In Press, Journal of Experimental Psychology: Learning, Memory, & Cognition

Structural Processing and Implicit Memory for Possible and Impossible Figures

Pepper Williams
Yale University
Michael J. Tarr
Brown University

Previous investigations have shown that participants are biased to respond "possible" to studied items when asked to decide whether objects could or could not exist in an object possibility test. The present study clarifies and extends the concept of bias in implicit memory research in two ways. First, we show that participants are biased to respond "possible" (rather than "impossible") on the object possibility test because structural processing is facilitated by prior study of possible, but not impossible, portions of objects. Second, we demonstrate that bias in this context is a form of, not an alternative to, implicit memory, by showing priming effects in response times when accuracy scores for studied and unstudied items are equated. We conclude by comparing proceduralist and memory-systems accounts of implicit memory effects, and suggest that the two approaches can be seen as complementary, rather than conflicting.

In the *object possibility test*, an implicit memory test introduced by Schacter, Cooper, and Delaney (1990)[1], participants decide if line drawings represent "possible figures," that could potentially exist as three-dimensional objects, or "impossible figures," that could not be instantiated in three dimensions (see Figure 1). Some of the figures on the test are also presented during an earlier study task, whereas other figures are completely new; as with other implicit memory tasks, memory for studied items can be inferred if participants perform differently on studied and unstudied test items.

An extensive series of investigations by Schacter, Cooper, and their colleagues have contrasted implicit memory effects for novel objects, as measured by object possibility tests, with explicit memory effects, as measured by old/new recognition tests. Performance on the two tasks has been compared across various encoding conditions (Schacter et al., 1990; Schacter & Cooper, 1993), participant populations (Schacter, Cooper, Tharan, & Reubens, 1991; Schacter, Cooper, & Valdiserri, 1992), stimulus transformations (Cooper, Schacter, Ballesteros, & Moore, 1992), and figure types (Schacter, Cooper, Delaney, Peterson, & Tharan, 1991). The principal finding from these studies has been that experimental manipulations often have significant effects on recognition performance, but do not significantly affect object possibility priming (on occasion, the opposite dissociation has been observed—significant effects of an independent variable on priming but not on recognition). For example, testing participants on pictures of objects that are larger or smaller than studied pictures significantly hinders recognition performance, but does not significantly affect the size of possibility priming effects (Cooper et al., 1992; Williams, 1995).

Based on these results, Schacter and Cooper have argued that separable memory systems underlie performance on these implicit and explicit memory tests. More specifically, Cooper and Schacter (1992, p. 145) proposed that recognition performance is mediated by an episodic memory system "dedicated to coding the semantic and visual information that creates distinctive representations of individual objects," whereas possibility priming is mediated by a *structural description system* (SDS) that "represents the global three-dimensional organization of parts of an object."

Ratcliff and McKoon (1995) have recently challenged this proposal. Although Ratcliff and McKoon's criticisms extend to the entire memory-systems approach to implicit memory research (Ratcliff & McKoon, 1997a, 1997b), the primary focus of their critique regarding the object possibility paradigm was Schacter and Cooper's findings involving performance on possible versus impossible figures. In their experiments, Schacter, Cooper, and colleagues have consis-

---

[1] Schacter, Cooper, and their colleagues dubbed this task the *object decision test*, but we prefer the more specific term object possibility test, since other decisions about objects have been usefully employed in other experimental paradigms (e.g. Kroll & Potter, 1984).

This research was supported by Air Force Office of Scientific Research Grant F49620-91-J0169 to MJT and a National Science Foundation Graduate Fellowship to PW. We thank Robert G. Crowder, William G. Hayward, Isabel Gauthier, and Vlada Aginsky for their detailed and thoughtful comments. Experiments 1 and 3a were first reported in a poster given at the 35th Annual Meeting of the Psychonomics Society, St. Louis, MO, Nov. 11–13, 1994. Requests for copies of the stimuli used in the experiments and other correspondence should be addressed to PW, who is now at the Psychology Department, University of Massachusetts-Boston, 100 Morrissey Blvd., Boston, MA 02125, email: *pepper.williams@yale.edu*. Version 7.0

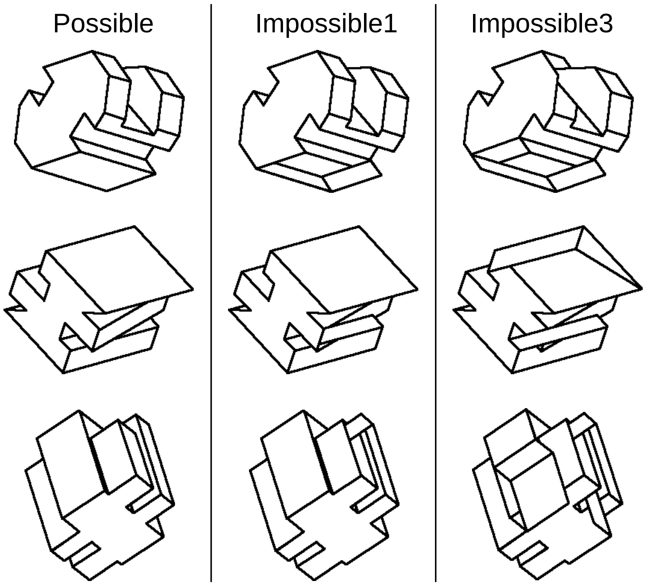

Figure 1. Examples of possible, impossible1, and impossible3 figures. Impossible figures could not actually exist as three-dimensional objects; impossible1 figures have one structural violation, while impossible3 figures have approximately three structural violations.

tently found significant priming for possible figures, but consistently failed to find significant priming for impossible figures (e.g. Schacter, Cooper, Delaney, Peterson, & Tharan, 1991). According to Cooper and Schacter (1992, p. 144), "the failure of priming to be exhibited in the case of impossible objects reflects constraints on the computational capabilities of the structural description system." In other words, since global representations of impossible objects are, by definition, not computable, the SDS cannot encode global information about impossible figures. Therefore, "impossible" decisions cannot be primed.

As an alternative to this interpretation, Ratcliff and McKoon (1995) suggested that possibility decisions are affected by study encounters in two different ways, as diagrammed in Figure 2. First, participants tend to respond "possible" more often to objects they have studied than to unstudied objects (Figure 2a). The same tendency is evidenced for both possible and impossible test items, so on its own, this "bias effect" would lead to positive priming for possible test figures, but negative priming for impossible test figures (since "possible" responses are incorrect when given to impossible items). *Such a pattern should still be classified as an implicit memory effect*, in the sense that test response rates are different for studied and unstudied items. However, this pattern varies from that predicted by the traditional definition of priming (e.g., Tulving & Schacter, 1990), in that performance is not necessarily *improved* by prior study (see Ratcliff, McKoon, & Verwoerd, 1989 and Ratcliff, Albritton, & McKoon, 1997 for examples of other implicit memory tasks that show bias effects).

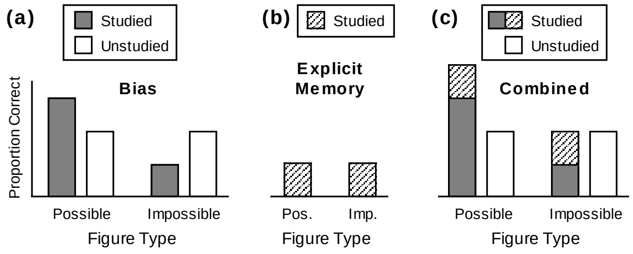

Figure 2. An illustration of Ratcliff and McKoon's (1995) two-process model of object decision priming.

The second component of Ratcliff and McKoon's (1995) model involves explicit memory for test items. Participants may remember "some particular configuration of corners, angles, or twists from an object that is associated with information about whether the object is possible or impossible" (Ratcliff & McKoon, 1995, p. 758). This information would help participants to make correct possibility decisions for all studied objects, boosting performance for both possible and impossible test items (Figure 2b). As shown in Figure 2c, the two components of Ratcliff and McKoon's model can combine to produce the pattern found by Schacter, Cooper, and their colleagues: robust priming for possible test figures combined with null priming for impossible test figures.

To test their model, Ratcliff and McKoon (1995) designed object possibility experiments that included manipulations intended to eliminate the availability and/or usefulness of explicit memory processes. Possibility decisions, they predicted, would therefore only be affected by the bias component of their model, resulting in better accuracy rates for studied than unstudied possible figures, but worse accuracy for studied than unstudied impossible figures (the pattern shown in Figure 2a). This prediction was confirmed by experiments that forced participants to carry an extra memory load (limiting the resources available for explicitly remembering the objects), make possibility decisions very quickly (precluding relatively slow explicit processes from working), or identify possible and impossible figures that looked very similar to each other (making explicit memory for the objects less useful). The experiments reported in the present paper make use of this last strategy: objects were studied in either a possible or an impossible version (Figure 1; the fact that there were actually two different impossible versions will be explained below), then the same object was presented at test in both its possible *and* its impossible form. Since possible and impossible versions are quite similar to each other, familiar-looking objects on the possibility test could not be assumed to be either possible or impossible, even if participants explicitly remembered that the object was seen in one version or another at study. As we will see, this "switching" of possible and impossible versions between study and test also allows effective tests of the hypotheses to be examined in these experiments.

Response biases found in connection with an explicit memory experiment are typically assumed to be caused by demand characteristics of the particular testing situation, and thus taken to be relatively uninteresting. On an implicit memory test, however, the fact that *responses* are biased differentially for studied and unstudied items could indicate that some cognitive *process* is being biased to operate differently on previously-encountered stimuli. This point was also made by Ratcliff et al. (1989, p. 379) in their study of bias in the perceptual identification test: "when we talk about bias, we do not mean a strategic postperceptual response bias. Rather, any bias effects could be perceptual or semantic in nature, depending on the task, and they could be the result of unconscious information processes."

If object possibility priming is caused by such a *perceptual bias*, then what is the nature of this bias? Ratcliff and McKoon (1995, p. 756) suggested that possibility decisions might be performed by constructing representations in an object recognition system, and that "bias because of previous study of an object could be based on modifications to processes engaged in building these representations." In the main, we agree with this assessment, but it begs another question: *exactly which processes are modified in order to produce the observed bias?* Visual object recognition begins with an array of points of light detected on the retinae, goes through a number of intermediate perceptual processes, and ends when a representation of the perceived object has been matched with a representation in memory. Ratcliff and McKoon (1995) made no hypothesis as to which of these processes is responsible for the perceptual bias found on the object possibility task.

Let us consider a simple model of object possibility test performance in which "structural evidence" is gathered for each test item. This model will both help us in evaluating how different types of perceptual bias would affect object possibility test performance, and help relate the concept of perceptual bias to the more traditional meaning of the term *bias* in signal detection theory (SDT). The guiding assumption here is that participants perform the task by attempting to analyze the three-dimensional structure of the test object, and that the successful analysis of a possible portion[2] of an object constitutes structural evidence (see Enns & Rensink, 1990, for evidence that such a *structure-extraction* process occurs early and automatically in visual perception). Since even impossible objects have some possible portions, participants will be able to glean some structural evidence from every test item, although more possible information must always be available from possible than from impossible items. In SDT terms, then, we can treat the possible information available from impossible items as the "noise," and the extra possible information available from possible test items as the "signal," as represented in the probability distributions shown in Figure 3 (compare this figure with Goldstein, 1996, Figure E.3). As in standard SDT, we assume that participants establish a criterion specifying a certain amount of evidence (i.e., somewhere along the X-axis in Figure 3), and respond "possible" if the amount of evidence obtained from a test item falls to the right of this criterion, or "impossible" otherwise.

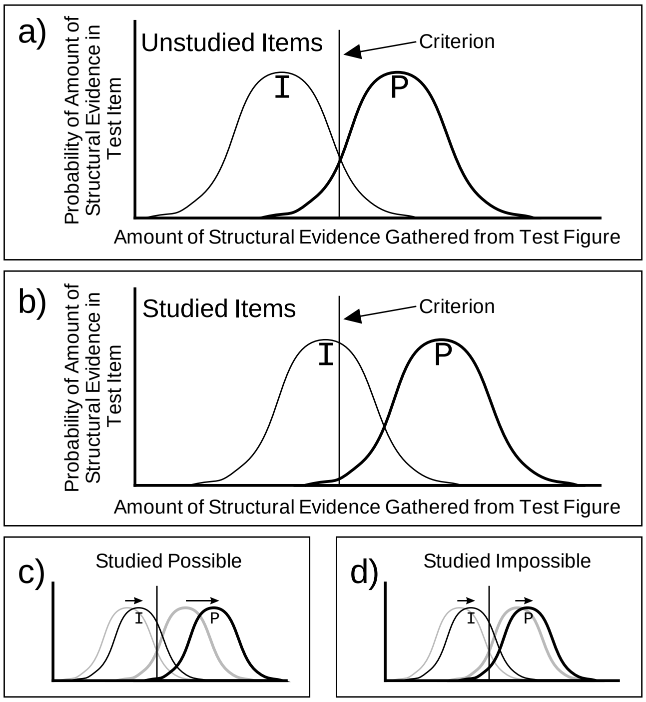

Figure 3. A signal-detection theory model of object possibility performance.

Figure 3a shows probability distributions and a criterion setting for unstudied test items. Ratcliff and McKoon (1995) showed that, when explicit memory processes are precluded, priming on the object possibility test is reflected in participants' response biases. In SDT terms, this implies a criterion shift for studied compared to unstudied items—that is, more studied than unstudied items fall to the right of the criterion. Usually, criterion shifts are caused by postperceptual processes in which participants strategically place their criterion at one value or another (depending, for example, on payoff matrices). Such mundane strategic processes seem unlikely to be able to account for priming on an implicit memory task, since they would require conscious registration on the part of participants as to whether an item is studied or unstudied (indeed, in Ratcliff and McKoon's experiments, such explicit memory was suppressed by various mechanisms). Instead, we propose that the criterion remains in the same absolute

---

[2] We use the term "portions" of objects to refer to groups of lines that can be parsed into groups of surfaces that indicate valid (possible portions) or invalid (impossible portions) 3-D structure. How, or even if, the visual system extracts surfaces and parts from line drawings are questions beyond the scope of this paper (see Waltz, 1975, Barrow & Tenenbaum, 1981, Biederman, 1987, and Bülthoff, Edelman, & Tarr, 1995, for discussions of these issues). All we mean to imply by the use of these terms is that some portions of impossible figures have structurally valid interpretations, while other portions of such figures are not interpretable in three dimensions.

location for all test items, but that the probability distributions for studied items shift to the right (Figure 3b), reflecting the action of a perceptual bias that causes more structural evidence to be gathered for studied than for unstudied items.

Our assumption that distributions shift as a result of prior study is somewhat analogous to the idea in the recognition-memory literature that different familiarity distributions arise for different types of stimuli (i.e., high- and low-frequency words). However, in recognition-memory experiments the signal is defined as the increase in familiarity resulting from prior study. In the context of the object possibility task, we propose that the signal (the extra structural evidence available from possible compared to impossible objects) is present in the test stimulus regardless of whether the item was studied or not. Nevertheless, we propose, prior study affects performance by shifting the evidence distributions for studied items compared to unstudied items. Note that if the distributions for both possible and impossible test objects shift in tandem, then overall accuracy (sensitivity) will not vary for studied and unstudied objects, as observed by Ratcliff and McKoon (1995).

So far, we have clarified the concept of perceptual bias, but still have not explained exactly what perceptual processes are biased in the object possibility task. At least two types of processes are viable options. First, consider a hypothetical perceptual bias in a process yielding information that would not distinguish between a possible and impossible object, such as the object's outline shape (Hayward, 1997, has shown that outline-shape information may be regularly used in object recognition processes). Although such a process would not directly increase the amount of structural evidence available from studied objects, the "perceptual fluency" (Jacoby, 1983) resulting from what we will call a non-diagnostic perceptual bias might be interpreted as evidence for the positive, "possible" response (much as perceptual fluency for previously-seen names in the experiments of Jacoby, Woloshyn, & Kelley, 1989, presumably leads to the "false fame" effect). We will refer to this as the *perceptual-fluency hypothesis*, but it should be understood that underlying the perceptual fluency is a perceptual bias in processing non-diagnostic information in studied objects.

Second, consider a perceptual bias in a process that operates only (or at least most efficiently) on possible portions of objects. The process we are thinking of here is the structure-extraction process itself. Since impossible portions of objects, by definition, do not have valid three-dimensional structures, these portions would be difficult for a structure-extraction mechanism to process during the study phase of an experiment. Therefore, the mechanism would only be biased to process possible portions of studied objects more efficiently during the test phase, resulting again in more structural evidence for studied than for unstudied test items. We will refer to this as the *structure-extraction hypothesis*.

Both the perceptual-fluency and structure-extraction hypotheses account for the distribution shifts (and resultant relative criterion shift) in Figure 3b, and thus both are consistent with Ratcliff and McKoon's, 1995, bias-plus-explicit-memory proposal. However, the two hypotheses can be distinguished by their predictions regarding the effects of *studying* different types of objects. Specifically, since only possible portions of objects contain valid structural information, the structure-extraction hypothesis predicts that priming effects should be larger following study of possible than following study of impossible objects. This is shown in Figures 3c-d: when an object was studied in its possible version, a large perceptual bias should result, and the distributions should shift relatively far to the right. For objects studied in their impossible versions, the perceptual bias should be relatively small, resulting in a relatively smaller distribution-shift.

The structure-extraction hypothesis also predicts at least a small difference in sensitivity between items studied in possible and impossible versions. Consider an object that was studied in its impossible version. Regardless of what version the object is tested in, only the possible portions of the test object will be primed. Thus the distribution-shifts in Figure 3d should be equivalent for possible and impossible test items, as shown. However, an object studied in its possible version should prime its possible test version more than its impossible test version, because the former contains more possible structure to be primed than the latter. As a result, the possible distribution will shift more than the impossible distribution, causing both a change in bias and sensitivity compared to unstudied items (Figure 3c).

In contrast to these predictions of the structure-extraction hypothesis, an outline-shape or other perceptual bias that does not operate differentially on possible and impossible portions of objects will be equivalent for possible and impossible studied objects, resulting in the same amount of perceptual fluency and therefore the same degree of distribution-shifting. This hypothesis thus predicts that items studied in possible and impossible versions should show equivalent bias changes relative to unstudied items, and that sensitivity should be identical for all items (studied possible, studied impossible, and unstudied).

The data collected by Ratcliff and McKoon (1995) seem to support the perceptual-fluency hypothesis, since priming effects for possible and impossible studied figures were statistically indistinguishable in their experiments. However, Ratcliff and McKoon (1995) were not specifically looking for differences between study conditions, and did not design their experiments in a way that would be conducive to such a finding. Given unlimited time, participants agreed that Ratcliff and McKoon's possible figures were possible only 78% of the time, and that their impossible figures were impossible only 76% of the time. This relatively low level of agreement indicates that the structural integrity of Ratcliff and McKoon's possible objects may not have been much greater than that of their impossible figures—in other words, the structural evidence obtainable from possible objects may not have been substantially greater than the structural evidence obtainable from impossible objects. This would not be a problem for demonstrating different biases on studied and unstudied objects (since a perceptual bias would affect whatever structural evidence was available from any studied figure). However, compressing the range of structural information available from possible vs. impossible objects would lessen Ratcliff and McKoon's chances of detecting an effect of figure impossibility on the magnitude of priming effects.

To provide a more sensitive test of the structure-extraction hypothesis, we developed a new set of objects in which the structural integrity of possible and impossible figures was less disputable. In Experiment 1, we had participants study three types of objects: *possible* objects, that could clearly exist in three dimensions; *impossible1* objects, that each had one distinct impossible portion; and *impossible3* objects, that each had three impossible portions (Figure 1). As reported in the Methods section of this experiment, 92% of participants agreed that our possible objects were possible, and 94% that our impossible3 objects were impossible. Other aspects of the experiment were similar to Ratcliff and McKoon's (1995) Experiment 6: the study task was followed by an object possibility test in which objects were shown very briefly (45 ms each), and by matching possible and impossible objects, we "disabled" the explicit memory component of Ratcliff and McKoon's model. (That is, objects that were studied in an impossible version were later tested in both possible and impossible3 versions, so explicit memory that an object was impossible at study would not be very helpful in determining whether the object is possible or impossible at test.)

If priming effects on the object possibility task are caused by a perceptual bias operating solely on possible portions of studied figures, then the magnitude of these priming effects should depend on the amount of possible structural information available to be encoded from studied figures. This structure-extraction hypothesis thus predicts the largest priming effects (i.e., the largest bias to respond "possible") for objects studied as possible figures, somewhat less priming for objects studied as impossible1 figures, and the least amount of priming for objects studied in their impossible3 forms. The alternate hypothesis is that object possibility priming is caused by a non-diagnostic perceptual bias and perceptual fluency. This hypothesis predicts that, as in Ratcliff and McKoon's (1995) Experiment 6, all types of studied figures should lead to approximately equivalent tendencies to respond "possible."

In keeping with the SDT-based model proposed in Figure 3, we chose to analyze the results in terms of signal-detection measures of discrimination (d_L) and bias (C_L)[3], by treating correct "possible" responses (to possible test items) as hits and incorrect "possible" responses (to impossible test items) as false alarms. Although not absolutely necessary to test the two hypotheses in question, this data transformation allowed us to evaluate predictions of the hypotheses more easily, because results for possible and impossible test items were combined into a single measure of response bias. The perceptual-fluency hypothesis predicts equivalent values of C_L for different types of studied figures, whereas the structure-extraction hypothesis predicts the lowest value of C_L (negative values indicate a bias to respond "possible") for possible studied figures, a smaller bias for impossible1 studied figures, and the smallest bias for impossible3 studied figures. For sensitivity, the perceptual-fluency hypothesis predicts equivalent values of d_L for different types of studied figures, whereas the structure-extraction hypothesis again predicts a graded effect: the largest d_L should be found for studied possible objects, and progressively smaller values should be found for impossible1 and impossible3 objects.

Norris (1995) has recently criticized the use of SDT measures in evaluating other "bias models," such as Morton's (1969) logogen model (Ratcliff and McKoon's, 1997b, counter model is a close cousin of the logogen model). Norris' main criticism is that the logogen model and its relatives violate a fundamental assumption of SDT, that a single criterion be used for all stimuli in a test sequence. Therefore, standard interpretations of SDT statistics may not be applicable when evaluating the predictions of these models, even though the models make reference to the concept of bias. In the model we have proposed here, our assumptions (i.e., signal and noise distributions of evidence strength with a single response criterion) are tied directly to those of standard SDT, and thus our application of SDT statistics is appropriate.

Experiment 1

Method

Participants. Thirty-eight Yale undergraduates from the introductory psychology pool served as participants, but data from two participants were discarded because the participants performed at below-chance levels on the object possibility test.

Materials. Line drawings of 40 possible objects were created on a Macintosh computer using a commercial drawing program (Figure 1). Two different impossible versions of each object were created by adding or removing lines in such a way that the resulting figures appeared as if they could not be instantiated in three dimensions. Impossible1 figures were designed to be impossible in only one portion of the figure, leaving the rest of the drawing possible, whereas impossible3 figures have approximately three impossible portions. Each experiment utilized 36 of the 40 objects; exactly which objects were used varied somewhat from experiment to experiment.

---

[3] As described by Snodgrass and Corwin (1988), the signal detection model based on logistic distributions, which uses d_L and C_L as measures of sensitivity and bias, is functionally equivalent and considerably easier to calculate than the more traditional model based on normal distributions, which uses d' and β.

Agreement about the possible/impossible nature of each figure was assessed by giving 30 additional participants unlimited time to classify objects as possible or impossible, with instructions to take as much time as needed and be as accurate as possible. Each participant classified three-quarters of the figures; the number of participants rating each individual figure varied from 14 to 28. On average, possible figures were classified "possible" by .92 of participants (range: .79 to 1.00; standard deviation: .075), impossible1 figures were classified "impossible" by .88 of participants (.59 to 1.00; .098), and impossible3 figures were classified "impossible" by .94 of participants (.80 to 1.00; .059). Although the agreement rate was somewhat lower for impossible1 figures than for the other two versions, we note that participants in Experiments 1-3 saw only possible and impossible3 figures on the object possibility test, so agreement on figures for which participants actually made possibility decisions was very high.

Stimuli were shown on a Macintosh computer screen, in black lines on a white background, and were approximately 9 cm x 9 cm in Experiment 1 and approximately 7.5 cm wide by 7.5 cm high in Experiments 2, 3a, and 3b. Participants viewed the figures from approximately 50 cm away from the screen, resulting in visual angles of approximately 10.3° in Experiment 1 and 8.6° in Experiments 2, 3a, and 3b. Thirty-six of the 40 objects were selected for use in Experiment 1.

Procedure. Participants were initially instructed only that they would be making judgments and decisions about novel objects. In the study phase, 9 possible, 9 impossible1, and 9 impossible3 figures were shown in a random order for 5 s each. Participants were requested to decide what direction they thought each object faced, choosing among the 8 options of straight up, up and to the right, straight right, and so on. They entered their choice using the numeric keypad. No mention of a memory test was made, and the possible/impossible nature of the figures was not explained until the study task was completed.

Immediately following the study task, participants received directions for the object possibility task. They were informed that some figures in this task would represent "valid, possible three-dimensional objects that could exist in the world," whereas others would represent "impossible objects that could not actually exist in the real three-dimensional world." Participants were asked to decide whether each object was possible or impossible, and to respond as quickly and as accurately as possible. The test trials were preceded by 8 practice figures, to which participants made a possibility decision and were then given a second viewing of the figure, along with feedback. Very few participants expressed any difficulty understanding the possible/impossible distinction.

Test trials were computer-paced. Each trial began with a 2.5 s blank screen, then a 500 ms fixation cross. The stimulus then appeared for 45 ms, followed by a pattern mask (consisting of a crisscross pattern of lines the same width as those in the stimuli), which was shown for 500 ms. Participants pressed the "z" key if they thought the figure was possible, or the "m" key if they thought it impossible. The sequence for the next trial then began. Pilot testing indicated that the 45 ms exposure duration would allow participants to make accurate possibility decisions on our objects about 65-75% of the time, leaving room for both positive and negative priming effects. Experiments reported by Schacter and Cooper (1993) employed exposure times as small as 17 ms and as large as 100 ms to get baseline levels of performance similar to ours, whereas Ratcliff and McKoon (1995) displayed their test figures for between 150 and 250 ms. This rather wide range is presumably due to variations in the figure sets, participant populations, and computer monitors used in the different studies.

Participants viewed 72 critical test figures: 36 objects which were each shown in both their possible and impossible3 versions (having each object appear as both a possible and an impossible test figure allowed us twice as many observations per participant, and we hoped that effects of the first test exposure on the second would be minimal). The objects were organized into four groups that were rotated through the four study conditions (studied possible, studied impossible1, studied impossible3, and unstudied) between participants, so that each object participated equally often in each condition. Each participant viewed the stimuli in a different random order.

Participants were cautioned that two or more objects would sometimes look very similar, but that they should make each decision independently, because sometimes they would see two possible versions, sometimes two impossible versions, and sometimes one of each. Four filler objects (that had not been studied and were not subsequently analyzed) were shown twice in their possible version or twice in their impossible3 version, and the practice trials also included a possible and an impossible3 figure that were shown twice. These precautions were taken to ensure that participants could not assume that once they had seen a possible version of an object, the next similar object would be an impossible version (or vice-versa). Participants were not told that any test objects had been previously studied.

Design and statistical analyses. The dependent variable was accuracy, which was analyzed in terms of sensitivity (d_L) and bias (C_L) in the main analyses. Independent variables of interest included Studied Figure (studied possible, studied impossible1, studied impossible3, or unstudied) and Test Figure (possible or impossible3), both of which were manipulated within-participants. Throughout this paper, separate signal detection measures were calculated for each subject, and means were then calculated across subjects. Reported ANOVAs treat participants as the random factor. Separate analyses using items as the random factor yielded the same patterns of results and generally smaller p values, in all experiments. An α level of .05 was adopted for the entire study;

Table 1
Object possibility performance: Experiment 1

	Studied Figure			
Tested Figure	Possible	Impossible1	Impossible3	Unstudied
Possible	0.80	0.76	0.73	0.70
Impossible3	0.66	0.70	0.69	0.75
Sensitivity (d_L)	2.06	2.11	1.78	1.95
Bias (C_L)	-0.33	-0.14	-0.12	0.12

specific p values are reported only in exceptional cases.

Results

Each participant was tested on both the possible and the impossible3 version of each test object, so a preliminary analysis was run to determine whether order effects were present in the data. The majority of first test exposures occurred in the first half of the test, whereas the majority of second test exposures occurred in the second half, so both of these factors—Test Order (first or second) and Trial Number (trials 1-40 vs. trials 41-80)—were included in ANOVAs, along with the factor of Test Figure. Neither main effect approached significance, nor did they interact with each other or with Test Version (possible or impossible3), all $F < 1.06$. Since participants saw each test figure for only 45 ms, it is not surprising that the first test exposure had no significant effect on the second in this experiment. All other analyses were collapsed over these factors.

Accuracy scores in Experiment 1 reveal a clear bias to respond "possible" more often to studied than to unstudied objects (see Table 1). For possible test objects, participants were more accurate if the object was studied than if it was unstudied, whereas for impossible test figures, participants were less accurate if they had seen the object during the study task. The significance of this effect was confirmed by a significant main effect of Studied Figure on bias (C_L) scores, $F(3, 105) = 5.643$, $MS_e = 0.222$. Planned one-tailed t- and sign tests indicated that participants were significantly more biased to respond "possible" in each of the studied conditions (studied possible, studied impossible1, and studied impossible3) than they were in the unstudied condition, all $t(35) > 1.91$; all $z > 2.00$. Additionally, a two-tailed t-test indicated that objects studied in possible versions showed a larger bias effect than objects studied in impossible3 versions, $t(35) = 2.25$. Participants did not show significant differences in sensitivity (d_L) for different Studied Figure conditions, $F(3, 105) = 0.580$, $MS_e = 1.267$. Furthermore, the largest single difference between sensitivity scores, possible-studied vs. impossible3-studied objects, was not significant by either a t- or sign test ($t(35) = 1.08$; $z = 0.17$).

Discussion

Experiment 1 replicated two important results from Ratcliff and McKoon's (1995) Experiment 6. First, studying impossible versions of objects clearly affected performance for these objects on the object possibility task, demonstrating that priming can be found for impossible figures, despite the many failures to find such effects in Schacter and Cooper's studies (e.g. Schacter, Cooper, Delaney, Peterson, & Tharan, 1991). Second, positive priming effects for possible test figures were balanced by negative priming effects for impossible test figures, resulting in a bias to respond "possible" to all studied objects.

A third finding from Experiment 1 is that, unlike in Ratcliff and McKoon's (1995) experiments, studying different types of figures led to different degrees of bias. Ratcliff and McKoon (1995, Experiment 6) found approximately equal amounts of priming for test objects that were studied as possible and impossible figures. In contrast, our data indicate that for the present stimulus set, the tendency to respond "possible" was stronger for possible than for impossible1 or impossible3 studied figures. As stated in the General Introduction, we assume that Ratcliff and McKoon's failure to find such a difference was due to the fact that the possible and impossible versions of their objects were not sufficiently differentiated. That is, if object possibility priming is due to a bias to extract structural information more efficiently from studied objects, and if Ratcliff and McKoon's possible objects did not contain significantly more structural information than their impossible objects, then it would not be surprising that they failed to detect significant differences between studied possible and impossible objects.

The present pattern of effects is particularly striking because this hypothesis successfully predicted a violation of the encoding specificity principle (Tulving & Thomson, 1973): possible studied figures led to the largest priming effects both when these figures were most similar to test items (possible test figures) and when they were least similar to test items (impossible3 test figures). (The results of Ratcliff and McKoon's Experiment 6 also hint at the latter effect—priming from possible studied figures to impossible test figures (.08) was slightly larger than priming from impossible studied figures to impossible test figures (.06)).

For unstudied objects, participants demonstrated a bias to respond "impossible," whereas in every other condition in Experiments 1, 2, and 3a, participants were biased to respond "possible." We speculate that this result reflects a kind of unconscious bias-correction process. Participants may want, overall, to respond "impossible" about as often as they respond "possible." However, they see so many studied objects, for which they gather extra evidence of possibility, that they end up experiencing a reverse bias, to respond "impossible," to unstudied objects, for which the structural evidence is comparatively weaker.

One potential objection to our interpretation of Experiment 1 is based on the observation that the three versions of our objects differed in ways other than their possibility. The most obvious such factor is complexity. Possible figures appear to be less complex than impossible1 figures, which in turn seem less complex than impossible3 figures; therefore, differences in complexity rather than impossibility could have been responsible for our finding that studying possible figures led to larger biases than studying impossible figures. To address this issue, we asked 36 additional participants to rate the complexity of possible, impossible1, and impossible3 versions of each of our 40 objects, on a 1 to 7 scale. Complexity did vary strongly by figure version, $F(2, 78) = 529.77$, $MS_e = .0965$: possible figures were rated least complex (mean, 2.77; range, 1.81 - 4.17), followed by impossible1 figures (3.73; 2.36 - 5.33), followed by impossible3 figures (5.02; 4.11 - 5.86).

If the smaller priming effects for impossible compared to possible figures were due to the difference in complexity between figure types, then the priming effects for individual figures should have also varied by figure complexity: larger priming effects, for both possible and impossible3 test figures, should be associated with less complex studied figures. To test this prediction, we first normalized complexity scores for the 108 studied objects (possible, impossible1, and impossible3 versions of each of the 36 test objects) by subtracting the median score for each figure type from each figure's complexity score. We then computed the correlations between normalized complexity and priming for possible and impossible3 test figures. The correlation for impossible3 test figures reached significance, $r = .28$, $p < .01$; the correlation for possible test figures was only .05.

This analysis indicates that for impossible3 test figures, studied figures rated as less complex primed possibility decisions slightly more than studied figures rated as more complex. However, just as impossible figures are more complex than possible figures, figures that are more complex are likely to be "more impossible." Therefore, it could be that this correlation simply reflects a tendency across individual figures that mimics the pattern of effects across study conditions: the more possible an object was, the more information about the object was available to prime possibility decisions. Furthermore, complexity only accounts for eight percent of the variance in priming for impossible3 test figures, and less than a quarter of one percent of the variance for possible test figures. It thus seems unlikely that differences in figure complexity *per se* can account for the differences we found between possible, impossible1, and impossible3 figures (but see Carrasco & Seamon, 1996, for further consideration of this issue).

More serious objections to our interpretation of the present results involve the magnitude of priming effects for impossible studied figures and the effects on sensitivity (d_L). First, since impossible1 figures possess more possible portions than do impossible3 figures (see Figure 1), the structure-extraction hypothesis predicted that studying the former versions would lead to a larger bias than studying the latter. Although results of the bias measure (C_L) were numerically in this direction, the difference in bias between possible and impossible1 objects (.19) was much larger than the difference between impossible1 and impossible3 objects (.02). This result could be considered an artifact, since a shift of three percentage points or so in the accuracy rate of impossible3 test objects primed by impossible3 studied versions would have brought all the points into line with the structure-extraction predictions. Alternatively, the insignificant difference between bias resulting from study of impossible1 and impossible3 objects could be taken as evidence for the perceptual-fluency hypothesis, which predicted equivalent biases for each of the three studied-object conditions.

Second, differences between d_L values in the different study conditions were small and insignificant, whereas the structure-extraction hypothesis predicted larger sensitivity for studied-possible compared to studied-impossible objects. This finding will be considered in the discussion of Experiment 2, since results in that experiment were similar.

In sum, the significantly larger bias for possible than impossible3 studied items supports the structure-extraction hypothesis, while the near-equivalent bias for impossible1 and impossible3 studied items and the equivalent sensitivities for all study conditions support the perceptual-fluency hypothesis. We note that the burden of proof in this experiment was on the structure-extraction hypothesis, since it predicted positive effects, whereas the perceptual bias hypothesis predicted null effects. Therefore, we feel that overall, Experiment 1 supports the structure-extraction hypothesis. Nevertheless, we sought converging evidence for this hypothesis in Experiment 2.

Experiment 2

According to the structure-extraction hypothesis, participants are biased to process possible portions of objects differently when the objects were studied from when unstudied. Since impossible portions are not structurally interpretable, a structure-extraction bias should not affect these portions of objects. Therefore, we reasoned, occluding a possible portion (which could support bias if not occluded) of a studied object should cause a reduction in bias, compared to occluding an impossible portion (which would not support bias in any case). This prediction was tested in Experiment 2. In the study phase of the experiment, participants viewed impossible1 figures, with some portion of each figure occluded by a black rectangle (see Figure 4). For one half of the objects, the occluder covered the impossible portion of the figure (occluded-impossible stimuli), while for the other half, the occluder covered an equally large possible portion of the figure (occluded-possible stimuli). The test phase of the experiment was exactly the same as that in Experiment 1: participants viewed, for 45 ms, the possible and impossible3 versions of studied and unstudied objects, and decided whether each figure was possible or impossible.

For the same reasons as in the previous experiment, we expected a larger bias to respond "possible" (i.e., a lower C_L score) for studied than unstudied objects. The pattern of results on the two occlusion conditions should distinguish between the structure-extraction and perceptual-fluency hy-

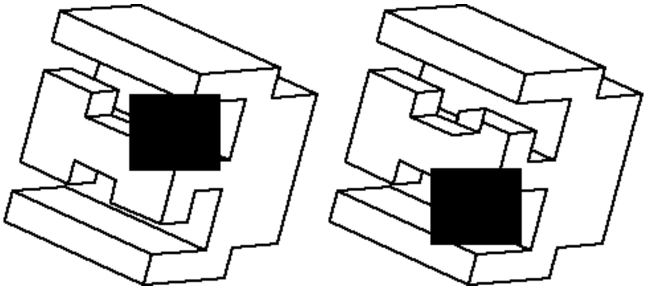

Impossible portion occluded

Possible portion occluded

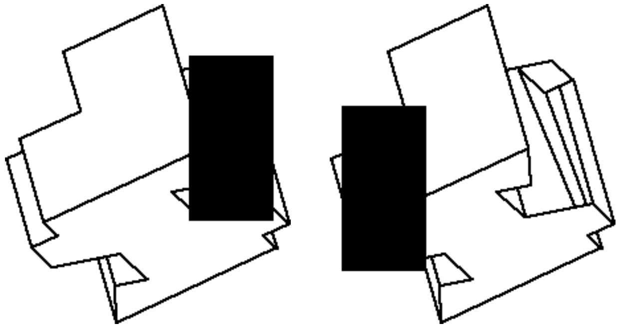

Figure 4. Examples of occluded objects used in Experiment 2.

potheses. The former predicts greater priming (in the form of a larger bias to respond “possible”) from the occluded-impossible than from the occluded-possible stimuli, as explained above. On the other hand, if priming is caused by non-diagnostic information (perceptual information that does not distinguish between possible and impossible objects) and perceptual fluency, priming should be equivalent in the two conditions, since the occluders were exactly the same size in both types of studied stimuli.

Method

Participants. Forty-two members of the Brown University community participated in the experiment in exchange for $6.00 each. Data from three participants were discarded because they performed at below-chance levels on the object possibility test.

Materials, procedure, and design. These were identical to those used in Experiment 1, except for the following. First, since there were only three Studied Object conditions (occluded-possible, occluded-impossible, and unstudied), the 36 objects used were broken up into three groups and rotated between-participants, so that each object participated equally often in each condition. Second, in the study phase of the experiment, participants saw each figure once while deciding what direction it faced (as in previous experiments), then saw each figure again in a second block of trials, where the task was to decide whether each object looked most like a building, a tool, or a spaceship. Stimuli were shown for 5 s each in both study tasks. Third and most importantly, participants viewed the impossible1 version of each studied object, and a portion of each figure was occluded by a black rectangle (Figure 4). For one half of the studied objects, the rectangle covered the portion of the figure that made it impossible, while for the other half of the studied objects, a possible portion was covered. To generate the occluded-impossible version of an object, we drew the smallest possible rectangle that completely covered the impossible portion of the object’s impossible1 version. The occluded-possible version was then created by re-positioning the same rectangle over a possible portion of the impossible1 version of the object. Across participants, the occluded-possible and occluded-impossible versions of each object were seen equally often.

Table 2
Object possibility performance: Experiment 2

	Studied Objects		
Tested Figure	Occluded Impossible	Occluded Possible	Unstudied
Possible	0.77	0.76	0.71
Impossible3	0.53	0.59	0.64
Sensitivity (d_L)	1.36	1.55	1.61
Bias (C_L)	-0.56	-0.39	-0.16

Results

Once again, participants were biased to respond “possible” more often to studied than to unstudied objects (Table 2), as confirmed by a significant main effect of Studied Object on C_L scores, $F(2, 76) = 9.70$, $MS_e = 0.165$, but were not any more sensitive for studied than for unstudied objects, $F(2, 76) = 1.04$ for d_L scores. The crucial question in the present experiment was whether participants would show a greater bias for objects studied with impossible parts occluded than for objects studied with possible parts occluded. Planned *t*- and sign tests revealed the difference between biases in these conditions to be significant, $t(38) = 1.90$; $z = 1.67$.

Discussion

Participants in Experiment 2 showed a larger bias to respond “possible” to figures studied with impossible portions occluded than to figures studied with possible portions occluded. This result runs counter to the perceptual-fluency hypothesis: since the occluder concealed the same surface area of the line drawing when covering impossible and possible portions (see Figure 4), this hypothesis predicted equivalent biases to respond “possible” in the two conditions. In contrast, the larger bias for occluded-impossible than occluded-possible studied figures is consistent with the structure-extraction hypothesis, since figures with possible portions occluded have less valid three-dimensional structure to be processed than do figures with impossible portions occluded.

As is evident in Figure 4, the visible (non-occluded) portions of occluded-impossible figures overlap precisely with possible test figures, whereas the visible portions of occluded-possible figures do not precisely overlap with impossible3 test figures. Overall, then, it could be argued that the correspondence between occluded-impossible studied figures and test figures is slightly less than the overlap for occluded-possible studied figures with test figures. This might seem to create a problem for our interpretation of the findings in Experiment 2. However, if image consistency was the sole determining factor in amount of priming, then priming effects should have been larger for possible test figures, which overlapped more with studied figures, than for impossible3 test figures; instead, we found larger priming effects for impossible3 than for possible test figures. Furthermore, impossible3 test figures were more strongly primed by occluded-impossible than by occluded-possible figures, *even though the latter figures were more similar to these test figures*. As in Experiment 1, the structure-extraction hypothesis correctly predicted this violation of the encoding specificity principle. For possible test figures, the structure-extraction hypothesis correctly predicted that occluded-impossible figures would lead to more priming than occluded-possible figures, although the magnitude of this difference was quite small.

Although the response bias (C_L) effects in Experiments 1 and 2 support the structure-extraction hypothesis, results concerning sensitivity (d_L) effects in the two experiments do not support the predictions of this hypothesis. As explained in the General Introduction and diagrammed in Figure 3c-d, objects studied in possible versions should prime possible test objects more than impossible test objects, while objects studied in impossible versions should prime possible and impossible test objects equally. As a result, sensitivity should be higher for studied-possible than for studied-impossible items. In fact, however, no significant effects of studied version were observed in either Experiment 1 or Experiment 2.

Since test stimuli were presented for only 45 ms and were masked, direct perceptual information (that is, information about the two-dimensional contours that composed the test figures) would have been quite fragmented. According to the structure-extraction hypothesis, priming in the object possibility task does *not* affect the perception of these two-dimensional lines; rather, priming operates on a later stage in object processing, in which the lines are assembled into three-dimensional structures. Therefore, it could be that when an object's possible version was studied and its impossible3 version tested, impossible portions of the line drawing were sometimes not perceived in the test flash, and three-dimensional structure was "filled in" in the structure-extraction system. To the extent that this happened, it would have decreased the difference in priming effects between possible and impossible3 test items, and consequently decreased any difference in sensitivity (d_L) between the studied-possible and studied-impossible3 conditions.

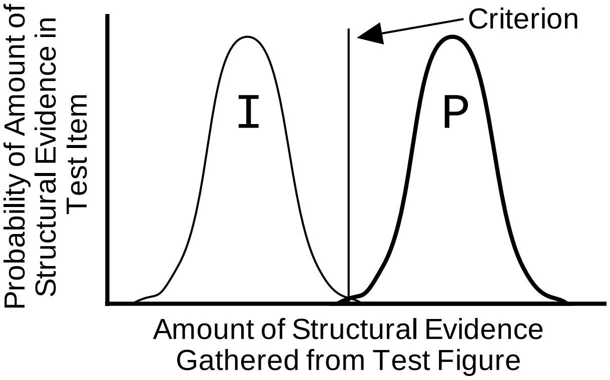

Figure 5. Hypothetical probability distributions for impossible and possible test items given long object possibility test performance.

The structure-extraction hypothesis made strong predictions: that both bias (C_L) and sensitivity (d_L) would vary for different study conditions in Experiments 1 and 2. The perceptual-fluency hypothesis, on the other hand, predicted no difference between study conditions on either dependent measure. As we stated in the conclusion of Experiment 1, we feel that confirming a positive effect (significant differences in C_L) should be weighed more heavily than confirming a null effect (no significant differences in d_L), and therefore that these experiments best support the structure-extraction hypothesis. If the post-hoc argument given above for why sensitivity differences were not found proves unsatisfactory, a third hypothesis may be needed that predicts both a difference in bias but no difference in sensitivity.

Experiment 3a

Regardless of one's stance on the two hypotheses being tested in Experiments 1 and 2, the results of these experiments, as well as the results of Ratcliff and McKoon's (1995) experiments, show that when given brief test exposures, participants demonstrate a tendency to respond "possible" to studied objects more often than to unstudied objects on the object possibility test. What would happen if, instead of flashing test objects on the screen, participants were given much longer exposures to test objects? This was the experimental question posed in Experiment 3a, which used exactly the same study and test procedures as Experiment 1, except that participants were allowed up to five seconds to make test responses. Unlike in previous experiments, there should be very little uncertainty about the possibility or impossibility of any test item given this much time to make decisions. Therefore, in terms of the SDT model introduced previously, we can predict that the probability distributions of structural evidence for possible and impossible3 test figures should overlap much less (Figure 5), and that participants should easily be able to place their criterion so that they will almost always be able to correctly distinguish between possible and impossible objects. Furthermore, we predicted that there should be no distribution-shifting for studied and unstudied objects (as there was in Experiments 1 and 2), because given virtually unlimited viewing time, the amount of structural evidence ultimately available should be solely dependent on the currently-shown test object.

To summarize, probability distributions for possible and impossible objects should be clearly separable and should not vary for studied and unstudied items. In SDT terms, then, sensitivity (d_{L}) should be much higher than in Experiments 1 and 2, and neither sensitivity nor bias (C_{L}) should vary depending on whether an object was studied or not. However, if object possibility priming is caused by a perceptual bias to process studied items more efficiently than unstudied items, we predicted that participants would be *faster* to respond to studied than to unstudied objects, at least for possible test items. On the other hand, encoded structural information or perceptual fluency could hamper participants' ability to reject impossible test objects, leading to slower responses for studied than unstudied items on these items. In any case, patterns of effects for response times in these experiments should help us constrain models of object possibility performance, since such models should posit that similar mechanisms are responsible for test performance given either short or long stimulus exposures.

Method

Participants. Thirty-six Yale undergraduates from the introductory psychology pool participated in the experiment.

Materials, procedure, design, and analyses. These were identical to those used in Experiment 1, except that a) object possibility test figures remained on the screen until participants responded, or until 5 s had passed, and b) response time, as well as accuracy, served as a dependent measure. Geometric means of participants' response times were used to minimize the effects of outliers.

Results

As in Experiment 1, each test object was shown twice: once in its possible version and once in its impossible3 version. Therefore, preliminary analyses including the factors of Test Order, Trial Number (see the Results section of Experiment 1 for an explanation of these terms), and Test Version were performed on accuracy scores and response times. There were no significant main effects or interactions involving Test Order or Trial Number on accuracy, all $F < 1.25$. For response time, the effect of Trial Number was significant, $F(1,35) = 6.91$, $MS_{\mathrm{e}} = 238850$, but the effect of Test Order was not, $F(1,75) = 2.47$, $MS_{\mathrm{e}} = 91016$, and no interactions approached significance (all $F < 1.49$). These analyses indicate that participants simply got faster as the test went on (responding more quickly on the second than on the first half of the test), so other analyses were collapsed over these two factors.[4]

Table 3
Object possibility performance: Experiment 3a

	Studied Figure			
Tested Figure	Possible	Impossible1	Impossible3	Unstudied
Possible				
Accuracy	0.98	0.96	0.97	0.97
RT	1393	1481	1559	1581
Impossible3				
Accuracy	0.94	0.92	0.93	0.94
RT	1678	1773	1700	1717
Sensitivity (d_{L})	5.10	4.75	4.97	5.04
Bias (C_{L})	-0.15	-0.17	-0.13	-0.14

Note: Trials for which subjects failed to respond in the allotted 5 s (0.7% of possible trials and 2.0% of impossible trials) were counted as neither correct nor incorrect.

Showing test figures for an essentially unlimited amount of time had the desired effect of boosting accuracy scores, and measured sensitivity, to near-ceiling levels in all conditions (Table 3). As is evident from the consistently large d_{L} and consistently negative C_{L} scores in Table 3, participants were very accurate, but slightly biased to respond "possible" overall. Importantly, however, these trends were equivalent for both studied and unstudied objects: the main effect of Studied Figure (possible, impossible1, impossible3, or unstudied) was significant neither on sensitivity (d_{L}), $F(3, 105) = 1.58$ nor on bias (C_{L}), $F(3, 105) < 1$.

Although memory for studied items did not affect participants' response tendencies, we did find robust priming of object possibility response times. Planned *t*- and sign tests indicate that possible test items were significantly primed by both possible and impossible1 studied figures, $t(35) =$ 3.87 and 2.03 and $z = 4.00$ and 1.86, respectively, but no significant priming effects were observed for impossible3 test figures, all $t(35) < .838$; all $z < 0.67$). For the omnibus ANOVA, the main effects of Studied Figure, $F(3, 105) = 3.00$, $MS_{\mathrm{e}} = 61506$, and Test Figure, $F(1, 35) = 5.82$, $MS_{\mathrm{e}} = 563755$, were significant, while the interaction term was marginally significant, $F(3, 105) = 2.30$, $MS_{\mathrm{e}} = 58672$, $p = .082$.

---

[4] Because of the long display durations used in this experiment, it might seem surprising that the first test exposure did not prime the second. However, when the first test figure was possible, the second test figure was impossible, and we failed to find any priming at all for impossible test figures. When test figures were seen in the reverse order (impossible followed by possible), we would also not expect a significant amount of priming, since impossible3 studied figures did not appreciably facilitate object decisions for possible test figures.

Discussion

Experiment 3a demonstrates that previous exposure to objects can have a significant effect on object possibility decisions even when response biases in accuracy, as measured by C_L, are equivalent for studied and unstudied stimuli. Specifically, we found implicit memory effects on participants' response times: responses to possible test objects were faster when objects were studied than when they had not been studied. These priming effects were largest for objects studied in their possible versions and decreased monotonically for impossible1 and impossible3 studied figures.

Interestingly, significant priming was not observed for impossible3 test figures. Furthermore, participants took considerably longer to make "impossible" than to make "possible" decisions. These findings could have important implications for detailed models of object possibility performance. However, a less interesting explanation is also plausible: it could be that participants were able to rapidly analyze impossible information, but were not willing to trust these first impressions, and subsequently rescanned test figures in order to confirm the structural violations.

In an attempt to rule out this explanation, Experiment 3b employed the materials and design of Experiment 3a with a go/no-go procedure. Half of the participants responded only to possible objects and not to impossible objects; the other half responded to impossible but not to possible objects. In a go/no-go procedure, participants know they are only responsible for detecting a target, since they do not have to make an alternate response. Therefore, participants in the "Go-Impossible" condition should respond as soon as they can gather information sufficient to make an "impossible" response. In addition, we stressed in the instructions that participants would be best off responding on the basis of their "first impression," and told them not to go back to check whether this first impression was correct. We also encouraged quick responses by allowing participants only three seconds in which to respond, rather than the five seconds allowed in Experiment 3a. This procedure should thus be more diagnostic for the relative speeds of processes engaged in making "possible" and "impossible" responses.

Experiment 3b

Method

Participants. Seventy-two college-age participants from Case Western Reserve University and the Brown University community were recruited for the experiment, and participated either for class credit or for $6.00.

Materials, procedure, and design. These were identical to those used in Experiment 3a, except in the way participants were asked to respond during the object possibility test. Thirty-six participants pressed the space bar if a test object was possible, or allowed the trial to time out if the test object was impossible (trials timed out after three seconds). A second group of 36 participants received the opposite instructions, pressing the space bar for impossible objects and doing nothing for possible objects. In addition, all participants were prodded to respond as quickly as possible with the following instructions: "Previous research has shown that many people who do this task take longer to respond to test figures than they have to—even after they have determined that the test figure is possible/impossible [this word varied for the two groups of participants], many people go back and re-examine the figure to make absolutely sure that it could/could not be a real 3D object. This previous research has also indicated that this 'extra analysis' is often detrimental to performance. In this experiment, we are trying to test this idea by asking you to respond to each trial based on your 'first impression.' In other words, please respond as soon as you can if the object is possible/impossible, and do not go back to check whether this first impression is correct. Don't worry if it seems like you might be getting some trials wrong—just keep on responding as soon as you can make a decision about the test object."

Results

Response times and accuracy rates for Experiment 3b are reported in Table 4. The top two rows of figures show performance for participants asked to respond to possible objects, which we will refer to as the Go-Possible condition; the bottom two rows show performance for participants in the Go-Impossible condition. Accuracy and response time could be accurately measured only for "go" responses, since a timed-out trial could reflect either a negative response or indecision on the part of the participant. Furthermore, the accuracy rates reported are underestimates, since time-outs due to indecision are normally counted as neither correct nor incorrect in calculating accuracy. Results mirrored those from Experiment 3a. Participants in the Go-Possible condition responded significantly faster than participants in the Go-Impossible condition, pooled $t(70) = 4.49$. Accuracy rates did not vary across Studied Figure conditions for either Go-Possible or Go-Impossible participants, both $F < 1$, while response times were primed in the Go-Possible but not in the Go-Impossible condition. Also paralleling Experiment 3a, response time priming was largest for objects studied as possible figures, smaller for objects studied as impossible1 figures, and smallest for objects studied as impossible3 figures. The priming effect of possible studied figures on possible test figures was significant by a planned t-test, $t(35) = 2.66$.

Discussion

The results of Experiment 3b replicate those of Experiment 3a: priming can be found on object possibility response times even in the absence of differential accuracy rates for studied and unstudied objects. Furthermore, Experiment 3b indicates that participants were not slow to respond "impossible" in Experiment 3a simply because they were wary of making such responses and went back to rescan the test images.

Table 4

Object possibility performance: Experiment 3b

	Studied Figure			
"Go" Response	Possible	Impossible1	Impossible3	Unstudied
Possible				
Accuracy	0.96	0.96	0.97	0.96
RT	958	991	1039	1038
Impossible				
Accuracy	0.90	0.92	0.90	0.91
RT	1312	1282	1297	1299

Participants in the present experiment had every incentive to respond as quickly as they could, yet they still made “possible” responses more quickly than “impossible” responses. We also failed again to find any response time priming for impossible test figures. These two results may have important implications for detailed models of object possibility task performance, as discussed below.

General Discussion

Schacter and Cooper’s object possibility paradigm (Schacter et al., 1990) generated a considerable amount of excitement in the cognitive psychology community, because it extended the implicit memory movement beyond the domain of verbal stimuli. The most widely-cited finding from Schacter, Cooper, and their colleagues’ studies was that object possibility decisions about possible objects could be primed, but that decisions about impossible objects could not. This result has been replicated many times (e.g., Schacter, Cooper, Delaney, et al., 1991) and is consistent with Schacter and Cooper’s notion that priming on the possibility task is dependent on information stored in a memory system that codes three-dimensional structural descriptions of objects. Ratcliff and McKoon (1995) redefined the priming demonstrated in object possibility studies as a bias to respond “possible” to studied items, rather than an ability to more accurately make decisions about possible objects. To some investigators, this redefinition may depreciate the theoretical value of possibility priming, since in psychophysical studies and recognition experiments, response biases are often considered artifactual products of participants’ idiosyncratic motivations and strategies.

One objective of the present study has been to counter this notion, by showing that response biases are consequences, rather than causes, of implicit memory in the object possibility paradigm. Thus Experiments 3a and 3b demonstrated that robust priming effects can be evidenced in participants’ response times to make possibility decisions when accuracy rates for studied and unstudied objects are equated. A second objective was to explore the underlying basis of the implicit memory effects found in Experiments 3a and 3b and previous studies (e.g. Schacter et al., 1990; Ratcliff & McKoon, 1995). Ratcliff and McKoon (1995) found statistically indistinguishable priming effects following study of possible and impossible objects, supporting the perceptual-fluency hypothesis outlined in our Introduction. In contrast, our Experiments 1 and 2 (as well as Experiments 3a and 3b), which used more discriminable possible and impossible objects than were tested in Ratcliff and McKoon’s (1995) experiments, revealed significant differences between studied-object conditions, supporting the structure-extraction hypothesis outlined earlier. Table 5 shows the magnitude of priming effects resulting from study conditions in the four experiments reported here. In the remainder of the General Discussion, we review the implications of our findings, then conclude with more speculative remarks on the cognitive foundations of object possibility priming.

Table 5

Summary of test exposure times, dependent measures, and priming effects from Experiments 1-3b.

Experiment	**Test Exposure**	**Dependant Measure**	**Priming for Study Conditions**		
			Possible	Impossible1	Impossible3
1	45 ms	C_L	.45*	.26*	.24*
			Occ. Impos.	Occ. Pos.	
2	45 ms	C_L	.40*	.23*	
			Possible	Impossible1	Impossible3
3a	≤ 5000 ms	RT	188*	100*	22
3b	≤ 3000 ms	RT	80*	47	-1

Note. Asterisks (*) indicate priming scores that were significantly different from 0 by *t*-tests. Priming in Experiments 3a and 3b is reported for possible test objects only; no significant priming effects were observed for impossible test objects in these Experiments. Occ. Impos. = Occluded Impossible; Occ. Pos. = Occluded Possible.

Priming and bias

Participants in every one of our experiments evidenced an overall bias to respond “possible” on the object possibility task. A survey of previous object possibility studies (e.g. Schacter, Cooper, Delaney et al., 1991; Ratcliff & McKoon, 1995) reveals that this overall bias has also been present in nearly every other object possibility experiment reported. This should come as no surprise: since, by definition, we never see impossible figures in the real world, our visual systems should be strongly biased to perceive valid three-dimensional structure when we view a line drawing of a novel object. Furthermore, our visual systems may be biased to discount portions of objects that seem “impossible” as being due to accidents of viewpoint, occlusion by other objects, or other normally-occurring optical illusions.

Although this overall bias to respond “possible” comes as no surprise, a stronger bias to respond “possible” to some object X when compared to another object Y strikes us as considerably more interesting. Such a finding indicates that the visual system is more likely to believe it is perceiving valid three-dimensional structure in X than it is in Y. When X represents studied objects and Y unstudied objects, as in the experiments by Ratcliff and McKoon (1995) and our Experiments 1 and 2, the inevitable conclusion must be that some kind of encoded information about studied figures is leading them to be perceived differently from unstudied figures. This is the essence of an implicit memory effect: remembered information has caused a variation in performance between studied and unstudied objects, even though the test task does not require participants to explicitly recollect previous events.

In Experiments 3a and 3b, participants were given sufficiently long test exposures to respond correctly a vast majority of the time, so accuracy rates (and response biases, as measured by C_L) did not vary significantly between studied and unstudied objects. Nevertheless, both experiments revealed reliable implicit memory effects: response times were facilitated for studied as compared to unstudied objects (Table 5). These findings reinforce the idea that differences between studied and unstudied items on the object possibility task are not *caused* by a response bias to respond "possible." Instead, as suggested here and in Ratcliff et al. (1989), implicit memory effects may be caused by a perceptual bias in an information processing system involved in performance of the implicit memory task. Implicit memory is revealed in response tendencies when brief test exposures are given, or in response times given long test exposures.[5]

The basis of perceptual bias

We suggested that two different types of perceptual processes might be biased in a way that would produce the pattern of response biases found by Ratcliff and McKoon (1995). In addition to replicating Ratcliff and McKoon's results with a new set of objects, Experiments 1 and 2 were designed to tease apart these alternatives. In both experiments, participants studied various types of figures that contained more or less valid three-dimensional structural information. If object possibility decisions are primed directly by processes involved in the extraction of three-dimensional structure, then studying figures containing more extractable structure should lead to more priming. On the other hand, if possibility decisions are primed by perceptual fluency resulting from a non-diagnostic perceptual bias (i.e., bias in a process that does not extract information distinguishing between possible and impossible objects), then priming should be approximately equivalent in the different conditions. Both experiments utilized brief test exposures (45 ms), so we expected priming to occur in the form of a tendency to respond "possible" to studied more often than to unstudied objects, as in the experiments reported by Ratcliff and McKoon (1995).

In Experiment 1, participants studied possible, impossible1, and impossible3 objects (Figure 1), each of which possess progressively larger areas in which the three-dimensional structure of the object is compromised. This experiment revealed larger priming effects (a larger bias to respond "possible") when participants studied possible objects than when they studied either type of impossible object, supporting the structure-extraction hypothesis. In Experiment 2, participants studied impossible1 objects that were partially occluded by a black rectangle (Figure 4). Results showed that covering a possible portion of a studied figure led to a smaller bias effect than did covering an impossible portion of the figure, even though roughly the same amount of contour was obscured in both cases. This finding again supports the structure-extraction hypothesis: occluding a possible portion of an object, which could have provided valid structural information to be primed, led to a smaller priming effect than occluding an impossible portion of an object, which would not have provided any valid structural information.

Participants in Experiment 2 always studied impossible1 objects, but in the occluded-impossible condition, the entire impossible portion of the object was covered. Interestingly, the priming effect of this study condition was considerably larger than the effect of impossible1 studied objects in Experiment 1, and almost as large as the effect of possible studied objects in Experiment 1 (Table 5). This may indicate that participants were able to fill in three-dimensional structure behind an occluder (and in place of missing or fragmented perceptual information following flashed test exposures, as hypothesized above), but that impossible portions actually disrupted structural processing during the long stimulus exposures employed during the study task. In other words, the structure-extraction mechanism may be able to compensate for missing perceptual information, but be unable to do anything at all with conflicting information coming from perceptually intact impossible portions. However, further studies (preferably employing a within-participants occlusion manipulation) would be needed to confirm this conclusion, since there were potentially confounding methodological changes between Experiments 1 and 2 (for example, the ratio of studied to unstudied test items was 3:1 and 2:1 in the two experiments, respectively).

Returning now to Experiments 3a and 3b, we note that the simple model of object possibility performance assumed in Figures 3 and 5 can be adapted to account for the response time priming observed in these experiments. The underlying assumption of this model is that decisions are made based on the amount of structural evidence gathered from possible and impossible test objects. Figure 6 adds a time-course element to the model, positing that evidence is accumulated at different rates for different objects. It seems reasonable to assume that either a structure-extraction or perceptual-fluency bias would lead evidence to be gathered more quickly from

---

[5] This phenomenon is akin to that of perceiving Ronald James' famous dalmation picture (viewable in most introductory textbooks): the first time one sees it, the picture looks like a random collection of dots, but once one perceives the dalmation, it is impossible not to see the dog every time the picture is viewed. Similarly, in the present experiments, participants seem to have automatically perceived studied objects differently from unstudied objects, regardless of whether or not extra information was needed to perform the task.

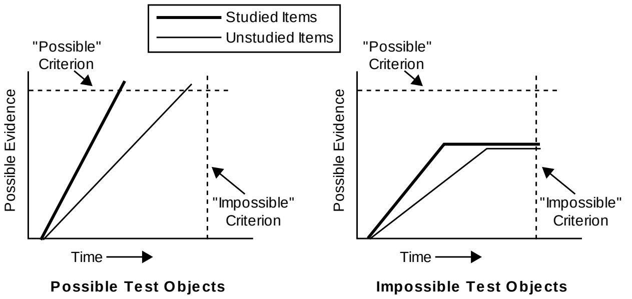

Figure 6. A time-course model of object possibility processing.

studied than from unstudied items. Given ample time to analyze test objects, participants in Experiments 3a and 3b may have established a “possible” threshold specifying an amount of evidence, but an “impossible” threshold specifying a period of time (Figure 6). If the evidence threshold is reached before the time threshold is crossed, a “possible” response is made; if the time threshold is reached before reaching the evidence threshold, an “impossible” response is made.

As shown in Figure 6, the evidence threshold will be crossed more quickly for studied than unstudied possible test objects (resulting in response time priming), but since the time threshold remains the same for studied and unstudied items, “impossible” responses are made no more quickly for studied than for unstudied items. Furthermore, the model correctly predicts that “impossible” responses will always be slower than “possible” responses. When responses must be made based on limited perceptual evidence (i.e., in Experiments 1 and 2 and past object possibility studies), the time criterion is ignored, and participants make possible/impossible judgments based solely on the amount of structural evidence that was gatherable from the briefly-presented test figure. Figure 6 makes it clear that there will be more structural evidence available for studied than unstudied objects at any point in the time course and regardless of whether the object is possible or impossible.

This model is attractive because it can rely on a single perceptual mechanism, extraction of possible three-dimensional structure, that other evidence suggests is already at work in normal everyday perception (Enns & Rensink, 1990; this model could also be made to work with a non-diagnostic perceptual bias and perceptual fluency, but such a solution would be decidedly less elegant than one utilizing a structure-extraction mechanism). We would like to re-emphasize that this model is not an *alternative* to Ratcliff and McKoon’s (1995; Ratcliff & McKoon, 1997a) proposal that object possibility priming is due to a perceptual bias for studied items; rather, it extends Ratcliff and McKoon’s analysis by offering specific mechanisms through which perceptual bias could operate.

Finally, while many of our experimental results directly support the structure-extraction bias hypothesis, at least two results, the lack of a significant difference between priming following impossible1 and impossible3 items in Experiment 1 and the fact that sensitivity did not vary by study condition in either Experiment 1 or Experiment 2, are inconsistent with this hypothesis. Since multiple perceptual processes are involved in performing any complex visual task, it could very well be that more than one process is biased by previous study in the object possibility paradigm. Thus we cannot rule out the notion that perceptual fluency resulting from some non-diagnostic perceptual bias contributes to priming effects in the object possibility task. Or, it could be that yet another hypothesis may eventually be proposed that accounts for all the findings presented here and in past research reports. For the present, though, we conclude that a structure-extraction bias is at least one major cause of priming effects on the object possibility task.

Architectural issues

We have asserted that object possibility priming in Experiments 1–3b is largely based on a structure-extraction bias—that is, facilitated processing of the three-dimensional structure of the possible portions of studied objects. How might such facilitation be implemented? One possibility is through Schacter and Cooper’s proposed structural description system (SDS). Complete representations of the structure of possible objects, along with representations of impossible objects that are either incomplete (depicting only possible portions) or filled in with inferred valid structural information, could be represented in the SDS and could provide the basis for the observed priming effects (the results of Experiment 2 suggest that incomplete representations of impossible objects would be more likely).

Ratcliff and McKoon (1995, p. 765) offered an alternative account of how structural information might be encoded in order to produce priming: connections in a network model of object recognition might be altered in such a way that repeated presentations of a stimulus would “increase the probability of constructing a representation of a possible object” under degraded viewing conditions (i.e., short test exposures). If such alterations also led to speeded construction of structural representations under more optimal viewing conditions, this account would also be able to explain response time priming in our Experiments 3a and 3b.

As we see it, the fundamental difference between Schacter and Cooper’s (1995) and Ratcliff and McKoon’s (1995) viewpoints is similar to the propositionalist-proceduralist distinction characterized by Kolers and Roediger (1984). Processing an image of an object almost certainly involves extracting information about the object’s three-dimensional structure (Marr, 1982), and the present results indicate that possibility priming occurs because the structure-extraction process performed when test stimuli have been studied is different from that performed when they have not been studied. The question is, do participants remember the *process* of extracting the structure of an object from a two-dimensional image, or do they remember a *coded description* of the resulting structure?

Ratcliff and McKoon's (1995) proposal conforms to the proceduralist doctrine that memory effects are caused by stimulus-specific modifications to information processing systems (Crowder, 1993; Kolers & Roediger, 1984). In contrast, by referring to the SDS as a "memory system," Schacter and Cooper imply that the SDS is primarily responsible for representing the output of the structure-extraction system. This standpoint is propositionalist in the sense that it assumes that perceptual input is recoded (in much the same way that sentences are often assumed to be recoded into propositions) and placed in a cognitive system expressly dedicated to memory storage.

Both approaches have strengths and weaknesses. A proceduralist explanation appears to be more parsimonious. If we assume that humans possess a cognitive system for extracting three-dimensional structure from two-dimensional images, then by incorporating representations of object structure into the extraction system itself, we avoid postulating an additional storage system. On the other hand, a propositionalist theory allows us to more naturally describe the relationship between distinct sets of represented information. For instance, if encoded structural information affects possibility decisions more than recognition judgments, whereas encoded information about image size has the opposite effect (Cooper et al., 1992; Williams, 1995), it makes sense to think of these two sources of information as residing in different memory systems.

Dedicated storage systems also appeal to our intuitions about how explicit memory tasks are performed: when asked to recognize whether we have seen an object before or not, we often have the sense of processing it and then attempting to match the output of this processing to relevant stored representations. In contrast, the proceduralist approach inherently links representation with processing, and is thus consistent with the intuition that on implicit memory tasks, studied items simply seem easier to process than do unstudied items. For example, an account of how encoded information affects the object possibility task might include a neural network that performs structural extraction (e.g., Hummel & Stankiewicz, 1996) on test images and in which the weights between units may be modified by previous encounters with studied stimuli. Such a network could become increasingly efficient at deriving structural descriptions for studied objects as compared to unstudied objects (this is a similar account to that offered by Ratcliff and McKoon, 1995).

Based on these considerations, we propose that Schacter and Cooper (1995) are correct in asserting that object possibility priming is based on an SDS, but that this cognitive system is best characterized as a "structure-*describing* system," rather than a structural *description* system. Its primary function is to extract the three-dimensional structure of objects from two-dimensional images, and memory is encoded as a by-product of the system's information-processing duties. Schacter and Cooper must also be correct in asserting that old/new recognition performance is often mediated by information retained elsewhere, such as in an episodic memory system (Tulving, 1983). For example, the structure-extraction system will not be the storage site for associations between novel and familiar objects, yet this type of information will be used in making old/new recognition judgments.

In our view, conceiving of the SDS as a system primarily dedicated to processing, rather than storing, structural information makes Schacter and Cooper's (1995) theory much more compatible with Ratcliff and McKoon's (1995) model of possibility priming. We have proposed that the structure-extraction system will only be able to process information about possible portions of studied figures, for the simple reason that only these portions have extractable three-dimensional structures. Therefore, participants demonstrate a bias to respond "possible" to studied items when test exposures are very short (Ratcliff & McKoon's, 1995, Experiments 2-7, and our Experiments 1 and 2) because processing of possible but not impossible information about studied figures is facilitated. Furthermore, when differential accuracy rates for studied and unstudied items are eliminated by providing longer test exposures, participants demonstrate shorter response times to studied than to unstudied items (our Experiments 3a and 3b), again due to more efficient processing of studied structural information.

Schacter and Cooper's primary goal has been to describe the relationship between memory demonstrated on recognition and object possibility tests (Cooper & Schacter, 1992). Their multiple-memory systems theory is helpful in understanding this relationship, since different kinds of information are often more or less useful on each test. Ratcliff and McKoon's focus, on the other hand, has been on how object possibility priming should be modeled (McKoon & Ratcliff, 1995). Their bias-plus-explicit-memory model is helpful in providing a precise description of the observed behavior. In the present paper, we have elaborated the bias aspect of Ratcliff and McKoon's proposal. By proposing and testing the hypotheses stemming from a more specific "bias model" of object possibility performance (Figures 3, 5, and 6), we found evidence against the notion that priming on the object possibility task is merely a result of perceptual fluency for previously-encountered objects. Instead, we have argued, item-specific changes in an information-processing system for extracting three-dimensional structure from two-dimensional images is at least one significant source of object possibility priming effects.

References

Barrow, H. G., & Tenenbaum, J. M. (1981). Interpreting line drawings as three-dimensional surfaces. *Artificial Intelligence*, *17*, 75-116.

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Re-*

view, 94, 115-117.

Bülthoff, H. H., Edelman, S. Y., & Tarr, M. J. (1995). How are three-dimensional objects represented in the brain? *Cerebral Cortex*, *5,* 247-260.

Carrasco, M., & Seamon, J. G. (1996). Priming impossible figures in the object decision test: The critical importance of perceived stimulus complexity. *Psychonomic Bulletin & Review*, *3*, 344-351.

Cooper, L. A. & Schacter, D. L. (1992). Dissociations between structural and episodic representations of visual objects. *Current Directions in Psychological Science*, *1*, 141-146.

Cooper, L. A., Schacter, D. L., Ballesteros, S., & Moore, C. (1992). Priming and recognition of transformed three-dimensional objects: Effects of size and reflection. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 43-57.

Crowder, R. G. (1993). Systems and principles in memory theory: Another critique of pure memory. In A. F. Collins, S. E. Gathercole, M. A. Conway, & P. E. Morris (Eds.), *Theories of memory* (pp. 130-161). East Sussex: Erlbaum.

Enns, J. T., & Rensink, R. A. (1990). Sensitivity to three-dimensional orientation in visual search. *Psychological Science, 1,* 323-326.

Goldstein, E. B. (1996). *Sensation and Perception*. Pacific Grove, CA: Brooks/Cole.

Hayward, W. G. (1997). Effects of outline shape in object recognition. In press, *Journal of Experimental Psychology: Human Perception & Performance*.

Hummel, J. E., & Stankiewicz, B. J. (1996). An architecture for rapid, hierarchical structural description. In T. Inui & J. McClelland (Eds.), *Attention and Performance XVI*.

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

Jacoby, L. L., Woloshyn, V., & Kelley, C. (1989). Becoming famous without being recognized: Unconscious influences of memory produced by dividing attention. *Journal of experimental psychology: General*, *118*, 115-125.

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

Kroll, J. F. and M. C. Potter (1984). Recognizing words, pictures, and concepts: A comparison of lexical, object, and reality judgments. *Journal of Verbal Learning & Verbal Behavior*, *23*, 39-66.

Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco: Freeman.

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

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

Norris, D. (1995). Signal detection theory and modularity: On being sensitive to the power of bias models of semantic priming. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 935-939.

Ratcliff, R., Albritton, D.W., & McKoon, G. (1997) Bias in auditory priming. In Press, *Journal of Experimental Psychology: Learning, Memory, and Cognition*.

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. (1997a). Bias effects in implicit memory tasks. In press, *Journal of Experimental Psychology: General*.

Ratcliff, R. & McKoon, G. (1997b). A counter model for implicit priming in perceptual word identification. In press, *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.

Schacter, D. L. & Cooper, L. A. (1993). Implicit and explicit memory for novel visual objects: Structure and function. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*, 995-1009.

Schacter, D. L. & Cooper, L. A. (1995). Bias in the priming of object decisions: Logic, assumption, 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., Tharan, M., & Rubens,

A. B. (1991). Preserved priming of novel objects in patients with memory disorders. *Journal of Cognitive Neuroscience*, *3*, 118-131.

Schacter, D. L., Cooper, L. A., & Valdiserri, M. (1992). Implicit and explicit memory for novel objects in older and younger adults. *Psychology and Aging*, *7*, 299-308.

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

Snodgrass, J. G., & Corwin, J. (1988). Pragmatics of measuring recognition memory: Applications to dementia and amnesia. *Journal of Experimental Psychology: General*, *117*, 34-50.

Tulving, E. (1983). *Elements of Episodic Memory*. Oxford, England: Clarendon.

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

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

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

Williams, P. (1995). Effects of picture-plane orientation, size, and color transformations on object decision priming. Paper presented to the 3rd Annual Pre-Psychonomics Workshop on Object Perception and Memory, Los Angeles, CA, November 9, 1995.