Journal of Experimental Psychology: General
1996, Vol. 125, No. 4, 403–421

0096-3445/96/$3.00

Bias Effects in Implicit Memory Tasks

Roger Ratcliff and Gail McKoon
Northwestern University

A major focus of recent research in memory has been performance on implicit tasks. The phenomenon of most interest has been *repetition priming,* the effect that prior exposure to a stimulus has on later perception of the stimulus or on a later decision about the stimulus. Picture naming, word identification, and word production in stem- and fragment-completion tasks all show repetition priming effects. The separation of implicit from explicit memory systems provides one account of this data, but a different theoretical view is proposed here: Repetition-priming effects come about because the processes that perform a task are biased by prior exposure to a stimulus. The processing of the prior stimulus leaves behind by-products, temporary modifications of the processes, which influence later processing. The aim of this article is to demonstrate the potential of this view for developing new theories and for prompting new empirical questions.

When a prior experience affects perception even though the experience is not consciously recollected, "implicit" memory is said to be implicated. Implicit memory has become a central focus of research in memory, with a large number of recent experiments examining the unconscious priming effects of prior experience on perceptual as well as other kinds of tasks. Our goal for the experiments described in this article was to set the stage for development of theories to explain these priming effects. Although this area of research only recently has acquired the label *implicit memory,* interactions of perception and memory have long been studied. Despite this history and despite large quantities of new research, little theoretical understanding about process or representation has been gained. Recent accounts that attribute priming effects to a multiplicity of memory systems offer no explanation of the mechanisms responsible for priming in implicit tasks. In this article, we offer a specific hypothesis about why priming effects come about and a specific model for priming in one implicit memory task. Both hypothesis and model are grounded in the information-processing tradition (e.g., Broadbent, 1958; LaBerge & Samuels, 1974; Morton, 1969; Posner, 1978). We begin by reviewing older data about implicit priming effects.

Suppose you are shown an unambiguous version of the old woman–young wife ambiguous picture (Boring, 1930; see Neisser, 1967, p. 142), for example, the old woman version; then suppose two weeks elapse and you are shown the ambiguous version of the picture. Will you again perceive it to be the old woman? Subjects in experiments conducted by Leeper in 1935 did see the ambiguous version of the picture the same way on the second exposure as the first (with high probability). This result was considered particularly noteworthy because, on the second viewing, the picture was shown unexpectedly, without warning, in the middle of a classroom lecture, for only 1 s. Leeper (1935) also investigated the perception of fragmented pictures, pictures that showed an object such as a violin by displaying only tiny bits of the original picture of the whole violin (see Neisser, 1967, p. 60). Leeper let about 3 weeks pass between first and second exposures to the fragmented pictures and, as with the old woman–young wife, subjects were likely to identify a picture in the same way on the second presentation as on the first. Leeper's experiments followed earlier work by Rubin (1915, 1921; cited in Woodworth, 1938), who had examined perception of figure–ground relationships. Rubin made up colored nonsense shapes, which were displayed against a black surround. Subjects were asked to judge which part was figure and which part was ground, and they tended to give the same response on a second presentation as on the first, with 30–45 min intervening between the two presentations.

This early work can be described as investigating the effects of memory on perception. The results of the experiments were interpreted as showing that once one kind of sensory organization had been achieved for a picture, this organization persisted to a second encounter. The results were considered especially interesting because the effects persisted for such long periods of time. Interest in this work was carried forward into the beginnings of modern cognitive psychology. In 1960, Epstein and Rock recognized the role in perception of memory traces for specific objects and argued that "the problem of how a memory trace can determine a perceptual outcome" (p. 228) should be faced. But little was done to address this problem, such that Neisser, in reviewing Leeper (1935) and Epstein and Rock's work,

Roger Ratcliff and Gail McKoon, Psychology Department, Northwestern University.

This research was supported by National Institute of Mental Health Grants HD MH44640 and MH00871, National Institute for Deafness and Other Communication Disorders Grant R01-DC01240, and National Science Foundation Grant SBR-9221940.

We thank Doug Hintzman, Mike McCloskey, Roddy Roediger, and David Rubin for thoughtful comments on this research. We also are grateful to Alan Baddeley and Dennis Norris for their comments and for the hospitality of the Applied Psychology Unit in Cambridge, England, where parts of this article were written.

Correspondence concerning this article should be addressed to Roger Ratcliff, Psychology Department, Northwestern University, Evanston, Illinois 60208.

raised the question of why such phenomena had been "so little studied" (Neisser, 1967, p. 142).

Research into the effects of memory on perception enjoyed some resurgence with Kolers's work in the early 1970s on how the encoding of stimuli changes as a function of prior encounters. In his well-known experiments (1973, 1974, 1975a, 1975b, 1976; Kolers & Ostry, 1974), participants read texts that were geometrically inverted. Even a year later, they were still able to read those same texts better than new, previously unseen, inverted texts. Kolers argued that repeated exposures to a stimulus "modify the encoding operations"; by his view, "memory then is not traces that are matched to a stimulus (or vice versa) but procedures, operations, ways of encoding the stimulus" (1975a, p. 700). Kolers (1975a) also drew a distinction between the effects that repeated exposures might have in facilitating execution of the same actions versus switching processing to different, more skilled actions. It is the former kind of item-specific facilitation that we focus on in this article, not the acquisition of skilled actions in, for example, motor learning (Cohen & Squire, 1980; Corkin, 1965; Milner, 1962) or mirror reading (Cohen & Squire, 1980).

In the 1980s and 1990s, there has been renewed interest in the interactions of memory and perception. This time so much research has been generated that one author suggests it might be called a "golden age" (Schacter, 1992, p. 559). What Wallach said in 1949 could truly be said by many cognitive psychologists today: "I have recently become impressed with the extent to which memory traces participate in simple perceptual processes" (1949, p. 13; see also Wallach, O'Connell, & Neisser, 1953). For example, in experiments analogous to Leeper's (1935), Tulving, Schacter, and Stark (1982) found that subjects are likely to respond to a fragment of a word (selected letters of the word) with the same word they have been exposed to previously. In the perceptual (word) identification task (Broadbent, 1967; Jacoby, 1983a, 1983b; Jacoby & Dallas, 1981; Morton, 1968; Neisser, 1954; Ratcliff, McKoon, & Verwoerd, 1989; Winnick & Daniel, 1970), briefly presented words are more likely to be correctly identified if they were read in a prior phase of the experiment. In correctly deciding that a drawing of a three-dimensional (3-D) object depicts a possible object that could actually occur in the real world, subjects are more accurate if they have viewed the object before. If subjects are asked to name a picture of a familiar object, they are faster if they have seen the same picture before (e.g., Biederman & Cooper, 1991), even when as much as a week elapses between exposures (Cave & Squire, 1992).

Strong new impetus for research in this domain came from findings with amnesic subjects (e.g., Corkin, 1965, 1968; Milner, 1962; Warrington & Weiskrantz, 1968). Across a number of tasks, it has been found that amnesic subjects who cannot even recollect that they were previously in an experiment still show effects of previous encounters with the stimuli. Amnesic subjects who are severely impaired on tests like recognition and recall show preserved learning of motor skills and perceptual skills, and they are sometimes claimed to be equivalent to nonamnesic subjects in their tendency to identify a specific stimulus the same way on a second encounter as on the first encounter (see discussion in Ostergaard & Jernigan, 1993). Initially these findings led to a distinction between a procedural memory system that is responsible for the learning of motor skills, perceptual motor skills, and cognitive skills (e.g., Cohen & Squire, 1980) and a declarative memory system that allows explicit reference to previously encountered stimuli or events. More recently, a major focus in research with amnesic subjects has been the perceptual and cognitive effects of prior exposure to specific stimuli (e.g., Gardner, Boller, Moreines, & Butters, 1973; Graf, Squire, & Mandler, 1984; Jacoby & Witherspoon, 1982; Rozin, 1976; Shimamura, 1986), that is, the repetition priming effects that are the topic of this article.

There have also been new insights into the nature of the interactions between memory traces and perceptual processes. For one, there are the new labels that describe these interactions: "Implicit memory" is memory for an earlier encounter with a test stimulus, a memory that does not involve conscious recollection of the previous encounter but can affect perception of the stimulus or a decision about it. An *implicit task* is one that asks subjects to perceive, identify, or make a decision about a stimulus and does not explicitly ask them for conscious recollection of any previous exposure to the stimulus as they do so. *Priming* is the change in probability or speed of responses to perceptual objects that results from recent previous encounters with the objects. (In the implicit memory domain, priming refers to facilitation due to repetition of the same item; in other domains, it refers to facilitation between associated items, which has quite different characteristics from repetition priming; for discussion see Ratcliff & McKoon, 1988.) Complementing the findings with amnesic patients, the performance of nonamnesic subjects on implicit tasks has been found to dissociate from performance on tasks that explicitly require conscious recollection, tasks such as recall and recognition (cf. Tulving & Schacter, 1990).

On the basis of these findings, Moscovitch, Schacter, Squire, Tulving, and others (Moscovitch, 1992; Schacter, 1994; Squire, 1987; Tulving & Schacter, 1990) have claimed that the implicit memory shown by priming effects reflects the operation of a special memory system or systems (or modules, Moscovitch, 1992) separate from a memory system used for conscious recollection. Squire (1992, 1994) classifies memory into two broad categories, declarative (explicit) and nondeclarative (implicit), with each category having several subcategories, and argues that "multiple forms of memory are supported by different brain systems and have different characteristics" (Squire, 1994, p. 225). Schacter and Tulving have argued that "human memory can be classified into five major categories, plus a number of subcategories" (Schacter & Tulving, 1994, p. 32) and that there are at least three different "perceptual representation systems": a visual word-form system, an auditory word-form system, and a structural description system for 3-D objects (Schacter, 1994; Schacter & Tulving, 1994). Moscovitch (1994) divides memory into different modules for different computations performed in different classes of tasks (with some of the modules coinciding with the perceptual representation systems described by Schacter, 1994).

A major problem that arises with these classification schemes is that simply postulating multiple memory systems does not explain priming effects. There has been little effort to provide the explanations that are needed—little exploration of what mechanisms underlie performance on implicit tasks or how those mechanisms might be affected by prior experiences. If perceptual representation systems are key components of information processing, then it is surprising how much is *not* known about how they work. This contrasts with the large body of research and theory that focuses on the conscious recollection processes of recall and recognition, research that has led to the development of models that explain and integrate a wide variety of empirical results.

Parenthetically, it must be noted that proposals of implicit memory systems are sometimes accompanied by hedges. For example, Schacter (1994) has said that the implicit–explicit distinction is descriptive and does not refer to or imply the existence of distinct underlying memory systems. But later in the same paragraph, he has argued that data do suggest separate memory systems, and the rest of the article continues as though separate systems exist. Whether the terms *explicit* and *implicit* are to be taken as descriptive of tasks or as labels for memory systems is not a trivial issue. Richardson-Klavehn and Bjork (1988) list several reasons against using the terms interchangeably as descriptors and names of memory systems. For example, a major problem caused by mixing uses of the terms is the resulting tendency to assume that an implicit task is measuring implicit memory or that an episodic task is measuring episodic memory, uncontaminated by any other form of memory. In this article, we treat the multiple-systems view as a serious proposal, assuming that it is not just a descriptor of tasks and that its proponents really do intend to say that separate memory systems exist.

If the postulation of separate memory systems does not give an account of the mechanisms by which perception and memory interact to produce priming, then the question arises as to whether an account can be formulated from a different approach. A candidate is the information-processing approach, as articulated by LaBerge and Samuels (1974), Posner (1969, 1978), and Morton (1969, 1970) and at least dating back to Broadbent (1958). Posner (1978) summarized a large body of research by describing the processing of information over time and the progression of information from perception through stages of coding, culminating in a semantic representation and understanding of the stimulus. A similar framework was developed by LaBerge and Samuels (1974) to describe the processes involved in reading. According to their model, perception of a word led to a progression of codes—visual, phonetic, and semantic—over the time course of processing of the word. Morton (1969, 1970) also offered a detailed model for how a word could be identified through processes that accumulate perceptual information over time. Figure 1 shows a schematic version of LaBerge and Samuels's model.

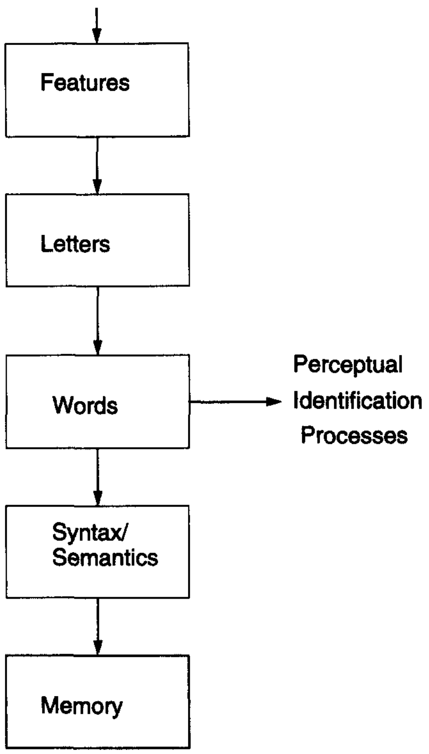

Figure 1. A generic information-processing model.

Models in the information-processing tradition, such as the models just described, do not address the interactions between perception and memory that were demonstrated in Leeper's (1935) early experiments and others described in Neisser (1967). In fact, most current models in this tradition, such as connectionist models for word identification (McClelland & Rumelhart, 1981; Seidenberg & McClelland, 1989) and connectionist models for identification of pictured objects (Edelman & Weinshall, 1991; Hummel & Biederman, 1992), do not address the priming phenomena that have been brought to such prominence by the implicit memory enterprise.

Nevertheless, it is possible to think of priming effects as due to modifications of the earlier stages of processing in a scheme like that of LaBerge and Samuels (1974; Figure 1). The focus of information-processing models, unlike implicit memory systems, is on processing: how information is transformed over time to allow stimuli to be identified and understood and to allow decisions to be made about them. By this view, priming effects can be conceptualized as the "by-products" of information processing (Morton, 1970, p. 216).

In keeping with this view, our goal was to develop hypotheses about how priming in a specific task derived from the mechanisms necessary to explain performance in the task. Our general proposal is that priming reflects the operation of bias: Prior presentation of an item biases processing of the item on subsequent presentations toward a particular response, away from other responses. A key aspect of bias is that it entails costs as well as benefits. A bias toward a particular response will be beneficial only if it is the correct response. In this article, we demonstrate bias across a range of experimental paradigms, and we suggest that in all of them, despite differences in their underlying cognitive processes, bias can be understood in terms of temporary modifications to the processes of perception, identification, and decision. We also describe a model for the specific task of word identification and show how bias operates in that task.

The bias proposal is compatible with the transfer-appropriate processing view that has been presented by Roediger and his colleagues (e.g., Roediger, 1990; Roediger & Blaxton, 1987; Roediger & McDermott, 1993; Roediger, Weldon, & Challis, 1989). *Transfer-appropriate processing* means that the processing of a stimulus will be facilitated to the extent that it overlaps with processing of the stimulus on a previous presentation. For example, changes in perceptual form (e.g., font changes, modality changes) from first to second presentation reduce the amount of priming because these changes mean there are differences in early perceptual processes, as conceptualized in a scheme like that shown in Figure 1. The transfer-appropriate processing proposal suffers from circularity in that the amount of overlap can be defined only by changes in amount of priming, although the proposal can be bolstered by intuitively plausible assumptions. The bias proposal may also help to overcome the problem of circularity by connecting the transfer-appropriate processing idea to the mechanisms of information-processing models.

The bias hypothesis and the transfer-appropriate processing hypothesis both hold that multiple memory systems are not needed to explain currently available data from implicit tasks; they assert that it is sufficient to have a single information-processing system subsume a variety of processes. Proponents of a single system are not required, as Schacter has argued, to "maintain that both explicit and implicit remembering are based on newly-created episodic representations within a unitary memory system" (Schacter, 1990, p. 548). His statement oversimplifies, making no allowance for different processes to operate at different points in the processing system or for different processes to operate for different tasks. Taking a middle ground, Moscovitch (1992) has suggested there are both multiple processes and multiple memory systems, and Shimamura (1993, p. 281) has suggested that the different positions are "simply the result of scientists working from different perspectives." However, they miss a critical point, which is that a focus on multiple memory systems "foregrounds" memory and puts into the background an understanding of the mechanisms responsible for priming effects, leaving it unlikely that these mechanisms will be explored as explanations of repetition priming.

The step that needs to be taken, from all perspectives, is theoretical examination of questions such as What is the representation into which a stimulus is encoded? What are the encoding processes? How is a decision made about the correct identification of the stimulus? and How does a repetition of the same stimulus facilitate or inhibit correct responses to it? It is to begin to address these questions that we examine bias empirically across a range of tasks and describe how it might be treated by information-processing models.

We have examined bias in some of the implicit paradigms that are currently most popular. Experiments with stem completion, picture naming, and fragment completion are reported in this article, and similar results from experiments with object decision and perceptual word identification are reviewed. Bias represents a strong prediction of a specific pattern of results. Perhaps surprisingly, given the wide variety of stimuli and tasks in the list of paradigms, the data were all consistent with the prediction that prior presentation biased the responses to a stimulus toward a particular response, leading to costs as well as benefits. In the sections that follow, we present the various paradigms and their data. Then in the General Discussion, we suggest some theoretical implications and review a quantitative model for bias in perceptual word identification.

Object Decision

A typical object-decision experiment (Schacter, Cooper, & Delaney, 1990; Schacter, Cooper, Delaney, Peterson, & Tharan, 1991; Schacter, Cooper, & Treadwell, 1993) has two phases, a study phase and a test phase. In the study phase, subjects are presented with a series of line drawings of 3-D objects, some that are possible objects and some that are impossible (for examples, see Figure 2). For each object, the subjects are asked to judge whether it is left facing or right facing. In the test phase, drawings of 3-D objects are flashed for brief amounts of time, and subjects are required to judge whether each object depicts a possible object that

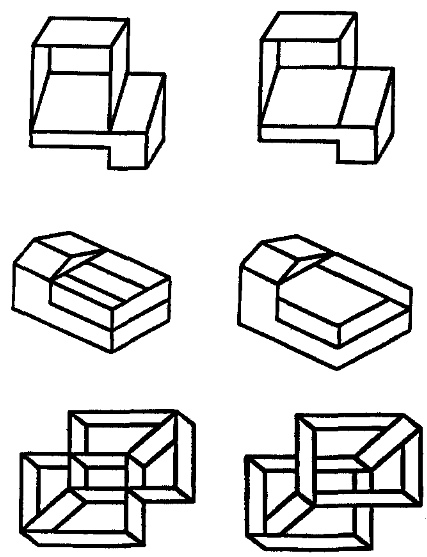

Figure 2. Examples of possible and impossible objects. Reprinted from "Bias in the Priming of Object Decisions," by R. Ratcliff and G. McKoon, 1995, *Journal of Experiemntal Psychology: Learning, Memory, and Cognition, 21,* p. 755. Copyright 1995 by the American Psychological Association.

could actually occur in the real world or an impossible object that could not occur. Subjects are not told before the study phase that there will be a later test phase, and they are not told to use memory for objects from the study phase in making their decisions in the test phase. The results found by Schacter and his colleagues showed priming for possible figures but not for impossible figures: In the test phase, subjects were more accurate in their decisions about possible objects if the objects had been presented in the study phase than if they had not been, but there was no significant effect of prior study on decisions about impossible objects. This pattern of results, combined with neuropsychological data and with findings of dissociations between performance on the object-decision task and performance on recognition, was taken as evidence for a perceptual "structural description system" that encodes possible 3-D figures but not impossible ones (Schacter et al., 1990; Schacter et al., 1991; Schacter et al., 1993). The structural description system is one of the implicit perceptual representation systems proposed by Tulving and Schacter (1990). Object-decision priming is held to be one of the more striking pieces of evidence for these implicit memory systems because the stimuli are nonverbal and because they are novel to the subjects.

In Ratcliff and McKoon (1995), we proposed that performance in Schacter et al.'s (1990, 1991) experiments was based on a bias to respond "possible" to an object that was previously studied (see also McKoon & Ratcliff, 1995; Schacter & Cooper, 1995). This hypothesis leads to the prediction that prior study should affect impossible as well as possible objects, by increasing the probability of calling an impossible object possible. We suggested that the reason this is not generally observed is that there is an offsetting influence of retrieval of episodic information about some feature or configuration of features of an impossible studied object that cues the object's impossibility. In other words, the bias to respond "possible" to previously studied impossible objects is offset by episodic information that indicates their impossibility.

To support this explanation of priming in object decision, we separated the episodic retrieval process from the bias process by manipulating retrieval conditions. One method we used was to eliminate episodic retrieval by imposing a deadline on the object-decision process. Subjects were given a deadline that required them to respond to the test objects very quickly. Another method was to eliminate episodic retrieval by imposing a memory load of seven digits for subjects to keep in mind while performing the object decisions. With both methods, results showed that prior study biased subjects to respond "possible" to both possible and impossible objects and that the effect of prior study was as large for impossible objects as for possible objects. Table 1 shows the probabilities of "possible" responses, comparing objects that had been previously studied to those that had not. Without a deadline or memory load, the data replicate the finding of priming for possible but not impossible objects. With a deadline or memory load, there was a substantial increase in the probability of "possible" responses from the no-study condition to the previously studied condition for impossible as well as possible objects.

Prior study should, according to the bias hypothesis, affect performance not only for a previously studied stimulus itself but also performance for similar stimuli. This was demonstrated for the object-decision task by using pairs of objects that were very similar and differed only in the few lines that it took to make one of them possible and the other impossible (see Figure 2). These stimuli were used in an experiment in which, for each object, subjects studied the possible version or the impossible version or neither and were tested on either the possible or impossible version. The effect of study was the same, whichever version was studied and whichever version was tested. In other words, prior study of either version increased the probability of responding "possible" for both the possible and impossible test objects, and there were no significant differences in the size of the increase as a function of which version had been studied or which tested.

The first conclusion that we drew from our results was that there is no compelling reason to postulate an implicit structural description system that stores only possible, not impossible, objects. Either with procedures designed to eliminate episodic retrieval or when similarity made episodic information unreliable, impossible objects showed as large an effect of prior study as possible objects. The second conclusion was that bias gives a good description of priming for object decisions. Object decisions show both costs and benefits in that prior study hurts responses to impossible test objects as much as it helps responses to possible test objects.

The object-decision task is a categorical one: Subjects are asked to categorize each test item as belonging to the

Table 1
Object Decision

		Pr (Possible)		
Study form	Test form	Replication	200-ms deadline	Memory load
Possible	Possible	0.67	0.64	0.66
No study	Possible	0.58	0.54	0.61
Impossible	Impossible	0.42	0.48	0.42
No study	Impossible	0.41	0.33	0.34

Note. Pr = probability. Adapted from "Bias in the Priming of Object Decisions," by R. Ratcliff and G. McKoon, 1995, *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* pp. 758, 759. Copyright 1995 by the American Psychological Association.

possible category or the impossible category. Thus, bias appears as a tendency to respond with one of these categories. For a test object with which they are familiar because of prior study, subjects choose the possible category. In other implicit memory tasks, subjects are asked to give the specific name of the test item as a response, for example, to name a word that is flashed in perceptual identification. We now examine bias in some of these tasks.

Experiment 1: Picture Naming

The perception of pictures of real and familiar objects, like the perception of drawings of novel 3-D objects, is assumed by Schacter and his colleagues to depend on the (implicit) structural description system (Schacter, 1994; Schacter et al., 1993; Tulving & Schacter, 1990). The speedup due to prior exposure in the time it takes a participant to name a pictured object has been considered a fairly pure measure of implicit memory (Cave & Squire, 1992). Mitchell and Brown (1988) found a speedup in naming time even when as much as six weeks intervened between exposures. Cave and Squire (1992) found that both amnesic and nonamnesic subjects showed facilitation after a 1-week interval.

To look for a bias mechanism operating in picture naming, we used a similarity manipulation. We constructed pairs of pictures of objects such that the two objects were quite similar to each other; examples are shown in Figure 3. The two objects of a pair were easily named as different common objects, but they were very similar visually. Subjects participated in two experimental sessions. For the first session, they were shown some of the objects and asked to name them as quickly as possible. For the second session, which took place about a week later, they were again asked to name objects quickly. Some of the objects were the same objects as in the first session, some were the very similar pair mates of objects from the first session, and some were new. When the exact same object was named in both sessions, we expected to replicate previous findings of facilitation. But when an object in the second session was very similar to but different from an object from the first session, we expected to see a cost to performance: Object-naming response times should show inhibition.

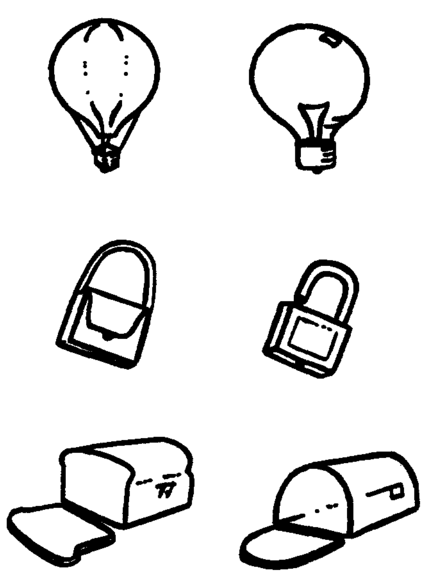

Figure 3. Examples of similar pictures.

Method

Materials. Thirty pairs of pictures of objects were constructed. The 2 objects of a pair were drawn to be as visually similar as possible. The 2 objects were not related to each other in any other way, for example, associatively or semantically. There were also 5 pictures used for practice items. When the objects were displayed for naming on a personal computer (PC) screen, they ranged from about 1.7 cm to 6.4 cm in width and from about 1.8 cm to 5.6 cm in height.

Design and subjects. There were six conditions in the experiment, formed by crossing two variables: Either one or the other of the two objects of a pair was tested in the second session, and either that same object, the other object, or neither was presented in the first session. The six conditions were combined with six sets of pictures and groups of subjects in a Latin square design. There were 30 subjects, but one was discarded due to problems with recording responses. Each subject participated in two 10-min sessions, with about 7 days intervening between the two sessions. As with all the experiments reported here, the subjects participated in order to receive credit in an introductory psychology class.

Procedure. We designed our procedure to follow that of Cave and Squire (1992). Pictures of objects were presented for naming in both sessions of the experiment. The procedure was the same in both. The pictures were displayed on a PC screen. Each picture was preceded by a 500-ms warning signal (a row of plus signs), and then the picture was displayed. The picture remained on the screen until the subject either named the object aloud or said the word "no" to indicate that no name came to mind. Naming latencies were recorded by voice key. There was a 4.5-s pause between pictures to allow the experimenter to record the subject's response. Five practice pictures began each session. The order of the experimental pictures in the two sessions was random, with the randomization changed after every second subject. Subjects were instructed to name the objects as quickly as possible.

Results

The mean naming latency for each condition for each subject was calculated, and the means of these means are displayed in Table 2. For about 11% of the items, subjects were not able to give a name or they gave some inappropriate name; these were excluded from the analyses. For all of the results of analyses of variance (ANOVAs) reported in this article, $p < .05$.

Bias. The naming latencies displayed in Table 2 show a facilitation effect from first to second session. Objects in the first session were named with a mean latency of 789 ms, and when the same objects were presented in the second session, the mean was 720 ms. The latencies also show a clear bias effect. In the second session, the naming latency for an object that was not tested in the first session was 809 ms. If the same object had been tested in the first session, this

Table 2
Picture Naming

Study condition	Response time (ms)
Target picture	720
Similar picture	879
No study, Session 2	809
No study, Session 1	789

latency decreased to 720 ms. If the similar but different object had been in the first session, the latency increased to 879 ms. The cost of prior exposure to a similar object was much the same size as the benefit of prior exposure to exactly the same object. These data show the same bias pattern as the object-decision data but with different stimuli, a different measure, and a different procedure.

Across the three conditions in the second session, the naming latencies were significantly different, $F(2, 56) = 22.2$. Both the increase and the decrease relative to no previous exposure were shown significant by planned comparisons, $F(1, 56) = 7.0$ and $F(1, 56) = 11.0$, respectively. The standard error of the means was 19 ms.

In addition to the bias effect shown in correct responses (Table 2), there were also intrusion errors. For an object tested in Session 2, if subjects had seen its similar pair mate in Session 1, then they tended to give the name of the similar pair mate in error in Session 2. The probability of this error was .097, larger than the probability of giving the name of the similar pair mate in error when there had not been prior study of the pair mate, which was .015.

Episodic memory. In the standard object-decision task, without a deadline or memory load, we found that subjects' fast responses showed bias and that their slower responses did not show bias; we interpreted this difference as reflecting a fast, biased component of processing and a slower component of processing based on episodic memory (Ratcliff & McKoon, 1995). This raises the question of whether or how episodic information could be affecting processing in the picture-naming task.

First, episodic memory did not affect picture naming in the same way as it did the object-decision task. The bias pattern observed in the correct naming latencies held across the entire distribution of response latencies. Both the fastest and the slowest halves of the responses showed the bias pattern (although of course its size was reduced for the fast responses, to about ± 40 ms). This rules out any hypothesis by which episodic memory affected slow responses differently than fast responses.

Another possible hypothesis would be that episodic information, as well as implicit information, affected all responses, both faster and slower ones. However, this hypothesis is not tenable. Episodic information would bias responses toward the previously encountered stimulus; this would give facilitation if the test item was the same as the previously encountered stimulus but inhibition if they were different. Implicit information, on the other hand, should give only facilitation because it always provides an improvement in processing (at least, in the most straightforward view of implicit information). If these two sources of information, one episodic and the other implicit, combine, then there should be more facilitation than inhibition, a prediction contradicted by the data.

Third, it might be claimed that there was no contribution at all to performance in our experiments from implicit memory; performance was entirely based on episodic memory. This would be an unlikely argument, because our procedures and facilitation effects mirror those of studies that have been unambiguously interpreted as showing implicit memory. It is also an argument for which a resolution might be sought in dissociation logic: If our results dissociated from the results of an explicit memory task, then we could be sure they were due to implicit memory. However, as is pointed out in the General Discussion, dissociation logic is not conclusive. Instead, the key to resolving the issue of what processes are involved in picture naming in particular and implicit tasks in general lies in the development and testing of models, such as those for picture naming that we review in the General Discussion.

Experiment 2: Picture Identification—Forced Choice

In Experiment 1, we found bias when subjects were asked to name pictures of common objects. In another similar implicit task, perceptual identification, subjects are asked to name common words. While typical picture-naming and perceptual identification experiments have much in common, there are also large differences. In picture naming, a picture is presented to the subject for an unlimited amount of time, and the dependent measure is the time taken to name the object in the picture. In perceptual identification, a word is flashed for only a few milliseconds and the dependent measure is either the probability of naming the flashed word correctly or the probability of choosing which of two alternative words matches the flashed word. Our aim in Experiment 2 was to bring the picture-identification paradigm closer to the word-identification paradigms so that we could be sure that general claims about bias applied to them in similar ways. We designed Experiment 2 to use picture stimuli with a perceptual word-identification procedure and dependent measure. Subjects were shown pictures and asked to study them for a later memory test. In the test phase, there was a series of target pictures. Each target picture was flashed briefly, and then subjects were asked to choose which of two alternative pictures matched the one that had been flashed.

To test for bias due to prior study, we used the same similarity manipulation as in Experiment 1. There were three conditions: A target picture was presented in the study phase, its very similar pair mate was presented in the study phase, or neither of them was presented in the study phase. In the test phase, the target picture was flashed briefly and then the target and its pair mate were presented for forced choice. We expected to see facilitation in the probability of a correct response to the target when the target had been studied but inhibition when its pair mate had been studied.

Method

Materials, subjects, and design. The same 30 pictures were used as in Experiment 1. For the forced-choice test, the two pictures of a pair were displayed on the PC screen side by side. In order to make the pictures difficult to identify when they were flashed in the test phase, we used a mask made up of a number of crisscrossing lines, measuring about 7.5 cm horizontally and about 7.0 cm vertically. The mask was a collection of lines from the various stimulus objects randomly placed on the screen to cover the area in which the pictures were presented. There were three conditions in the experiment: The target picture was studied, its pair mate was studied, or neither of them was studied. Which of the two pictures of a pair was designated the target was decided randomly. The three conditions were combined with the 30 pictures and 15 subjects in a Latin square design.

Procedure. The stimuli were presented on a PC screen, and keypress responses were recorded on the PC's keyboard. The study phase consisted of the presentation of 4 practice pictures followed by the 20 pictures of the experimental design. Each picture was displayed for 2 s, separated from the next picture by a 200-ms blank interval. Subjects were asked to study the pictures for a later memory test.

After the study phase, subjects were given instructions for the test phase and then the test phase began immediately. The sequence of events for each test was as follows: First, a row of pluses was displayed for 700 ms on the PC screen as a fixation point; then the target picture was flashed on the screen for 50 ms (i.e., three raster scans of the PC screen); it was followed by the mask displayed for 400 ms; then the two pictures for the forced choice were displayed until the subject pressed a response key; then there was a blank interval of 300 ms. There were 5 practice items, followed by the 30 experimental items. Subjects were instructed to press the ?/ key if the right-hand picture of the forced-choice alternatives matched the picture that had been flashed and the *Z* key if the left-hand picture matched.

The order of the pictures in the study and test lists, which of each pair of pictures was designated the target, and which response (left or right) was correct were all chosen randomly, with a new randomization used after every second subject.

Results

The mean probability correct was calculated for each subject for each condition. When neither a target picture nor its pair mate had appeared in the study list, the probability of a correct response was .76. Previous study of the target increased the probability correct to .87. But, consistent with the bias hypothesis, this benefit was associated with a cost: Previous study of the similar pair mate decreased the probability of a correct response to the target to .58. The differences among the three conditions were significant, $F(2, 28) = 15.4$, and the standard error of the means was .042. As with Experiment 2, the bias pattern was observed across both fast and slow responses.

The cost (.76 down to .58) was about twice as large as the benefit (.76 up to .87). This is due to a ceiling effect: For some subjects, performance was about perfect in the no-study condition and in the condition in which the target was studied. Given our equipment, we could not flash the target for less than three raster scans without reducing performance of most subjects to near chance, and so we were not able to eliminate the ceiling effect. Despite this equipment limitation, the results showed both facilitation and inhibition on forced choice in perceptual identification for pictures.

Experiment 3: Stem Completion

The object-decision and picture-naming tasks are assumed to tap one of the implicit perceptual representation systems, the structural description system (Schacter et al., 1990, 1991, 1993). Stem completion, fragment completion, and perceptual identification are assumed to tap another of the perceptual representation systems, the visual word-form system (Schacter, 1994; Schacter & Tulving, 1994). Experiments 3 and 4 were designed to show bias operating in stem and fragment completion.

In stem completion, subjects are given several letters and asked to provide the first word that comes to mind that begins with those letters. Priming in stem completion occurs when subjects' responses tend to be words to which they were previously exposed. To show bias in stem completion, we used the same strategy as with picture naming: a similarity manipulation. We constructed a list of pairs of words like "abstain–absent" and "scope–scoop." The words in a pair began with the same three letters, and the fourth letter of one of the words (designated the target word) also occurred in the other word but in a position later than fourth. The experiment was conducted in two phases: Subjects were first exposed to some of the words in the study phase, and then in the second phase, they were given the first four letters of each target word (the "stem") and they were asked to complete the stem with the first word that came to mind. For target words that had been presented in the study phase of the experiment, we expected to replicate previous results of priming (e.g., Graf et al., 1984; Warrington & Weiskrantz, 1968): Correct completions should be produced faster and with higher probability than completions of targets that had not been presented in the first phase. But we expected this benefit to performance to be offset by costs. When a word that was similar to the target but not a correct completion of the stem was presented in the study phase, correct completions of the target should be inhibited.

Method

Materials. Forty-eight pairs of words were constructed such that the first three letters of each of the words in a pair were the same, the fourth letters were different, and the fourth letter of one of the words (the target) appeared in a position later than fourth in the other word. The 48 target words were chosen from words used by Graf et al. (1984).

Procedure and equipment. The stimuli were presented on a PC screen, and keypress responses were recorded on the PC's keyboard. Stem-completion responses were recorded by voice key. Each subject was tested in one 20-min session. The procedure followed that used by Graf and Mandler (1984).

In the first phase of the experiment, a list of 32 words was presented on the PC screen. The words were presented one at a time, and subjects were asked to judge for each word how much they liked the word. Judgments were indicated by pressing keys on the PC keyboard, pressing the *1* key for "dislike extremely," the *2*, *3*, and *4* keys for intermediate judgments, and the *5* key for "like extremely." Each word remained on the PC screen until a key was pressed, and then after a 400-ms pause, the next word was presented. Subjects were encouraged to take their time, making a careful judgment.

The second, test phase of the experiment began immediately after the first. For each test item, a string of four letters was presented on the PC screen. Subjects were instructed to say aloud, as quickly as they could, a word that began with those four letters. Each string remained on the screen until subjects responded; then, after a 4,500-ms pause, the next four-letter string was presented. Subjects were instructed to respond "no" if they could not think of any word that began with the four letters. Subjects' responses were recorded by an experimenter who sat behind the subject. There were 48 test strings, preceded by 3 practice strings.

Design and subjects. There were three conditions in the experiment: A target word was presented in the study phase, the other word of its pair was presented in the study phase, or neither word was presented in the study phase. In all three conditions, the first four letters of the target word were presented in the test phase. The three conditions were combined with groups of subjects (8 per group) and sets of items (16 per set) in a Latin square design. The order of presentation of words in the study phase and test strings in the test phase was random, with the randomization changed after every second subject.

Results

The numbers of different kinds of responses given by the subjects and the mean times required to give those responses are shown in Table 3 (means were calculated across all responses and not across subject means). First, consider the case where the response was the target word, a response that correctly matched the four-letter stem presented for completion (first columns of data in Table 3). As was expected from previous research (Graf & Mandler, 1984; Graf, Squire, & Mandler, 1984; Warrington & Weiskrantz, 1970), the target word was more likely to be given as a response if it was encountered in the study phase of the experiment. The time required to give the target response was also faster if it was encountered previously. The facilitation was offset by costs, of almost equal size in the case of response times. The time required to give the target response was increased, and the likelihood of giving the target response was decreased if a word similar to the target was studied. It is as though study of the similar word blocked production of the target. An ANOVA showed a significant overall effect of study condition on the speed of target-word responses, $F(2, 46) = 6.6$, with a standard error in the means of 83 ms, and a planned test showed slower responses in the similar-word study condition than in the no-study condition, $F(1, 46) = 4.8$. An ANOVA also showed a significant effect of study condition on the probability that a target response was produced, $F(2, 46) = 42.5$, with a standard error in the mean probability of .023, but a planned test comparing the similar-word and no-study conditions showed no significant difference, $F(1, 46) = 1.1$. As with Experiments 1 and 2, the bias pattern held for both fast and slow responses, although it decreased in size (to about ± 35 ms) for the fastest half of responses.

While correct target responses were the data of main interest, Table 3 also shows that subjects occasionally responded to a four-letter stem with the word similar to the target, even though it was an incorrect completion of the four-letter stem, and that they were more likely to do this if the similar word had been presented in the study phase. Overall, the data are consistent in showing that study of the similar (but incorrect) word inhibited the production of correct completions.

Experiment 4: Fragment Completion

Fragment completion, like stem completion, is an implicit task for which responses show priming, and so it has been assumed to depend on the visual word-form system (Schacter, 1994). In fragment completion, subjects are given some of the letters of a word and asked to produce the additional letters that will make a legitimate word. Typically, the letters that are given are not the first consecutive letters of a word, so this is a more difficult task than stem completion, and response times can be on the order of tens of seconds. This means that effects of bias on response time, if they were of the size found in stem completion, would be lost in the increased variability (noise) that necessarily arises with such long response times. So, to make it more likely that bias effects would be observed, we designed a procedure by which a time limit was imposed on fragment completions. A fragment was presented for 4 s, and then a test word was presented; subjects were asked to indicate "yes" or "no" as to whether the word was a correct completion of the fragment. Subjects were instructed to spend the 4 s trying to generate a word to complete the fragment.

Table 3
Stem Completion

Study condition	Target word RT (ms)	Target word *n*	Similar word RT (ms)	Similar word *n*	Some other word RT (ms)	Some other word *n*	No response RT (ms)	No response *n*
Target word	1,156	243	1,335	4	1,634	125	4,811	11
Similar word	1,561	136	1,513	30	1,699	189	5,390	28
No study	1,363	149	1,366	10	1,555	195	5,299	29

Note. RT = response time.

To discourage them from comparing the test word letter by letter to the fragment in order to make certain of a correct yes–no response, we required them to give their yes–no response within 800 ms of presentation of the test word.

To look for bias, we used a similarity manipulation, as in Experiments 1, 2, and 3. We created pairs of words and a single fragment for each pair, for example, *TRAMWAY*, *FRAMEWORK*, and _R_ MW____. One word of each pair was designated the target, and it was a correct completion of the fragment (TRAMWAY is a correct completion of _R_MW____). The other word contained the letters of the fragment and so was somewhat similar to it, but was not a correct completion of the fragment. We expected that prior study of the correct completion would lead to facilitation but that prior study of the similar but incorrect alternative word would lead to inhibition.

Method

Materials. There were 54 pairs of words used in the experiment, each with a corresponding fragment. One word of each pair, the target, was a legitimate completion of the fragment; the other word was not, even though it did contain the letters of the fragment. Some of the fragments and their legitimate completions were chosen from the materials used by Tulving et al. (1982), and others were chosen from materials used by Gibson and Watkins (1988). The set of materials we used is given in the Appendix.

Design and subjects. There were six conditions in the experiment, derived from crossing two variables: Either the target, the similar word, or neither was presented to subjects in a study list, and either the target or the similar word was used in the test phase as the word about which subjects were to make a yes–no decision according to whether or not it was a correct completion of the fragment. The fragment in the test phase was always correctly completed by the target word. These six conditions were combined with six sets of pairs (nine per set) and six sets of subjects (six per set) in a Latin square design.

Procedure. The experiment began with the presentation of the study list. Each of the 36 study words was shown on a PC screen one at a time for 5 s (with a 100-ms blank pause between each word). Subjects were instructed to learn the words for a later (unspecified) memory test. This study procedure was copied from Tulving et al. (1982). The study phase was followed by an unrelated lexical decision experiment that took about 10 min, and then the fragment completion test.

The test for each of the 54 items began with presentation of the target fragment on the PC screen for 4 s. Subjects were instructed to spend the 4 s trying to generate a completion for the fragment. At the end of the 4-s interval, the fragment was erased from the screen and a word, either the target word (a correct completion) or the similar word (an incorrect completion), was displayed to the right of where the fragment had been. Subjects were instructed to respond "yes" or "no" according to whether the test word was a correct completion for the fragment, pressing the *?/* key for "yes" and the *Z* key for "no." Subjects were also instructed to make their response within 800 ms, and, to encourage compliance with this instruction, we displayed their response time on the screen after their response. If their time was longer than 800 ms, the message "TOO SLOW" was displayed for 1 s. Subjects were asked to respond within 800 ms in order to discourage a strategy of checking each letter of the word against the fragment. Our success in discouraging such a strategy is shown by the high error rates. The orders of the items in the study list and in the test list were randomized, with a new randomization used for each second subject.

Results

The 5% of responses that were slower than 800 ms were discarded from the analyses (the pattern of results did not differ when these responses were included). Then the mean percent of "yes" responses was calculated for each subject in each condition, and means of these means are shown in Table 4.

The first aspect of the data to note is that when subjects had not previously studied either the target or the similar word, they were more likely to respond "yes" (correctly) to the target than they were to respond "yes" (incorrectly) to the similar word.

The critical aspect of the data concerns the effect of prior study: Subjects were more likely to respond "yes" to a test word that they had studied before than to a test word they had not studied, compared to the no-study baseline, and this was true for both the targets and the similar words. For example, for a target test word, prior study increased the likelihood of subjects' deciding it was a correct completion by 4%, and for the similar test word, prior study increased the likelihood of deciding that it was a correct completion (incorrectly) by 7%. In other words, prior study of a test word biased subjects to respond that it was a correct completion. Similarly, prior study of a different word than the test word biased subjects against the test word (see also Smith & Tindell, 1996, for similar results).

The significance of these bias effects was demonstrated by ANOVA, with test word and study condition as factors. The main effect of test word was, of course, significant, $F(1, 35) = 26.1$, as was the interaction of test word and study condition, $F(2, 70) = 8.0$. Planned tests showed the difference between performance on the target and similar test words larger for the study-target-word condition than the no-study condition, $F(1, 70) = 8.1$, and smaller for the study-similar-word condition than the no-study condition, $F(1, 70) = 10.0$, as would be predicted by the bias hypothesis. The standard error of the means was .026.

Visual Word Identification

The third task in which priming is assumed to show the operation of the visual word-form system is perceptual identification (Schacter, 1994). In the first phase of a visual

Table 4
Fragment Completion

	Test word, Pr (yes)	
Study condition	Target word	Similar word
Target word	.83	.61
Similar word	.76	.73
No study	.79	.66

Note. Pr = probability.

perceptual identification experiment, subjects study a list of words. Then, in the test phase, target words are flashed, and subjects are asked either to name the flashed word or to choose which of two alternative words matches the flashed word. A target word is flashed for only a few milliseconds, so identification is difficult.

Ratcliff and McKoon (1996; see also Ratcliff et al., 1989) showed that the bias hypothesis provided a good description of perceptual identification data. When subjects were asked to name the flashed target, prior study of the target increased the probability of a correct response. Also, prior study of a similar word increased the probability that the similar word was given as an incorrect response. Table 5 shows the costs and benefits of prior study for forced choice: Prior study of the target word increased the probability that it was correctly selected as matching the flashed word, and prior study of a similar word decreased the probability of correctly selecting the target. There are also two other important effects shown in Table 5. First, the bias pattern appeared only when the forced-choice alternatives were similar to each other, not when they were dissimilar. In Table 5, results for the three study conditions with dissimilar forced-choice alternatives are combined because they were not significantly different from each other. What this means is that if the word "died" was flashed as a target, prior study of "died" increased the probability of correctly choosing "died" from the alternatives "died" and "lied" but it did not affect the probability of choosing "died" from the alternatives "died" and "sofa." Even when similar and dissimilar forced-choice alternatives were mixed in the ratio of either 4:1 or 1:4, the same patterns of data were obtained (Ratcliff & McKoon, 1996, Experiment 1), arguing against a strategy based on guessing. The second important result shown in Table 5 is that the bias effect was as large when the target was flashed for only 10 ms and performance was near chance as it was when the target was flashed for a longer time and performance was above chance. It is also noteworthy that (like all the other paradigms except object decision) the patterns of results were the same for fast responses as for slow responses.

Auditory Word Identification

The auditory word-form system is the third of the proposed implicit perceptual systems (Schacter & Church, 1992). We have found a bias effect with the task used to demonstrate this system, just as with the other tasks and systems. The task is an auditory analog of the visual perceptual-identification task. In our experiments (Ratcliff, Allbritton, & McKoon, in press), subjects listened to a study list of single words, followed by a list of test items. Subjects heard each target test word in noise so that it was difficult to identify (following procedures used by Schacter & Church, 1992). For each target, either the target itself had been presented at study, a very similar sounding but different word had been studied, or the target had not been studied. The test was forced choice. Table 6 shows the bias effect in the data. When the two choices were similar to each other, correct choices of the target were facilitated by prior study of the target but inhibited by prior study of a similar word, with the facilitation and inhibition effects being the same size. When the two test choices were dissimilar to each other, there was no effect of prior study on performance. Overall, the pattern of results was the same as that observed with visual perceptual identification.

Table 5
Probability Correct in Forced Choice for a Flashed Target Perceptual Identification

	Similar alternatives			Dissimilar alternatives
Flash time (ms)	Study target	Study similar	No study	3 study conditions grouped
10	.585	.408	.507	.551
20	.667	.564	.635	.675
40	.804	.686	.745	.880

Note. Data are from Ratcliff and McKoon (1996).

Table 6
Probability Correct in Forced Choice for Auditory Perceptual Identification

Alternatives	Target word
Similar	
Study target	.653
Study similar	.550
No study	.604
Dissimilar	
3 study conditions grouped	.758

Note. Data are from Ratcliff, Allbritton, and McKoon (1996).

General Discussion

The view that there exist multiple memory systems is currently the dominant view in the study of implicit memory. It is obvious to ask whether this view is contradicted by the data presented here. The answer is that there is no direct contradiction. This is because instantiations of the multiple-systems approach make no relevant predictions. As the various proposals currently stand, there are no detailed accounts of the processing mechanisms that underlie performance in implicit tasks, and so there is no means of predicting bias effects.

One way to proceed, given this situation, would be to build a theory that would explain bias effects in the context of multiple memory systems. However, an important and prior issue must be faced: Is there truly a need for multiple systems to be a part of the explanation of implicit priming effects? Or are there alternative models that do not require multiple systems but still can explain all the kinds of data that have been said to require multiple systems, as well as the new data presented in this article? In the sections below, we address the issues raised by these two questions. First, we argue that the several lines of evidence used to support the existence of multiple systems are not fully compelling, and second, we illustrate ways that implicit priming effects might be explained without recourse to multiple systems.

Multiple Memory Systems

Research guided by the multiple memory systems approach has concentrated mainly on separating memory systems from each other. It is for this reason that there is no account of implicit processing mechanisms sufficiently detailed to predict bias effects. In 1984, Tulving proposed that memory systems could be distinguished using five criteria: their functions and the kinds of information they represent; the laws and principles by which they operate; their neural substrates; their ontogenetic and phylogenetic development; and their formats for representing information. These criteria are now thought to be too narrow, and Schacter and Tulving (1994) have proposed three new criteria: First, a memory system is able to perform a large number of tasks of a particular class, and if the system is damaged, deficits are observed in tasks of this class but not the tasks of other memory systems. Second, for each system, there should be a property list that allows the system's identity to be determined and its relationship to other systems to be specified (see examples in Tulving, 1983, and the critique of such lists by McKoon, Ratcliff, & Dell, 1986; also Ratcliff & McKoon, 1986; Tulving, 1986). Third, convergent dissociations should provide experimental separation of the memory systems. It is hard to see how these three criteria will provide a precise way of discriminating among memory systems, and none of these criteria distinguish memory systems by making reference to mechanisms and processes. Thus, they offer no means of answering the question of how to explain bias.

There have been a few isolated instances of speculation in the multiple memory systems literature about the cause of priming effects. Squire (1992, p. 211) has suggested "an increased facility for detecting or identifying words or other stimuli" as a function of prior experience, but an increased facility would not predict that there would sometimes be costs to performance. Schacter (1994, p. 237) has suggested that "visual priming may make it easier for the perceptual representation system mechanisms involved with visual word form representation to extract visual information from the test cue." If it is easier to extract visual information, then there should be no costs, only benefits (or at least more benefits than costs).

We set aside these specific suggestions, because they are contradicted by data and because they do not derive directly from any explicit multiple memory systems proposal. We return to the more general issue: If the multiple memory systems approach cannot explain bias, that is, if it cannot explain in detail how prior experience affects performance on implicit tasks, then what can it explain? There have been a number of answers to this question, and we review them here; all are problematic.

Task dissociations and functional independence. These occur when performance on an implicit task is affected by different variables from performance on a task that requires conscious recollection. A dissociation is explained as the implicit task tapping a different memory system from that used for conscious recollection. However, it has been repeatedly pointed out that dissociations can occur between tasks that tap the same memory system (Hintzman, 1984, 1990; Kolers & Roediger, 1984; McKoon, Ratcliff, & Dell, 1986; Nosofsky, 1988; Ratcliff & McKoon, 1986; Roediger, 1990; Shimamura, 1993). If dissociations can occur with a single system, then the hypothesis that there exist multiple memory systems does not help explain them.

Stochastic independence. This is used to describe a finding that performance on an implicit task does not significantly correlate with performance on an explicit task; that is, how good performance is with an item on one task does not predict how good performance is with that item on another task. The postulation of multiple memory systems does not provide an explanation of stochastic independence because intercorrelation patterns among tasks do not always line up with a split between implicit memory and explicit memory (e.g., Perruchet & Baveux, 1989; Witherspoon & Moscovitch, 1989) and because different subsets of stimuli can show different relationships between implicit and explicit tasks (Hintzman & Hartry, 1990; Shimamura, 1985). Moreover, as with functional independence, stochastic independence can appear with different tasks that tap a single memory system (Hintzman, 1987; Nosofsky, 1988). One further problem with the use of stochastic independence to distinguish memory systems is that there has often not been enough power to find dependence even if it were present (Ostergaard, 1992).

Preserved memory performance in amnesic patients. This is often cited as the most compelling kind of data explained by multiple memory systems. The explanation is simple: Performance is preserved because an implicit memory system is preserved. But the explanation is circular: Implicit memory is what is spared in amnesia; what is spared in amnesia is shown by performance on implicit tasks, and implicit tasks are those that tap implicit memory. In a particularly clear example of this circularity (cited in Ostergaard & Jernigan, 1993), Gabrieli, Milberg, Keane, and Corkin (1990) gave patient H.M. an explicit cued recall task. When his performance on this task was as good as control subjects', Gabrieli et al. concluded that the particular cued recall task they used must have been a task that tapped implicit memory. Circularity is also often shown when the performance of subjects with amnesia on an implicit task is compared to the performance of nonamnesic subjects (see Ostergaard & Jernigan, 1993). If the relevant implicit system is preserved under amnesia, then the patients with amnesia should perform as well as nonamnesic subjects. When the patients with amnesia do not perform as well as nonamnesic subjects, it is said to be because the nonamnesic subjects used conscious recollection. Ironically, as the problems outlined here show, it is to avoid the circularity of explaining the data from amnesic patients that an account of the mechanisms of performance in implicit tasks is most urgently needed.

Summary. The hypothesis that there exist multiple memory systems offers a description of classes of empirical effects. Sometimes, but not always, there is a dissociation between performance on a task that taps one system and performance on a task that taps another system. Often, but not always, there is no dissociation between performance on a task that taps one system and performance on another task that taps the same system. Often, but not always, amnesic patients perform as well on implicit tasks as nonamnesic subjects. In none of these cases is it obvious how the hypothesis of multiple memory systems offers clear or falsifiable predictions. To this list of problems, the data reported here add the problem of explaining the bias pattern of priming effects.

Information-Processing Approaches

In the absence of explanations of bias effects from the multiple systems framework, we advocate a more traditional approach: information-processing models that attempt to describe cognitive processes in sufficient detail to give quantitative as well as qualitative accounts of the data. For picture naming, we point to computational vision models that might be able to correctly predict bias. For word identification, we describe a new model that gives correct predictions for bias and also accommodates other findings from the word-identification literature. It is important to note that the bias effect should not be viewed as a simple response bias arising from a change in criterion setting at the final output stage of processing. The models we discuss locate bias in modifications to processes in their internal structures.

For object decision, stem completion, and fragment completion, there are no existing models to use to predict bias. For the object-decision task, it is unclear what it is about the stimuli that allows subjects to make the possible versus impossible judgment. Progress on modeling must await extensive empirical investigation of what dimensions of the stimuli enter decision processes to affect performance and be affected by prior experience. For stem completion, good sources of inspiration might be models that have been developed for cued recall (e.g., Raaijmakers & Shiffrin, 1981); the processes might be similar in that the letters of a stem provide a cue to a word, which could be any word in the lexicon. Fragment completion might be more like free recall (e.g., Raaijmakers & Shiffrin, 1981): Candidate words from the lexicon are generated and matched against the letters of the fragment.

Picture naming. As an example of a computational vision model that might make specific predictions about bias in picture naming, we describe Edelman and Weinshall's model (1991; see also Edelman & Poggio, 1992; Poggio & Girosi, 1990). The model is a two-layer connectionist network designed to assign names to 3-D wire frame objects. Two-dimensional (2-D) views of the objects are presented to the network. Each node in the input layer corresponds to the end of a line in the 2-D view or the crossing point of two lines. Nodes in the input layer are connected to nodes in the second layer, the representation layer. Learning takes place via the presentation to the network of a series of different 2-D views of a 3-D object, one 2-D view immediately following another, each showing the next view in a rotation of the object. When one of the 2-D views is presented to the network, nodes of the input layer are activated, and activation spreads from them to nodes in the representation layer. Using a winner-take-all scheme, one of the nodes in the representation layer gains maximum activation and comes to represent the 2-D view that was presented to the network. The weights on the connections from this representation node to the input nodes are strengthened, and the thresholds of all active representation nodes are raised so that they will be less likely to be activated by other input patterns. When the next 2-D view of the object is presented to the network, a different node in the representation layer gains maximum activation, and a link between the two nodes in the representation layer is strengthened. After the series of views of the object has been presented to the network, a pattern of activation values across the nodes of the representation layer has been built up, and this pattern corresponds to the object. Once one object has been learned in this manner, the patterns for other objects can be learned.

When, at a later time, a 2-D view of one of the learned objects is presented to the network, a pattern of activation arises in the representation layer. This pattern is matched against the pattern that was learned from the series of 2-D views of the object. The goodness of the match is taken to determine the time required to name the object. To model priming, it could be assumed that presentation of the 2-D view modifies the network in the same manner as on the earlier learning trials. The additional learning would lead to a better match when the same 2-D view of the object was presented again, giving facilitation to naming speed. The additional learning could also lead to a worse match for a 2-D view of a very similar object with a different name. Thus, this model is a promising candidate to explain bias in picture naming. Eventual success would be determined by simulations designed to show that the numerical amounts of facilitation and inhibition could be correctly predicted.

Hummel and Biederman (1992) implemented a different neural network model for recognition of pictures of objects. The model has seven layers of elements that take as input a line drawing of an object and, at output, activate a unit that labels the object with its name. In Hummel and Biederman's implementation, the network was trained to assign an output unit to one view of each of a set of objects. Then other, different, views of the objects were tested. The activation value in an object's output node was not affected by spatial translations, size modifications, or mirror reversal, but it was affected by rotation. This is the same pattern of results as has been found with human subjects, the results that Schacter (1994) and Cave and Squire (1992) have argued demonstrate multiple memory systems.

In Hummel and Biederman's (1992) simulations of the model, there was only the one initial learning stage, with no need for learning during later tests of the network. If there were learning at the later tests, that is, if there were changes to connection weights at test, the model could potentially predict bias in the same way as the Edelman and Weinshall (1991) model.

Although neither currently deals with priming explicitly, both Hummel and Biederman's (1992) and Edelman and Weinshall's (1991) models can potentially be adapted and applied. If they can accommodate the data, they will have explicated a set of mechanisms that can produce priming, and they will do so without postulating multiple systems. If they cannot predict the correct patterns of data, the reasons for their failures may point the way to extensions of the models that encompass bias effects as well as explain the original databases of the models.

Perceptual identification of words. Words are probably the most tractable of the stimuli that have been used in implicit tasks, and the processes by which they are identified have been among the most studied of cognitive psychology. Experiments manipulating various characteristics of words can be implemented in straightforward ways, and much is known about how the various characteristics affect performance. Also, there are a number of well-developed models. Thus, when choosing a paradigm in which to begin to attempt to model bias, we were led to perceptual identification.

We first attempted to account for bias with existing models, the interactive activation model (McClelland & Rumelhart, 1981) and Seidenberg and McClelland's (1989) distributed connectionist model. When we found that neither these models nor previously suggested applications of them (cf. Rueckl, 1990) could accommodate bias (Ratcliff & McKoon, 1996), we developed a new model, a counter model, that combines elements of the logogen model (Morton, 1969, 1970, 1979) with the decision mechanism of a random walk (Laming, 1968; Link & Heath, 1975; Ratcliff, 1978; Stone, 1960). We give an overview of the model here, and it is described in detail in Ratcliff and McKoon (1996).

The counter model was designed to explain masked word identification and priming in masked word identification. Like the logogen model, the counter model assumes one decision counter for each word in the lexicon. The counters are organized by similarity such that the counters for similar words are close together in a cohort and the counters for dissimilar words are far apart (see Andrews, 1992; Coltheart, Davelaar, Jonasson, & Besner, 1977; Goldinger, Luce, & Pisoni, 1989). The accumulation of counts in the counters is the mechanism by which a word-identification decision is made. Counts are accumulated at a constant rate such that, for each unit of time, one count is accumulated to one (and only one) counter. Because of the impoverished stimulus conditions of perceptual identification experiments, only some counts are determined by the perceptual features of the stimulus. Other counts that are not determined by the stimulus are random noise, and they are called null counts. Counts that correspond to perceptual features are accumulated into an appropriate counter. A count from a feature of the letter *d,* for example, might be accumulated into the counter for *died.* The null counts can be accumulated to any counter.

The assumption of the model that explains priming is that counters can become attractors. At the time of the perceptual identification test, prior exposure to a word causes the word's counter to attract a few more counts than it otherwise would, stealing them away from the counters of other similar words. It is assumed that this attractive force is quite weak and so its influence extends only through the neighborhood of similar words and not to faraway, dissimilar words.

When a target word is flashed, the accumulation of counts begins. If the flash time is extremely short, as it is in perceptual identification experiments, then counts will continue to be accumulated after presentation of the target word has terminated. Accumulation of counts into counters continues until the total number of counts in one counter exceeds the maximum of the others by a criterial amount, k (this is a generalization of a random walk process, Laming, 1968; Link & Heath, 1975; Ratcliff, 1978; Stone, 1960).

When the decision required is a forced choice, the flashed target is immediately followed by two alternatives, one of them the target word. At that point, the decision process is restricted in that the accumulation of counts is restricted to the counters for those two words. Every count is accumulated by one or the other of these two counters. Counts that correspond to features that are a part of both words are accumulated by the target's counter with probability .5 if neither of the choices was previously studied. Null counts are also accumulated by the target's counter with probability .5 if neither choice was studied. Counts that correspond to features that are a part of one of the words but not the other are accumulated to the appropriate counter. Counts of this last kind are labeled *diagnostic.* There are fewer diagnostic counts if the alternatives are similar to each other, and also the proportion of diagnostic counts decreases as flash time decreases.

Prior study of a word causes the counter for that word to steal nondiagnostic counts away from the counters of other words similar to it. In forced choice, if the two alternatives are similar, the theft gives a benefit to the target if it was previously studied. But if the alternative was previously studied, it steals counts from the target. The increase in the probability that a nondiagnostic count is accumulated in the counter of a word that was previously studied is small. In quantitative fits to data (described in Ratcliff & McKoon, 1996), an increase of only .01 (from .50 to .51) was sufficient. With the iterative and additive accumulation of counts toward the criterion, this small increase is sufficient to yield the bias effects found in the data. When the proportion of diagnostic counts is very low (i.e., when flash time is very short), null counts still allow a decision to be made, and attraction of the null counts still biases responses toward a previously studied word. However, bias only occurs when the forced choice is between two similar words, one of which was previously studied. When the two forced-choice alternatives are dissimilar, there is no effect of prior study because the attractive force that arises due to previous study only extends through a word's immediate neighbors.

When the task is to name the flashed target word aloud instead of forced choice, there are no alternatives to restrict the decision process to only two counters. Null counts are randomly distributed among all counters, and the probability that any one of them is accumulated in the flashed target's counter is very low. Of those counts that are not null, some correspond to features that distinguish the target word from words similar to it, and these are accumulated by the target. Other counts correspond to features that distinguish the target and the other words in its cohort from all the other words in the lexicon; each word in the cohort has an equal chance of accumulating these counts, if no word in the cohort was previously studied. Unlike forced choice where all counts are accumulated by one of only two counters, in naming counts can be accumulated by any counter in the lexicon, so there may never be sufficient counts in any one counter to exceed the criterion number, and a stopping rule is needed. The stopping rule used in the tests of the model by Ratcliff and McKoon (1996) was based on an evaluation of the numbers of counts in counters at three discrete points: when a total of 35, 100, or 400 counts had entered the system. At each of these points, if no counter had accumulated a significant number of counts (e.g., at the 100-count point, if no counter had accumulated at least 12 counts), then processing was terminated (see Ratcliff & McKoon, 1996, for details and also a discussion of how this discrete rule could be made more continuous). Terminating processing with the stopping rule corresponds to a subject in an experiment saying "no" when no response can be made.

In naming, as in forced choice, prior study leads to theft. The counter of a studied word steals counts from other counters in its cohort, and the counters for the words in the studied word's cohort steal counts away from counters outside the cohort. Ratcliff and McKoon (1996) showed that only a small increase in the probability of the studied word's accumulating a count was needed to correctly predict data. Combining the two sources of theft, the increase was .015, about the same size as the increase due to prior study for forced choice.

As demonstrated by Ratcliff and McKoon (1996), the assumptions of the counter model are sufficient to give quantitative as well as qualitative explanations of the bias patterns in forced choice and naming, why bias appears in forced choice for similar but not dissimilar alternatives, and why the size of bias effects does not decrease as the duration of the flashed word decreases. The model, as its origin in the logogen model would suggest, is designed to explain bias in the context of a more general model for word identification. As such, it must explain not only bias but also standard effects in the word-identification literature. Ratcliff and McKoon (1996) showed its successful account of word frequency and neighborhood effects.

The counter model provides an explanation of the processes that give priming in an implicit task, and in doing so, it addresses the dual problems of explaining bias in word identification and answering the question of whether separate memory systems offer potentially useful explanations of priming data. The counter model does not have multiple memory systems. Instead, it is designed to be part of the prototypical flow of information processes shown in Figure 1. The counter model would be located at the more perceptual end of processing. Other tasks that involved other levels of information such as semantic or elaborated episodic information would reside at later levels and be tapped by different tasks such as recognition or recall, and there would be models of processing at those levels (e.g., Gillund & Shiffrin, 1984; Hintzman, 1986; Murdock, 1982; Ratcliff & McKoon, 1988). It would be expected that different variables would affect the different processes; for example, the characteristics of a word that made it easier to identify than some other word would not correlate with the aspects of meaning that made it easier to memorize. Therefore, dissociations would be expected. It would also be anticipated that amnesic patients could have impairments in some parts of the system, but not other parts.

Conclusion

Object identification, picture naming, stem completion, fragment completion, visual word identification, and auditory word identification all show bias patterns in priming. The finding that a large collection of implicit tasks shows a bias pattern suggests a communality of explanations for the tasks. We have proposed that these explanations can reside in traditional information-processing theories that emphasize the mechanisms by which information is processed over time, and we have supported the proposal with an explicit counter model for one task, perceptual identification (Ratcliff & McKoon, 1996). However, there will be no single mechanism for bias across all the tasks; rather it is likely that there will be related but different mechanisms for bias depending on the particular processes needed to perform a task. We see nothing to rule out mechanistic models like the counter model in favor of the alternative approach of postulating multiple memory systems. Instead, there is much to learn from both perspectives and from competitive development of theories. Even neuropsychological research that has attempted to locate the brain structures that embody memory systems can profitably interact with research that has attempted to examine the processes involved in perceptual and other information-processing tasks.

In this light, what is surprising is the extent to which implicit memory research and the information-processing tradition of modeling have proceeded in isolation from each other. Articles from one domain rarely cite the other. Our hope is that the generality of the bias results demonstrated in this article will demand an explanation in terms of explicit modeling and so empower an interchange that will bring the two domains together.

Finally, it should be pointed out that, from an ecological perspective, biases that reflect prior exposure are likely to be a good thing: We are likely to encounter objects that are the same as those we have previously encountered and not so likely to encounter objects that are similar but require a different response such that, on average, perceptual processes that produce biases will be adaptive. From the perspective of cognitive theory, biases need explanation, and we urge interactions between two approaches: the postulation of multiple memory systems and the development of information-processing models.

References

Andrews, S. (1992). Frequency and neighborhood effects on lexical access: Lexical similarity or orthographic redundancy? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 234–254.

Biederman, I., & Cooper, E. E. (1991). Evidence for complete translational and reflectional invariance in visual object priming. *Perception, 20,* 585–593.

Boring, E. G. (1930). A new ambiguous figure. *American Journal of Psychology, 42,* 444–445.

Broadbent, D. E. (1958). *Perception and communication.* London: Permagon.

Broadbent, D. E. (1967). Word frequency effect and response bias. *Psychological Review, 74,* 1–15.

Cave, C. B., & Squire, L. R. (1992). Intact and long-lasting repetition priming in amnesia. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 509–520.

Cohen, N. J., & Squire, L. R. (1980). Preserved learning and retention of pattern analyzing skill in amnesia: Dissociation of knowing how and knowing that. *Science, 210,* 207–209.

Coltheart, M., Davelaar, E., Jonasson, J. T., & Besner, D. (1977). Access to the internal lexicon. In S. Dornic (Ed.), *Attention and Performance, VI* (pp. 535–555). Hillsdale, NJ: Erlbaum.

Corkin, S. (1965). Tactually guided maze-learning in man: Effects of unilateral cortical excisions and bilateral hippocampal lesions. *Neuropsychologia, 3,* 339–351.

Corkin, S. (1968). Acquisition of motor skill after bilateral medial temporal excision. *Neuropsychologia, 6,* 255–265.

Edelman, S., & Poggio, T. (1992). Bringing the grandmother back into the picture: A memory-based view of object recognition. *International Journal of Pattern Recognition and Artificial Intelligence, 6,* 37–61.

Edelman, S., & Weinshall, D. (1991). A self-organizing multiple-view representation of 3D objects. *Biological Cybernetics, 64,* 209–219.

Epstein, W., & Rock, I. (1960). Perceptual set as an artifact of recency. *American Journal of Psychology, 73,* 214–228.

Gabrieli, J. D. E., Milberg, W., Keane, M. W., & Corkin, S. (1990). Intact priming of patterns despite impaired memory. *Neuropsychologia, 28,* 417–428.

Gardner, H., Boller, F., Moreines, J., & Butters, N. (1973). Retrieving information from Korsakoff patients: Effects of categorical cues and reference to the task. *Cortex, 9,* 165–175.

Gibson, J. M., & Watkins, M. J. (1988). A pool of 1,086 words with unique two-letter fragments. *Behavioral Research Methods, Instruments, & Computers, 20,* 390–397.

Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review, 91,* 1–67.

Goldinger, S. D., Luce, P. A., & Pisoni, D. B. (1989). Priming lexical neighbors of spoken words: Effect of competition and inhibition. *Journal of Memory and Language, 28,* 501–518.

Graf, P., & Mandler, G. (1984). Activation makes words more accessible, but not necessarily more retrievable. *Journal of Verbal Learning and Verbal Behavior, 23,* 553–568.

Graf, P., Squire, L., & Mandler, G. (1984). The information that amnesic patients do not forget. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10,* 164–178.

Hintzman, D. L. (1984). MINERVA 2: A simulation model of human memory. *Behavior Research Methods, Instruments, & Computers, 16,* 96–101.

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

Hintzman, D. L. (1987). Recognition and recall in MINERVA2: Analysis of the "recognition failure" paradigm. In P. E. Morris (Ed.), *Modelling cognition* (pp. 215–229). London: Wiley.

Hintzman, D. L. (1990). Human learning and memory: Connections and dissociations. *Annual Review of Psychology, 41,* 109–139.

Hintzman, D. L., & Hartry, A. L. (1990). Item effects in recognition and fragment completion: Contingency relations vary for different subsets of words. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 955–969.

Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review, 99,* 480–517.

Jacoby, L. L. (1983a). Perceptual enhancement: Persistent effects of an experience. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9,* 21–38.

Jacoby, L. L. (1983b). Remembering the data: Analyzing interactive processes in reading. *Journal of Verbal Learning and Verbal Behavior, 22,* 485–508.

Jacoby, L. L., & Dallas, M. (1981). On the relationship between autobiographical memory and perceptual learning. *Journal of Experimental Psychology: General, 3,* 306–340.

Jacoby, L. L., & Witherspoon, D. (1982). Remembering without awareness. *Canadian Journal of Psychology, 36,* 300–324.

Kolers, P. A. (1973). Remembering operations. *Memory & Cognition, 1,* 347–355.

Kolers, P. A. (1974). Remembering trivia. *Language and Speech, 17,* 324–336.

Kolers, P. A. (1975a). Memorial consequences of automatized encoding. *Journal of Experimental Psychology: Human Learning and Memory, 1,* 689–701.

Kolers, P. A. (1975b). Specificity of operations in sentence recognition. *Cognitive Psychology, 7,* 289–306.

Kolers, P. A. (1976). Reading a year later. *Journal of Experimental Psychology: Human Learning and Memory, 2,* 554–565.

Kolers, P. A., & Ostry, D. J. (1974). Time course of loss of information regarding pattern analyzing operations. *Journal of Verbal Learning and Verbal Behavior, 13,* 599–612.

Kolers, P. A., & Roediger, H. L. (1984). Procedures of mind. *Journal of Verbal Learning and Verbal Behavior, 23,* 425–449.

LaBerge, D., & Samuels, S. J. (1974). Toward a theory of automatic information processing in reading. *Cognitive Psychology, 6,* 293–323.

Laming, D. R. J. (1968). *Information theory of choice reaction time.* New York: Wiley.

Leeper, R. (1935). A study of a neglected portion of the field of learning—The development of sensory organization. *Journal of Genetic Psychology, 46,* 41–75.

Link, S. W., & Heath, R. A. (1975). A sequential theory of psychological discrimination. *Psychometrika, 40,* 77–105.

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review, 8,* 375–407.

McKoon, G., & Ratcliff, R. (1995). How should implicit memory phenomena be modeled? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 777–784.

McKoon, G., Ratcliff, R., & Dell, G. S. (1986). A critical evaluation of the semantic/episodic distinction. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 295–306.

Milner, B. (1962). Les troubles de la memoire accompagnant des lesions hippocampiques bilaterales [Memory impairments associated with bilateral hippocampic lesions]. *Physiologie de l'hippocampe* (pp. 257–272). Paris: Centre National de la Recherche Scientifique.

Mitchell, D. B., & Brown, A. S. (1988). Persistent repetition priming in picture naming and its dissociation from recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 213–222.

Morton, J. (1968). A retest of the response bias explanation of the word frequency threshold effect. *British Journal of Mathematical and Statistical Psychology, 21,* 21–33.

Morton, J. (1969). Interaction of information in word recognition. *Psychological Review, 76,* 165–178.

Morton, J. (1970). A functional model for memory. In D. A. Norman (Ed.), *Models of human memory* (pp. 203–254). New York: Academic Press.

Morton, J. (1979). Facilitation in word recognition: Experiments causing change in the logogen model. In P. A. Kolers, M. E. Wrolstad, & H. Bouma (Eds.), *Processing visible language 1* (pp. 259–268). New York: Plenum.

Moscovitch, M. (1992). Memory and working-with-memory: A component process model based on modules and central systems. *Journal of Cognitive Neuroscience, 4,* 257–267.

Moscovitch, M. (1994). Memory and working with memory: Evaluation of a component process model and comparisons with other models. In D. L. Schacter & E. Tulving (Eds.), *Memory Systems 1994* (pp. 269–310). Cambridge, MA: MIT Press.

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

Neisser, U. (1954). An experimental distinction between perceptual process and verbal response. *Journal of Experimental Psychology, 47,* 399–402.

Neisser, U. (1967). *Cognitive psychology.* New York: Appleton-Century-Crofts.

Nosofsky, R. M. (1988). Exemplar-based accounts of relations between classification, recognition, and typicality. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 700–708.

Ostergaard, A. L. (1992). A method for judging measures of stochastic dependence: Further comments on the current controversy. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 413–420.

Ostergaard, A. L., & Jernigan, T. L. (1993). Are word priming and explicit memory mediated by different brain structures? In P. Graf & M. E. J. Masson (Eds.), *Implicit memory: New directions in cognition, development, and neuropsychology* (pp. 327–349). Hillsdale, NJ: Erlbaum.

Perruchet, P., & Baveux, P. (1989). Correlational analyses of explicit and implicit memory performance. *Memory & Cognition, 17,* 77–86.

Poggio, T., & Girosi, F. (1990). Regularization algorithms for learning that are equivalent to multilayer networks. *Science, 247,* 978–982.

Posner, M. I. (1969). Abstraction and the process of recognition. In G. Bower & J. T. Spence (Eds.), *Psychology of learning and motivation* (Vol. 3). New York: Academic Press.

Posner, M. I. (1978). *Chronometric explorations of mind.* Hillsdale, N. J.: Erlbaum.

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

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

Ratcliff, R., Allbritton, D. W., & McKoon, G. (in press). Bias in auditory priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition.*

Ratcliff, R., & McKoon, G. (1986). More on the distinction between episodic and semantic memories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 312–313.

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

Ratcliff, R., & McKoon, G. (1995). Bias in the priming of object decisions. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 754–767.

Ratcliff, R., & McKoon, G. (in press). A counter model for implicit priming in perceptual word identification. *Psychological Review.*

Ratcliff, R., McKoon, G., & Verwoerd, M. (1989). A bias interpretation of facilitation in perceptual identification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 378–387.

Richardson-Klavehn, A., & Bjork, R. A. (1988). Measures of memory. *Annual Review of Psychology, 39,* 475–543.

Roediger, H. L., III. (1990). Implicit memory: Retention without remembering. *American Psychologist, 45,* 1043–1056.

Roediger, H. L., III, & Blaxton, T. A. (1987). Effects of varying modality, surface features, and retention interval on priming in word fragment completion. *Memory & Cognition, 15,* 379–388.

Roediger, H. L., III, & McDermott, K. B. (1993). Implicit memory in normal human subjects. In F. Boller & J. Grafman (Eds.), *The Handbook of Neuropsychology* (Vol. 8, pp. 63–131). New York: Elsevier Science Publishers B. V.

Roediger, H. L., III, Weldon, M. S., & Challis, B. H. (1989). Explaining dissociations between implicit and explicit measures of retention: A processing account. In H. L. Roediger, III, & F. I. M. Craik (Eds.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (pp. 3–41). Hillsdale, NJ: Erlbaum.

Rozin, P. (1976). The psychobiological approach to human memory. In M. R. Rosenzweig & E. L. Bennet (Eds.), *Neural mechanisms of learning and memory* (pp. 1–48). Cambridge, MA: MIT Press.

Rubin, E. (1915). *Synsoplevede figurer, studier i psykologisk analyse* [Visually experienced objects: Studies in psychological analysis]. Copenhagen, Denmark: Gyldendal, Nordisk Forlag.

Rubin, E. (1921). *Visuell wahrgenommene Figuren: Studien in psychologischer analyse* [Visually perceived figures: Studies in psychological analysis]. Copenhagen, Denmark: Gyldendalske boghandel.

Rueckl, J. G. (1990). Similarity effects in word and pseudoword repetition priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16,* 374–391.

Schacter, D. L. (1990). Perceptual representation systems and implicit memory: Toward a resolution of the multiple memory systems debate. In A. Diamond (Ed.), *Annals of the New York Academy of Sciences: Vol. 608. Development and neural bases of higher cognitive functions* (pp. 543–572). New York: New York Academy of Sciences.

Schacter, D. L. (1992). Understanding implicit memory: A cognitive neuroscience approach. *American Psychologist, 47,* 559–569.

Schacter, D. L. (1994). Priming and multiple memory systems: Perceptual mechanisms of implicit memory. In D. L. Schacter & E. Tulving (Eds.), *Memory systems 1994* (pp. 233–268). Cambridge, MA: MIT Press.

Schacter, D. L., & Church, B. A. (1992). Auditory priming: Implicit and explicit memory for words and voices. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 915–930.

Schacter, D. L., & Cooper, L. A. (1995). Bias in the priming of object decisions: Logic, assumptions, and data. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 768–776.

Schacter, D. L., Cooper, L. A., & Delaney, S. M. (1990). Implicit memory for unfamiliar objects depends on access to structural descriptions. *Journal of Experimental Psychology: General, 119,* 5–24.

Schacter, D. L., Cooper, L. A., Delaney, S. M., Peterson, M. A., & Tharan, M. (1991). Implicit memory for possible and impossible objects: Constraints on the construction of structural descriptions. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17,* 3–19.

Schacter, D. L., Cooper, L. A., & Treadwell, J. (1993). Preserved priming of novel objects across size transformation in amnesic patients. *Psychological Science, 4,* 331–335.

Schacter, D. L., & Tulving, E. (1994). What are the memory systems of 1994? In D. L. Schacter & E. Tulving (Eds.), *Memory Systems 1994* (pp. 1–38). Cambridge, MA: MIT Press.

Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review, 96,* 523–568.

Shimamura, A. P. (1985). Problems with the finding of stochastic independence as evidence for multiple memory systems. *Bulletin of the Psychonomic Society, 23,* 506–508.

Shimamura, A. P. (1986). Priming effects in amnesia: Evidence for a dissociable memory function. *Quarterly Journal of Experimental Psychology, 38A,* 619–644.

Shimamura, A. P. (1993). Neuropsychological analyses of implicit memory: History, methodology, and theoretical interpretations. In P. Graf & M. E. J. Masson (Eds.), *Implicit memory: New directions in cognition, development, and neuropsychology* (pp. 265–285). Hillsdale, NJ: Erlbaum.

Smith, S. M., & Tindell, D. R. (1996). *Memory blocks in word fragment completion caused by involuntary retrieval.* Manuscript submitted for publication.

Squire, L. R. (1987). *Memory and Brain.* New York: Oxford University Press.

Squire, L. R. (1992). Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychological Review, 99,* 195–231.

Squire, L. R. (1994). Declarative and nondeclarative memory: Multiple brain systems supporting learning and memory. In D. L. Schacter & E. Tulving (Eds.), *Memory Systems 1994* (pp. 203–231). Cambridge, MA: MIT Press.

Stone, M. (1960). Models for choice reaction time. *Psychometrika, 25,* 251–260.

Tulving, E. (1983). *Elements of episodic memory.* New York: Oxford University Press.

Tulving, E. (1984). Multiple learning and memory systems. In K. M. J. Lagerspetz & P. Niemi (Eds.), *Psychology in the 1990's* (pp. 163–184). Amsterdam: Elsevier.

Tulving, E. (1986). What kind of a hypothesis is the distinction between episodic and semantic memory? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 307–311.

Tulving, E., & Schacter, D. L. (1990). Priming and human memory systems. *Science, 247,* 301–306.

Tulving, E., Schacter, D. L., & Stark, H. A. (1982). Priming effects in word-fragment completion are independent of recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 8,* 336–342.

Wallach, H. (1949). Some considerations concerning the relation between perception and cognition. *Journal of Personality, 18,* 6–13.

Wallach, H., O'Connell, D. N., & Neisser, U. (1953). The memory effect of visual perception of three-dimensional form. *Journal of Experimental Psychology, 45,* 360–368.

Warrington, E. K., & Weiskrantz, L. (1968). New method of testing long-term retention with special reference to amnesic patients. *Nature (London), 217,* 972–974.

Warrington, E. K., & Weiskrantz, L. (1970). The amnesic syndrome: Consolidation or retrieval? *Nature, 228,* 628–630.

Winnick, W. A., & Daniel, S. A. (1970). Two kinds of response priming in tachistoscopic recognition. *Journal of Experimental Psychology, 84,* 74–81.

Witherspoon, D., & Moscovitch, M. (1989). Stochastic independence between two implicit memory tasks. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 22–30.

Woodworth, R. S. (1938). *Experimental Psychology.* New York: Henry Holt & Company.

Appendix

Stimulus Materials for Fragment Completion Study

The following word pairs were used in Experiment 4. For each pair, the top word was the target, the bottom word was the distractor, and the uppercase letters in the top word were the fragment.

	PhoENiX	sUrVivaL	fLaMInGo	aFfOrDed
	appendix	juvenile	climbing	daffodil
tRaMWay	liQUiDs	hiGhWAy	oUtFOx	uRetHRa
framework	aqueduct	megawatt	buffalo	brethren
leNGtHs	sCiSsoRs	meMBraNE	fRaNtiC	kIdNeY
dinghy	accessory	semblance	arsenic	dignity
hYbRiD	pEtUNIa	luKewaRM	rHEtORic	SIliCoN
myriad	debunking	hookworm	theories	sicilian
CROqUet	rePrieVE	cOntEXts	maVeriCK	qUaRtET
croupier	captive	anorexia	livestock	laureate
sAPPhIre	eVEninGs	lEtTUCe	liQUOr	rHoMBuS
appetite	average	kentucky	sequoia	chambers
sPItTlE	fiLTraTE	sQUaSh	doWAGer	rUFFiAn
opiates	alternate	equals	sewage	bouffant
oUtSidER	mYstIqUe	cYlinDeR	cOcOnUT	nIrVAnA
crusader	symposium	bystander	workout	giveaway
disCOVer	caLYpsO	bLaRNEy	cAsHmeRE	stYLisH
alcoves	ballyhoo	clarinet	catheter	daylight
hELPFul	eQualiZE	chIMNeY	iNsoMNiA	hiBiSCuS
leapfrog	squeeze	alimony	pneumonia	biscuits
frIDaYs	aArdVArk	vICeROy	bEGoNIa	pEnDuLuM
acidity	caravan	hickory	legions	feudalism

Received June 26, 1995
Revision received December 22, 1995
Accepted May 29, 1996 ■