

# Visual antipriming: Evidence for ongoing adjustments of superimposed visual object representations

CHAD J. MARSOLEK

*University of Minnesota, Minneapolis, Minnesota*

DAVID M. SCHNYER

*Boston VA Healthcare System, Boston, Massachusetts  
and Boston University School of Medicine, Boston, Massachusetts*

REBECCA G. DEASON

*University of Minnesota, Minneapolis, Minnesota*

MAUREEN RITCHEY

*University of Notre Dame, Notre Dame, Indiana*

and

MIEKE VERFAELLIE

*Boston VA Healthcare System, Boston, Massachusetts  
and Boston University School of Medicine, Boston, Massachusetts*

A fundamental question of memory is whether the representations of different items are stored in localist/discrete or superimposed/overlapping manners. Neural evidence suggests that neocortical areas underlying visual object identification utilize superimposed representations that undergo continual adjustments, but there has been little corroborating behavioral evidence. We hypothesize that the representation of an object is strengthened, after it is identified, via small representational changes; this strengthening is responsible for repetition priming for that object, but it should also be responsible for antipriming of other objects that have representations superimposed with that of the primed object. Functional evidence for antipriming is reported in young adults, amnesic patients, and matched control participants, and neurocomputational models. The findings from patients dismiss explicit memory explanations, and the models fit the behavioral performance exceptionally well. Putative purposes of priming and comparisons with other theories are discussed. Priming and antipriming may reflect ongoing adjustments of superimposed representations in neocortex.

A fundamental question of knowledge representation is whether the representations of different items are stored in separate or overlapping manners. Many cognitive models and theories posit localist/discrete representations (see, e.g., Grainger & Jacobs, 1996; Hintzman, 1986; Jacoby, 1983; Logan, 1990; Medin & Schaffer, 1978; Poggio & Edelman, 1990; Shiffrin & Steyvers, 1997), with different items stored in separate traces. Many other cognitive mod-

els and theories posit superimposed representations (see, e.g., Anderson & Hinton, 1981; Eich, 1982; Knapp & Anderson, 1984; Masson, 1995; McClelland & Rumelhart, 1985; Murdock, 1982; Seidenberg & McClelland, 1989), with different items stored in distributed and overlapping traces. In this article, we report a new effect, "visual antipriming," which provides functional evidence for superimposed representations of visual objects and a continual-adjustment account of repetition priming effects.


Neural implementational evidence suggests that familiar visual shapes are represented in a distributed and partially overlapping manner in ventral temporal cortex. In neurophysiological studies, different whole shapes activate inferior temporal visual neurons to differing degrees (Gross, Rocha-Miranda, & Bender, 1972). Different neurons are maximally sensitive to different moderately complex features of whole objects (Tanaka, 1993); these features are simple enough that they are present in multiple whole objects but also complex enough that they are not present in all whole objects. Similarly, neuroimaging

---

This work was supported in part by the Center for Cognitive Sciences at the University of Minnesota, the National Institutes of Health (HD07151, K23MH64004, and P50NS26985), and the Medical Research Service of the Department of Veterans Affairs. We thank Chris Azorson, Jeff Bowers, Wilma Koutstaal, Randy O'Reilly, and Carmen Westerberg for valuable discussion and comments, Wilma Koutstaal for the use of her object stimuli, and Scott Domansky, Carissa Dotseth, David Frasch, Debbie Hackman, Patrick Hoben, Lindsay D. Nicholls, Shanna Trudeau, and Reed Waterworth for valuable assistance with data collection and analysis. Correspondence should be addressed to C. J. Marsolek, Department of Psychology, University of Minnesota, 75 East River Road, Minneapolis, MN 55455 (e-mail: chad.j.marsolek-1@umn.edu).

studies indicate that patterns of activation from different classes of shapes (e.g., houses, bottles, shoes, faces, chairs) overlap in human ventral temporal cortex (Haxby et al., 2001; Ishai, Ungerleider, Martin, Schouten, & Haxby, 1999). In addition, neurocomputational modeling indicates functional utilities for learning and storing partially superimposed representations (Hinton, McClelland, & Rumelhart, 1986). The overlap between representations not only increases the efficiency of storage, but can also be used to store systematic relationships between inputs and appropriate outputs that enable the system to efficiently generalize to novel inputs (e.g., identifying a previously unseen piano as a piano; McClelland & Rumelhart, 1985). Interestingly, though, this utility is attained at a price. In order for new categories to be learned in the system (e.g., a representation for a new gadget), trials of learning the new information must be interleaved with trials of relearning the old familiar representations (e.g., piano; McClelland, McNaughton, & O'Reilly, 1995). Otherwise, many trials of learning the new information all at once (without the interleaved relearning of old information) can cause old representations to be irretrievably lost (McClelland et al., 1995; McCloskey & Cohen, 1989). We hypothesize that this kind of interference may not be limited to situations of learning new object representations, but may also occur continually from the relearning of old familiar representations that occurs when we identify objects.

The results reported in this article are important for providing behavioral evidence that corroborates the neural and computational evidence for superimposed representations. Our theory (Marsolek, 2003; see also McClelland et al., 1995) is that partially superimposed representations in neocortex represent familiar visual object shapes and that these traces undergo continual adjustments. The representation of an object is strengthened after every use (i.e., after every experience of identifying that object), which entails small representational changes analogous to small weight changes that occur in a training trial in neurocomputational models (Bowers, Damian, & Havelka, 2002; Marsolek & Burgund, 1997; McClelland & Rumelhart, 1985; Rueckl, 2003; Stark & McClelland, 2000). These weight changes may correspond to synaptic modification effects like long-term potentiation and long-term depression evidenced in human inferior and middle temporal cells (Chen et al., 1996). The results of such changes are enhanced neuronal representation of the encoded object (Kobatake, Wang, & Tanaka, 1998; Sigala & Logothetis, 2002). Our theory is that weight changes strengthen the preexisting representation of the identified object (e.g., piano), and this is responsible for repetition priming effects (i.e., object identification is facilitated by the recent experience of viewing that object; Bartram, 1974). However, the same weight changes simultaneously interfere with storage of any other representations with which the primed object representation is superimposed (e.g., desk; see Figure 1). Strengthening the representation for one object (piano) entails weight changes that increase the likelihood that one of its features (e.g., right-front leg) will subsequently reactivate that object's representation but also decrease the likeli-

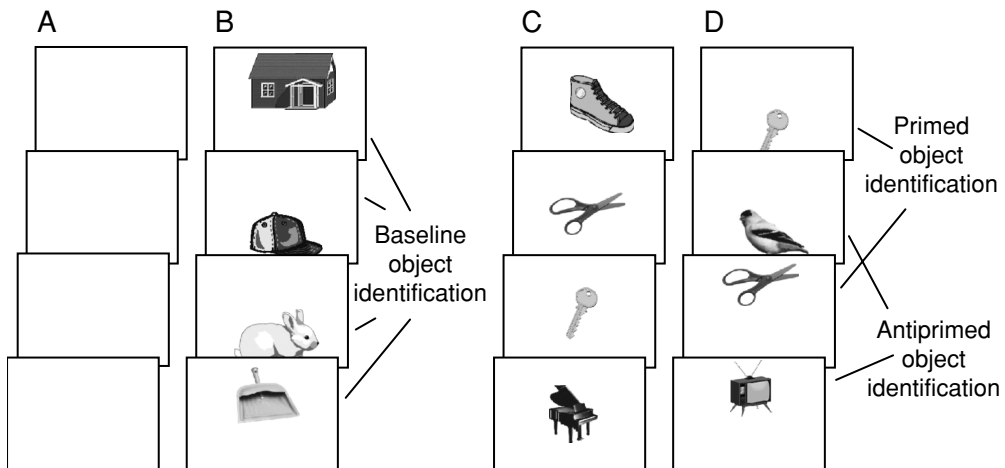


**Figure 1.** Schematic depiction of superimposed representations of two object shapes. The four units to the right of the piano should be understood to be the same four units as those depicted to the right of the desk, with their differing patterns of activation reflecting which object is being represented. Each of the four units is maximally activated by a moderately complex feature of visual objects. The middle two units are activated by two of the features in the piano, and the bottom two units are activated by two of the features in the desk. Because the piano and the desk contain a common feature, their distributed representations are superimposed.

hood that that feature will subsequently activate the representation for a different object that also contains that feature (desk). In other words, some of the changes that enhance the ability to reactivate the representation for the primed object (piano) simultaneously impair the ability to activate the representations of other objects (desk, table) that share features or representational subpatterns with the primed object (right-front leg, but note that the critical elements of representational overlap are not yet well understood).

*Antipriming* is our label for the hypothesized detrimental effect of previous weight changes on the subsequent identification of objects that are different from the primed objects but that have representations that are superimposed with those of the primed objects. Theoretically, the relevant objects need only be superimposed in their representations for antipriming to occur; they need not be discernibly very similar (e.g., chairs and houses are not very similar visually, yet they activate overlapping areas of ventral temporal cortex; Haxby et al., 2001; Ishai et al., 1999). The present study provides functional evidence for visual antipriming through behavioral, neurocomputational, and neuropsychological investigations.

In our paradigm, participants first stare at a blank display while hearing and judging the names of objects (Figure 2A). Then they see and identify a set of briefly presented objects that do not correspond with the previously heard names (Figure 2B); this task provides a *baseline* measure of visual object identification performance (against which both priming and antipriming will be measured) that occurs following a period in which no visual object identification took place, but processing of object



**Figure 2.** Schematic depiction of examples of trials in the four phases (A–D) of the behavioral experiments. The objects were presented in color in the experiments. Baseline identification of briefly presented objects is measured during the second phase (B), following a period during which no objects were viewed (A). Priming is measured during the fourth phase (D) as more accurate identification of briefly presented objects (compared against baseline) that are repeated from the preceding phase (C). Antipriming is measured during the fourth phase (D) as less accurate identification of briefly presented objects (compared against baseline) that are new with respect to the preceding phases, following a period of viewing other visual objects (C).

names did. Next participants view and judge a new set of visual objects (Figure 2C). Presumably, the visual representations of these objects are strengthened during this phase. Finally, participants again see and identify briefly presented objects (Figure 2D), but in this final phase, half of the objects are *primed* (i.e., they were viewed in the immediately preceding phase), and the other half are *antiprimed* (i.e., they are different from all previously presented objects). Note that the main difference between performance in the antiprimed condition and the baseline condition is that the period preceding antiprimed performance involves viewing of objects whereas the period preceding baseline performance involves no viewing of objects. Perceptual identification of very briefly presented objects is used in the baseline, primed, and antiprimed conditions to increase the degree to which visual processing underlies performance in those conditions. In addition, unlike in a methodologically similar study (Ratcliff & McKoon, 1996), in this study, none of the objects was chosen to be highly visually similar. Under our conditions, priming is measured as more accurate performance in the primed condition than in the baseline condition, and antipriming is measured as less accurate performance in the antiprimed condition than in the baseline condition. Such an antipriming effect should be observed if the antiprimed test objects have representations that are superimposed with representations of the visual objects that were strengthened in the preceding phase.

## EXPERIMENT 1

We first tested whether the hypothesized antipriming effect would be observed in healthy young adults.

## Method

**Participants.** Twenty undergraduate students at the University of Minnesota participated for course credit.

**Materials.** The experimental stimuli were 250 colored line drawings (5.4° in height), each depicting a unique familiar object, and 250 auditory recordings of the entry-level names of those objects (Jolicœur, Gluck, & Kosslyn, 1984). An additional 9 drawings were used for practice. The drawings of objects were a subset of those developed for and used in a previous study of object priming (Koutstaal et al., 2001). None of the objects was highly similar to any other object in the set. A personal computer controlled presentation of stimuli. Full counterbalancing assured that each object represented each experimental condition an equal number of times across participants.

**Procedure.** The experiment had four phases. First, participants heard auditory recordings of the names of 50 objects while staring at a blank computer display (Figure 2A). For each name, they judged whether they liked or disliked what it represented (considering the meanings associated with the objects, not how they sounded or looked).

Second, for the baseline measure, participants viewed 100 drawings of objects presented briefly (for 15 msec) and centered 4.3° above or below the center of the display, with presentation location pseudorandomly intermixed (Figure 2B). These objects were all different from the objects whose names were presented in the preceding task. For each object, participants tried to identify it and then spoke aloud an acceptable name for it as accurately and quickly as possible. Acceptable names for an object included synonyms (e.g., either “phone” or “telephone” would be acceptable for a telephone) or near-synonyms to its entry-level name (e.g., either “sofa” or “couch” would be acceptable for a sofa). Response times were measured using a voice-activated key. (Objects were presented slightly above or below the center because pilot results indicated that identification accuracy was at ceiling when objects were presented centrally.)

Third, participants viewed 50 drawings of objects presented centrally for 3 sec each (Figure 2C). These objects were all different from the objects that were presented in the preceding phases (either visually or as an auditory name). For each trial, participants judged whether they liked or disliked what the object represented (again


considering the meanings associated with the objects, not how they sounded or looked).

Fourth, participants viewed 100 drawings of objects presented in the same manner as in the second phase (Figure 2D). Half of the objects that were presented above fixation and half of the objects that were presented below fixation were primed, in that they were the same as the objects presented visually in the immediately preceding phase. The other half of the objects that were presented above fixation and the other half of the objects that were presented below fixation were antiprimed, in that they were different from the visual objects presented previously (they also did not correspond with the auditory names of objects presented previously) and they were identified at a time that followed a period of viewing visual objects. For the antiprimed and primed test objects (pseudorandomly intermixed), participants tried to identify each one and then spoke aloud an acceptable name for it as accurately and quickly as possible. Response times were measured using a voice-activated key.

## Results and Discussion

Separate ANOVAs were conducted using identification accuracy and mean and median response times for correctly identified objects as dependent measures. In all analyses, the independent variable was test presentation condition (baseline, antiprimed, or primed).

The main effect of test presentation condition was significant in accuracy rates [ $F(2,38) = 105.1$ ,  $MS_e = 45.0$ ,  $p < .0001$ ; Figure 3]. Most important, antiprimed objects were identified less accurately (51.0%) than objects presented in the baseline condition (55.8%) [ $F(1,38) = 5.12$ ,  $MS_e = 45.0$ ,  $p < .05$ , for the simple effect contrast].



**Figure 3.** Results indicating significant antipriming and priming in Experiment 1 (with young adults). Mean object identification accuracy rates ( $\pm SEM$ ) are plotted. Antiprimed objects were identified significantly less accurately than baseline, and primed objects were identified significantly more accurately than baseline. \* $p < .05$ . \*\* $p < .0001$ .

Significant priming also was observed, replicating many previous studies of repetition priming for familiar objects (see, e.g., Bartram, 1974; Biederman & Cooper, 1991; Cave & Squire, 1992; Jacoby, Baker, & Brooks, 1989; Marsolek, 1999; Mitchell & Brown, 1988; Snodgrass & Feenan, 1990; Srinivas, 1993; Weldon & Roediger, 1987); primed objects were identified more accurately (79.7%) than objects presented in the baseline condition (55.8%) [ $F(1,38) = 127.1$ ,  $MS_e = 45.0$ ,  $p < .0001$ , for the simple effect contrast]. In addition, the magnitude of the antipriming effect was significantly smaller than the magnitude of the priming effect [ $t(19) = 7.32$ ,  $p < .0001$ ], which is important for comparisons with effects observed in the following neurocomputational models.

The same pattern of results was observed in the analyses of response times for correct responses, but significance was not observed in the relevant simple effects. The main effect of test presentation condition was marginally significant in mean response times [ $F(2,38) = 3.14$ ,  $MS_e = 5,754.6$ ,  $p < .06$ ], and significant in median response times [ $F(2,38) = 3.80$ ,  $MS_e = 8,696.5$ ,  $p < .05$ ]. However, the differences between antiprimed and baseline conditions (means of 1,167 msec vs. 1,139 msec, medians of 1,145 msec vs. 1,117 msec, respectively) were not significant [ $F(1,38) = 1.34$ ,  $p > .25$ , for means;  $F(1,38) < 1$ , for medians], and the differences between primed and baseline conditions (means of 1,107 msec vs. 1,139 msec, medians of 1,065 msec vs. 1,117 msec, respectively) also did not achieve significance [ $F(1,38) = 1.81$ ,  $p > .15$ , for means;  $F(1,38) = 3.10$ ,  $p > .08$ , for medians]. Most important, however, these results indicated no trade-offs between speed and accuracy of response.

## NEUROCOMPUTATIONAL MODELS

We next investigated whether extant neurocomputational models of repetition priming that utilize superimposed representations (Stark & McClelland, 2000) also exhibit antipriming of the sort observed in Experiment 1. This prediction has not been investigated previously. We examined parallel distributed processing models that are biologically inspired but simplified and relatively abstract with respect to neural details (McClelland & Rumelhart, 1985). These are autoencoder networks with units that accept external input from the environment and internal input from weighted connections to other units during an iterative settling process with patterns of activation updated in cycles. These models can be trained to identify patterns using a biologically plausible Hebbian learning rule (Hebb, 1949), which guides small changes in the weights on the connections in each training trial. Importantly for present purposes, these models learn and store superimposed representations. After these models are trained to identify a set of patterns, repetition priming can be simulated very successfully by applying an additional training experience for some of the previously learned patterns, which enhances subsequent identification of those patterns (Stark & McClelland, 2000). If such models accurately simulate antipriming as well, they will

provide a concrete specification of the hypothesized origin of both effects.


## Method

The Hebbian rule trained models of pattern identification and priming examined by Stark and McClelland (2000) and based on previous work (McClelland & Rumelhart, 1985) were replicated and extended for this study. Ten models (simulating 10 participants) of 64 fully interconnected units were examined, with 64-element orthogonal patterns serving as proxies for different visual objects. The models first were trained to identify a set of input patterns to a level of performance that was good, but not at a ceiling level that would prohibit possible subsequent improvement. Then, their accuracy in identifying those patterns was measured to obtain a baseline level of performance (analogous to the baseline performance in Experiment 1). Next, one additional trial of training for half of the patterns was administered, followed by tests of accuracy in identifying the primed patterns (known patterns that were retrained in the preceding phase) and the antiprimed patterns (known patterns that were not included in the preceding phase of retraining other patterns).

All aspects of the procedure were the same as those of the replicated study (Stark & McClelland, 2000) with the following exceptions. First, we measured both antipriming and priming (not priming only). Second, because the behavioral experiments in our study involved identification of only familiar objects, we used a set of familiar patterns throughout the study (instead of comparing priming among familiar, pseudofamiliar, and unfamiliar patterns). Third, because the behavioral experiments in our study involved identification of very briefly presented objects and accuracy was the primary dependent variable, we set the number of settling cycles to 10 during the simulated experiment (half as many as the 20 cycles used during initial training). We measured accuracy of performance at that point (normalized dot product between the network's pattern of activation after settling and the target output for that pattern) as the dependent variable (as opposed to the number of cycles needed to achieve a particular strength of response).

## Results and Discussion

Replicating results from Stark and McClelland (2000) and results from other modeling studies of priming (Bowers et al., 2002; Marsolek & Burgund, 1997; Rueckl, 2003), primed patterns were identified significantly more accurately (normalized dot product = .730) than baseline patterns [.695;  $F(1,18) = 2,922.8, MS_e = 2E-6, p < .0001$ ; Figure 4]. More important, antiprimed patterns were identified significantly less accurately (normalized dot product = .689) than baseline patterns [.695;  $F(1,18) = 78.9, MS_e = 2E-6, p < .0001$ ; Figure 4]. The main effect of test presentation status was significant [ $F(2,18) = 2,321.2, MS_e = 2E-6, p < .0001$ ]. Also, the magnitude of the antipriming effect was significantly smaller than the magnitude of the priming effect [ $t(9) = 49.7, p < .0001$ ]. These results were very similar to the behavioral results obtained in Experiment 1, as confirmed with a simple quantitative analysis. Pattern identification performance in the models was fitted to object identification accuracy in Experiment 1 with a least squares linear fit (accuracy = normalized dot product  $\times 706.85 - 435.96$ ). The fit was extremely good ( $R^2 = .999$ ). In line with our theory, this study indicates that the weight changes that enable neurocomputational models with superimposed representations to simulate repetition priming also cause them to exhibit antipriming.



**Figure 4. Results indicating significant antipriming and priming in neurocomputational models simulating Experiment 1. Mean pattern identification accuracy ( $\pm SEM$ ) is plotted. Antiprimed objects were identified significantly less accurately than baseline, and primed objects were identified significantly more accurately than baseline. NDP, normalized dot product.  $**p < .0001$ .**

## EXPERIMENT 2


We next tested whether antipriming would be observed when no primed objects were presented with the antiprimed objects during the final phase of the experiment. Such a finding would indicate that antipriming is not caused by a bias or attention toward the primed objects made at the cost of disregarding antiprimed objects. The experiment was conducted in the same manner as Experiment 1, except an additional group of 30 participants were tested, and the primed trials in the final phase of the experiment were deleted.

## Results and Discussion

As in Experiment 1, antiprimed objects were identified significantly less accurately (52.0%) than objects presented in the baseline condition (55.7%) [ $F(1,29) = 6.47, MS_e = 31.7, p < .05$ ; Figure 5]. Also as before, no effect in the analyses of response times approached significance. Thus, antipriming is not caused by a bias or attention toward primed objects made at the cost of disregarding antiprimed objects.

## EXPERIMENT 3

We also tested another alternative explanation for the antipriming effect observed in the previous experiments. Because antiprimed object identification is measured at



**Figure 5.** Results indicating significant antipriming in Experiment 2 (with young adults), when primed objects were not included in the final phase of the experiment. Mean object identification accuracy rates ( $\pm$ SEM) are plotted. Antiprimed objects were identified significantly less accurately than baseline.  $^*p < .05$ .

a time that is later in the experimental session than when baseline object identification is measured, it is possible that the putative antipriming effect is due to greater participant fatigue or boredom during the final phase of our paradigm than during the earlier phase. If this were the case, poorer performance in the antiprimed condition than in the baseline condition should be observed even if the third phase of the experiment involved hearing object names instead of viewing drawings of objects. This is because roughly the same time should elapse between the second and fourth phases and roughly the same effort is required during that time, when comparing the two procedures. According to our theory, however, no difference in performance should be observed in the antiprimed condition compared with the baseline condition if the third phase of the experiment involved no viewing of drawings of objects (viewing of objects is needed for object representations to be strengthened, which causes visual antipriming). The purpose of Experiment 3 was to test these predictions. It was conducted in the same manner as Experiment 1, with an additional group of 20 participants, the only exception being that, rather than viewing drawings of objects in the third phase of the experiment, the participants instead heard the names of the objects.


## Results and Discussion

Unlike in Experiment 1, the main effect of test presentation condition was not significant in accuracy rates [ $F(2,38) = 1.44$ ,  $MS_e = 49.2$ ,  $p > .20$ ; Figure 6], mean response times [ $F(2,38) = 2.52$ ,  $MS_e = 4,056.5$ ,  $p > .09$ ],

or median response times [ $F(2,38) = 1.52$ ,  $MS_e = 5,683.1$ ,  $p > .20$ ]. Most important, in contrast with Experiment 1, antiprimed objects were identified at a rate (55.3%) that was numerically greater than objects presented in the baseline condition (53.8%), although this difference was not significant ( $F < 1$ , for the simple effect contrast). These results run counter to the hypothesis that the antipriming effects observed in the previous experiments were due to a fatigue effect. Also unlike in Experiment 1, (cross-modality) primed objects were identified at a rate (57.5%) that was not significantly greater than objects presented in the baseline condition (53.8%) [ $F(1,38) = 2.86$ ,  $MS_e = 49.2$ ,  $p > .09$ , for the simple effect contrast]. The lack of significant positive priming in this experiment and the highly significant (within-modality) positive priming in Experiment 1 indicate that priming from visual object representations was largely responsible for the strong positive priming in Experiment 1. With very brief visual presentations of test objects, our paradigm is propitious for measuring only priming and antipriming that take place within the visual modality.

## EXPERIMENT 4

A critical test is whether explicit or declarative memory may be responsible for the antipriming effect observed in Experiment 1. This possibility stands in contrast with our explanation that antipriming reflects small structural changes to superimposed object representations in visual



**Figure 6.** Results indicating no "antipriming" effect in Experiment 3 (with young adults), in which participants, instead of viewing visually presented objects, heard names of objects in the third phase of the experiment. Mean object identification accuracy rates ( $\pm$ SEM) are plotted. "Antiprimed" objects were not identified less accurately than baseline.

neocortex. Participants may find the object identification task difficult enough that they utilize their explicit memories for recently viewed objects to help their performance. This could impair identification of antiprimed objects because none of them is in explicit memory. Thus, we investigated whether explicit memory impaired patients with medial temporal amnesia would exhibit antipriming.

## Method


**Participants.** Nine patients with amnesia due to bilateral medial temporal lobe lesions (followed at the Memory Disorders Research Center, Boston University School of Medicine) and 10 control participants without history of neurological illness participated in this experiment. The control participants were matched to the patients in age, education, and IQ. The patients' etiologies, demographic information, and performance on standardized tests are reported in Table 1.

**Procedure.** The procedure was the same as in Experiment 1, except that during the second and fourth phases of the experiment, drawings of objects were presented for 33 msec each instead of 15 msec each; pilot data indicated that the patients and the older matched control participants required longer durations than the young adults to obtain similar levels of correct identification. Also, only 9 patients could be tested (and response times for one of the patients were lost due to equipment malfunction), disallowing complete counterbalancing of stimuli in that group. Full counterbalancing of stimuli was accomplished for the 10 control participants however.

## Results and Discussion

The main analyses were the same as in Experiment 1 with the addition of a participant group variable, so that  $2 \times 3$  mixed-effects repeated measures ANOVAs were performed.

The main effect of test presentation condition was significant in accuracy rates [ $F(2,34) = 66.8$ ,  $MS_e = 29.1$ ,  $p < .001$ ; Figure 7]. Most important, replicating Experiment 1, antiprimed objects were identified less accurately (33.5%) than baseline objects (40.2%) [ $F(1,34) = 14.6$ ,  $MS_e = 29.1$ ,  $p < .01$ ], and primed objects were identified more accurately (53.6%) than baseline objects (40.2%) [ $F(1,34) = 58.6$ ,  $MS_e = 29.1$ ,  $p < .0001$ ]. In addition,



**Figure 7.** Results indicating significant antipriming and priming in amnesic patients and matched control participants in Experiment 4. Mean object identification accuracy rates ( $\pm SEM$ ) are plotted. Antiprimed objects were identified significantly less accurately than baseline, and primed objects were identified significantly more accurately than baseline, in both amnesic patients and matched control groups. \* $p < .01$ . \*\* $p < .001$ .

the main effect of group was significant [ $F(1,17) = 8.03$ ,  $MS_e = 483.5$ ,  $p < .05$ ]; amnesic patients were less accurate at identifying objects (33.7%) than control participants (50.2%). Finally, the interaction between group and test presentation condition was significant [ $F(2,34) = 5.14$ ,  $MS_e = 29.1$ ,  $p < .05$ ]; this interaction was caused by a significant difference in the magnitudes of the priming effect in patients (7.8%) versus controls (18.5%) [ $F(1,34) =$

**Table 1**  
Etiologies, Demographics, and Neuropsychological Performance Values for the Patients ( $n = 9$ ) and Means for the Patients and the Control Participants ( $n = 10$ )

Patient	Etiology	Age	Ed.	WAIS-III VIQ	WMS-III			
					GM	AD	VD	WM
1	Anoxia	75	18	113	75	80	72	102
2	Anoxia	48	14	111	59	72	52	93
3	Encephalitis	49	14	92	45	56	55	96
4	Thalamic stroke	63	12	84	73	67	84	99
5	Anoxia	52	12	83	52	56	55	81
6	Anoxia	44	14	90	45	53	52	91
7	Anoxia	20	12	91	45	46	56	102
8	Encephalitis	61	12	106	69	77	68	111
9	Anoxia	51	17	134	70	67	75	88
Patient Mean		51	13.9	100	59	64	63	96
Control Mean		55.4	13.7	105				

Note—Age and education are in years; WAIS-III, Wechsler Adult Intelligence Scale, 3rd edition; VIQ, verbal IQ; WMS-III, Wechsler Memory Scale, 3rd edition (Wechsler, 1997); GM, general memory index; AD, auditory delay index; VD, verbal delay index; WM, working memory index.

9.87,  $MS_e = 29.1$ ,  $p < .01$ ], as the magnitudes of the antipriming effect did not differ for patients (8.0%) and controls (5.5%) ( $F < 1$ ). It is important to note, however, that the amnesic patients still demonstrated a significant level of priming. Replicating previous studies of patients (see, e.g., Cave & Squire, 1992), primed objects were identified more accurately (41.6%) than baseline objects (33.8%) [ $F(1,51) = 13.5$ ,  $MS_e = 20.2$ ,  $p < .001$ ]. Similarly, for the control participants, primed objects were identified more accurately (64.4%) than baseline objects (45.9%) [ $F(1,51) = 84.7$ ,  $MS_e = 20.2$ ,  $p < .0001$ ]. Furthermore, the patients identified antiprimed objects less accurately (25.8%) than baseline objects (33.8%) [ $F(1,51) = 14.3$ ,  $MS_e = 20.2$ ,  $p < .0001$ ]; the control participants also identified antiprimed objects less accurately (40.4%) than baseline objects (45.9%) [ $F(1,51) = 7.49$ ,  $MS_e = 20.2$ ,  $p < .01$ ].

In the analyses of response times for correct responses, only the main effect of group approached significance [ $F(1,16) = 3.72$ ,  $MS_e = 65,400.7$ ,  $p < .08$ , for means] or was significant [ $F(1,16) = 7.51$ ,  $MS_e = 55,780.1$ ,  $p < .05$ , for medians]. The patients generally identified objects more slowly than the control participants (means of 1,133 msec vs. 998 msec; medians of 1,129 msec vs. 952 msec). No other effects approached significance, thus the results indicated no trade-offs between speed and accuracy.

An initially surprising aspect of these results was that the magnitude of priming was smaller for the patients than for the control participants. However, priming measures with amnesic patients often are normalized (by dividing the difference between primed and unprimed performance by the unprimed measure) to control for the typically differing levels of unprimed performance. In typical studies of priming, what is labeled the unprimed condition during the test phase is what we label antiprimed in our paradigm. Thus, we calculated normalized priming scores in the typical manner [(primed – antiprimed performance)/antiprimed performance] and with these scores we found no difference between priming in the patients (.81) and in the control participants (.71) [ $F(1,17) < 1$ ].

Perhaps more interesting, however, is the following possible reason that the (nonnormalized) priming magnitude was reduced in patients compared with that in the control participants. The patients generally were slow in their processing, as has been reported in previous studies of amnesia (Corkin, 1984; Postle & Corkin, 1998). In the present experiment, the patients exhibited lower accuracy and greater response times overall than did the control participants. Reduced speed and accuracy in neural systems that operate independently of the medial temporal areas that are damaged in the patients may be a secondary aspect of their performance that could be responsible for the reduced priming. According to our theory, priming for one object at test is caused mostly by the weight changes that occur after that one object is identified during visual encoding. In other words, priming for one object is caused mostly by one trial's worth of changes. By contrast, antipriming for one object at test is caused by the weight changes that occur after each of several differ-

ent objects is identified during visual encoding—that is, those objects that have representations that are superimposed with the representation of the test object. In other words, antipriming for one object is caused by many trials' worth of changes. This may be important because, when neurocomputational models are challenged or damaged, they tend to retain relatively specific knowledge (like that from changes in one trial) more poorly than relatively generalized knowledge (like that from many previous accumulated changes; McClelland & Rogers, 2003). Thus, magnitudes of priming in particular (and not the accompanying magnitudes of antipriming) may be detrimentally affected when such systems are slowed down.

To test this possibility, the networks that were used in the previous modeling study were reexamined in another simulated priming and antipriming experiment. The only difference was that the number of settling cycles per trial was set to 5 instead of 10 during pattern identification. Restricting the number of settling cycles that are allowed before measuring the response in a trial may simulate performance when secondary effects of brain damage cause slow identification. Supporting this reasoning, overall identification accuracy in the 5-cycle models (normalized dot product = .463) was about two thirds that of the identification accuracy in the 10-cycle models (.705), mirroring the nearly two thirds overall accuracy reduction in patients (33.7%) compared with that in the control participants (50.2%). Yet, as previously found with the 10-cycle models, primed patterns were identified significantly more accurately (.479) than baseline patterns (.456) [ $F(1,18) = 7,059.1$ ,  $MS_e = 3.7E-7$ ,  $p < .0001$ ], and antiprimed patterns were identified significantly less accurately (.453) than baseline patterns (.456) [ $F(1,18) = 142.2$ ,  $MS_e = 3.7E-7$ ,  $p < .0001$ ]. Most relevant here, the reduction in priming magnitude between 10-cycle and 5-cycle models (.035 vs. .023) was significant [ $t(9) = 12.7$ ,  $p < .0001$ ], paralleling the reduced magnitude of priming in patients compared with that in the control participants.

Finally, we tested this explanation for reduced priming in the patients in the following way. Taking advantage of individual differences in patient performance, we correlated their overall identification accuracy levels with their magnitudes of priming. This correlation was significant ( $r = .68$ ,  $p < .05$ ), indicating that the less accurate the patient's responses, the lesser the magnitude of priming observed for that patient. This finding and the simulation of cognitive slowing and reduced priming in amnesic patients may be crucial for helping to effectively interpret memory effects in both patients and control participants.

Most important, the amnesic patients were significantly impaired in standard explicit memory measures (Table 1), yet they exhibited significant antipriming and priming. Thus, in line with our theory, intact explicit memory is not necessary for antipriming.

## GENERAL DISCUSSION

Representations of different items of well-established knowledge may be localist/discrete or superimposed/over-

lapping. This distinction has been understood for decades as fundamental, with major implications for theories of memory (see, e.g., Rumelhart & Norman, 1984). For representations of visual object shapes in ventral neocortex, neural evidence is that representations are superimposed. Neurocomputational modeling indicates important utilities for this kind of representation, but these utilities are accompanied by certain costs. In this article, we report functional evidence for superimposed representations of visually presented objects. We report visual antipriming effects that are predictable from the hypothesized detrimental effects of using superimposed representations to store different visual object shapes and from continual adjustments of such representations with each use. By our theory, the representation of an object is strengthened via small representational changes after that object is identified, and this is responsible for repetition priming. However, those same structural changes simultaneously interfere with storage of other representations with which the strengthened representation is superimposed, and this is what causes visual antipriming. We measure priming as more accurate performance in the repetition primed condition than in our baseline condition, and we measure antipriming as less accurate performance in the antiprimed condition than in our baseline condition (within which little priming or antipriming should be measured).

We should note that our paradigm involves a novel baseline condition that differs from that used in most priming studies. Typically, primed (previously presented) and unprimed (not previously presented) objects are presented during a test phase, and priming is calculated as more effective processing of primed objects than unprimed objects. By our theory, this is a comparison between a primed condition and an antiprimed (not unprimed) condition. The baseline condition in our paradigm, obtained following a period in which participants stared at a blank display and did not view any objects, is necessary to separately measure the positive effect of priming and the negative effect of antipriming. It is worth noting that even in our paradigm, the baseline measure is not completely clean, as real-world object identification before the experiment begins may cause some degree of antipriming of the baseline objects. Nevertheless, any such (uncontrollable) antipriming in our baseline measure should work against our finding a significant antipriming effect, yet we did in fact find antipriming to be significant and reliable across experiments.

Another important aspect of our paradigm is that the task differs between the preceding phase (liking judgments) and the final phase in which antiprimed and primed performance is measured (identification of visually presented objects). Thus, any rapid learning of a particular response to a particular stimulus (Logan, 1990) that may occur during the preceding phase could not be responsible for any antipriming or priming. This is important because rapid response learning can underlie behavioral and neuroimaging priming effects when the tasks are the same (Dobbins, Schnyer, Verfaellie, & Schacter, 2004). Such rapid response learning is not posited within the object

identification systems that we hypothesize to underlie priming and antipriming of objects. Instead, a form of strengthening and maintaining of partially superimposed representations is posited in these systems.

### **Purpose for Priming?**

Long-term repetition priming is a well-established and frequently studied aspect of human perception and memory. Forms of priming have been in evidence for over a century (Schacter, 1987), densely amnesic patients exhibit robust priming despite very poor performance in explicit memory tasks (Warrington & Weiskrantz, 1974), and participants who exhibit priming need not be aware that memory is exhibited in performance (Bowers & Schacter, 1990). However, why priming exists or what purpose it may have is debatable.

One possibility is that priming does not serve a purpose per se. Instead, it may be a by-product of repeated information processing. This by-product could take the form of persisting lowered thresholds for activation of the relevant representations (Morton, 1979), or it may reflect a bias that operates in a postperceptual decision stage of identifying objects that “can be understood simply as a by-product of ordinary object perception and identification processes” (Rouder, Ratcliff, & McKoon, 2000, p. 13).

Other theories have posited some functional purpose or utility for priming. A potential utility is that it may be evolutionarily adaptive to effectively reprocess recently encountered stimuli, because recently encountered stimuli may be more likely to be encountered again (Roediger, 2003; Schacter, 1994; Squire, Knowlton, & Musen, 1993). This is reasonable but difficult if not impossible to test. In addition, priming may reflect sharpening of the representations that underlie perceptual identification (Desimone, 1996; Wiggs & Martin, 1998). Repeated familiar objects activate a smaller number of neurons and elicit smaller magnitude neuroimaging signals than objects that have not been processed recently. Thus, activations elicited by primed objects may be more efficient than activations elicited by unprimed objects, which provides a benefit, although deriving this conclusion from the results can be debated (Henson & Rugg, 2003). Unless the sharpening of one representation is hypothesized to detrimentally affect other representations, this theory does not explain antipriming.

Our theory posits a different purpose for priming. Priming effects are caused by structural changes in visual representations after objects are identified, enabling storage and maintenance of superimposed representations. Superimposed representations have important benefits; therefore, the structural changes that maintain them have an important purpose. In addition, our theory thus leads us to hypothesize an ongoing cycle. Strengthening one representation weakens other representations with which it is superimposed. This creates a purpose for the subsequent strengthening of a weakened representation at the next opportunity (when its object is next encountered and identified), thus starting this kind of cycle over again.

### Contrasts With Bias Theory

A theory that appears on the surface to be similar to our theory posits that priming reflects a bias effect (Ratcliff & McKoon, 1996; Rouder et al., 2000). The important theoretical differences between our theory and the bias theory are the following.

First, antipriming requires superimposed representations, since it is the negative effect of strengthening one object representation on subsequent processing of other objects with which the original object is superimposed. In contrast, in the Rouder et al. (2000) model of bias in object identification, different object shape representations are not superimposed; instead, completely separate modules are used to process the shapes of different objects.

Second, antipriming involves changes to visual shape representations. In contrast, in Rouder et al.'s (2000) model, prior encounters with stimuli do not affect shape representations, but rather postperceptual decision processing. The shape-processing module for an object has an output node that is connected to a subsequent decision node for that object. During object identification, the final activation of a decision node for an object is calculated as a ratio of its activation relative to the summed activations of the decision nodes for all other objects. Bias from an encounter with an object (e.g., a lightbulb) is implemented as a strengthening of the connection between the output node and the decision node for that object. A benefit of bias is measured when the strengthened connection causes that object (the lightbulb) to be subsequently reidentified faster than before. A cost of bias occurs only between two highly similar objects: A previous encounter with one object (a lightbulb) can reduce the final activation of the decision node for another object at test (e.g., a hot air balloon) only to the extent that the two objects are similar enough that the previously unrepresented test object (the hot air balloon) activates the decision node for the previously presented object (the lightbulb). Such activation of the decision node for the previously presented object (the lightbulb) increases the summed activations of the decision nodes for the nontested objects, thereby decreasing the final activation of the previously unrepresented test object's decision node (the hot air balloon), which is calculated relative to that sum. By this theory, neither the positive nor the negative effect of bias involves changes to shape representations per se, and any negative effect of bias occurs only between highly similar objects (such as a lightbulb and a hot air balloon).

Thus, an important feature of the experiments reported in this article is that none of the objects were selected to be highly visually similar to any other in the set. The antipriming effects that we observed between visually dissimilar objects are not predicted by the Rouder et al. (2000) model, which therefore should be revised to accommodate our results.

Perhaps most important, our theory can account for both our antipriming effects (with visually dissimilar objects) and the empirical evidence that Ratcliff and McKoon (1996) offer in support of their bias theory. Using highly similar visually presented objects, they found that after

objects (e.g., a lightbulb) were named, object naming was faster (a benefit) for test objects that had been processed previously (the lightbulb) than for test objects that had not been processed previously and were dissimilar to the previously processed objects (e.g., a mailbox). In contrast, object naming was slower (a cost) for test objects that had not been processed previously and were very visually similar to previously processed objects (e.g., a hot air balloon presented after previous processing of a lightbulb) than for test objects that had not been processed previously and were dissimilar to the previously processed objects (e.g., a mailbox). This finding was considered to reflect a bias in the interpretation of the test stimulus in part because the processing cost approximated the processing benefit in magnitude (see also Ratcliff & McKoon, 1995, 1997). It is also possible, however, that the processing cost reflects a relatively strong form of antipriming caused by high levels of superimposition of representations for visually similar objects and strengthening of the representations for the primed objects.

Alternatively, the results from Ratcliff and McKoon (1996) may be due to another memory source. Because they used the same task (object naming) in both the encoding and the test sessions, rapid learning of a particular response (i.e., a name) to a particular stimulus (i.e., a shape) could have occurred during the encoding session (see Logan, 1990), and this may have been responsible for the observed benefit and the observed cost in the test session. Our recent evidence indicates that such rapid response learning depends on the medial temporal structures that support explicit memory (Schnyer, Dobbins, Nicholls, Schacter, & Verfaellie, 2006). Thus, the bias effects described by Ratcliff and McKoon may be due to different memory processes than those that underlie visual antipriming, which does not depend on the medial temporal structures that support explicit memory (see Experiment 4).

### Conclusion

Our approach to uncovering the nature of knowledge representation and the nature of repetition priming and related memory effects involves integrating evidence from neural, functional, and computational levels of analysis. This approach is proving to be of great utility in many areas of cognitive neuroscience. The integration reported in this article converges on the conclusion that visual object representations are distributed and superimposed, not localist/discrete, and that both priming and antipriming of visual objects reflect ongoing adjustments of these superimposed representations in neocortex.

### REFERENCES

- ANDERSON, J. A., & HINTON, G. E. (1981). Models of information processing in the brain. In G. E. Hinton & J. A. Anderson (Eds.), *Parallel models of associative memory* (pp. 23-62). Hillsdale, NJ: Erlbaum.
- BARTRAM, D. J. (1974). The role of visual and semantic codes in object naming. *Cognitive Psychology*, *6*, 325-356.
- BIEDERMAN, I., & COOPER, E. E. (1991). Evidence for complete translational and reflectional invariance in visual object priming. *Perception*, *20*, 585-593.

- BOWERS, J. S., DAMIAN, M. F., & HAVELKA, J. (2002). Can distributed orthographic knowledge support word-specific long-term priming? Apparently so. *Journal of Memory & Language*, **46**, 24-38.
- BOWERS, J. S., & SCHACTER, D. L. (1990). Implicit memory and test awareness. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **16**, 404-416.
- CAVE, C. B., & SQUIRE, L. R. (1992). Intact and long-lasting repetition priming in amnesia. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **18**, 509-520.
- CHEN, W. R., LEE, S., KATO, K., SPENCER, D. D., SHEPHERD, G. M., & WILLIAMSON, A. (1996). Long-term modifications of synaptic efficacy in the human inferior and middle temporal cortex. *Proceedings of the National Academy of Sciences*, **93**, 8011-8015.
- CORKIN, S. (1984). Lasting consequences of bilateral medial temporal lobectomy: Clinical course and experimental findings. *Seminars in Neurology*, **4**, 252-262.
- DESIMONE, R. (1996). Neural mechanisms for visual memory and their role in attention. *Proceedings of the National Academy of Sciences*, **93**, 13494-13499.
- DOBBINS, I. G., SCHNYER, D. M., VERFAELLIE, M., & SCHACTER, D. L. (2004). Cortical activity reductions during repetition priming can result from rapid response learning. *Nature*, **428**, 316-319.
- EICH, J. M. (1982). A composite holographic associative retrieval model. *Psychological Review*, **89**, 627-661.
- GRAINGER, J., & JACOBS, A. M. (1996). Orthographic processing in visual word recognition: A multiple read-out model. *Psychological Review*, **103**, 518-565.
- GROSS, C. G., ROCHA-MIRANDA, C. E., & BENDER, D. B. (1972). Visual properties of neurons in inferotemporal cortex of the macaque. *Journal of Neurophysiology*, **35**, 96-111.
- HAXBY, J. V., GOBBINI, M. I., FUREY, M. L., ISHAI, A., SCHOUTEN, J. L., & PIETRINI, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, **293**, 2425-2430.
- HEBB, D. O. (1949). *The organization of behavior*. New York: Wiley.
- HENSON, R. N. A., & RUGG, M. D. (2003). Neural response suppression, haemodynamic repetition effects, and behavioural priming. *Neuropsychologia*, **41**, 263-270.
- HINTON, G. E., MCCLELLAND, J. L., & RUMELHART, D. E. (1986). Distributed representations. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 1: Foundations* (pp. 77-109). Cambridge, MA: MIT Press.
- HINTZMAN, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, **93**, 411-428.
- ISHAI, A., UNGERLEIDER, L. G., MARTIN, A., SCHOUTEN, J. L., & HAXBY, J. V. (1999). Distributed representation of objects in the human ventral visual pathway. *Proceedings of the National Academy of Sciences*, **96**, 9379-9384.
- JACOBY, L. L. (1983). Perceptual enhancement: Persistent effects of an experience. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **9**, 21-38.
- JACOBY, L. L., BAKER, J. G., & BROOKS, L. R. (1989). Episodic effects on picture identification: Implications for theories of concept learning and theories of memory. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 275-281.
- JOLICEUR, P., GLUCK, M. A., & KOSSLYN, S. M. (1984). Pictures and names: Making the connection. *Cognitive Psychology*, **16**, 243-275.
- KNAPP, A., & ANDERSON, J. A. (1984). A signal averaging model for concept formation. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **10**, 616-637.
- KOBATAKE, E., WANG, G., & TANAKA, K. (1998). Effects of shape-discrimination training on the selectivity of inferotemporal cells in adult monkeys. *Journal of Neurophysiology*, **80**, 324-330.
- KOUTSTAAL, W., WAGNER, A. D., ROTTE, M., MARIL, A., BUCKNER, R. L., & SCHACTER, D. L. (2001). Perceptual specificity in visual object priming: Functional magnetic resonance imaging evidence for a laterality difference in fusiform cortex. *Neuropsychologia*, **39**, 184-199.
- LOGAN, G. D. (1990). Repetition priming and automaticity: Common underlying mechanisms? *Cognitive Psychology*, **22**, 1-35.
- MARSOLEK, C. J. (1999). Dissociable neural subsystems underlie abstract and specific object recognition. *Psychological Science*, **10**, 111-118.
- MARSOLEK, C. J. (2003). What is priming and why? In J. S. Bowers & C. J. Marsolek (Eds.), *Rethinking implicit memory* (pp. 41-64). Oxford: Oxford University Press.
- MARSOLEK, C. J., & BURGUND, E. D. (1997). Computational analyses and hemispheric asymmetries in visual-form recognition. In S. Christman (Ed.), *Cerebral asymmetries in sensory and perceptual processing* (pp. 125-158). Amsterdam: Elsevier.
- MASSON, M. E. J. (1995). A distributed memory model of semantic priming. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 3-23.
- MCCLELLAND, J. L., MCNAUGHTON, B. L., & O'REILLY, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, **102**, 419-457.
- MCCLELLAND, J. L., & ROGERS, T. T. (2003). The parallel distributed processing approach to semantic cognition. *Nature Reviews Neuroscience*, **4**, 310-322.
- MCCLELLAND, J. L., & RUMELHART, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General*, **114**, 159-188.
- MCCLOSKEY, M., & COHEN, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 24, pp. 109-165). New York: Academic Press.
- MEDIN, D. L., & SCHAFFER, M. M. (1978). Context theory of classification learning. *Psychological Review*, **85**, 207-238.
- MITCHELL, D. B., & BROWN, A. S. (1988). Persistent repetition priming in picture naming and its dissociation from recognition memory. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **14**, 213-222.
- MORTON, J. (1979). Facilitation in word recognition: Experiments causing change in the logogen models. In P. A. Kolers, M. E. Wrolstad, & H. Bouma (Eds.), *Processing of visible language* (Vol. 1, pp. 259-268). New York: Plenum.
- MURDOCK, B. B., JR. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review*, **89**, 609-626.
- POGGIO, T., & EDELMAN, S. (1990). A network that learns to recognize three-dimensional objects. *Nature*, **343**, 263-266.
- POSTLE, B. R., & CORKIN, S. (1998). Impaired word-stem completion priming but intact perceptual identification priming with novel words: Evidence from the amnesic patient H.M. *Neuropsychologia*, **36**, 421-440.
- RATCLIFF, R., & MCKOON, G. (1995). Bias in the priming of object decisions. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 754-767.
- RATCLIFF, R., & MCKOON, G. (1996). Bias effects in implicit memory tasks. *Journal of Experimental Psychology: General*, **125**, 403-421.
- RATCLIFF, R., & MCKOON, G. (1997). A counter model for implicit priming in perceptual word identification. *Psychological Review*, **104**, 319-343.
- ROEDIGER, H. L., III (2003). Reconsidering implicit memory. In J. S. Bowers & C. J. Marsolek (Eds.), *Rethinking implicit memory* (pp. 3-18). Oxford: Oxford University Press.
- ROUDER, J. N., RATCLIFF, R., & MCKOON, G. (2000). A neural network model of implicit memory for object recognition. *Psychological Science*, **11**, 13-19.
- RUECKL, J. G. (2003). A connectionist perspective on repetition priming. In J. S. Bowers & C. J. Marsolek (Eds.), *Rethinking implicit memory* (pp. 67-104). Oxford: Oxford University Press.
- RUMELHART, D. E., & NORMAN, D. A. (1984). Representation in memory. In R. C. Atkinson, R. J. Herrnstein, P. H. Lindsay, & R. D. Luce (Eds.), *Handbook of experimental psychology* (pp. 511-587). New York: Wiley.
- SCHACTER, D. L. (1987). Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **13**, 501-518.
- SCHACTER, D. L. (1994). Priming and multiple memory systems: Perceptual mechanisms of implicit memory. In D. L. Schacter & E. Tulving

- ing (Eds.), *Memory systems 1994* (pp. 233-268). Cambridge, MA: MIT Press.
- SCHNYER, D. M., DOBBINS, I. G., NICHOLLS, L., SCHACTER, D. L., & VERFAELLIE, M. (2006). Rapid response learning in amnesia: Delineating associative learning components in repetition priming. *Neuropsychologia*, **44**, 140-149.
- SEIDENBERG, M. S., & MCCLELLAND, J. L. (1989). A distributed, developmental model of visual word recognition. *Psychological Review*, **96**, 523-568.
- SHIFFRIN, R. M., & STEYVERS, M. (1997). A model for recognition memory: REM—retrieving effectively from memory. *Psychonomic Bulletin & Review*, **4**, 145-166.
- SIGALA, N., & LOGOTHETIS, N. K. (2002). Visual categorization shapes feature selectivity in the primate temporal cortex. *Nature*, **415**, 318-320.
- SNODGRASS, J. G., & FEENAN, K. (1990). Priming effects in picture fragment completion: Support for the perceptual closure hypothesis. *Journal of Experimental Psychology: General*, **119**, 276-296.
- SQUIRE, L. R., KNOWLTON, B., & MUSEN, G. (1993). The structure and organization of memory. *Annual Review of Psychology*, **44**, 453-495.
- SRINIVAS, K. (1993). Perceptual specificity in nonverbal priming. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **19**, 582-602.
- STARK, C. E. L., & MCCLELLAND, J. L. (2000). Repetition priming of words, pseudowords, and nonwords. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **26**, 945-972.
- TANAKA, K. (1993). Neuronal mechanisms of object recognition. *Science*, **262**, 685-688.
- WARRINGTON, E. K., & WEISKRANTZ, L. (1974). The effect of prior learning on subsequent retention in amnesic patients. *Neuropsychologia*, **12**, 419-428.
- WECHSLER, D. (1997). *Wechsler adult intelligence scale* (3rd ed.). San Antonio, TX: Harcourt.
- WELDON, M. S., & ROEDIGER, H. L., III (1987). Altering retrieval demands reverses the picture superiority effect. *Memory & Cognition*, **15**, 269-280.
- WIGGS, C. L., & MARTIN, A. (1998). Properties and mechanisms of perceptual priming. *Current Opinion in Neurobiology*, **8**, 227-233.

(Manuscript received November 25, 2005;  
revision accepted for publication April 1, 2006.)