of match of the compound against memory. Thus, if the prime and target were previously associated in memory (as semantic associates like *dog–cat* or as two words of a previously studied sentence), then the goodness of match will be larger than if they were not associated, and a larger value of match will lead to faster and more accurate responses. This compound-cue theory for priming is different than a spreading activation account of priming, because the prime is not assumed to affect (activate) any items in memory prior to the presentation of the target as would be assumed by spreading-activation theories. Instead, the compound as a whole is matched against all items in memory. The differences between these two accounts of priming, and the evidence that supports the compound-cue theory over spreading activation, are discussed by Ratcliff and McKoon (1988b). The compound-cue view of the retrieval processes for a prime and target can be applied to targets that represent inferences. When different primes are used with the same inference target, then the degree of match between these different compounds and the mental representation of the text gives a picture of the information relevant to the inference in the textual representation.

B. DEGREE OF ENCODING AND RETRIEVAL CONTEXT

The representations of inferences in memory can be studied by varying the retrieval contexts of the target test words. For inferences about predictable events, those with high degrees of encoding should match information encoded in memory under most retrieval conditions (that is, the same retrieval conditions as information that was explicitly presented at study). For inferences with lower degrees of encoding, their match with information in memory should be improved when they are combined in a compound cue with a prime that also matches information from the same text. This pattern is shown in the data in Table IV.

TABLE IV

SPEEDED RECOGNITION OF TARGET WORDS IN TWO RETRIEVAL CONTEXTS

	Neutral prime word		Text word	
Target word	Predicting	Control	Predicting	Control
dead	ready	ready	actress	actress
Error rate (%)	34	23	48	21
sew	ready	ready	housewife	housewife
Error rate (%)	59	40	64	40
cow	ready	ready	granddaughter	granddaughter
Error rate (%)	66	40	63	51
goat	ready	ready	granddaughter	granddaughter
Error rate (%)	41	41	38	36

In the experiments illustrated in Table IV (McKoon & Ratcliff, 1986, 1989a,c), the target word (e.g., *dead*) was tested in two retrieval contexts, one with the neutral prime word *ready* and one with a word from the text (e.g., *actress*). The prime words were presented for a brief time (200 msec), and then the target word was presented immediately. The prime word did not require any response; only the target word required a recognition decision. (The data shown in Table IV for the neutral priming conditions are those from Tables I, II, and III.[1])

For inferences with support from well-known information, like the inferences about sewing and cows, the test word matches the mental representation of the predicting text to such a degree that error rates are high even when the test word is presented by itself (in the neutral priming condition). For an inference without such support at encoding, like the inference about death, the difference between the predicting and control conditions is larger when the test word is combined with an explicitly stated word from the text (*actress*) than when it appears by itself. Finally, a nontypical category member (*goat*) does not match the encoded inference so that error rate is not increased even with a prime from the text.

Another example of the potentially large effects of retrieval context is given in Table V, with a different kind of inference (McKoon & Ratcliff, 1988a). The predicting text is more relevant to the features of color in the meaning of *tomato* than the control text. To test whether such contextually relevant meaning was differentially encoded with the two texts, a sentence-verification procedure was used (McKoon & Ratcliff, 1988a). Subjects read a list of two texts and then were presented with a list of sentences for verification. Some of the sentences could be judged true or false according to general knowledge (e.g., *Tomatoes are red*). Others could be judged only according to information from a text (e.g., *The still life would require accuracy*). Retrieval context was varied with the test sentence that immediately preceded the target test sentence. The priming sentence was either neutral (some sentence true by general knowledge but not related to any studied text) or it was from one of the studied texts.

---

[1] The interpretation of the data that are given here is based on the assumption that speeded recognition data reflect information encoded into the mental representation of a text during reading. Potts, Keenan, and Golding (1988) have suggested an alternative interpretation, that recognition responses are due to processes that occur at the time of the recognition text and not at encoding. These processes are supposed to compute the compatibility of a test word against textual information in memory. However, this interpretation has been ruled out in two ways (McKoon & Ratcliff, 1989b, in press). First, it has been shown that compatibility (as measured by subjects' ratings) does not predict recognition performance (for individual items). Second, there are words that are compatible with their predicting texts that are not inhibited in recognition.

TABLE V

SENTENCE VERIFICATION AND EFFECTS OF RETRIEVAL CONTEXT

Predicting text: *This still life would require great accuracy. The painter searched many days to find the color most suited to use in the painting of the ripe tomato.*

Control text: *The child psychologist watched the infant play with her toys. The little girl found a tomato to roll across the floor with her nose.*

Verification sentence: *Tomatoes are red.*

	Neutral prime		Text prime	
	Predicting	Control	Predicting	Control
	neutral	neutral	still life requires accuracy	psychologist watched infant
	tomatoes are red	tomatoes are red	tomatoes are red	tomatoes are red
Response time (msec)	1049	1065	1155	1265

The data in Table V show that when the sentence *Tomatoes are red* is presented in the neutral priming condition, it does not match the encoded information from the predicting text sufficiently to give a significant speed-up in response time relative to the control text. However, when the sentence is presented in the context of a prime from its text, then the combined information does match better in the predicting than in the Control condition, and response time is significantly speeded. Surprisingly, this speed-up is observed even when the prime sentence is from some text other than the one about tomatoes (see McKoon & Ratcliff, 1988a). Apparently, the context of "studied in this experiment" that comes from a prime from any text is a context that can combine with the target to increase the match between the target and the mental representation of studied textual information.

These patterns of data demonstrate the necessity of examining inferences under various retrieval conditions. For the inferences about death and the color of tomatoes, consideration of the neutral priming conditions alone would have led to the conclusion that they were not encoded in memory. It is only under some retrieval conditions that evidence for these inferences can be observed. Furthermore, differences between inferences can be seen only across retrieval contexts. Specifically, it is the fact that inferences with support from well-known information are less sensitive to retrieval conditions that allows them to be judged to have a higher degree of encoding in memory than inferences without such support.

V. The Time of Availability of Inferences

The final dimension of inference to be discussed in this article concerns the mechanisms by which inferences are generated. Little information has been obtained about the actual psychological processes involved in forming inferences, partly because research has concentrated on whether inferences are generated and not on how they are generated and partly because of the lack of incisive experimental methods. However, one dimension which can be examined is the time required for the construction of different kinds of inferences.

In the previous sections, it was suggested that inferences are encoded to a higher degree if they can be based on well-known information. It may be that well-known information contributes to degree of encoding by virtue of its speed and ease of availability. The more information that is available and the more quickly it is available, the more information will be encoded into the representation of the text and the more strongly encoded will be the information that forms an inference.

The time of availability of implicit information was examined with the materials exemplified in Table VI (McKoon & Ratcliff, 1989a). The predicting text predicts the event represented by the test word *hurt;* this event is predicted only at the point of the final word of the text. The control text is exactly the same as the predicting text, except that the final word is

TABLE VI

TIME AVAILABILITY OF IMPLICIT INFORMATION

Predicting text: *The diver prepared to do a double somersault into the pool; he jumped, spun, and hit the cement.*
Control text: *The diver prepared to do a double somersault into the pool; he jumped, spun, and hit the water.*
Immediate recognition of *hurt:*
  Predicting: 768 msec and 18% errors
  Control: 748 msec and 15% errors

Predicting text: *After shopping for hours, the grandmother headed for her favorite chair.*
Control text: *After shopping for hours, the grandmother headed for her favorite store.*
Immediate recognition of *sit:*
  Predicting: 758 msec and 16% errors
  Control: 693 msec and 13% errors

Predicting text: *After shopping for hours, the grandmother headed for her favorite chair.*
Control text: *After shopping for hours, the grandmother finally found the perfect chair.*
Immediate recognition of *sit:*
  Predicting: 741 msec and 27% errors
  Control: 762 msec and 27% errors

changed so that the test word is not predicted. These texts were presented to subjects one word at a time (250 msec per word), and the test word was presented immediately after the final word of the text. Thus, the time between presentation of the text word that allowed the inference and presentation of the test word was only 250 msec. The test word was presented for recognition, so that subjects had to decide whether or not it had appeared in the text just read. For test words that expressed inferences, the correct response was "no," so evidence that the inference had been generated would be inhibition, either slowed responses or increased errors. Inferences that were not supported by well-known information (inferences like *hurt*) showed no significant inhibition of responses (the first set of data in Table VI). However, the well-known information that connects *chair* to *sit* did lead to significant inhibition, and this inhibition was apparent whether the association was appropriate for the meaning of the text as a whole (as in the second predicting text in Table VI) or inappropriate for the meaning of the text as a whole (as in the third predicting text).

These data give us one piece of information about the processes that construct inferences. Well-known information based on semantic associations is available during text comprehension within a short amount of time (the 250 msec for which *chair* was presented plus some part of the response time that could be used for further processing). But inferences that cannot be based on such information are not so quickly available. The speed with which the well-known information is available brings to mind another kind of information that is also available quickly, the referent of an anaphor. Several experiments (Corbett, 1984; Dell, McKoon & Ratcliff, 1983) have shown that the relation between an anaphor and its referent can be calculated within about the same time as the relation between *chair* and *sit* in the experiments in Table VI. Perhaps, for some types of anaphors, the referents are quickly and easily available from short-term memory just as well-known information is quickly and easily available from long-term memory.

The data in Table VI give only one piece of information about the processes that construct inferences. But this information is consistent with other results about inferences that indicate the importance of the ease of availability of features that make up the meanings of inferences.

VI. Conclusions

The collection of data summarized in this article is consistent with the view that inferences are variable on several dimensions. Some inferences may be so completely encoded during reading that they behave as though they had been explicitly stated. Others are only partially encoded and appear only in the retrieval context of other information from the text. And still others must be generated after reading, using appropriate cues and enough time for processing. These variable degrees of encoding may depend on the ease with which the information that supports the inference is available from memory. In addition, the variable nature of encoding emphasizes the importance of testing for inferences under a variety of retrieval conditions.

This collection of data is also consistent with a more general view of text comprehension that we have described previously. According to this view, inferences mainly establish local coherence among immediately available pieces of information, and there is only minimal encoding of other kinds of inferences (McKoon & Ratcliff, 1986; Ratcliff & McKoon, 1988a). The inferences necessary to connect propositions by argument repetition include the inferences that connect two instances of the same concept (Ratcliff & McKoon, 1978) and the inferences that connect an anaphor to its referent (Clark & Haviland, 1974; Corbett, 1984; Corbett & Chang, 1983; Dell *et al.,* 1983; McKoon & Ratcliff, 1980). These inferences are generated quickly and automatically, perhaps because they are available from working memory similarly to the way well-known information is available from long-term memory. In contrast, inferences that are not necessary for coherence, such as elaborative inferences about death from falling off a fourteenth-story roof, inferences about schema information (Alba & Hasher, 1983; Seifert, McKoon, Abelson, & Ratcliff, 1986), and inferences about the instruments of verbs (Corbett & Dosher, 1978; McKoon & Ratcliff, 1981) are not encoded completely; they may be partially encoded or not encoded at all. This minimal-coherence view arises directly from research that emphasizes the variable nature of inference processes, and the data in this article demonstrate the potential gains from such research. However, even though this approach will add to our knowledge of inference processes, much more needs to be done to increase our knowledge of the factors described in this article, to examine other factors and their relationships to those discussed here, and to begin to develop models of inference processes.

ACKNOWLEDGMENTS

This research was supported by NSF grant 85-16350 to Gail McKoon and NSF grant 85-10361 to Roger Ratcliff.

REFERENCES

Alba, J. W., & Hasher, L. (1983). Is memory schematic? *Psychological Bulletin,* **93,** 203–231.

Clark, H. H. & Haviland, S. E., (1974). What's new? Acquiring information as a process in comprehension. *Journal of Verbal Learning and Verbal Behavior,* **13,** 512–521.

Corbett, A. T. (1984). Pronominal adjectives and the disambiguation of anaphoric nouns. *Journal of Verbal Learning and Verbal Behavior,* **17,** 683–695.

Corbett, A. T., & Chang, F. R. (1983). Pronoun disambiguation: Accessing potential antecedents. *Memory and Cognition,* **11,** 283–294.

Corbett, A. T., & Dosher, B. A. (1978). Instrument inferences in sentence encoding. *Journal of Verbal Learning and Verbal Behavior,* **17,** 479–491.

Craik, F. I. M., & Tulving, E. (1975). Depth of processing and the retention of words in episodic memory. *Journal of Experimental Psychology: General,* **104,** 268–294.

Dell, G. S., McKoon, G., & Ratcliff, R. (1983). The activation of antecedent information during the processing of anaphoric reference in reading. *Journal of Verbal Learning and Verbal Behavior,* **22,** 121–132.

Dosher, B. A. (1984). Discriminating preexperimental (semantic) from learned (episodic) associations: A speed–accuracy study. *Cognitive Psychology,* **16,** 519–555.

Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review,* **19,** 1–65.

Gronlund, S. D., & Ratcliff, R. (1989). The time-course of item and associative information: Implications for global memory models. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* **15,** 846–858.

Hintzman, D. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review,* **93,** 411–428.

McKoon, G. (1988). *Word identification and elaborative inference: The relation between theory and empirical measurement.* Submitted.

McKoon, G., & Ratcliff, R. (1980). The comprehension processes and memory structures involved in anaphoric reference. *Journal of Verbal Learning and Verbal Behavior,* **19,** 668–682.

McKoon, G., & Ratcliff, R. (1981). The comprehension processes and memory structures involved in instrumental inference. *Journal of Verbal Learning and Verbal Behavior,* **20,** 671–682.

McKoon, G., & Ratcliff, R. (1986). Inferences about predictable events. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* **12,** 82–91.

McKoon, G., & Ratcliff, R. (1988a). Contextually-relevant aspects of meaning. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* **14,** 331–343.

McKoon, G., & Ratcliff, R. (1989a). Semantic association and elaborative inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* **15,** 326–338.

McKoon, G., & Ratcliff, R. (1989b). Assessing the occurrence of elaborative inference with recognition: Compatibility checking vs. compound cue theory. *Journal of Memory and Language,* **28,** 547–563.

McKoon, G., & Ratcliff, R. (1989c). Inferences about contextually-defined categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* **15,** 1134–1146.

McKoon, G., & Ratcliff, R. (1990). Textual inferences: Models and measures. In K. Rayner, D. Balota, & F. D'Arcais (Eds.), in press.

Murdock, B. B. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review,* **89,** 609–626.

Potts, G. R., Keenan, J. M., & Golding, J. M. (1988). Assessing the occurrence of elaborative inferences: Lexical decision versus naming. *Journal of Memory and Language*, **27**, 399–415.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, **85**, 59–108.

Ratcliff, R. (1988). Continuous versus discrete information processing: Modeling the accumulation of partial information. *Psychological Review*, **95**, 238–255.

Ratcliff, R., & McKoon, G. (1978). Priming in item recognition: Evidence for the propositional structure of sentences. *Journal of Verbal Learning and Verbal Behavior*, **17**, 403–417.

Ratcliff, R., & McKoon, G. (1982). Speed and accuracy in the processing of false statements about semantic information. *Journal of Experimental Psychology: Human Learning and Memory*, **8**, 16–36.

Ratcliff, R., & McKoon, G. (1988a). Memory models, text processing, and cue-dependent retrieval. In F. I. M. Craik & H. L. Roediger (Eds.), *Varieties of memory and consciousness: Essays in honor of Endel Tulving*. Hillsdale, NJ: Erlbaum.

Ratcliff, R., & McKoon, G. (1988b). A retrieval theory of priming in memory. *Psychological Review*, **95**, 385–408.

Ratcliff, R., & McKoon, G. (1988c). Similarity information versus relational information: Differences in the time course of retrieval. *Cognitive Psychology*, **21**, 139–155.

Seifert, C. M., McKoon, G., Abelson, R. P., & Ratcliff, R. (1986). Memory connections between thematically similar episodes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **12**, 220–231.

Singer, M. (1978). *The role of explicit and implicit recall cues*. Paper presented at the meeting of the American Psychological Association, Toronto.

Tulving, E. (1974). Cue-dependent forgetting. *American Scientist*, **62**, 74–82.