Journal of Experimental Psychology:
Learning, Memory, and Cognition
1988, Vol. 14, No. 2, 331-343

0278-7393/88/$00.75

Contextually Relevant Aspects of Meaning

Gail McKoon and Roger Ratcliff
Northwestern University

A series of six experiments investigated whether inferences about contextually relevant aspects of meaning were encoded into memory during reading. In all the experiments, subjects studied short paragraphs. Then, test sentences were presented that expressed relevant aspects of meaning that had not been explicitly stated in the paragraphs. For example, for a paragraph about searching for the correct color to paint a picture of a tomato, a relevant aspect of meaning would be that tomatoes are red. The test sentences were presented either immediately following the relevant paragraph or after a delay. With immediate testing, it was argued that the facilitation obtained in verification latency could result from processes occurring either when the context was read or when the test sentence was verified. With delayed testing, evidence was found to support the hypothesis that contextually relevant aspects of meaning are incorporated into the memory representation of the paragraph, but such evidence was obtained only when the retrieval environment encouraged the use of newly learned information in the decision process on the test sentence.

It is well accepted that general knowledge is used during reading to understand the meaning of a text. Definitions of words, reasons for looking at a menu, expectations of "happy endings," and expectations of significant *F* values are all thought to be combined with the information stated in a text in order to achieve comprehension. This article examines under what conditions such general knowledge is incorporated into the memory representation of a text, when it is not explicitly stated in the text.

A basic issue is whether different aspects of the meaning of a word are encoded into memory as a function of the context in which the word is read. Figure 1 shows two paragraphs, both of which mention the noun *tomato*. Different aspects of the meaning of *tomato* are relevant to the two paragraphs; the first paragraph has more to do with tomatoes being red, the second more to do with tomatoes being round. The experimental concern is whether the appropriate aspect of meaning, *tomatoes are red* for the first paragraph or *tomatoes are round* for the second paragraph, is encoded into memory with the paragraph.

Several previous studies have been concerned with a similar question. Barsalou (1982) showed that subjects processed properties of nouns differently in different contexts. The subjects were presented (for 6 s) with a context sentence with one noun underlined; then, they were asked to verify whether the noun possessed a particular property (not stated in the context sentence). Verification time (for some kinds of nouns and properties) was faster when the property was relevant to the context. The same finding was obtained by Tabossi and Johnson-Laird (1980) and Tabossi (1982). They presented a context sentence for 4, 5, or 6 s (in different experiments), a blank interval for 1 s, and then a yes–no question about a property of one of the nouns from the context sentence. Questions about properties relevant to the context sentence were answered more quickly than questions about irrelevant properties.

These findings clearly show that some combinations of context and property information are easier or faster (or both) to process than others. Thus, they lead to theoretically interesting questions about the kinds of factors that govern effects of context relevance (Barsalou, 1982; Tabossi, 1982). But it is difficult to assign an explanation of these effects to any one type or component of processing. Although faster verification times for context-relevant properties might reflect an automatic activation process (in Posner & Snyder's, 1975, sense of automatic), the long presentation times for the context sentences suggest the involvement of strategic processes. Furthermore, the context-relevance effect might be due either to processes occurring at the time of reading the context sentence or to processes occurring at the time of verifying the property. In the first case, the relevant property would be activated during reading and, because it is already activated, would be faster to verify. In the second case, the relevant property would not be activated during reading; instead, it would be faster to verify because it was easier to integrate with the context sentence. Finally, the finding of an effect of context relevance gives no information about the memory representation of the context sentence. Even if the relevant property was activated during reading of the context sentence, it might or might not be encoded into memory with the context sentence.

This last issue, whether implicit properties are included in the memory representation of text information, is the issue of concern in the present article. Although Barsalou (1982), Tabossi and Johnson-Laird (1980), and Tabossi (1982) did not address this issue, their work does show that different

An earlier discussion of this work was presented at the meetings of the Psychonomic Society in November 1982.

This research was supported by National Institutes of Health Grant HD18812 and National Science Foundation Grant 85-16350 to Gail McKoon and National Science Foundation Grant BNS-8510361 to Roger Ratcliff.

Correspondence concerning this article should be addressed to Gail McKoon, Psychology Department, Northwestern University, Evanston, Illinois 60201.

Version 1

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.

Target Test Sentences

Matching: Tomatoes are red. (True by general knowledge)
Mismatching: Tomatoes are round. (True by general knowledge)

Filler Test Sentences

The still life would require great accuracy. (True by paragraph)
Balloons are heavy. (False by general knowledge)

Version 2

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

Target Test Sentences

Matching: Tomatoes are round. (True by general knowledge)
Mismatching: Tomatoes are red. (True by general knowledge)

Filler Test Sentences

The little girl played with her toys. (True according to paragraph)
Balloons are heavy. (False by general knowledge)

Figure 1. Example of Paragraphs Used in Experiment 1.

aspects of the meanings of words are sensitive to context and that an immediate verification task demonstrates this sensitivity. Thus, in our first experiment, we used immediate verification to validate our materials. Experiments 2–5 used a delayed verification task to investigate whether the aspects of meaning shown to be context sensitive in Experiment 1 were integrated with the context into the memory representation of the context. If they were so integrated, then aspects of the meanings of nouns that matched the context paragraph would be verified faster than aspects that did not match.

In Experiment 1, with immediate verification, short paragraphs like those shown in Figure 1 were presented to subjects. Immediately after each paragraph, test sentences were presented for a verification (true or false) judgment. Sentences that were designed to test the experimental hypothesis expressed some aspect of the meaning of a noun in their paragraph, an aspect that had not been stated explicitly. This aspect of meaning either matched or mismatched the meaning of the paragraph. As expected, we found that matching test sentences were verified faster than mismatching test sentences.

If relevant aspects of meaning are integrated with the paragraphs into memory, then the matching effect should also be obtained when testing is delayed. In Experiments 2–5, with delayed verification, the same paragraphs as in Experiment 1 were presented to subjects, but they were presented in study lists of three paragraphs, and sentences for verification were not given until after the three paragraphs of a list had been read, so that there was a delay between reading and test. The target verification sentences were the same as in Experiment 1: They either matched or mismatched their paragraph in meaning. Matching targets in Experiments 2–5 were verified faster than mismatching targets when tested in the context of studied material, confirming the hypothesis that matching aspects of meaning were integrated into memory.

The other hypotheses of Experiments 2–6 concerned the *retrieval* context in which the test sentences were verified. We were interested in the question of whether the newly learned information of the studied paragraphs was "compartmentalized," separate from general knowledge (Potts & Peterson, 1985). If this were the case, then a matching effect might obtain when the information relevant to a paragraph, *tomatoes are red*, was tested in a retrieval context of newly learned information, but not when it was tested in a retrieval context of general knowledge. In a retrieval context of newly learned information, the decision process for a target test sentence would be likely to include newly learned information, and this newly learned information would give the matching effect. In a retrieval context of general knowledge, newly learned information would be less likely to be included in the decision about the target; the decision would be more likely to depend on general knowledge, and so a matching effect would be less likely.

Retrieval context was manipulated by the test sentence that immediately preceded the matching or mismatching target. Across the experiments, this sentence was either a true sentence from the same paragraph as the target sentence, a true sentence from another paragraph in the list of three paragraphs just read, or a sentence that was not related to any of the three paragraphs but that was known to be true by general knowledge. With a sentence from a studied paragraph preceding the target, matching targets should be verified faster than mismatching targets. But with a sentence from general knowledge preceding the target, the matching effect might be reduced or eliminated. This effect of retrieval context was demonstrated in Experiments 2–6.

Experiment 1

Experiment 1 was designed to show a matching effect when test sentences were presented immediately after the paragraphs were read. Subjects each read a total of 60 paragraphs, 40 used in the experimental design and 20 fillers. Presentation of each paragraph was immediately followed by presentation of three test sentences. Subjects were instructed to respond "true" or "false" to each sentence, and they were instructed that some sentences were true or false according to the paragraph just read, whereas others could be verified only according to general knowledge.

Examples of the experimental paragraphs are shown in Figure 1. For each item, there were two versions of the paragraph and both filler and target test sentences. One target test sentence matched the meaning of the first version but not the second version, and the other matched the meaning of the second version but not the first. The target test sentences were always true according to general knowledge, and were always presented first, immediately after a paragraph was read. Two filler test sentences followed. For filler paragraphs, the first of the three test sentences was always false.

Method

Subjects. The subjects were 16 Dartmouth undergraduates who participated in the experiment for extra credit in an introductory psychology course.

Materials. A total of 40 pairs of paragraphs like that shown in Figure 1 were written. One version of a pair matched one aspect of meaning of a particular noun (e.g., *tomato*); the other version matched another aspect of meaning of the same noun. The aspects of meaning were chosen so that sentences testing them would be clearly true, and would not test obviously infrequent or unlikely properties. Otherwise, the nouns and properties were not chosen in any special way (a complete set of materials is available from the authors). The paragraphs varied in number of words from 20 to 30. For each pair of paragraphs, there were two true target test sentences that mentioned the critical noun in the subject position (varying in number of words from 3 to 7). One of the targets matched the aspect of meaning of the noun in one of the paragraphs and the other matched the aspect of meaning in the other paragraph. The predicates of the targets (e.g., *are red*) were not mentioned in the paragraphs. Thus, a subject could know these test sentences were true only according to general knowledge and not according to information given in the paragraph. For each paragraph of a pair, there were also two other filler test sentences; these were true or false according to information in the paragraph or false according to general knowledge.

In addition to the 40 pairs of paragraphs used in the experimental design, there were 20 filler paragraphs of the same general style and length. Each of these 20 had three test sentences. In total, a subject read 60 paragraphs and was tested with 40 sentences true according to general knowledge (those like *tomatoes are red* used for the design of the experiment), 40 sentences false according to general knowledge, 50 sentences true according to a paragraph, and 50 sentences false according to a paragraph.

Procedure. Each subject was tested in one 50-min session. Presentation of all materials and collection of data was controlled by a microcomputer driven by Dartmouth's time-sharing computer system.

Each subject read 5 practice paragraphs, and then one version of each of the 40 pairs of paragraphs and the 20 filler paragraphs. Presentation of each paragraph began with an instruction to the subject to press the space bar on a CRT keyboard to initiate the paragraph. Then, the paragraph was presented on the CRT screen for 8 s, it was removed from the screen, and three test sentences were presented one at a time. Each test sentence remained on the CRT screen until the subject responded "true" or "false"; then it disappeared from the screen and the next test sentence was presented after a 200-ms pause. After the third test sentence, the instruction to press the space bar to begin the next paragraph appeared.

Subjects were instructed to respond as quickly and accurately as possible, pressing the *?/* key on the CRT keyboard for true and the *Z* key for false. If any response was incorrect, the word ERROR was presented for 2 s before presentation of the next test sentence. Subjects were informed that it would be necessary to respond to some test sentences according to general knowledge and to others according to information given in the paragraph just read.

Design. Three variables were combined in a Latin square design. One was experimental condition; there were two possible paragraphs of a pair and two possible test sentences, yielding four conditions (two matching and two mismatching). Which paragraphs and which test sentences were grouped together (i.e., whether *tomatoes are red* was assigned to the first group and *tomatoes are round* to the second group or vice versa) was arbitrarily determined. The other variables were groups of subjects (four groups with 4 subjects per group) and sets of pairs of paragraphs (four sets with 10 pairs per set). For the 40 paragraphs of the experimental design, the critical test sentence (true according to general knowledge) was always presented first after the paragraph was read. For the 20 filler paragraphs, the first test sentence was always false. Otherwise, a different random order of presentation of paragraphs and test sentences was used for every 2 subjects.

Results

All analyses were based on mean correct response times for each subject or each item in each condition.

If the meaning of a target test sentence about a noun matched the aspect of meaning of that noun in the paragraph read by a subject, response time for the sentence was faster than if the meanings did not match. Response times for targets that matched averaged 1,270 ms; the average time for targets that did not match was 1,390 ms. This difference was significant with subjects treated as a random variable, $F(1, 12) = 20.5$, $p < .01$, and with sentences treated as a random variable, $F(1, 36) = 15.5$, $p < .01$. Average standard error of these means was 26 ms. Which of the two groups of targets was tested (the group with *tomatoes are red* or the group with *tomatoes are round*) did not affect response times, Fs < 1, or interact with the effect of matching, Fs ≤ 1.3.

Subjects were not only faster but also more accurate when the aspect of meaning in the targets matched the meaning in the paragraph. There were 4.7% errors when the meaning matched and 9.7% when it did not match, a difference significant with subjects as a random variable, $F(1, 12) = 5.6$, $p < .05$, and with sentences as a random variable, $F(1, 39) = 27.8$, $p < .01$. Although targets had been arbitrarily assigned to the two groups in the experimental design, there were more errors for one group (5.3% vs. 9.1%). This difference was significant with materials as a random variable, $F(1, 39) = 14.9$, $p < .01$, and approached significance with subjects as a random variable, $F(1, 12) = 3.9$, $p = .07$. There was no interaction between the matching effect and groups of test sentences, Fs < 1.

Control Experiment

We would like to think that the targets that matched the studied paragraphs in meaning had faster response latencies because the meaning of the whole sentence matched the meaning in the paragraph. However, an alternative interpretation is that the matching effect depended only on the predicate of the test sentence. For example, the target *tomatoes are red* might have been responded to faster after the paragraph about painting because the paragraph mentioned words related to the predicate *are red*, words such as *color* and *painting*. To counter this alternative interpretation, a control experiment was performed.

This experiment used new versions of the paragraphs of Experiment 1, but the same test sentences. An example is shown in Figure 2. In the new versions of the paragraphs, the critical noun (e.g. *tomatoes*) was mentioned and the words related to the predicate of the test sentences were mentioned, but these words did not refer in any way to the critical noun. Thus, the predicates of the target test sentences matched or mismatched the paragraphs (depending on experimental condition), but the test sentences as whole sentences did not vary in the extent to which they matched the paragraphs. Except for this change in materials, the experiment was identical to Experiment 1.

The results showed no matching effect. When the predicates of the targets matched the paragraph that was read, mean response time was 1,225 ms (6% errors); when they did not match, 1,232 ms (8% errors), $Fs \leq 1.2$.

Clearly, the matching effect found in Experiment 1 was not the result of matching predicates. It is the whole sentence, *tomatoes are red*, that is important for the matching effect, not just the predicate, *something is red*.

Discussion

Experiment 1 shows that it is easier to verify a property of a noun in a context in which the property is relevant than in a context in which it is not relevant. This is the same conclusion reached by Barsalou (1982), Tabossi and Johnson-Laird (1980), and Tabossi (1982) in previous work. The current results extend theirs in two ways. First, it is somewhat less likely in the current experiment that the context-relevance effect was due to strategic processes. This is because a variety of test sentences was presented, including sentences about information explicitly stated in the paragraph and sentences not about property information, so that subjects would be less likely to try to anticipate a particular kind of test sentence. Also, subjects were required to read context paragraphs more quickly in the current experiment (approximately 3 words per second). Second, the control experiment shows that relevant context facilitates responses not to property information alone (e.g., something being red) but to property information tied to the appropriate noun (e.g., *tomatoes are red*). This control had not been included in the previous experiments.

The results of Experiment 1 validate the materials to be used in Experiments 2–5. The particular aspects of the meanings of the nouns chosen for use in Experiment 1 were verified more quickly in the presence of a relevant paragraph. Thus, these aspects of meaning are the kind of general knowledge that might be included in the representation in memory of a relevant text but not included in the representation of an irrelevant text.

Version 1

While eating his bacon and tomato sandwich, the painter searched for the colors most suited to the still life painting of the rustic English countryside.

Target Test Sentences

Matching: Tomatoes are red. (True by general knowledge)
Mismatching: Tomatoes are round. (True by general knowledge)

Filler Test Sentences

The painter was eating. (True by paragraph)
The painting would be a pencil sketch. (False by paragraph)

Version 2

While eating her bacon and tomato sandwich, the child psychologist watched the infant play with her toys. The little girl found a ball to roll across the floor with her nose.

Target Test Sentences

Matching: Tomatoes are round. (True by general knowledge)
Mismatching: Tomatoes are red. (True by general knowledge)

Filler Test Sentences

The psychologist was eating. (True by paragraph)
The little girl just sat quietly. (False by general knowledge)

Figure 2. Example of Paragraphs Used in Control Experiment.

Experiment 2

To test whether relevant aspects of meaning were encoded into memory, a study paragraph was separated from its test sentences by other paragraphs, other test sentences, or both. Test sentences were not presented immediately after a paragraph, but were delayed until a total of three paragraphs had been read.

Subjects were presented with a series of study–test trials. On each trial, they read three paragraphs and then verified 11 test sentences. Among the test sentences were those critical to testing the experimental hypotheses. These target sentences either matched or mismatched the aspect of meaning of a noun mentioned in one of the paragraphs (see Figure 3). In addition, each of these sentences was immediately preceded in the test list by a priming test sentence. A priming sentence was either a true sentence from the paragraph that had mentioned the noun or a sentence not related to any of the paragraphs but true according to general knowledge. For example, in Figure 3, the proposition *tomatoes are red* is not explicitly stated in either paragraph, but it matches the aspect of meaning of tomatoes in the first paragraph. It is primed in the test list either by a sentence from the paragraph (the *painter* sentence in the matching condition, the *psychologist* sentence in the mismatching condition) or by the general knowledge sentence about newspapers.

This experiment was designed to test two hypotheses. The first was that aspects of meaning relevant to a text are included in the memory representation of that text so that matching targets have faster response times than mismatching targets. They would be faster perhaps because the matching information was represented in a new trace in the mental representation of the text in addition to the already existing general knowledge trace, or because the existing trace was tagged as information connected to the text.

The second hypothesis was that the kind of information included in the decision about a target test sentence depends on retrieval context, where the retrieval context is determined by the immediately preceding test sentence. If the preceding test sentence is a test sentence about newly learned information, then newly learned information will tend to be included in the decision about the target. This tendency will lead to a matching effect because, although both matching and mismatching targets include newly learned information (the concept *tomatoes*), only the predicate of the matching target has

Matching Version

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.

Target Test Sentence

Tomatoes are red. (True)

Priming Test Sentences

The still life would require great accuracy. (True)
Newspapers are reading material. (True)

Filler Test Sentences

The painter searched for many days. (True)
Balloons are heavy. (False)

Mismatching Version

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

Target Test Sentence

Tomatoes are red. (True)

Priming Test Sentences

The child psychologist watched the infant. (True)
Newspapers are reading material. (True)

Filler Test Sentences

The little girl played with her toys. (True)
Balloons are heavy. (False)

Figure 3. Example of Paragraphs Used in Experiment 2.

been encoded in memory with the text. In contrast, if the preceding test sentence prime is a test sentence true by general knowledge, then the decision process on the target will tend to be based on general knowledge. To the extent that general knowledge is verified independently of recently studied information, then the matching effect should be reduced.

Putting the two hypotheses together gives an interaction between type of priming test sentence and type of target. With a prime from newly learned information, a matching effect is expected; but with a prime from general knowledge, the matching effect should be reduced or absent. Thus, the prediction is that the matching effect will depend on retrieval context.

Method

Subjects. The subjects were 24 Yale undergraduates who participated in the experiment for extra credit in an introductory psychology course.

Materials. The same 40 pairs of paragraphs that were used in Experiment 1 were used in this experiment. However, for each pair of paragraphs, only one (arbitrarily chosen) critical target test sentence was used. The meaning in the chosen target sentence matched the meaning in one of the paragraphs but mismatched the meaning in the other paragraph. For each target, there were three possible priming test sentences (see the example in Figure 3). One of the priming sentences was true according to one of the paragraphs of the pair, another was true according to the other paragraph of the pair, and the third was true according to general knowledge. This last priming sentence contained only concepts that had never been mentioned in any paragraph in the experiment. For each paragraph, there were also two other filler test sentences; these were true or false according to information in the paragraph or false according to general knowledge. These filler test sentences did not repeat any of the words or concepts mentioned in the targets or the priming sentences.

There were 20 filler paragraphs, each with three test sentences; these were the same as were used in Experiment 1.

Over the course of the experiment (excluding practice), a subject read one version of each of the 40 pairs of paragraphs and all 20 of the filler paragraphs, for a total of 60 paragraphs. Each of the 40 experimental paragraphs was tested with the target test sentence, which was true according to general knowledge. Of these targets, 20 were primed with a sentence true according to the paragraph that was read, and 20 were primed with a sentence true according to general knowledge. Filler test sentences for the experimental paragraphs and test sentences for the filler paragraphs combined to yield an additional 50 test sentences true according to a paragraph that was read, 50 test sentences false according to a paragraph that was read, and 40 test sentences false according to general knowledge.

Procedure. Each subject was tested in one 50-min session. Presentation of all materials and collection of data was controlled by a microcomputer driven by an Apple computer.

Each subject was presented with 2 practice trials and 20 experimental trials. Each trial began with an instruction to the subject to press the spacebar on a CRT keyboard to initiate the trial. Then, three paragraphs were presented (two experimental and one filler), one at a time, for 8 s each. After the third paragraph, a row of asterisks was displayed for 1 s to warn the subject that the test list was about to start. A test list consisted of 11 sentences, presented one at a time. Each test sentence remained on the CRT screen until the subject responded "true" or "false"; it then disappeared from the screen and the next sentence appeared after a 50-ms pause. After the eleventh test sentence, the instruction to press the space bar to begin the next paragraph appeared.

Subjects were given the same instructions about responding to the test sentences as in Experiment 1. After incorrect responses, the word ERROR was presented for 2 s before presentation of the next test sentence.

Design. Three factors were combined in a Latin square design. One of these factors was groups of subjects (four groups with 6 subjects per group) and another was sets of pairs of paragraphs (four sets with 10 pairs per set). The third factor was the four experimental conditions: a target test sentence was primed either by a sentence true according to the paragraph that was read or by a sentence true according to general knowledge, and the target either matched or mismatched the paragraph that was read.

A different random order of presentation of paragraphs and test sentences was used for every 2 subjects. Order of presentation of test sentences was subject to two constraints: a target test sentence could not appear earlier than the third position in the test list, and the filler test sentences for the experimental paragraphs could not appear earlier in the test list than the target. No test sentence was presented more than once in the test list.

Results

All analyses were based on mean correct response times for each subject or each test sentence in each condition. Means of these means are shown in Table 1. For the target test sentences, only correct responses preceded by correct re-

sponses to the priming test sentence were included in the analyses in an effort to ensure that both were in memory.

As shown in Table 1, if a target was primed by a sentence from the paragraph that was read, response time depended on whether the target matched the paragraph that was read. But if the target was primed by a sentence true according to general knowledge, there was little matching effect. For response times, this interaction was significant with subjects as a random variable, $F(1, 23) = 4.6$, $p < .05$, and with test sentences as a random variable, $F(1, 39) = 6.2$, $p < .05$. Analysis of variance also showed that the effect of matching was significant with subjects as a random variable, $F(1, 23) = 11.7$, $p < .01$, but not with test sentences as a random variable, $F(1, 39) = 2.7$, $p > .10$. The effect of which kind of prime was used (from the paragraph or general knowledge) was significant in both analyses, $F(1, 23) = 14.3$, $p < .01$, with subjects, and $F(1, 39) = 16.4$, $p < .01$, with sentences. Average standard error of the means was 16.8 ms. For errors, analysis of variance showed that target test sentences were responded to less accurately when they did not match the paragraph read, $F(1, 23) = 5.1$, $p < .05$, with subjects as the random variable and $F(1, 39) = 19.1$, $p < .05$, with test sentences as the random variable. There were no other significant effects for errors on the targets.

Priming test sentences true according to general knowledge were responded to more quickly than priming test sentences true according to the paragraph that was read; the means were 1,057 ms (3% errors) and 1,317 ms (3% errors).

For all filler test sentences (including the priming test sentences), mean response times were 1,012 ms (3% errors) for sentences true by general knowledge, 1,160 ms (7% errors) for sentences false by general knowledge, 1,170 ms (4% errors) for sentences true according to a paragraph, and 1,280 ms (19% errors) for sentences false according to a paragraph.

Discussion

Experiment 2 showed the predicted pattern of results. If the retrieval context was newly learned information, responses on targets showed a matching effect, indicating that relevant aspects of meaning were encoded with the appropriate text. If the retrieval context was general knowledge, then the matching effect was much reduced (and nonsignificant).

We think of the matching effect as facilitation, but we have no neutral priming condition in the experiment, so we cannot be sure that we have observed facilitation in the matching condition and not inhibition in the mismatching condition. But whichever the effect actually is, we would attribute it to information that is encoded in the memory representation of the studied text (e.g. information about the color of tomatoes in the "painting" text), so we refer to the effect simply as a matching effect.

Responses to targets were much faster when the prime tested general knowledge information than when it tested newly learned information. We suspect that this difference reflects the difficulty of the material; the newly learned information was not very well learned. There may be other differences between the general knowledge and newly learned primes, such as the degree of specificity of the concepts. However, these other differences would have the same effect as difficulty, namely, to change the kind or amount of information included in the decision on the target in such a way as to increase response times. Also, it should be noted that newly learned information is not necessarily much more difficult than general knowledge information; intensive study could make new information easier than some kinds of well-learned information. However, in the experiments in this article, the general knowledge test sentences were all easy, and the newly learned information was read only once.

The effect of retrieval context found in Experiment 2 is reminiscent of similar effects demonstrated by Potts and Peterson (1985); see also Potts & St. John (1986). They manipulated context by the overall content of test lists, rather than on an item-to-item basis the way it was done in Experiment 2, and they used newly learned information for target sentences rather than information already known to the subjects. They interpreted their results as reflecting the compartmentalization of information in memory: New information is separated in memory from general knowledge. Our interpretation is somewhat different; instead of separating facts in memory, we stress the flexibility of the decision process. The information involved in a decision may vary, depending on the kind of information required to produce a response and on the retrieval context.

This explanation of the results of Experiment 2 in terms of retrieval context leads to consideration of a third priming condition, one in which the prime is a sentence true according to a paragraph other than the paragraph relevant to the target. Given the hypothesized retrieval context effect, this prime should cause newly learned information to be included in the decision about the target, and so should lead to a matching effect on the target. Experiments 3 and 4 found such an effect.

Table 1
Reaction Time (RT; in Milliseconds) and Percentage of Error (%E): Experiments 2–4

	Prime					
	Same paragraph		Different paragraph		General knowledge	
Condition	RT	%E	RT	%E	RT	%E
Experiment 2						
Matching	1,155	8			1,049	5
Mismatching	1,265	12			1,065	8
Experiment 3						
Matching	1,167	1	1,157	3	1,145	1
Mismatching	1,235	2	1,256	9	1,162	3
Experiment 4						
Matching	1,074	2	1,124	4	1,082	1
Mismatching	1,143	4	1,185	6	1,110	2

Experiments 3 and 4

The procedure for these experiments was the same as for Experiment 2; on each trial, subjects read a list of paragraphs and then verified a series of test sentences. Target test sentences either matched or mismatched the meaning of a studied paragraph, and they were primed by one of three kinds of

priming test sentences: a sentence from the same studied paragraph, a sentence from a different studied paragraph, or a sentence from general knowledge.

When newly learned information enters the decision process on the target, that is, when the prime is from a studied paragraph, there should be a matching effect. If one of the studied paragraphs matches the target in meaning, then response times for the target should be faster than if none of the studied paragraphs matches the target in meaning. This matching effect should be independent of whether the prime is from the same or a different paragraph than the target, because either prime will invoke newly learned information. In contrast, when the prime is true according to general knowledge, the decision on the target should rely less on newly learned information and the matching effect should be reduced.

Although the matching effect should be independent of whether the prime came from the same or a different paragraph as the target, overall response times, for matching and mismatching targets, may be faster with the prime from the same paragraph. Same-paragraph primes should be more closely connected in memory to target-relevant information than different-paragraph primes. With some delay between study and test, these connections should lead to faster response times for the target (McKoon & Ratcliff, 1980a, 1980b, 1981; Ratcliff & McKoon, 1978, 1981a, 1981b; Seifert, McKoon, Abelson, & Ratcliff, 1986). In Experiment 3, there were only three paragraphs to study and only 13 test sentences, so on average, test sentences followed study relatively quickly, and same-paragraph primes did not lead to faster response times than different-paragraph primes. In Experiment 4, the delay was longer and the amount of intervening material was greater, with six paragraphs to study and 26 test sentences, and same-paragraph primes did lead to faster response times.

Method

Subjects. The subjects were 24 Yale undergraduates participating for extra credit in an introductory psychology course.

Materials. In general, the materials were the same as in Experiment 2, except that the number of experimental pairs was increased to 48 and the number of filler paragraphs was increased to 24. Over the course of the experiment (excluding practice), a subject read one member of each of the pairs and all of the fillers, and was tested on 96 test sentences true and 48 test sentences false according to general knowledge, and 108 test sentences true and 60 test sentences false according to paragraphs that were read.

Procedure. The procedure was basically the same as in Experiment 2. In Experiment 3, there were 24 trials (plus 2 practice trials) and 13 test sentences on each trial. In Experiment 4, the length of the study and test list was doubled: On each of 12 trials there were 6 paragraphs to study and 26 test sentences to verify.

Design. In each experiment, three variables were combined in a Latin square design, with six groups of subjects (4 subjects per group), six sets of materials (eight pairs of paragraphs per set), and six experimental conditions. The six conditions represented the type of prime (sentence from same paragraph as target, sentence from different paragraph, or sentence from general knowledge) crossed with type of target sentence (matching or mismatching).

The order of presentation of paragraphs and test sentences was determined in the same way as for Experiment 2.

Results

The data were analyzed as in Experiment 2, and means are shown in Table 1.

In Experiment 3, the results predicted by the retrieval context hypothesis were obtained. When a target test sentence was primed by a sentence from a paragraph that had just been studied, responses for matching sentences were faster than responses for mismatching sentences. This was true whether the prime was from the same paragraph as the target or from a different paragraph. But if a target was primed by a sentence true according to general knowledge, response times for matching and mismatching sentences were virtually identical. This interaction was significant with subjects as the random variable, $F(2, 36) = 3.7$, $p < .05$, and with items as the random variable, $F(2, 84) = 4.9$, $p < .01$. There were also significant main effects, between the matching and mismatching conditions, $F(1, 18) = 23.4$, $p < .01$, and $F(1, 42) = 17.6$, $p < .01$, and among the priming conditions, $F(2, 36) = 5.4$, $p < .01$, and $F(2, 84) = 3.7$, $p < .05$. The standard error of the means was 13.5 ms. For error rates, the main effect of priming condition was significant, $F(2, 36) = 6.8$, $p < .01$, and $F(2, 84) = 5.3$, $p < .01$, and the main effect of matching versus mismatching was significant, $F(1, 18) = 5.8$, $p < .01$, and $F(1, 42) = 6.3$, $p < .01$. The interaction was not significant, $Fs < 2.4$.

With respect to the matching effect, Experiment 4 replicated Experiment 3; response times for matching test sentences were faster than response times for mismatching test sentences when they were primed by a sentence from a studied paragraph but not when they were primed by a sentence true according to general knowledge.

The only difference between Experiments 3 and 4 was in the same-paragraph versus different-paragraph prime conditions. In Experiment 3, the delay between study and test was relatively short, and the same-paragraph primes did not lead to faster response times than did the different-paragraph primes. However, in Experiment 4, with a longer average delay between study and test, same-paragraph primes did lead to faster response times.

In Experiment 4, the interaction between matching and type of prime did not reach significance, but because this interaction had been found in both Experiments 2 and 3 (and other unpublished experiments), post hoc tests were performed. These showed a significant matching effect for the primes from studied paragraphs, but not for general knowledge primes: For example, for the different paragraph priming condition, $F(1, 36) = 4.6$, $p < .05$, with subjects as the random variable, and $F(1, 84) = 4.8$, $p < .05$ with materials as the random variable; for the general knowledge priming condition, $F(1, 36) = 0.98$ and $F(1, 84) = 1.0$.

The main effect of matching was significant, $F(1, 18) = 11.6$, $p < .01$, with subjects as the random variable, and $F(1, 42) = 11.8$, $p < .01$, with materials as the random variable.

Primes from the same paragraph as the target sentence led to faster response times than primes from a different paragraph. The main effect of type of prime was significant, $F(2, 36) = 6.1$, $p < .01$, with subjects as the random variable, and $F(2, 84) = 5.1$, $p < .01$, with materials as the random variable.

A post hoc test showed the expected difference between same- and different-paragraph primes to be significant, $F(1, 36) = 7.6$, $p < .01$, with subjects as the random variable, and $F(1, 84) = 6.6$, $p < .01$, with materials as the random variable. Standard error of the response time means was 20.0 ms.

The only significant effect in the error analyses was the main effect of priming condition in the analysis using subjects as the random variable, $F(2, 36) = 4.0$, $p < .05$; all other Fs were less than 2.4.

For Experiment 3, results for filler test sentences were as follows: true by newly-learned information, 1,285 ms and 3.6%; false by newly learned information, 1,408 ms and 19.8%; true by general knowledge, 1,091 ms and 2.9%; false by general knowledge, 1,289 ms and 8.0%. For Experiment 4, the results were as follows: true by newly learned information, 1,247 ms and 4.8%; false by newly learned information, 1,406 ms and 25.6%; true by general knowledge, 1,056 ms and 4.2%; false by general knowledge, 1,251 ms and 10.2%.

Discussion

The results of Experiments 3 and 4 confirmed predictions based on the hypotheses that relevant aspects of the meaning of a noun are encoded into the mental representation of a text, and that the decision process for a test sentence can be based on different kinds of information, depending on the type of priming test sentence that immediately precedes it. When the prime was from a studied paragraph, matching targets were verified faster than mismatching targets, and when the prime was from general knowledge, matching targets were not verified significantly faster than mismatching targets.

It is important to distinguish the three separate effects found in Experiments 3 and 4. The first is the effect of retrieval context, which determines the extent to which newly learned information will be included in the decision process on the target. The second is the effect of matching; matching targets are integrated into the studied text and therefore verified faster than mismatching targets. Finally, there is an effect that is due to within-paragraph connections in encoded information; a prime from the same paragraph as the target leads to facilitation on the target relative to a prime from a different paragraph.

Combining these three effects leads to the interactions shown in Table 1. A matching effect is obtained, but only in the appropriate retrieval context. Within the appropriate context, the matching effect is independent of type of prime (same or different paragraph) as it should be if matching information is encoded with the studied texts. However, same-paragraph primes do facilitate both matching and mismatching primes relative to different paragraph primes because both kinds of targets include information from the studied paragraph (e.g., the word *tomatoes*).

The results of Experiments 3 and 4 allow an alternative possibility about the locus of the matching effect to be rejected. We have assumed that the matching effect represents processes that occur at the time of encoding of the text. An alternative possibility would be that the matching effect occurs at test time. In this case, presentation of a prime from the same studied paragraph as the target would activate the paragraph, and in conjunction with the activated paragraph, a matching target would be easier to process than a mismatching target. However, this possibility can be rejected because a matching effect is also obtained with a prime from a different studied paragraph. A prime from a different paragraph would not be expected to activate the appropriate paragraph for the target, and so could not provide a context in which the target was easier to process. Thus, the locus of the matching effect must be at encoding; during reading, the aspects of the meaning of a word that are appropriate for the text are encoded into memory with the text.

Further predictions can be generated about matching and retrieval context, and these were tested in Experiments 5 and 6. Experiment 5 tested whether a relevant aspect of meaning (such as *tomatoes are red* for the painting paragraph) would prime a test sentence that had been explicitly stated, as it should do if it was encoded into the mental representation of the paragraph. Experiment 6 tested whether the retrieval context effects found with targets that were never explicitly stated in the studied paragraphs are also found with targets that are explicitly stated.

Experiment 5

In Experiment 5, test sentences representing matching and mismatching aspects of meaning were used as primes for test sentences representing information explicitly stated in studied paragraphs. The procedure was the same as that used in Experiments 2, 3, and 4, except that the target and prime were reversed. The target was a sentence that had been explicitly stated in a studied paragraph and the prime was not explicitly stated in the studied paragraph. Examples of the materials are shown in Figure 4.

In this experiment, the primes represented implicit information. In the previous experiments, responses to such test sentences were not necessarily affected by newly learned information. That is, they were not affected by newly learned information unless the previous test sentence required the use of newly learned information. This result might suggest that in Experiment 5, the primes would not engage newly learned

Paragraph:

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.

Target Test Sentence

The still life would require great accuracy. (True)

Condition 1: Matching Prime
Tomatoes are red. (True)
Condition 2: Mismatching Prime
Tomatoes are round. (True)
Condition 3: Explicitly Stated Prime From Same Paragraph
The painter searched for many days. (True)
Condition 4: Explicitly Stated Prime From Another Studied Paragraph

Figure 4. Example of Paragraphs Used in Experiment 5.

information and so they would not facilitate responses to target test sentences about newly learned information. However, predictions about the effects of these primes require a specific model of priming.

There are two general classes of models available, those based on spreading activation (e.g., Anderson, 1983; Collins & Quillian, 1969; Meyer & Schvaneveldt, 1976) and the retrieval theory proposed by Ratcliff and McKoon (in press). From a spreading activation theory, it is difficult to see why the primes representing implicit information should facilitate responses to explicitly stated information. Because these prime sentences do not, by themselves, access newly learned information (as shown by the results of Experiments 2, 3, and 4), they would not be expected to prime newly learned information. *Tomatoes are red* would not be expected to prime the explicit information *The still life would require great accuracy.*

The retrieval theory (Ratcliff & McKoon, in press) makes a different prediction. According to this theory, priming effects result from the formation of a compound cue made up of prime and target. The response to the target reflects the joint strength or familiarity of the compound; if prime and target are associated in memory, then the strength will be greater and so the response will be facilitated. In contrast to spreading activation, the prime does not affect long-term memory in advance of presentation of the target; instead the prime interacts with the target in short-term memory to form the compound cue. According to this theory, the primes representing implicit information should facilitate responses to their targets to the extent that the implicit and explicit information are associated in memory. Whether the implicit information is presented as prime and the explicit information as target, in Experiment 5, or the reverse (Experiments 2, 3, and 4), priming effects should still be obtained. The only difference between implicit information as a prime and implicit information as a target is that a prime is given somewhat less weight in the compound cue than a target (because the response is actually made to the target).

Experiment 5 also included two control conditions, in which both the primes and targets were explicitly stated in studied paragraphs. In one of the control conditions, the prime was an explicitly stated sentence from the same paragraph as the target, and in the other control condition, it was an explicitly stated sentence from a different studied paragraph. Examples are shown in Figure 4. The same-paragraph prime was expected to lead to faster response times for the target than was the different-paragraph prime. These control conditions were designed to show that the experiment had sufficient power to detect differences in target response times attributable to priming.

Method

Subjects. The subjects were 16 Yale undergraduates who participated for extra credit in an introductory psychology course.

Materials. There were 48 experimental paragraphs, one member of each of the pairs used in the preceding experiments. For each paragraph, the target sentence was the sentence used as a prime in the preceding experiments. There were three possible primes for this target: the matching or the mismatching sentence used in the preceding experiments, or a sentence from the paragraph. For each experimental paragraph, there was also one filler statement true according to general knowledge and one filler statement false according to general knowledge. The same filler paragraphs were used as in the previous experiments.

Overall (excluding practice), a subject was tested on 84 sentences true according to a studied paragraph, 66 sentences false according to a studied paragraph, 72 sentences true according to general knowledge, and 66 sentences false according to general knowledge.

Procedure. The procedure was the same as in Experiment 4. There were 12 trials (plus 2 for practice) with 6 paragraphs studied on each trial and 24 sentences tested on each trial.

Design. Three factors were combined in a Latin square: four groups of subjects (4 subjects per group), four sets of materials (12 paragraphs per set), and four experimental conditions. Across the experimental conditions, the target was always the same: a sentence true and explicitly stated in the paragraph. The prime was the matching sentence, the mismatching sentence, a true explicit sentence from the same paragraph as the target, or a true explicit sentence from a different paragraph.

Of the six paragraphs studied on a trial, four were experimental paragraphs (presented in randomly chosen positions) and two were fillers. A test list was constructed by placing the target for each experimental paragraph in a randomly chosen position (but not in the first position), and then placing the prime appropriate for its condition in the immediately preceding test position. Then filler sentences were placed randomly in the remaining test positions, under the same constraints as in Experiment 2.

Results

The data were analyzed as in Experiment 2. For the primes true according to studied paragraphs (sentences explicitly stated in the paragraphs), a prime from the same paragraph led to faster response times for the target than a prime from a different paragraph; response times averaged 1,288 ms (6.3% errors) and 1,352 ms (7.1% errors). Similarly, for the primes true according to general knowledge (sentences stated only implicitly in the paragraphs), a matching prime led to faster response times than a mismatching prime: 1,316 ms (11% errors) and 1,363 ms (6.5% errors), respectively. Two factors were entered into an analysis of variance: The primes were either from general knowledge or studied paragraphs and the primes were either appropriate (same-paragraph or matching) or inappropriate (different-paragraph or mismatching) for the target. The only significant effect was the advantage given to response times on the targets by the appropriate primes, $F(1, 15) = 6.8$, $p < .05$, with subjects as the random variable, and $F(1, 44) = 6.3$, $p < .05$, with materials as the random variable. The standard error of the means was 30 ms. There were no significant effects in the analyses of errors.

Response times and error rates for fillers and primes were as follows: sentences true by general knowledge (not including primes), 1,117 ms and 6.0% errors; matching primes true by general knowledge, 1,177 ms and 3.1% errors; mismatching primes true by general knowledge, 1,249 ms and 11.5% errors; sentences true according to studied paragraphs, 1,333 ms and 8.8% errors; sentences false according to studied paragraphs, 1,293 ms and 8.9% errors; and sentences false according to general knowledge, 1,524 ms and 25.5% errors.

Discussion

The results of Experiments 2, 3, and 4 were interpreted as supporting the hypothesis that relevant aspects of meaning are incorporated into the memory representation of a text. If this hypothesis is correct, then according to the retrieval theory of priming proposed by Ratcliff and McKoon (in press), the implicit relevant information should prime information explicitly stated in the text. This is what was found in Experiment 5. Test sentences representing matching aspects of meaning speeded response times to target test sentences that had been explicitly stated.

Experiment 6

In accounting for the results of Experiments 2, 3, and 4, it was claimed that the kind of information used for the decision about one test sentence can affect the kind of information used for the decision about the next test sentence. When one decision process requires newly learned information from the studied paragraphs, the next decision process will also include newly learned information. In our experiments, newly learned information tends to make the decision process more difficult. Table 1 shows that response times to the target sentences were longer and error rates were higher when the prime was newly learned information. This was expected, given that the paragraphs were read only once, and that the general knowledge sentences are comparatively easy.

Table 2 shows the same effect of newly learned information on responses to filler test sentences. These are post hoc analyses of filler sentences from Experiments 3, 4, and 5, excluding data for those of the filler sentences true according to studied paragraphs that were preceded by a test sentence from the same studied paragraph. For filler sentences true by general knowledge, the same effect is shown as for the targets in Table 1: longer response times and higher error rates when the preceding test sentence was newly learned information.

Table 2
Response Times (RT; in Milliseconds) and Percentage of Error (%E): Retrieval Context Effects for Experiments 3–6

	Target			
	Studied paragraph		General knowledge	
Prime	RT	%E	RT	%E
Experiment 3[a]				
Studied paragraph	1,267	2.8	1,139	3.1
General knowledge	1,259	4.5	1,020	0.3
Experiment 4[a]				
Studied paragraph	1,206	4.6	1,082	5.2
General knowledge	1,204	1.4	993	2.6
Experiment 5[a]				
Studied paragraph	1,288	7.5	1,092	6.0
General knowledge	1,218	5.8	1,027	1.7
Experiment 6				
Studied paragraph	1,292	5.4	1,115	1.5
General knowledge	1,248	4.0	1,021	4.6

[a] Post hoc analysis.

For filler sentences true according to newly learned information, the decision process was already difficult (and perhaps subject to a ceiling effect), but even for these test sentences, averaged across the experiments shown in Table 2, response times were longer and error rates were higher when they were preceeded by other test sentences true according to newly learned information.

The data shown in Table 2 represent post hoc analyses of test sentences that were uncontrolled for test position. Also, the materials were not well controlled in that some of the items true by general knowledge were sentences that included nouns from the studied paragraphs (e.g., *Tomatoes are red*). Experiment 6 was designed to investigate the issue of retrieval context effects when this was not the case, when all test sentences true according to general knowledge did not include any nouns mentioned in the paragraphs. The conditions of the experiment were those shown in Table 2.

Method

Subjects. The subjects were 12 Yale undergraduates participating for extra credit in an introductory psychology course.

Materials. The same materials were used as in Experiment 3. For the priming and target test sentences from studied paragraphs, the priming sentences from Experiment 3 were used (e.g., in Figure 3, *The still life required great accuracy*). For the priming and target test sentences true according to general knowledge, the fillers from Experiment 3 were used (e.g., *Villages are small*). None of these contained any content words from studied paragraphs. Overall (excluding practice), a subject was tested on 60 sentences true according to general knowledge, 48 sentences false according to general knowledge, 72 sentences true according to studied paragraphs, and 60 sentences false according to studied paragraphs.

Procedure. The procedure was the same as in Experiment 2. There were 24 trials, with three paragraphs for study on each trial (two experimental and one filler) and 10 test sentences.

Design. There were four experimental conditions: A target test sentence was true either according to a studied paragraph or general knowledge, and a target and its prime were either both true by the same criterion or true by different criteria. When a prime and target were both from studied paragraphs, they were never from the same paragraph. The four conditions were combined in a Latin square with four groups of subjects (3 per group) and four sets of materials (12 per set).

On each trial, both the paragraphs in the study list and the sentences in the test list were ordered randomly, except that test sentences in the experimental conditions were not placed in first or second position in the test list. Otherwise, study and test lists were subject to the same constraints as in Experiment 2.

Results

For the four conditions of interest, means were calculated for those correct responses that were preceded by correct responses. Means of these means are shown in Table 2. The post hoc analyses of the other experiments shown in Table 2 also include only correct responses preceded by correct responses.

The data in Table 2 show a retrieval context effect for both test sentences true by general knowledge and test sentences true by newly learned information. Test sentences preceded

by a test sentence true by newly learned information have slower response times than test sentences preceded by a test sentence true by general knowledge. However, this retrieval context effect was large enough to be significant for only target test sentences true by general knowledge. The two factors examined by analyses of variance were the truth criterion for the target (general knowledge or newly learned) and whether the prime did or did not have the same truth criterion as the target. The main effect for the targets was significant; general knowledge targets had faster response times than newly learned targets, $F(1, 11) = 107.8$, $p < .01$, with subjects as the random variable and $F(1, 44) = 31.8$, $p < .01$ with materials as the random variable. There was no significant main effect of whether the prime and target were true by the same criteria ($Fs < 1.9$), but the interaction of the two variables was significant, $F(1, 11) = 10.3$, $p < .01$ and $F(1, 44) = 6.1$, $p < .01$. Post hoc analyses based on the interaction showed a significant sequential effect for the targets true by general knowledge, $F(1, 11) = 5.3$, $p < .05$, and $F(1, 44) = 6.8$, $p < .01$, but not for the targets true by newly learned information, $Fs < 1.5$. Standard error of the response time means was 29 ms. Analyses of variance of the error rates showed no significant effects, all $Fs < 2.4$.

Response times and error rates for filler test sentences were as follows: sentences true by newly learned information, 1,275 ms and 7.6% errors; sentences false by newly learned information, 1,383 ms and 23.8% errors; sentences true by general knowledge, 1,103 ms and 4.3% errors; and sentences false by general knowledge, 1,235 ms and 7.4% errors.

Discussion

The results of Experiment 6 were exactly as predicted from the retrieval context hypothesis, and as would be expected from the post hoc analyses of the earlier experiments. Target sentences true according to general knowledge had slower response times when preceded by test sentences true by newly learned information than when preceded by test sentences true according to general knowledge. There was also a retrieval context effect for targets true by newly learned information, but it was not significant. This may be because of a ceiling effect on response times or because newly learned information was already included in the decision process on the target to such an extent that no more could be added by an immediately preceding test sentence.

Summary and Discussion

The first experiment found a matching effect in verification of information about the meaning of a noun, when the information was presented for verification immediately after reading a text that mentioned the noun. When the tested information was relevant to the text, it was verified more quickly than when it was not. This matching result confirmed earlier findings (Barsalou, 1982; Tabossi, 1982; Tabossi & Johnson-Laird, 1980), but also extended them in important ways, by showing that faster verification times for relevant information could be obtained when strategic processing was unlikely and that faster verification times were tied to information specific to the target noun and were not due to facilitation of general predicate information.

It is important to realize that several different processes may be operating when information to be verified is presented immediately following a text, as was done in Experiment 1. Faster response times for contextually relevant information might be due to processes occuring during reading of the text information, during interpretation of the meaning of the information to be verified, or during integration of the information to be verified with the text. None of these processes necessarily require that contextually relevant information is encoded into memory during reading of the text.

Experiments 2–5 were designed to examine whether implicit information about the meanings of words is encoded into the memory representation of a text differentially. The results show that relevant information is encoded into memory, *but* it is not accessed under all retrieval conditions. Only in an appropriate retrieval environment will there be differential effects of relevant information over irrelevant information.

When the retrieval environment is appropriate, Experiments 2–5 show that relevant information about the meanings of words is verified faster than information that is not relevant. This matching effect is demonstrated in Table 1: the sentence *tomatoes are red* is verified faster when it matches the aspect of tomatoes relevant to a studied text (a text about the color of tomatoes) than when it does not match (a text about the shape of tomatoes).

In Experiments 2–4, the retrieval environment was determined by the test sentence that immediately preceded the matching or mismatching target test sentence. The preceding sentence tested either newly learned information from a studied paragraph or information from general knowledge. The idea was that the kind of information used on one decision would affect the kind of information used on the next decision. So the decision process on a target test sentence would be more likely to depend on newly learned information if it followed a test of newly learned information than if it followed a test of general knowledge.

Experiments 2–4 showed support for this idea about retrieval context in two ways. First, the matching effect on target test sentences was obtained only when the test sentence preceding the target required the use of newly learned information. This use of newly learned information led to the inclusion of newly learned information into the decision process about the target, and because this information included aspects of meaning relevant to the decision about the target, a matching effect was obtained. When the sentence that preceded the target did not require the use of newly learned information, the information was less likely to be used in the decision about the target and the matching effect was reduced to insignificance.

The second way in which the experiments demonstrated effects of retrieval context is shown in Table 2 and Experiment 6. With our materials, newly learned information was more difficult to verify than information from general knowledge. Therefore, test sentences that required the use of newly learned information tended to have slower and less accurate

responses. They also led to the inclusion of newly learned information into decisions about immediately following test sentences, so responses to these sentences were slowed. This effect was significant when the following test sentences were true by general knowledge, but not when they were true according to new information. These sentences would already be difficult and so perhaps subject to a ceiling effect.

The two major conclusions from the experiments in this article concern the matching effect and the retrieval context effect. There are also several other results that should be noted. One is that Experiments 4 and 5 replicated previous work with priming and text structure (e.g., McKoon & Ratcliff, 1980b). A test sentence was verified faster if it was primed by another sentence from the same studied paragraph than by a sentence from a different paragraph. This result was true both for target test sentences representing information explicitly stated in the paragraph and for target test sentences in which only the noun had been explicitly stated (sentences like *tomatoes are red*). Differential priming was not found in Experiment 3, but it was argued that this was due to the short delay, the small amount of information intervening between study and test, or both.

Priming effects like those between sentences of a paragraph have been explained by two general theories: spreading activation and the retrieval theory proposed by Ratcliff and McKoon (in press). Experiment 5 gave support to the retrieval theory. Priming test sentences were sentences that expressed implicit information and that did not (by themselves) access newly learned information. According to spreading activation theories, these would not be expected to activate newly learned information and so would not be expected to give priming. In the retrieval theory, on the other hand, a prime enters a compound cue with the target, and the response to the target is determined by the overall familiarity of the compound. Priming effects would therefore be expected to reflect relations between prime and target, and this is what the data showed.

Another result of importance for interpretation of the matching effect is that the effect was obtained both with primes from the same studied paragraph as the target and with primes from a different paragraph. If the effect had been obtained only with same-paragraph primes, then it could be interpreted as due to retrieval processes. Perhaps the prime activated the paragraph, and in the context of the paragraph, the target was easier to process (even though it had not been encoded during reading). However, this interpretation can be ruled out because it cannot account for the matching effect obtained when the prime was from a different paragraph. Instead, the matching effect must be due to encoding processes. During reading, contextually relevant aspects of the meanings of words must be incorporated into the representation of the context in memory.

Retrieval Context

In this article, we argue that the information entering a decision process depends on the environment in which that decision is made. For example, in an environment that stresses newly learned information, more new information will enter the decision process, and in an environment that stresses general knowledge, less new information will enter the decision process. In this respect, the results found in the experiments presented here are analagous to earlier results on encoding specificity (Tulving & Thompson, 1973).

The typical encoding specificity procedure involves words studied as cue–target pairs, and testing of the targets for recognition and cued recall. The major finding is that a large proportion of targets can be recalled even though they cannot be recognized, so long as recall is in the context of the cue originally studied with a target. For example, suppose the pair *glue–chair* is studied: In the context of previously studied information (the cue *glue*), *chair* may be recalled. But without that context, *chair* alone may not be recognized as having been studied. Similarly, in our experiments, the relevant aspect of meaning of *tomatoes* is more available than the irrelevant aspect of meaning when the context is previously studied information, but not when the context is general knowledge.

Encoding specificity is sometimes interpreted to mean encoding of specific meanings (cf. Crowder, 1976). The meaning of *chair* presented alone is assumed to be slightly different from the meaning of *chair* presented in the context of the cue *glue*. The experiments reported here go beyond the usual encoding specificity result in showing explicitly that different aspects of meaning are encoded: the encoded meaning of *tomatoes* is shown to be different in the context of a paragraph about the color of tomatoes than in a context about the shape of tomatoes. However, the test sentence *tomatoes are red* does not by itself make available what was encoded about tomatoes in the paragraph about color. If this were the case, then the match between encoding and test sentence would be good in any retrieval environment. Instead, the match requires the appropriate aspect of meaning plus some other aspects of meaning encoded with the newly learned information and available when newly learned information is used in the decision process. This is not surprising; we would not expect to be able to capture in one test sentence exactly the right aspects of meaning or all of the aspects of meaning that were encoded about a concept during reading of a text.

This discussion has important implications for the study of inference processes. We often want to know whether some inference is encoded into the memory representation of a text during reading. According to the results of the present experiments, we cannot assume that we can test for that inference in isolation. Instead we must take into account the retrieval environment, because the ability of any test item to access items in memory is modified by all the other information available in the test situation (see also Tulving, 1974). A test item must be considered not alone as a single cue to memory, but as part of a larger cue made up of the test item plus other available information.

Local Minimal Coherence

The second main conclusion of this article concerns the contrast between contextually relevant meaning and other kinds of inferences. McKoon and Ratcliff (1986) found that inferences about highly predictable events were encoded only minimally and that readers did not encode inferences con-

necting the goals of characters in stories with later actions of the characters (McKoon & Ratcliff, 1988). Seifert et al. (1986) found that inferences about connections between thematically related stories were not encoded automatically. These results about the inferences that readers do not make suggest a local and minimal view of comprehension processes (McKoon & Ratcliff, 1988), which can serve as a null hypothesis against which to contrast models that propose large amounts of inference processing during reading. By the local, minimal view, subjects automatically infer only enough information to connect together the concepts explicitly stated in the text, and only connect together those concepts in relatively local proximity in the text. In the absence of special strategies (Seifert et al., 1986), readers make predictive inferences only minimally and they do not make inferences about global connections among the ideas of a text.

Intuitively, these inferences that readers do not make may seem more compelling than inferences about the meanings of words. For example, the inference that someone will die if they fall off a 14th story roof (McKoon & Ratcliff, 1986) seems more compelling than the inference that rolling a tomato has to do with tomatoes being round. Yet it is the inference about tomatoes that is encoded into memory more than the inference about death. The difference between these inferences does not involve textual coherence; neither are required to connect together explicitly stated ideas. For example, the meaning of tomato could simply be encoded as "an edible fruit"; no further meaning is needed to establish coherence between concepts or ideas. However, one way in which the inferences about relevant aspects of meaning are different from the other kinds of inferences is that they involve semantic information about specific concepts. Perhaps it is the well-learned nature of the semantic information involved in these inferences that allows them to be encoded. In this case, a local, minimal view of text representation would have to be modified to recognize the contribution of semantic information to inference processing.

This speculation points to directions for future research. To understand the inference processes that occur during reading will require careful analytic investigation of a variety of kinds of inferences, in order to determine the factors that govern the encoding of implicit information. At the same time, this investigation must cover a variety of retrieval environments, so that interactions between inference processes and retrieval cues can be mapped out.

References

Anderson, J. R. (1983). *The architecture of cognition.* Cambridge: Harvard University Press.

Barsalou, L. W. (1982). Context-independent and context-dependent information in concepts. *Memory & Cognition, 10,* 82–93.

Collins, A. M., & Quillian, M. R. (1969). Retrieval from semantic memory. *Journal of Verbal Learning and Verbal Behavior, 8,* 240–247.

Crowder, R. G. (1976). *Principles of learning and memory.* Hillsdale, NJ: Erlbaum.

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

McKoon, G., & Ratcliff, R. (1980b). Priming in item recognition: The organization of propositions in memory for text. *Journal of Verbal Learning and Verbal Behavior, 19,* 369–386.

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. (1988). *Inferences about the goals of characters in stories.* Manuscript in preparation.

Meyer, D. E., & Schvaneveldt, R. W. (1976). Meaning, memory structure, and mental processes. *Science, 192,* 27–33.

Posner, M., & Snyder, C. (1975). Attention and cognitive control. In R. L. Solso (Ed.), *Information processing and cognition* (pp. 55–85). Hillsdale, NJ: Erlbaum.

Potts, G. R., & Peterson, S. B. (1985). Incorporation versus compartmentalization in memory for discourse. *Journal of Memory and Language, 24,* 107–118.

Potts, G. R., & St. John, M. (1986, November). *Incorporation versus compartmentalization in memory for discourse.* Paper presented at the meetings of the Psychonomic Society.

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. (1981a). Automatic and strategic components of priming in recognition. *Journal of Verbal Learning and Verbal Behavior, 20,* 204–215.

Ratcliff, R., & McKoon, G. (1981b). Does activation really spread? *Psychological Review, 88,* 454–462.

Ratcliff, R., & McKoon, G. (in press). A retrieval theory of priming in memory. *Psychological Review.*

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

Tabossi, P. (1982). Sentential context and the interpretation of unambiguous words. *Quarterly Journal of Experimental Psychology, 34,* 79–90.

Tabossi, P., & Johnson-Laird, P. N. (1980). Linguistic context and the priming of semantic information. *Quarterly Journal of Experimental Psychology, 32,* 595–603.

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

Tulving, E., & Thomson, D. M. (1973). Encoding specificity and retrieval processes in semantic memory. *Psychological Review, 80,* 352–373.

Received July 3, 1986
Revision received February 19, 1987
Accepted April 7, 1987 ■