CORRECTED APRIL 28, 2010; SEE LAST PAGE

Journal of Experimental Psychology:
Learning, Memory, and Cognition
2009, Vol. 35, No. 4, 837–848

0278-7393/09/$12.00 DOI: 10.1037/a0015143

# Visual Priming of Inverted and Rotated Objects

Barbara J. Knowlton and Sean P. McAuliffe
University of California, Los Angeles

Chase J. Coelho
Pennsylvania State University

John E. Hummel
University of Illinois, Urbana-Champaign

Object images are identified more efficiently after prior exposure. Here, the authors investigated shape representations supporting object priming. The dependent measure in all experiments was the minimum exposure duration required to correctly identify an object image in a rapid serial visual presentation stream. Priming was defined as the change in minimum exposure duration for identification as a function of prior exposure to an object. Experiment 1 demonstrated that this dependent measure yielded an estimate of predominantly visual priming (i.e., free of name and concept priming). Experiments 2 and 3 demonstrated that although priming was sensitive to orientation, visual priming was relatively invariant with image inversion (i.e., an image visually primed its inverted counterpart approximately as much as it primed itself). Experiment 4 demonstrated a similar dissociation with images rotated 90° off the upright. In all experiments, the difference in the magnitude of priming for identical or rotated–inverted priming conditions was marginal or nonexistent. These results suggest that visual representations that support priming can be relatively insensitive to picture-plane manipulations, although these manipulations have a substantial effect on object identification.

*Keywords:* priming, object recognition, implicit memory, rapid serial visual presentation, viewpoint-independent

Previous exposure to a stimulus results in robust facilitation in processing the stimulus relative to novel stimuli. This repetition priming effect is thought to occur through a modification of the stimulus representation, allowing it to be used more efficiently on subsequent trials (See Ochsner, Chiu, & Schacter, 1994, for a review). The nature of the representations supporting priming has been an important question in implicit memory research. The basic approach to investigating these representations is to manipulate the relationship between the prime and the probe stimuli and measure the resulting priming. To the extent that priming is sensitive to a given manipulation, it suggests that the representation supporting this priming contains information about the dimension being manipulated. Studies of visual object priming have shown that priming is insensitive to the retinal size, the left–right reflection, and the shading of objects, suggesting that the object representations supporting priming do not contain information about these transformations (Biederman & Cooper, 1991a, 1991b, 1992; Biederman & Gerhardstein, 1993, 1995; Stankiewicz & Hummel, 2002; Stankiewicz, Hummel, & Cooper, 1998). However, none of these transformations typically affects object identification, so it is not surprising that they also do not affect visual priming.

The question remains as to whether transformations that do affect object identification also affect visual priming. For example, object identification has been shown to be very sensitive to rotation in the picture plane, as assessed by naming latency and error rates (e.g., Jolicoeur, 1985; Tarr & Pinker, 1989). However it is unclear whether these effects are due to the sensitivity of the shape representation to orientation or whether orientation influences other aspects of the process of naming familiar objects (e.g., as described by Stankiewicz, 2002). Thus, it may be that explicit measures of object identification are dissociable from implicit measures of object shape representations measured by priming. In the domain of short-term priming, in which the prime appears less than 1 s before the probe, the amount of priming obtained has been shown to decrease sharply if the prime and the probe are presented in different orientations (Arguin & Leek 2003). In this study, priming was measured in reference to a baseline in which unrelated objects were presented immediately before probe objects. However, it is unclear whether longer term priming is similarly viewpoint dependent.

The present study was designed to investigate visual priming of objects. We used minimum exposure duration in a rapid serial visual presentation (RSVP) stream as a measure of explicit identification performance. Participants viewed images of objects embedded in RSVP streams of nonobject images, and the participants' task was to name the object in each stream (see, e.g., Figure 1). If the participant failed to identify the object in a given stream, they were later tested with another stream containing the same image, but with a longer stimulus onset asynchrony (SOA) between images in the stream and, therefore, a longer exposure duration. The dependent measure was the duration (in ms) of the

---

Barbara J. Knowlton and Sean P. McAuliffe, Department of Psychology, University of California, Los Angeles; Chase J. Coelho, Department of Psychology, Pennsylvania State University; John E. Hummel, Department of Psychology, University of Illinois, Urbana-Champaign.

Correspondence concerning this article should be addressed to Barbara J. Knowlton, Department of Psychology, Franz Hall, University of California, Los Angeles, CA 90095-1563.

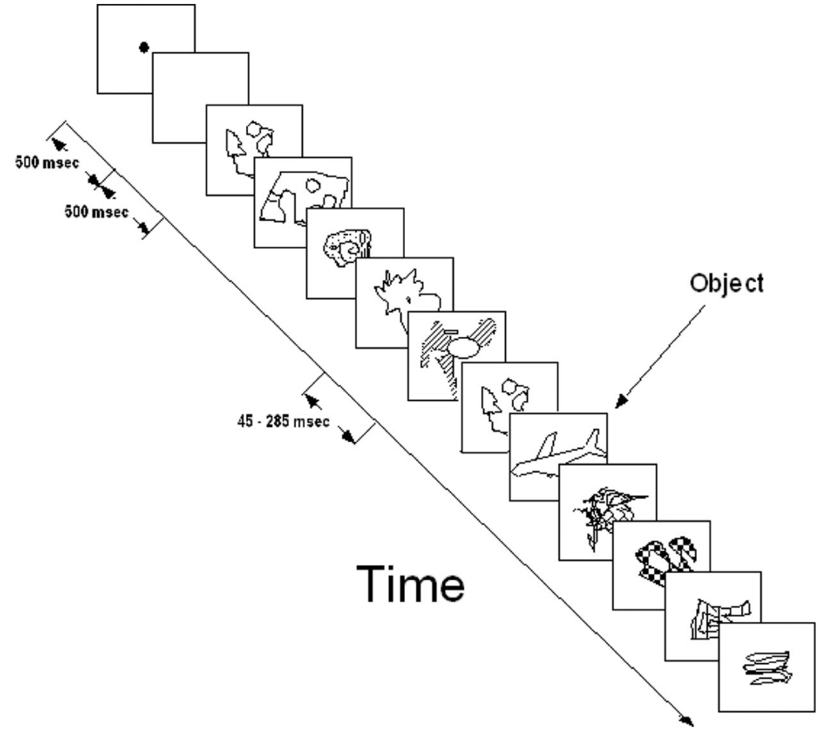

*Figure 1.* Illustration of the sequence of images during a probe trial.

shortest exposure at which the participant successfully identified the object. As elaborated below, and as supported by the findings in Experiment 1, this task was used because it provides a measure of perceptual processing speed devoid of factors such as the time required to access the object's name.

We operationalized priming as the change (relative to a baseline) in the minimum exposure duration required to recognize an object as a function of prior exposure to an image of that object (during a priming phase of the experiment). For example, if participants required, on average, $n$ ms of exposure to identify an unprimed object and $m$ ms to recognize the same object following a prime, then the priming was calculated as $n - m$.

To the extent that the view-sensitivities of human object identification reflect the view-sensitivities of the underlying perceptual representations, these two measures—identification and perceptual priming—should converge, showing comparable sensitivities to variations in viewpoint. For example, if both object identification and visual representation of shape are sensitive to orientation in the picture plane, then longer exposure durations would be required to identify objects presented upside down than to identify objects presented right-side up, and right-side up images ought to prime themselves more than they prime their upside-down counterparts (likewise, upside-down images ought to prime themselves more than they prime their right-side up counterparts).

But to the extent that the representation of shape is relatively invariant with picture-plane orientation and that the process of object identification takes orientation into account, perceptual priming may show greater invariance to viewpoint than does identification (e.g., with right-side up images priming upside-down images and vice versa, but with identification of primed upside-down images still requiring longer exposure durations than does identification of primed right-side-up images). Thus, the goal of the present study was to investigate the nature of the object shape representations supporting visual priming. By manipulating the prime and the probe orientation in the picture plane, we were able to test whether visual priming is invariant across transformations that have a large effect on object identification. The results will be informative about the nature of mental representations of object shape and the extent to which they are sensitive to previously experienced viewpoints.

Operationalized as a change (usually a decrease) in naming response time (RT) and errors as a function of repeated exposures to a stimulus, identification priming is subject to at least two sources of influence (see Biederman & Cooper, 1991a): (a) RTs and errors may decrease due to priming in the visual representation of object shape (visual priming) and (b) RTs and errors may decrease due to priming of nonvisual representations, such as the concept or the name of the object (nonvisual priming). Rather than using naming RT as the dependent variable, which reflects both visual and nonvisual aspects of processing, we used minimum exposure duration necessary for identification compared with a baseline condition in which participants had only read the object's name. The assumption behind this dependent measure is that perceptual priming should enable the participant to identify the object with a shorter exposure, relative to an unprimed object. However, any priming of the object's name or concept should be absent from this estimate of priming.

Following Biederman and Cooper (1991a, 1991b), a common way to separate visual and nonvisual sources of object priming is to use a same name different exemplar (SNDE) control condition, in which priming for a new view of an object (e.g., priming from one view of a jumbo jet to another view of the same jet) is compared with priming for the same view of a different object with the same basic-level name (e.g., from a jumbo jet [named *airplane*] to a Cessna [named *airplane*]). The logic behind this control is that SNDE objects (e.g., the Cessna) share the same name and most of the same conceptual information as the prime (the jumbo jet) but differ visually, due to their differing shapes. Priming for the SNDE provides an estimate of the magnitude of nonvisual sources of priming that by subtraction from the magnitude of priming from an image to itself, yields an estimate of purely visual priming (see Biederman & Cooper, 1991a). In Experiment 1, we used an SNDE control condition to confirm that our dependent measure does indeed yield an estimate of purely perceptual priming.

## Experiment 1

In Experiment 1, we used the SNDE control (e.g., Biederman & Cooper, 1991a, 1991b, 1992; Biederman & Gerhardstein, 1993; Stankiewicz et al., 1998) to estimate the magnitude of nonvisual priming, to assess whether our paradigm could provide a relatively pure measure of visual priming. In the SNDE control condition, a probe object is primed with a different exemplar of the same basic-level category. For example, if the probe depicts a Cessna (basic-level name *airplane*) the corresponding SNDE prime might be an image of a jumbo jet (basic-level name *airplane*). The idea behind this control is that the SNDE prime shares the name and much of the semantic information with the probe but differs from it in shape (Biederman & Cooper, 1991a; Jolicoeur, Gluck, & Kosslyn, 1984). As such, the SNDE prime is used to estimate semantic and name priming, free from visual priming. The difference between (a) the priming from one view of the Cessna to a different view of the Cessna and (b) the priming from a view of the jumbo jet to the same view of the Cessna, is taken as an estimate of the magnitude of visual priming (e.g., if Cessna View 1 primes the Cessna View 2 by, say, 250 ms and the jumbo jet View 2 primes the Cessna View 2 by 150 ms, then the remaining 100 ms is taken to be visual priming). In the context of the current paradigm—and more specifically, our suggestion that our dependent measure provides a measure of visual priming—the prediction is that our paradigm should result in zero priming in the SNDE condition compared with the baseline of reading the object name. It also may be the case that SNDEs differ in terms of their semantic representations. Thus, identical images may lead to greater semantic priming than do SNDEs. However, because the conceptual overlap in SNDEs is great, a lack of priming in the SNDE condition would point to a negligible contribution of semantic priming to the measure used in the current paradigm.

To maximize the likelihood of detecting any nonvisual priming in the SNDE condition, we inverted all probe images. Inverted images take longer to identify than do upright images and therefore afford a greater opportunity to observe priming.

Although different exemplars of the same basic-level class tend to have different shapes, they nonetheless tend to be perceptually similar (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). For example, a Cessna and a jumbo jet both have wings and a tail attached in similar places to a roughly cylindrical fuselage. To the extent that two SNDE exemplars have similar shapes, they could, at least in principle, prime one another visually. This visual priming would be misattributed to nonvisual sources (due to the logic of the SNDE control), leading the experimenter to underestimate the magnitude of visual priming in an experiment. To minimize this effect, we had a group of ten independent raters rank our SNDE pairs in terms of their visual similarity (i.e., the raters were naïve to the purpose of the experiment and did not participate in the experiment itself). The experiment included the 24 pairs that the raters ranked as most visually dissimilar.

### *Method*

*Participants.* Twelve undergraduates at University of California, Los Angeles (UCLA) participated for course credit. All had normal or corrected-to-normal vision. Ten additional undergraduates served as similarity raters for the SNDE pairs.

*Design.* This experiment had a two-factor design with prime–probe relationship (inversion, SNDE, and baseline) and object group as within-subjects variables. Because all probe items were presented upside down, in the inversion condition primes were presented in the canonical orientation (thus, inverted relative to probe images). Object group was a counterbalancing variable used to reduce the effect of variance in identification times for the different objects. Significance level was set at $p < .05$.

*Materials.* Stimuli were 24 pairs of black line drawings on a white background taken from Biederman and Cooper (1991a). Stimuli were presented on a Macintosh Performa 5200 color monitor (or, in the baseline condition, presented as names written in black letters on a white background). The nonobject images depicted collections of bounded surfaces, some with texture. Like most of the object images, the nonobject images were asymmetrical about the vertical axis. The entire experimental session was controlled by a program written in Macprobe (Aristometrics, Castro Valley, California). Participants viewed the display binocularly from a distance of approximately 70 cm, and object images subtended approximately 2.6° of visual angle.

The 24 SNDE pairs were rated as visually dissimilar by the 10 independent raters. The raters were naïve to the purpose of the experiment. We gave the raters 48 SNDE pairs and asked them to rank the similarity of each pair from 1 to 48 (i.e., with the most similar pair ranked as 1 and the second most similar pair as 2, etc.). The raters were specifically instructed to judge the images' visual similarity. The 24 pairs ranked most dissimilar (based on median similarity ratings) were used in the experiment (Appendix Figure A1). These 24 pairs were divided into three object groups of four pairs each. Across participants, the object groups were balanced across conditions (baseline, same exemplar prime, and SNDE exemplar prime).

*Procedure.* To reduce the variability in the names participants gave for the objects, all participants first read aloud all the object names from a piece of paper. The experimenter then read instructions to the participant. The experiment consisted of two phases. In the first (prime) phase, participants were instructed to name the object images (or to read object names, in the case of the baseline primes) that appeared sequentially on the screen. Prime displays began with a filled circle, which remained in the center of the screen for 500 ms. After a 500 ms blank screen, an image (in the priming conditions) or object name (in the baseline condition) appeared 4° to the left or right of fixation for 180 ms. In the prime phase, participants viewed upright images or SNDEs or read the names of objects. Because only 24 images would eventually appear in the probe phase, 36 other objects (12 in each condition) were presented in the prime phase. This minimized the influence of guessing in the probe phase by expanding the set of primed objects. Probe images were presented at fixation and prime images were presented off fixation, to avoid any possible retinotopic priming (McAuliffe & Knowlton, 2000).

During the probe phase, each trial presented an RSVP stream containing 1 object image among 17 nonobject images. The stream was presented at fixation, and the participant's task was to identify the object in the stream. The experimenter entered the first three letters of the participant's response into the computer, and the computer determined whether the response was correct or incorrect by comparing these letters with the first three letters of the correct response.

Figure 1 illustrates the sequence of images on a single probe trial. Each trial began with a fixation oval for 1,000 ms, followed by a blank screen for 500 ms. Seventeen distracter images were randomly chosen from a set of 24. The probe image appeared in a randomly chosen serial position between 7 and 13 (inclusive). Images subtended approximately 2.6° of visual angle. In the first block of the probe phase, all images were displayed for 45 ms. Any object the participant could not identify in the first block was presented again in the second block. Display time in the second block was 75 ms. This procedure was repeated, adding 30 ms to the display time on successive blocks, until all objects had been identified or until the display time reached 285 ms. Objects that could not be identified on or before the 285 ms block were treated as errors. In the probe phase, all images were inverted. Half the participants viewed one set of SNDEs in the probe phase (with the other half serving as primes in the SNDE condition). The assignment of images to prime and probe phases was reversed for the other half of participants. Thus, across all participants, each SNDE was equally likely to appear as a prime or probe. The experiment took about 10 min to complete.

## Results

Figure 2 shows the mean identification SOA in each condition. A Prime–Probe Relationship (inversion vs. SNDE vs. baseline) × Object Group (3 levels) repeated measures analysis of variance (ANOVA) revealed a significant effect of prime–probe relationship, $F(2, 99) = 12.10$, $MSE = 145.6$, $p < .001$, and a significant interaction, $F(4, 99) = 8.10$, $MSE = 145.6$, $p < .001$, but no main effect of object group ($F < 1$). The interaction between prime–probe relationship and object group was due to the fact that baseline identification SOA was higher for one of the object groups. Post hoc Bonferroni analysis revealed that the mean identification SOAs were significantly shorter in the inversion condition than in each of the other two conditions ($p < .01$ for SNDE; $p < .01$ for baseline), but mean identification SOA in these two conditions did not differ significantly from one other ($p > .6$). A power analysis revealed that we had an 80% chance of detecting a moderate effect ($d > 0.7$, approximately 10 ms). The error rates were 10.4% ± 4.4% for SNDE primes, 8.3% ± 3.5% for the baseline condition, and 4.2% ± 4.2% for identical primes.

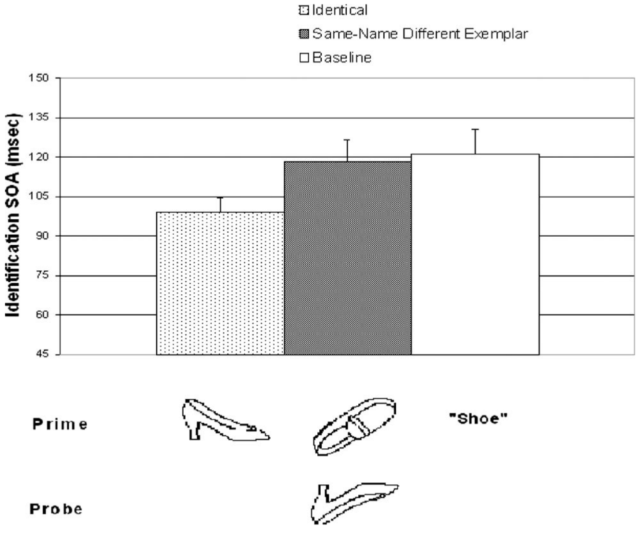

*Figure 2.* Mean identification stimulus onset asynchrony (SOA) in each condition of Experiment 1. Error bars indicate standard error of the mean.

## Discussion

The mean identification SOA for probes primed with visually dissimilar SNDEs was equivalent to the mean SOA for probes primed by word naming (i.e., baseline). Probes primed by their inversions were recognized at shorter SOAs than were probes primed either with SNDEs or with words. These results support the hypothesis that the priming observed with our dependent measure is primarily visual. That probes benefited equally from previous word naming and from exposure to an SNDE suggests that any semantic priming produced by an SNDE (beyond word naming) does not affect probe identification in this paradigm. This finding also underscores the advantage of this priming paradigm over the object naming task, in which semantic and linguistic processing as well as visual representations contribute to naming latency. In the following three experiments, we use this procedure to assess the effect of inversion and rotation on perceptual priming of object images.

# Experiment 2

In Experiment 2, we investigated the effects of inversion (reflection about the horizontal axis of the image plane) and left–right reflection (reflection about the vertical axis) on priming of upright and inverted objects. We manipulated the nature of the prime display (either an object name, in the baseline condition, or an object image, in the various priming conditions), the orientation of both the prime and probe images, and the relationship between them, which resulted in four prime–probe relationships for upright and for inverted objects: identical (i.e., with identical images presented in the prime and probe), inversion (i.e., forming the probe by reflecting the prime about the horizontal axis), left–right reflection (i.e., forming the probe by reflecting the prime about the vertical axis), and baseline (i.e., with an object name, rather than an image, presented as the prime, so that the probe depicted an object not seen during a priming trial). Minimum exposure durations for identification in this condition served as a baseline relative to which priming in the other conditions was measured. It is important to note that inverting an image is not the same thing as rotating it 180° about the line of sight; rather, an inversion is equivalent to a 180° rotation of a left–right reflection of the original image.

In the priming phase, object images (or names) were presented one at a time (i.e., not in an RSVP stream) off-fixation, and the participant's task was to name each one. In the probe phase, images were presented at fixation in an RSVP stream of nonobject images, and the participant's task was to name the object appearing in the stream. During the probe phase, initially brief SOAs were gradually increased on successive blocks until each object could be identified. The dependent variable was the shortest SOA at which an object image could be identified.

### Method

*Participants.* Sixteen undergraduates at University of California, Los Angeles participated for course credit. All had normal or corrected-to-normal vision.

*Design.* The experiment had a two factor design with probe orientation (upright vs. inverted) and prime–probe relationship (identical, inversion, left–right reflection, and baseline) as within-subjects variables, yielding eight conditions. As in Experiment 1, in the baseline condition the name of the object was presented during the priming phase. Significance level was set at $p < .05$ for all comparisons.

*Materials.* Stimuli included 48 black line drawings on a white background (Snodgrass & Vanderwart, 1980) presented as described in Experiment 1. Among these items, 24 were one image of the SNDE pairs that were presented in Experiment 1. The nonobject images used in the RSVP stream were identical to those used in Experiment 1.

*Procedure.* The procedure was similar to that used in Experiment 1. In all, there were 48 probe images plus six practice trials (given at the beginning of the probe phase) depicting objects not appearing in either the prime trials or the probe trials. Across participants, each object was equally likely to appear in each of the eight conditions. With 48 objects and eight conditions, there were 6 objects in each cell per participant. The experiment took about 15 min to complete.

### Results

Figure 3 shows the mean identification SOA in each condition. Priming in a given condition (identical, left–right reflection, or inversion) is operationalized as the difference between the identification SOA for probes in the baseline condition minus the identification SOA for probes in the priming condition. A repeated measures ANOVA (Probe Orientation [upright vs. inverted] × Prime–Probe Relationship [identical vs. left–right reflected vs. inverted vs. baseline]) performed on identification SOAs revealed significant main effects of probe orientation, $F(1, 15) = 62.55$, $MSE = 408.7$, Cohen's $d = 1.07$ and prime-probe relationship, $F(3, 45) = 16.37$, $MSE = 285.9$, but no significant interaction, $F(3, 45)=2.73$, $MSE = 298.6$. A post hoc Bonferroni analysis revealed that identification SOAs in the baseline condition were significantly longer than in each of the other conditions ($p$s $< .05$ for all conditions, Cohen's $d$ ranged from 0.98–1.32), but these three conditions did not differ significantly from one another ($p >$ .2 for all pairwise comparisons). For each of these comparisons for both probe orientations, power analyses revealed that we had an 80% chance of detecting a moderate effect ($d > 0.6$, approximately 9 ms).

In general, upright probes were identified at shorter SOAs than were inverted probes, but the advantage in priming enjoyed by identical probes (upright to upright or upside down to upside down) was not reliably greater than was the priming from inverted probes (upright to upside down or upside down to upright). Prim-

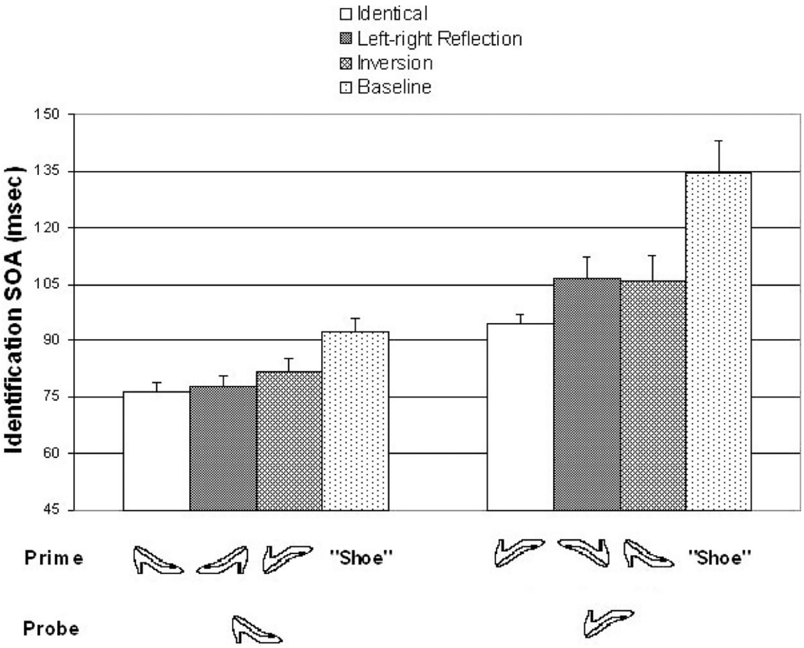

*Figure 3.* Mean identification stimulus onset asynchrony (SOA) in each condition of Experiment 2. Error bars indicate standard error of the mean.

ing was also insensitive to left–right reflection, in that the priming from an image to itself did not differ from the priming from an image to its left–right reflection. Only two conditions produced any errors. For inverted probes, the error rates were 3.1% ± 1.6% for baseline primes and 2.1% ± 1.4% for inverted primes.

We also conducted an item analysis to examine the amount of priming in each condition for the 48 objects across participants. The amount of priming for each object in each condition was determined as the difference between the mean primed versus the baseline identification SOA. One of the objects (the cow) was not correctly named by a large number of participants when it appeared in the inverted probe baseline condition; thus, it was not included in the item analysis. Paired *t* tests revealed no difference among any of the conditions, with either upright or inverted probes, $t$s(46) $< 0.68$, $p$s $> .5$. Thus, the item analysis supports the idea that similar levels of priming are obtained with primes that are upright and those that are inverted relative to the probe item.

### Discussion

Experiment 2 revealed a main effect of probe orientation on the time to identify object images, in that the mean SOA required to identify inverted probes (the four bars on the right of Figure 3) was longer than the mean identification SOA for upright probes (the four bars on the left of Figure 3). This result is not surprising and fits with an extensive literature showing that recognition of upside-down images takes longer than recognition of right-side up images (e.g., Jolicoeur, 1985, 1990; Tarr & Pinker, 1989, 1990).

More interesting is the fact that priming—change in identification SOA as a function of exposure to an object image during the priming phase of the experiment—did not show the same degree of sensitivity to inversion. Neither left–right reflection nor inversion (from an upright prime to an inverted probe or vice-versa) had a statistically reliable effect on the magnitude of visual priming for either upright or inverted probe images: As measured by the minimum exposure duration required for probe recognition, image identification profited as much from an inverted (relative to the probe) prime or a left–right reflected prime as it did from a prime that was identical to the probe. That is, the observed visual priming was invariant with left–right reflection and largely invariant with inversion (there were small but nonreliable effects of prime–probe inversion and reflection on the identification of upside-down probes). This relative invariance in visual priming, coupled with sensitivity to inversion as measured by identification SOA, suggests that sensitivity to orientation in recognition performance may not necessarily reflect sensitivity to orientation in the visual representations mediating identification (see also Stankiewicz, 2002).

The finding that visual priming is invariant with left–right reflection is not new (see, e.g., Biederman & Cooper, 1991b; Stankiewicz et al., 1998) and is straightforward to interpret in terms of current structural description theories of object recognition (e.g., Hummel, 1994, 2001; Hummel & Biederman, 1992). These theories represent both the *left-of* and *right-of* relations simply, as *beside*, with the result that any object image has exactly the same structural description as its left–right reflection. These models therefore predict that an image will prime its left–right reflection as much as it primes itself. By contrast, the finding that visual priming is predominantly invariant with left–right reflection is inconsistent with view-based models of object identification, including both those that postulate alignment or mental rotation to correct for rotation in the picture plane (e.g., Jolicoeur, 1985, 1990; Tarr and Pinker, 1989, 1990) and those that do not (e.g., Edelman, 1998; Edelman & Intrator, 2003; Poggio & Edelman, 1990). The reason is that all these models match images to views in memory according to the features they have in common (when *in common* means that a given feature is in the same or similar location in the image and the view in memory; see Hummel, 2000). A bilaterally asymmetric object image, such as those used in the experiments reported here, will tend to have few if any features in common with its left–right reflection.

Although the invariance of visual priming with left–right reflection is more consistent with structure-based theories than view-based theories, the observed relative invariance of visual priming with image inversion is inconsistent with both classes of theories. Theories based on view-matching without rotation or alignment (e.g., Edelman, 1998; Edelman & Intrator, 2003; Edelman & Poggio, 1990) predict little or no priming from an inverted image to its upright counterpart (or vice versa) because, like a left–right reflection, two such images will tend to have very few features in common. View-based theories that postulate alignment and/or mental rotation (e.g., Jolicoeur, 1985, 1990; Tarr, 1995; Tarr & Pinker, 1989, 1990) also predict little or no priming from an inverted image to its upright counterpart or vice versa. Such models hold that we recognize upside-down images by transforming them 180° to the upright prior to matching to memory for recognition. However, recall that a 180° rotation of an inverted image (either prime or probe) is not identical to the original uninverted image but is a left–right reflection of it. Therefore, transforming an inverted image will result in an upright image that has few features in common with—and therefore is not expected to prime or be primed by—its uninverted counterpart because inverting an image will alter some of the relations between parts (if Part A was above Part B in the upright image, this relation is not present in the inverted object).

Modern structural description theories are also at a loss to explain the observed invariance of visual priming with image inversion. Such theories predict that an inverted image will activate many of the same parts and relations in the object's (presumably upright) structural description in long-term memory, thereby priming that representation (see, e.g., Hummel, 1994; Hummel & Biederman, 1992). For the same reason, an upright image is predicted to prime, at least to a modest extent, its inverted counterpart. However, these theories incorrectly predict that an upright image will visually prime itself more than it will prime its inverted counterpart and that an inverted image will prime itself more than it will prime its upright counterpart.

## Experiment 3

The results of Experiment 2 suggest that visual priming is largely invariant with inversion (i.e., reflection about the horizontal axis). However, it is important to determine whether visual priming is invariant with picture plane orientation in general or whether vertical inversion is somehow a special case (e.g., Jolicoeur, 1985). Experiment 3 extended the results of Experiment 2 with the same prime–probe relationships (i.e., identity, reflection, and inversion) but changed the orientations of the prime and probe images. Instead of upside-down images, Experiment 3 used images rotated 90° off the upright. This method allows us to observe whether presenting an object in one noncanonical orientation (e.g., 90° clockwise) can prime the identification of that object in a different noncanonical orientation (90° counterclockwise).

### Method

*Participants.* Sixteen undergraduates at UCLA participated to fulfill a course requirement. All had normal or corrected-to-normal vision.

*Design.* This experiment used a two factor design with probe orientation (upright vs. 90° rotation) and prime–probe relationship (identical, inversion, left–right reflection, and baseline) as within-subjects variables. Significance level was set at $p < .05$.

*Materials.* The images were the same as those used in Experiment 2, except that half the images presented during the probe phase were rotated 90° from upright instead of inverted.

*Procedure.* The procedure was identical to that of Experiment 2, with the exception that the prime presentations were adjusted to maintain the correct prime–probe relationships (i.e., identical, left–right reflection, inversion, and baseline).

### Results

Figure 4 shows the mean identification SOA in each condition. Priming was operationalized in the same way as in Experiments 1 and 2. A Probe Orientation (upright vs. rotated) × Prime–Probe Relationship (identical vs. left–right reflection vs. inversion vs. baseline) repeated measures ANOVA revealed significant effects of probe orientation, $F(1, 15) = 108.39$, $MSE = 149.5$, Cohen's $d = 1.64$, and prime–probe relationship, $F(3, 45) = 42.25$, $MSE = 137.1$, but no interaction ($F < 1$). Post hoc Bonferroni analyses revealed that identification SOAs in the baseline condition were significantly longer than in each of the other three priming conditions ($p < .01$ for inversion, $p < .01$ for identical, and $p < .01$ for left–right reflection; Cohen's $d$ ranged from 1.67–1.96), but these three priming conditions did not differ significantly from each other ($p > .9$ for all pairwise comparisons). Power analyses performed on these pairwise comparisons showed that we had adequate power to detect a moderate effect of condition ($d > 0.6$; approximately 9 ms) with probability of 80% for both upright and rotated probes. Upright probes were identified at shorter SOAs than were rotated probes, but the advantage in priming enjoyed by identical probes (upright to upright, rotated 90° clockwise [or counterclockwise] to rotated 90° clockwise [or counterclockwise]) was not reliably greater than was the priming from inverted probes (upright to upside down [or vice versa] or rotated 90° clockwise to 90° counter-clockwise [or vice versa]). Only three conditions produced any errors. For upright probes, the error rates were 1% ± 1% for baseline primes. For 90° rotated probes, the error rates were 4.2% ± 1.8% for baseline primes and 2.1% ± 2.1% for inverted primes.

An item analysis by which we examined the amount of priming in each condition for each of the 48 objects showed no difference among any of the three priming conditions for either upright or inverted probes, $ts(47) < 1.57$, $ps > .1$.

### Discussion

Experiment 3 replicated Experiment 2 with probe images rotated 90° off upright in the place of the upside-down images used in Experiment 1. Participants required longer SOAs to identify images rotated 90° off upright than to recognize upright images (consistent with numerous previous findings demonstrating costs for recognizing objects rotated off the upright). However, visual priming was invariant with both left–right reflection and inversion. For example, a 90° counterclockwise rotation primed recognition of its left–right reflection (which would be a 90° clockwise rotation) as much as it primed itself, and it primed its inversion (which

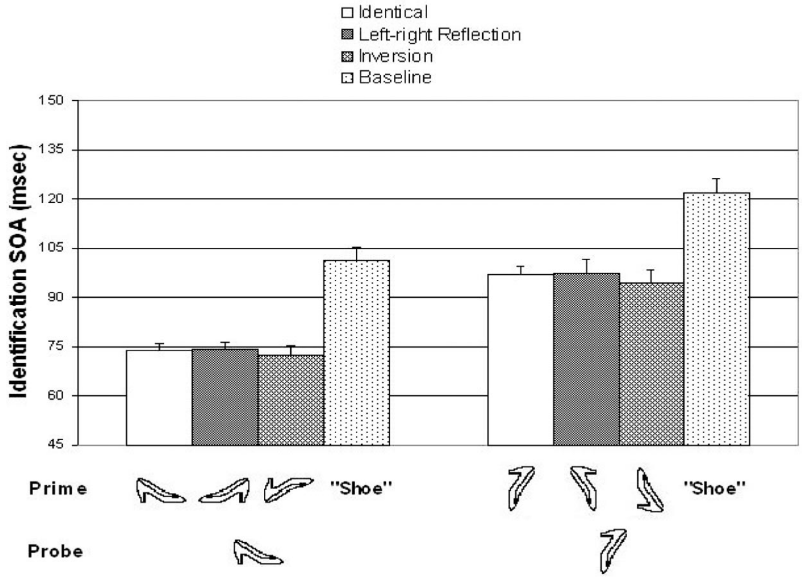

*Figure 4.* Mean identification stimulus onset asynchrony (SOA) in each condition of Experiment 3. Error bars indicate standard error of the mean.

would be an equivalent rotation of its left–right reflection) just as much as it primed itself. These results offer further support for the difference between visual priming and object recognition observed in Experiment 2.

## Experiment 4

The results of Experiments 2 and 3 suggest that visual priming may be nearly invariant with picture-plane orientation, but in these experiments, nonidentical prime–probe pairs were always related by a reflection either about the vertical axis (in the case of left–right reflection) or the horizontal axis (in the case of inversion). Experiment 4 was designed to extend these results with a different prime–probe relationship, namely a 90° rotation rather than a reflection about an axis. Prime and probe images were either upright or rotated 90° from upright. Experiment 4 used a total of three prime–probe relationships (i.e., identical, 90° rotations [clockwise or counterclockwise], and baseline).

The results of Experiments 2 and 3 suggest that object priming is relatively invariant with the picture-plane transformations tested here, even though these transformations had robust effects on identification. In all of these experiments, the prime images were presented off-center to avoid any possible retinotopic priming (McAuliffe & Knowlton, 2000). To test the possibility that larger orientation specific priming effects might be obtained when the retinal positions of the primes and probe matched, Experiment 4 included conditions in which primes and probes were presented in the same retinal location (foveally). In addition, more participants were included in the study to increase power.

### Method

*Participants.* Forty-eight undergraduates at UCLA participated for course credit. All had normal or corrected-to-normal vision.

*Design.* This experiment used a three factor design with probe orientation (upright vs. rotated) and prime–probe relationship (identical, rotation, and baseline) as within-subject variables, and location of the prime (center or off-center) as a between-subjects factor.

*Materials.* The images were the same as those used in Experiment 3.

*Procedure.* The procedure was identical to that of Experiment 3, with the following exceptions: During the prime phase, half the participants viewed images presented at fixation, and half viewed images presented off-center. For upright probe images, primes were upright (identical condition), rotated 90° clockwise or counterclockwise (rotation condition), or object names (baseline condition). For rotated probe images, primes were rotated (identical condition), upright (rotation condition), or object names (baseline condition). Across participants, each object was equally likely to appear in each of the six conditions. With 48 objects and six conditions, there were 8 objects in each cell per participant.

### Results

Figure 5 shows the mean identification SOA in each condition. There was no main effect of position of the prime on identification SOA of the probe ($F < 1$). As in Experiments 2 and 3, there was a main effect of probe orientation, with upright probe images identified at shorter SOAs, $F(1, 46) = 216.7$, $MSE = 221.2$, $p < .01$. There was also a main effect of priming condition, $F(2, 92) = 81.9$, $MSE = 155.2$, $p < .01$. The position of the prime interacted with the priming condition, $F(2, 92) = 3.41$, $MSE = 155.2$, $p < .05$. As can be seen in Figure 5A, when primes were presented at fixation, the magnitude of priming relative to baseline was greater than when probes were presented off-center (Figure 5B). Bonferroni corrected pairwise comparisons showed that all priming conditions differed from baseline (*p*s $> .01$, Cohen's *d*s $> 1.64$). However, for both upright and rotated probes, there was no significant difference between the priming conditions (*p*s $> .05$). No other interactions were statistically significant. A power analysis indicated that we had an 80% chance of detecting a moderate effect of condition (Cohen's $d = 0.4$, approximately 6 ms). As in the previous studies, error rates were relatively low. For participants receiving off-center primes, there were only errors in identifying probes that were rotated (7.2% ± 2.4% in the baseline condition and 1.0% ± 1.0% in the rotated prime condition). For participants receiving primes at fixation, the error rates for rotated probes were 4.7% ± 1.8% for the unprimed baseline, 3.1% ± 2.1% after rotated primes, and 0.7% ± 0.7% for identical primes. For upright probes, there were no errors after identical primes. For the rotated prime and baseline conditions, the error rates were 0.5% ± 0.5%.

An item analysis was conducted to assess the amount of priming in each condition for the 48 objects. When primes had been presented centrally, there was no significant difference in the amount of priming for identical and rotated primes, in either the upright condition or the rotated probe condition, $ts(47) < 0.44$, *p*s $> .6$. When probes were presented off-center, and when probes were presented in the upright orientation, there was no significant difference in the amount of priming in the upright or rotated probe conditions, $t(47) = 1.51$, $p > .1$. However, when probes were rotated, there was more priming when the prime was presented in the same rotated orientation as the probe item than when the prime was presented upright, $t(47) = 2.07$, $p < .05$, Cohen's $d = 0.2$.

### Discussion

Consistent with the results of Experiments 2 and 3, 90° rotation did not substantially affect the magnitude of visual priming but did affect identification SOAs for both primed and unprimed probe images. However, an item analysis revealed a small increase in priming under one condition (off-center primes, rotated probe images) when primes and probes were presented in the identical orientation. The item analysis had additional power that enabled us to detect this small effect. The small amount of orientation-sensitive priming observed in this condition is nevertheless much smaller than the amount of orientation-insensitive priming consistently observed in all conditions.

These results offer further support for the difference between priming and identification observed in the previous two experiments. The findings also demonstrate that the relative orientation invariance of object priming demonstrated in Experiments 2 and 3 generalizes beyond inversion and left–right reflection. These results are consistent with the ideas that visual priming of objects assessed in this paradigm is relatively view-invariant and that viewpoint specific priming is relatively small and variable in contrast to the magnitude of view-independent priming.

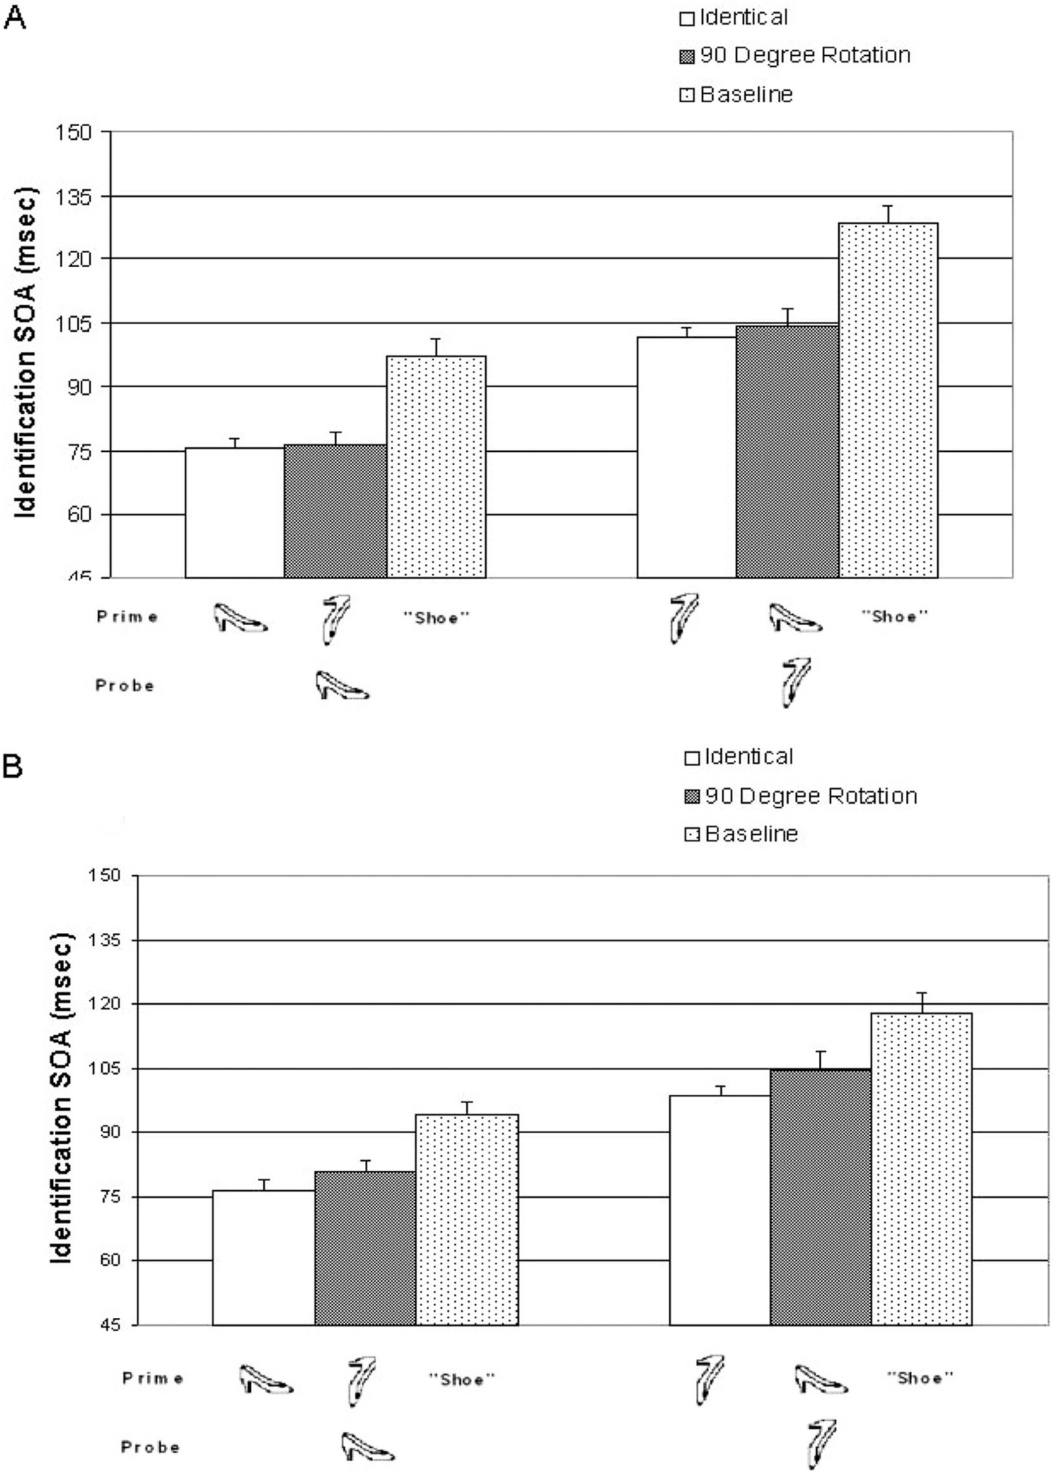

*Figure 5.* Mean identification stimulus onset asynchrony (SOA) in each condition of Experiment 4. Error bars indicate standard error of the mean. A: Primes presented at fixation. B: Primes presented off-center.

## General Discussion

Experiments 2–4 revealed a difference between the orientation sensitivity of human object identification operationalized in terms of identification SOA (the minimum exposure duration required to identify an object image in an RSVP stream) and the orientation sensitivity operationalized in terms of visual priming (change in identification SOA as a function of prior exposure to an image of the same object). Participants required longer SOAs to identify object images in unusual orientations in the picture-plane (upside down or rotated 90° from the upright) than to identify upright images: In this sense, object identification is sensitive to picture-place orientation, as has been reported elsewhere (e.g., Jolicoeur, 1985, 1990; Tarr & Pinker, 1989, 1990, among others). At the same time, however, visual priming was largely or completely invariant with rotation and inversion in the picture-plane: Images visually primed the inverted version or the rotated version of themselves nearly as much as they primed themselves. As elaborated below, this difference suggests that the view-sensitivity of object identification performance may not reflect any intrinsic view-specificity in the representation of shape itself. Instead, the representation of shape may be primarily view-invariant, but the process of object identification is view-sensitive because it relies on semantic information about viewpoint (as suggested by Stankiewicz, 2002).

In the experiments, we operationalized visual priming as a change in identification SOA, as a function of prior exposure to a prime stimulus (rather than defining it in terms of a change in naming RT, a more common measure of priming) to avoid any contributions to priming from repeatedly stating an object's name or accessing its concept. That is, our experimental paradigm was designed to provide as pure a measure of specifically visual priming as possible. Experiment 1 provided strong evidence that our dependent measure provided a pure estimate of perceptual priming. In contrast to widely used measures based on naming RT, the current paradigm revealed no priming whatsoever from an image of one object to an image of a different member of the same basic-level class. This result lends confidence to the conclusion that the observed view-invariant priming reflects priming of view-invariant visual representations rather than priming of an object's name or concept (which are trivially view-invariant).

The finding of view-invariant priming supports the idea that shape can be represented independently of viewpoint. Evidence from functional magnetic resonance imaging suggests that shape representations based on view-independent parts rather than on local features are supported by the lateral occipital complex (Hayworth & Biederman, 2006). In the present study, priming may be facilitating processing of these intermediate shape representations, which may be relatively invariant to specific viewpoint as they are based on relations between parts. Alternatively, given that the images used in the present study were all familiar objects, it is possible that shape representations do contain viewpoint information but that activation of one view through priming automatically fully activates all other possible viewpoints of that object, even rare, noncanonical views. However, such a conceptualization of view-based shape representations, in which all possible views are equally activated by presentation of an object image in any viewpoint, is difficult to distinguish from the concept of view-invariant representations.

The present results also do not eliminate the possibility that view-dependent representations of shape contribute to identification under some circumstances. In all experiments, there was a slight numerical advantage for priming from same view images (approximately 2–8 ms). In the item analysis in Experiment 4, significant viewpoint sensitive priming was obtained under one condition, and it is likely that we would find additional evidence of viewpoint specific priming in other experiments if we included a larger number of participants or if we used a measure of priming that was more sensitive. Nevertheless, it is important to keep in mind that any small decrease in priming due to rotating or inverting the probe stands in contrast to the large effects of viewpoint invariant priming and the effects of these manipulations on identification. The main conclusion from this study is not that all visual priming is viewpoint-insensitive but rather that the orientation of the prime has very little effect influence on the amount of long-term priming obtained. We interpret this finding to suggest that viewpoint-independent representation of object shape exists and makes a substantial contribution to object identification.

Support for the existence of both viewpoint-independent and viewpoint-dependent representations of shape was found by Burgund and Marsolek (2000). In this study, the authors found evidence of viewpoint dependent priming for object images rotated in depth when probe items were presented in the left visual field (i.e., to the right hemisphere), whereas view-invariant priming was obtained when items were presented in the right visual field (i.e., to the left hemisphere). This study used naming time to measure priming; thus, it is not clear whether the results reflect priming of semantic representation of objects in addition to shape. In the present series of experiments, probe images were presented centrally, so much of the priming observed was likely to have resulted from activation of shape representations that were most efficient in supporting performance, which may be based in the left-hemisphere and may be view-invariant. In addition, a small amount of reflection-sensitive priming has been reported with a procedure similar to the one used here when primes and probes are presented to the same retinal location, whereas priming is invariant to reflection when primes and probes are presented in different retinal locations (McAuliffe & Knowlton, 2000). Also, large orientation sensitivity effects have been shown in short-term object priming, in which the prime is presented within a second of the probe (Arguin & Leek, 2003). These results suggest that there are multiple representations of object shape and that these may vary to the extent that they are independent of viewpoint.

If long-term visual object priming can be largely supported by view-independent representations of object shape, then why is object identification so sensitive to viewpoint in many cases? Stankiewicz (2002) postulated that the visual system might use both shape and viewpoint as separate sources of information about object identity. The idea is that the representation of shape, per se, is relatively invariant with viewpoint but that the visual system takes information about viewpoint into account for the purposes of object recognition because it is diagnostic. Although viewpoint information is processed independently of shape, viewpoint information is clearly informative in making decisions about the identity of objects, in that many objects are encountered almost exclusively in an upright orientation. Information about the canonical orientation of an object would be part of the semantic representation of the object and could facilitate object identification in much the same the way that context is known to. For example, people know that horses stand upright on four legs, so information in the image suggesting an upright orientation (e.g., roughly parallel edges [the legs] below a roughly horizontal curved edge [the stomach]) would be consistent with the interpretation of the image as that of a horse. But one's ability to use such information in aid of identifying the image as a horse does not imply that it is integral to the representation of the horse's shape, per se. In the same way, being located on a track or in a barn could also increase one's confidence that the object one is viewing is a horse, even if that information is not an integral part of the representation of the horse's shape. From an engineering perspective, this approach would be an ideal way to design a visual system, provided that (a) both shape and viewpoint are predictive of object identity and (b) their contributions were relatively independent. On this account, the visual priming reported here may reflect changes in view-independent representations of object shape, whereas identification uses both these (view-independent) representations and information about viewpoint. Similarly, episodic representations of objects, as assessed by recognition memory, contain viewpoint information (Biederman & Cooper, 1991a). Thus, when identification is dependent on episodic memory of an item, such as identifying a recently presented novel item, it is likely that identification will be easier when the item is presented in the studied viewpoint (Lawson, 2004).

If object identification uses both view-independent representations and viewpoint information, then the problem for students of object recognition becomes one of understanding how the visual system comes to represent object shape independently of viewpoint, how it comes to represent viewpoint independently of shape, and how it combines the two in the service of object identification. Fortunately, the first part of this problem, arguably the hardest—representing shape independently of viewpoint—is not a new problem at all (see, e.g., Marr & Nishihara, 1978) but is simply an old problem revisited.

## References

Arguin, M., & Leek, E. C. (2003). Orientation invariance in visual object priming depends on prime-target asynchrony. *Perception and Psychophysics, 65,* 469–477.

Biederman, I., & Cooper, E. E. (1991a). Priming contour deleted images: Evidence for intermediate representations in visual object recognition. *Cognitive Psychology, 23,* 393–419.

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

Biederman, I., & Cooper, E. E. (1992). Size invariance in visual object priming. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 121–133.

Biederman, I., & Gerhardstein, P. C. (1993). Recognizing depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception and Performance, 19,* 1162–1182.

Biederman, I., & Gerhardstein, P. C. (1995). Viewpoint-dependent mechanisms in visual object recognition: A critical analysis. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 1506–1514.

Burgund, E. D., & Marsolek, C. J. (2000). Viewpoint-invariant and viewpoint-dependent object recognition in dissociable neural subsystems. *Psychonomic Bulletin and Review, 7,* 480–489.

Edelman, S. (1998). Representation is representation of similarities. *Behavioral & Brain Sciences, 21,* 449–498.

Ochsner, K. N., Chiu, C. Y., & Schacter, D. L. (1994). Varieties of priming. *Current Opinion in Neurobiology, 4,* 189–194.

Poggio, T., & Edelman, S. (1990, January 18) A network that learns to recognize three-dimensional objects. *Nature, 6255,* 263–266.

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8,* 382–439.

Snodgrass, J. G., & Vanderwart, M. (1980). A standardized set of 260 pictures: Norms for name agreement, familiarity, and visual complexity. *Journal of Experimental Psychology: Human Learning, 6,* 174–215.

Stankiewicz, B. J. (2002). Empirical evidence for independent dimensions in the visual representation of three-dimensional shape. *Journal of Experimental Psychology: Human Perception and Performance, 28,* 913–932.

Stankiewicz, B. J., & Hummel, J. E. (2002). The role of attention in scale- and translation-invariant object recognition. *Visual Cognition, 9,* 719–739.

Stankiewicz, B. J., Hummel, J. E., & Cooper, E. E. (1998). The role of attention in priming for left–right reflections of object images: Evidence for a dual representation of object shape. *Journal of Experimental Psychology: Human Perception and Performance, 24,* 732–744.

Tarr, M. J. (1995). Rotating objects to recognize them: A case study on the role of viewpoint dependency in the recognition of three-dimensional objects. *Psychonomic Bulletin & Review, 2,* 55–82.

Tarr, M. J., & Pinker. S. (1989). Mental rotation and orientation dependence in shape recognition. *Cognitive Psychology, 21,* 233–283.

Tarr, M. J., & Pinker. S. (1990). When does human object recognition use a viewer-centered reference frame? *Psychological Science, 1,* 253–256.

(*Appendix follows*)

# Appendix

## Image Pairs Used in Experiment 1

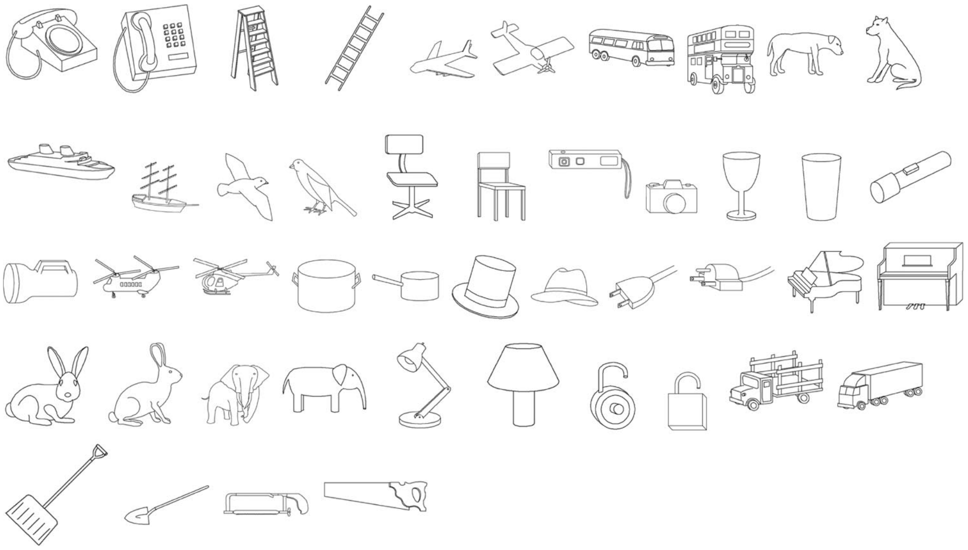

*Figure A1.* The 23 image pairs, in addition to the pair shown in Figure 2, used in Experiment 1. These were a subset of the image pairs used by Biederman and Cooper (1991a).

Received March 5, 2008
Revision received November 11, 2008
Accepted November 23, 2008 ■

## Correction to Knowlton et al. (2009)

In the article "Visual Priming of Inverted and Rotated Objects," by Barbara J. Knowlton, Sean P. McAuliffe, Chase J. Coelho, and John E. Hummel (*Journal of Experimental Psychology: Learning, Memory, and Cognition*, 2009, Vol. 35, No. 4, pp. 837–848), there was an error in the sixth sentence of the abstract. The sentence should read "Experiments 2 and 3 demonstrated that although identification was sensitive to orientation, visual priming was relatively invariant with image inversion (i.e., an image visually primed its inverted counterpart approximately as much as it primed itself)."