Psychonomic Bulletin & Review
2005, 12 (2), 321-327

When are viewpoint costs greater for silhouettes than for shaded images?

WILLIAM G. HAYWARD
Chinese University of Hong Kong, Hong Kong, China

ALAN C.-N. WONG
Vanderbilt University, Nashville, Tennessee

and

BRANKA SPEHAR
University of New South Wales, Sydney, New South Wales, Australia

Previous studies of object recognition have shown efficient recognition of silhouettes, suggesting that much of the information used to recognize objects resides in the outline. These studies, however, have used objects that contain many components, which provide redundant information. In this study, we examined recognition of silhouettes of less-complex objects, so that redundant information was reduced. We found that viewpoint generalization costs (the decrement of performance when recognizing nonstudied views) were greater for silhouettes than for shaded images, even when the same qualitative components were visible in the outline of both studied and nonstudied views. Thus, silhouettes do not always allow for view generalization as efficiently as do shaded images.

The information underlying human object recognition is still largely unknown. Whereas observers have little difficulty in judging an object to be the same over changes in pose, lighting, color, or even (in some circumstances) identity, each of these transformations results in profound alterations to the visual information available for such a decision. Much research has been directed at understanding the coordinate system describing object representations (e.g., Biederman & Gerhardstein, 1993; Tarr & Bülthoff, 1995), but relatively little has addressed the necessary and sufficient visual features for object recognition (though Biederman's [1987] recognition-by-components theory is one notable exception). The sheer variety of possible visual information that might be used by perceptual systems is one of the obvious difficulties in conducting such studies. One way to deal with this problem is to examine reduced sets of shape features and determine the extent to which they can account for object-recognition performance. In particular, a number of researchers have looked at the role of the outline shape, or bounding contour, of an image. Such information is both theoretically important (since it provides the basis for figure–ground segregation) and pragmatically useful (since it is easy to determine for any object view).

Studies in this area have found that recognition of silhouettes is largely indistinguishable from recognition of shaded images. For example, Lloyd-Jones and Luckhurst (2002) tested subjects' abilities to recognize silhouettes and shaded images of living and nonliving things and found that silhouettes were particularly useful for recognizing living objects. Hayward (1998) used both name priming and sequential matching tasks and found no difference in recognition of rotated objects between shaded image and silhouette depictions. Similarly, in a more systematic study using different object sets and more study–test manipulations, Hayward, Tarr, and Corderoy (1999) found that *viewpoint costs* (the relative decrement in performance following rotation of an object between study and test) were in general no greater for silhouettes than for shaded images. The implication for theories of object recognition from these results is that perceptual representations of objects might consist primarily, if not exclusively, of components of outline shape.

Tempering this conclusion, however, are several exceptions to the rule of equivalent recognition for shaded images and silhouettes. For example, Hayward et al. (1999) found greater viewpoint costs in a sequential matching task when both images were silhouettes than when either one or both of the stimuli were shaded images. In addition, Lloyd-Jones and Luckhurst (2002) found that although living objects were recognized well using outline information, nonliving things appeared to require internal features. These experiments show that outline shape does not account for object-recognition performance in all conditions. Equally, however, they do

This research was supported by a grant from the Australian Research Council to W.G.H. and B.S. and by Grant CUHK 4332/00H to W.G.H. from the Research Grants Council of Hong Kong SAR, China. Correspondence should be addressed to W. G. Hayward, Department of Psychology, Chinese University of Hong Kong, Shatin, NT, Hong Kong SAR, China (e-mail: whayward@psy.cuhk.edu.hk).

Copyright 2005 Psychonomic Society, Inc.

not fully specify the conditions under which outline shape information is insufficient in accounting for object-recognition performance. For example, it is not clear whether the differences between living and nonliving things observed by Lloyd-Jones and Luckhurst are caused by the nature of set membership or are correlated with existing differences in shape between members of the two categories (and, if so, what those differences are). The living/nonliving distinction is an important one because of the observation of a number of patients with temporal lobe lesions who can recognize living but not nonliving objects or vice versa (Humphreys & Forde, 2001; Warrington & McCarthy, 1983; Warrington & Shallice, 1984).

Previous studies of silhouette recognition have used a variety of objects, ranging from familiar objects (Hayward, 1998; Lloyd-Jones & Luckhurst, 2002) to novel objects composed of geons (Biederman, 1987) and nongeons in differing spatial relationships (Hayward, 1998; Hayward et al., 1999). However, these objects all have some characteristics in common; they are generally composed of a variety of components (normally at least six, and in some cases many more), and the components were generally joined in different possible spatial relationships. Therefore, an alteration of the image from shaded depiction to silhouette, while occluding some relevant information, would still leave considerable shape information available for the task. It is unclear whether this redundancy of information is necessary in order for silhouettes to be efficiently recognized. Surely, eliminating all discriminative information would have a deleterious effect on recognition performance. However, it is uncertain how the provision of minimal but sufficient amounts of information would affect the recognition of an object in silhouette form.

In this study, the stimulus set had restricted visual information. All of the objects had a single conical central component (albeit with different dimensions for each object) and two smaller components arranged midway up the cone, at 90º to each other (see Figure 1). The main features that contrast our object set with those of earlier studies are (1) its simplicity (in terms of number of distinctive features of each object) and (2) its consistency of spatial relations. The former is important because it increases the relative perturbation of total shape that occurs with occlusion of a single feature. For example, one component in the object set used here comprises a large subset of the total shape information about an object; if it is occluded or not visible in a silhouette, the total shape is largely transformed. The consistency of spatial relations is also important. As Johnston and Hayes (2000) have shown, objects with distinctive individual features but consistent spatial relations are more difficult to recognize than are objects with one set of features arranged in different spatial relationships. When different spatial relationships are used, the outline shape of each object becomes completely distinct from all others, and relatively early visual processes may be able to assist in recognition judgments. When there is a consistency of spatial relations, all outlines (even if they contain qualitatively distinct features) are relatively similar. In this experiment, we kept the spatial relations of parts as similar as possible across the objects, to maintain gross similarity (while leaving qualitative differences in each outline).

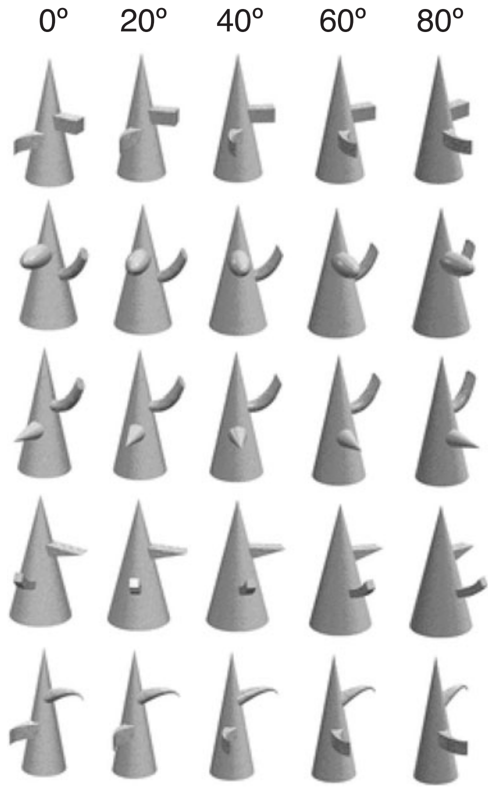

Figure 1. All the shaded images used in the experiment. Each row shows a different object. Silhouettes were black with the same smooth, antialiased outlines.

Note, however, that despite the reduced number of components, there remained discriminative information about the objects at all views, in both shaded image and silhouette depictions. Thus, sufficient information is available at all views for silhouettes to support accurate object recognition. The issue under investigation is whether such minimally sufficient conditions will produce impairments in viewpoint generalization of silhouettes.

We employed a sequential matching task, in which subjects determined whether an initial stimulus (hereafter termed *study*) depicted the same object as a second stimulus (*test*).[1] For shaded image depictions, no components were fully occluded across any of the test viewpoints. Therefore, on the basis of previous experiments (e.g., Hayward & Tarr, 1997; Tarr, Bülthoff, Zabinski, & Blanz, 1997; Tarr, Williams, Hayward, & Gauthier, 1998), we predicted that if shaded image versions of an object were displayed at both study and test, performance would show only small costs as the two views were rotated further from each other.

Two possible predictions can be made regarding performance with the test silhouettes. On the basis of earlier studies (Hayward, 1998; Hayward et al., 1999), one could predict that performance on trials where the test image was a silhouette would be identical to trials where it was a shaded image, unless the study image was also a silhouette (in which case greater viewpoint costs would be expected for silhouettes). On the other hand, if the lack of information redundancy was important, one would expect performance for test silhouettes to be worse than performance for shaded images, particularly since the study and test images were rotated relative to each other.

METHOD

Subjects

A total of 94 subjects from the Chinese University of Hong Kong participated in the experiment in return for either course credit or a small monetary reward.

Design

The experiment employed a sequential matching task in which the studied stimulus consisted of two images (either both the same view or two different views) depicting one object, and the test stimulus showed a single view of either the same or a different object. The studied stimuli consisted of two images, since we wanted to ensure that any results found were not due to the specific views that were tested. Therefore, we had three types of study trials: 20º view only, 60º view only, or both 20º and 60º views. In addition, we manipulated the manner in which the two study images were presented: simultaneously (for 1,000 msec) or sequentially (for 500 msec each, separated by an interstimulus interval [ISI] of 310 msec). The study view variable was manipulated within subjects, whereas study order was manipulated between subjects.

The experiment had a mixed design. Study image type was the principal between-subjects factor, with subjects studying either shaded images (48 subjects) or silhouettes (46 subjects). In addition, there were two main within-subjects variables: test image type (shaded image or silhouette) and test view (one of five possible views). The study view and study order variables (outlined earlier) were included as a procedural control that enabled us to test the generality of the results of principal interest. As discussed below, neither of these variables affected the results or the conclusions.

Stimuli

Five objects were created using 3-D modeling software (Carrara; Eovia, San Diego). Each object (displayed in Figure 1) consisted of a central cone and two other, smaller parts connected midway up the cone. The dimensions of the central cone were slightly different across objects. The smaller components were taken from a set of geons (Biederman, 1987); they were generally different across objects, although two geons were repeated so that subjects could not perform the test by simply paying attention to one component. The objects were rendered at five different views covering a total object rotation of 80º, with antialiasing and without cast shadows. Silhouettes were rendered with the same white background but no light source; thus, they had the same antialiased bounding contour, but all other internal pixels were black.

A mask was created by pairing together random features across the object set.

Procedure

Each subject participated in 4 practice trials and 300 experimental trials. The experimental trials comprised 1 trial for each of five objects, three combinations of study views (20º only, 60º only, or both 20º and 60º), two types of test trial (shaded image or silhouette), and five test views, resulting in 150 trials, along with 150 distractor trials. The trials were presented in mixed fashion, so subjects could not predict what the next type of trial would be like. Distractor trials were based on experimental trials, but one image was removed and an image of a different object put in its place. All independent variables were varied in distractor trials exactly as they were varied in experimental trials.

The experiment was controlled by RSVP software (available at http://www.tarrlab.org/RSVP). Each trial began with a 500-msec blank interval, followed by a 500-msec fixation cross in the center of the monitor. The two study views were then displayed either both simultaneously for 1,000 msec, one on each side of the fixation, or each sequentially in the center of the display for 500 msec, separated by an ISI of 310 msec. Following the second study image, a mask was displayed for 510 msec, and then the test image was displayed until the subject provided a response.

The subjects indicated whether study and test images depicted the same object or two different objects by pressing one of two keys on the keyboard. They heard an audible beep if they responded incorrectly. They were asked to respond as accurately and as quickly as possible.

RESULTS

Only trials in which study and test images depicted the same object, and in which subjects responded between 250 and 4,000 msec after the onset of the test image, were included in the analysis.[2] We analyzed both error rates for all trials adhering to these criteria and response latencies for correct trials.

Condition means are displayed in Figure 2. Note that there were five possible test views, which have been labeled 0º, 20º, 40º, 60º, and 80º. The 20º and 60º views were used for study. We conducted three-way analyses of variance (ANOVAs) on both response latencies and errors in which there was a between-subjects factor of studied image type (shaded image or silhouette) and two within-subjects factors: test image type (shaded image or silhouette) and test view.[3]

Analyses of errors were performed using both subjects (F_1) and items (F_2) as the random factor. Performance was better for test images that were shaded images than those that were silhouettes [$F_1(1,92) = 45.48$; $F_2(1,4) = 114.44$, $ps < .001$] and varied across viewpoints, with generally better performance for studied than for nonstudied views [$F_1(4,368) = 5.69$, $p < .001$], although this effect was not significant by items [$F_2(4,16) = 1.13$, $p > .05$]. There was no effect on performance on the basis of study image type [$F_1 < 1$; $F_2(1,4) = 2.16$, $ps > .05$].

The main effects were modulated by the two-way interactions, which were in turn modulated by the three-way interaction. The interaction of study and test image types [$F_1(1,92) = 5.8$; $F_2(1,4) = 14.4$, $ps < .05$] showed that the advantage for shaded images as test stimuli over silhouettes occurred because performance on test silhouettes was particularly poor when a shaded image had been studied. The interactions of study image type and test view [$F_1(4,368) = 7.93$; $F_2(4,16) = 10.09$, $ps < .001$] and test image type and test view [$F_1(4,368) =$

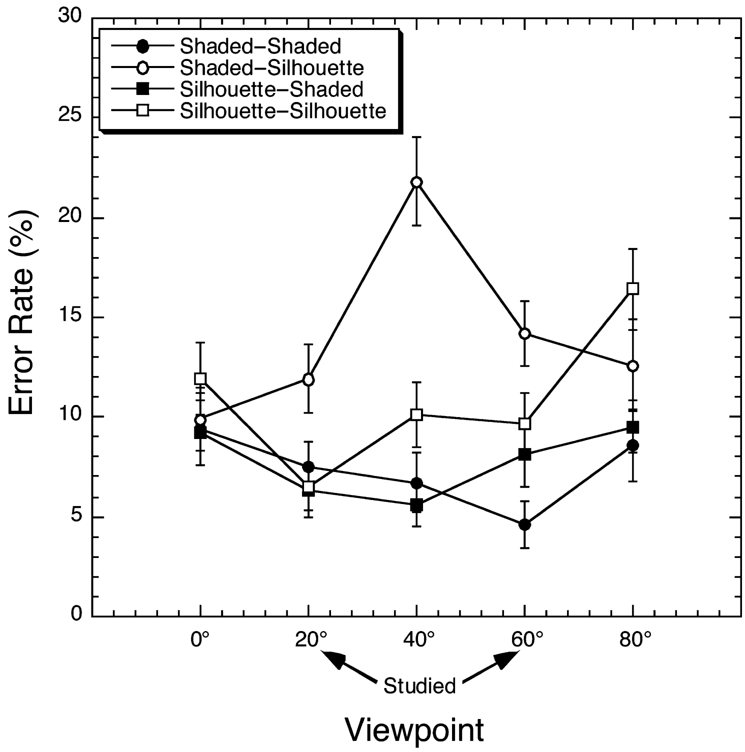

Figure 2. Error rates for the four conditions of the experiment. Error bars show standard errors.

7.06, $p < .001$], though not statistically significant by items [$F_2(4,16) = 1.9$, $p > .05$], both showed particularly poor performance at the unstudied 40° view when objects were studied as shaded images and tested as silhouettes; that the effect was particular to this condition was shown by the statistically reliable three-way interaction [$F_1(4,368) = 5.85$; $F_2(4,16) = 8.46$, $ps < .001$].

To more precisely localize the source of the three-way interaction, we conducted post hoc HSD tests between test image types (shaded vs. silhouette) separately for each viewpoint and each study image type. When shaded images were studied, test shaded images were recognized significantly more accurately than were test silhouettes at the 40° and 60° views ($p < .05$). On the other hand, when silhouettes were studied, test shaded images were recognized more accurately only at the 80° view ($p < .05$).

The analysis of response latencies shows a similar pattern, displayed in Figure 3. Again, both main effects of test image type [$F_1(1,92) = 42.98$; $F_2(1,4) = 38.08$, $ps < .01$] and test view [$F(4,368) = 8.99$; $F_2(4,16) = 4.21$, $ps < .05$] were statistically significant, showing better performance for shaded test stimuli and studied views. There was a main effect for study image type when items were the random factor, with better performance for shaded images than for silhouettes [$F_2(1,4) = 16.95$], although this effect was not statistically significant for the analysis by subjects [$F_1(1,92) = 1.61$, $p > .05$]. The study image type × test image type interaction [$F_1(1,92) = 7.08$, $p < .01$; $F_2(1,4) = 5.32$, $p = .08$] showed that the difference between test shaded images and test silhouettes was much larger when shaded images were studied than when silhouettes were studied. The interaction of study image type × test view [$F_1(4,368) = 6.71$; $F_2(4,16) = 11.53$, $ps < .001$] again seemed to show poorer performance at 40° when objects were studied as shaded images but tested as silhouettes (although performance was also quite poor at 80° following study of a silhouette). Although less definitive, this interpretation was also supported by a marginally significant interaction of test image type × test view by subjects [$F_1(4,368) = 2.18$, $p = .07$]; this interaction was not significant by items ($F_2 < 1$). The three-way interaction was significant by items [$F_2(4,16) = 3.03$, $p < .05$] but not by subjects [$F_1(4,368) = 1.44$, $p > .05$].

We computed the same post hoc HSD tests from the error data on latencies. When the study stimulus was a shaded image, recognition of shaded images was faster than was recognition of silhouettes at 20°, 40°, and 60° ($p < .05$). Conversely, there were no differences at any view between test shaded images and test silhouettes following study of silhouettes.

DISCUSSION

At issue in this article is the nature of silhouette recognition when an object rotates. Unlike previous studies, this study shows that silhouettes do not always support efficient view generalization. When the available features of an object are restricted, changing a stimulus from a shaded image to a silhouette results in a general performance loss, especially across changes in viewpoint. Efficient silhouette recognition appears to depend upon a redundancy of visual features.

Before we discuss the implications of the results presented here, we must clarify that these results do not show that recognition of silhouettes is *generally* difficult, even with the impoverished stimuli used in this experiment. In all conditions, performance was well above chance. In addition, we observed no impairment to recognition accuracy following study of silhouettes, as compared with study of shaded images.[4] Thus, even with a small set of objects that share a similar central volume and have two smaller volumes in similar spatial locations, the provision of outline shape allows subjects to respond accurately. However, although recognition of silhouettes is in general very good, subjects showed a differential impairment at generalizing across viewpoint for silhouettes, particularly after shaded images were studied. This differential impairment must be explained.

Hayward et al. (1999) showed that viewpoint generalization following study of shaded images and silhouettes was almost identical unless a silhouette was observed at both study and test. The results reported here extend and more precisely qualify this finding by showing that when the available information in an object is restricted, generalization to an unstudied silhouette is more impaired than generalization to an unstudied shaded image. Just as Hayward et al. reported, when participants studied a silhouette, we found greater viewpoint costs from the studied views when the test stimulus was a silhouette than when it was a shaded image (shown here at the 80º view). This finding is of general interest because it shows an improvement in recognition performance for *changing* the nature of the image (from silhouette to shaded image) between study and test. Unlike Hayward et al., however, we found here that the study of shaded images seems to produce markedly inferior generalization to silhouettes, at least at 40º. This is the only view at which one of the two smaller components disappears from the outline shape for all objects. Note also that because a silhouette contains a subset of information of a shaded image, good recognition of silhouettes following study of silhouettes shows that poor recognition of silhouettes following study of shaded images is not due simply to a lack of information in the image. The information is available in the image, but observers apparently fail to make appropriate use of it following study of shaded images.

This differential impairment of silhouettes at 40º following study of shaded images therefore suggests that subjects viewing shaded images encode information about internal, nonoutline features in the image. This finding is not surprising, but it was not necessarily supported by some of the results of Hayward (1998) and Lloyd-Jones and Luckhurst (2002). In particular, the results of our study suggest that attention is paid to the two

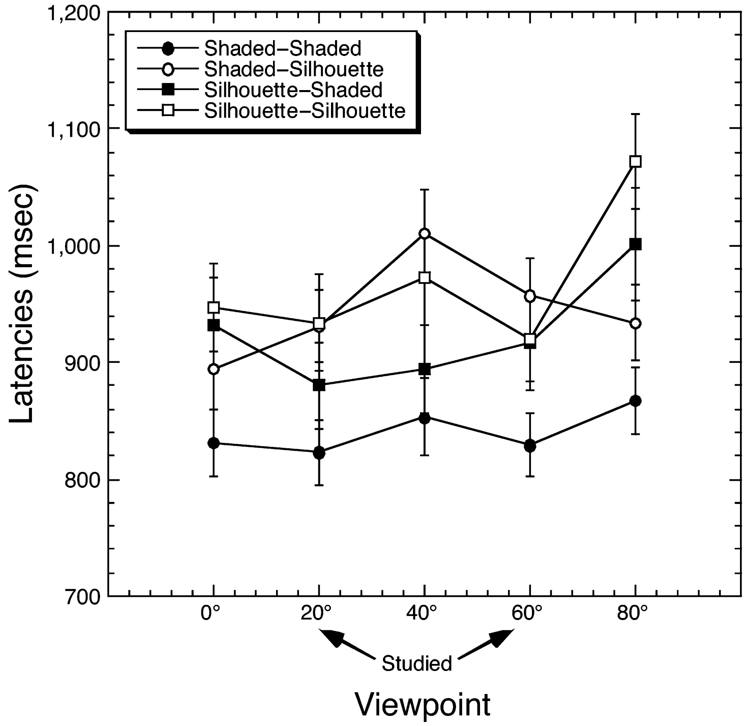

Figure 3. Response latencies for the four conditions of the experiment. Error bars show standard errors.

smaller components in the objects presented. In the 40° view, one of the components is always occluded in the silhouette, and so greater numbers of errors are found. At the 20° and 60° views, however, four of the five objects show both components in at least one view; when summed together, these components allow the test view to be more easily matched to the studied view (although the 60° view still shows a significant, if reduced, difference between recognition of shaded images and silhouettes following study of shaded images). On the other hand, if a silhouette is presented as the study item, the observer is forced to encode the outline; in this case, attention is not drawn to any of the object components that appear within the bounding contour of the image (since they are invisible). The observer is here required to encode information that will be useful in recognizing the object in either silhouette or shaded image form across all views.

Our results suggest differences in the way observers recognize silhouettes and shaded images. The robust performance for shaded images in all studies shows that representations can be derived from these stimuli that support efficient generalization across rotations of the object in 3-D space and that such representations are derived as easily for the simple objects of the present experiment as for the more complicated objects used previously. Performance at recognizing silhouettes, both in the study of Hayward et al. (1999) and in the data presented here, does not show view generalization to the same extent. Generalization to new views is possible if the object contains redundant information and its 3-D structure is known. But if this information is not available, the representation of the silhouette seems fragile and may support recognition only of views close to study.[5] Because of these differences, we speculate that perceptual representations of silhouettes consist of collections of features that are bound less strongly, and with a less clear structural relationship, than are features encoded for shaded images. Such representations of silhouettes support efficient recognition in some circumstances and across some views but will be less successful in general than representations of shaded images.

This view holds implications for the results of Lloyd-Jones and Luckhurst (2002), as well as neuropsychological reports of living/nonliving dissociations. Lloyd-Jones and Luckhurst report that silhouettes of living objects support efficient recognition, whereas silhouettes of inanimate objects do not. Although this difference may relate to differences in types of visual processing (Cree & McRae, 2003; Humphreys & Forde, 2001; Warrington & Shallice, 1984), it may be augmented by the availability of shape features in each type of object. The organic nature of living objects may result in silhouettes that provide better information for representations supporting 3-D view generalization by providing redundant cues (e.g., from the head, neck, torso, legs, feet, and tails of animals). Many inanimate objects, on the other hand, have fewer cues in the silhouette; even if these cues are sufficient for recognition (as shown by Hayward, 1998), they may be insufficient to derive a representation necessary for 3-D view generalization.

If, as proposed here, shape information necessary for making living object identification resides primarily in the outline shape (in redundant quantities), object-recognition procedures would be expected to make use of this regularity. On the other hand, recognition of nonliving objects may require the analysis of function, as proposed by Warrington & Shallice (1984), which would require attention to the relations between components and their potential interactions with human users. We predict that patients with living-object deficits might show particular impairments with silhouette studies such as those conducted here, since they may have a specific deficit in the perceptual processes used to combine shape fragments in the absence of compelling structural information.

REFERENCES

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, **94**, 115-147.

Biederman, I., & Gerhardstein, P. C. (1993). Recognizing depth-rotated objects: Evidence and conditions for 3D viewpoint invariance. *Journal of Experimental Psychology: Human Perception & Performance*, **19**, 1162-1182.

Cree, G. S., & McRae, K. (2003). Analyzing the factors underlying the structure and computation of the meaning of chipmunk, cherry, chisel, cheese, and cello (and many other such concrete nouns). *Journal of Experimental Psychology: General*, **132**, 163-201.

Hayward, W. G. (1998). Effects of outline shape in object recognition. *Journal of Experimental Psychology: Human Perception & Performance*, **24**, 427-440.

Hayward, W. G., & Tarr, M. J. (1997). Testing conditions for viewpoint invariance in object recognition. *Journal of Experimental Psychology: Human Perception & Performance*, **23**, 1511-1521.

Hayward, W. G., Tarr, M. J., & Corderoy, A. K. (1999). Recognizing silhouettes and shaded images across depth rotation. *Perception*, **28**, 1197-1215.

Humphreys, G. W., & Forde, E. M. E. (2001). Hierarchies, similarity, and interactivity in object recognition: "Category-specific" neuropsychological deficits. *Behavioral & Brain Sciences*, **24**, 453-509.

Johnston, M. B., & Hayes, A. (2000). An experimental comparison of viewpoint-specific and viewpoint-independent models of object representation. *Quarterly Journal of Experimental Psychology*, **53A**, 792-824.

Lloyd-Jones, T. J., & Luckhurst, L. (2002). Outline shape is a mediator of object recognition that is particularly important for living things. *Memory & Cognition*, **30**, 489-498.

Tarr, M. J., & Bülthoff, H. H. (1995). Is human object recognition better described by geon structural descriptions or by multiple views? *Journal of Experimental Psychology: Human Perception & Performance*, **21**, 1494-1505.

Tarr, M. J., Bülthoff, H. H., Zabinski, M., & Blanz, V. (1997). To what extent do unique parts influence recognition across changes in viewpoint? *Psychological Science*, **8**, 282-289.

Tarr, M. J., Williams, P., Hayward, W. G., & Gauthier, I. (1998). Three-dimensional object recognition is viewpoint-dependent. *Nature Neuroscience*, **1**, 275-277.

Warrington, E. K., & McCarthy, R. (1983). Category specific access dysphasia. *Brain*, **106**, 859-878.

Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain*, **107**, 829-854.

NOTES

1. Although testing visual working memory, many previous studies have shown that results with this paradigm match those of more ecologically valid, longer term identification paradigms (e.g., Hayward, 1998; Hayward & Tarr, 1997; Biederman & Gerhardstein, 1993; Tarr, Williams, Hayward, & Gauthier, 1998).

2. A total of 1.48% of the trials were excluded because responses were too short or too long.

3. The independent variables of study order (simultaneous or sequential) and study view (20° only, 60° only, or both 20° and 60°) showed few significant results and did not affect the conclusions presented in this article. For completeness, five-way analyses that include these variables are reported in the Appendix.

4. However, shaded images were recognized significantly faster than were silhouettes when analyzed by items.

5. Note that silhouettes should not be regarded as qualitatively distinct from other object depictions in this regard. Any stimulus manipulation that affects perception of 2-D and 3-D structures (e.g., illumination differences; stereoscopic depth manipulation) or stimulus features (e.g., occlusion) can be expected to have effects similar to those discussed here.

APPENDIX
Analyses of Study Order and Study View

To evaluate the role of the two study manipulations, study order (simultaneous or sequential) and study view (20º, 60º, or both 20º and 60º), we entered these variables into five-way ANOVAs (performed as subjects' analyses), along with the three variables tested previously (study image type, test image type, and test view), where study order was manipulated between subjects and study view was manipulated within subjects. Results for test image type and test view, and interactions between these factors, were as reported in the Results section, with the exception that the marginal effect of the interaction between test type and test view for recognition latencies disappeared ($p > .1$). There was a main effect of study order for latencies, with faster performance following sequential presentation than following simultaneous presentation [$F(1,77) = 8.97$, $p < .01$], but there was no effect for errors ($F < 1$). Conversely, the effect of study view, with better performance for studying both views than either view in isolation, was significant for errors [$F(2,180) = 5.86$, $p < .01$] but not latencies ($F < 1$). More important, there were few significant interactions of any type involving these two variables. The interaction of study order × test object type was marginally significant for errors [$F(1,90) = 3.95$, $p = .05$], with a slightly larger decrement in silhouette test performance following sequential study than following simultaneous study, but this effect was not significant for latencies ($F < 1$). The interaction of study order × study view was significant for latencies [$F(2,154) = 3.17$, $p < .05$] but not significant for errors [$F(2,180) = 1.23$, $p > .05$]; in latencies, responses following sequential study were slower following study of the 20º view than in either of the other study conditions, whereas performance following simultaneous study was not affected by the specific study view. Not surprisingly, the interaction of study view and test view was significant for both errors [$F(8,720) = 6.63$, $p < .001$] and latencies [$F(8,616) = 2.73$, $p < .01$], with better test performance for those views that were studied. The three-way interaction of study view × study order × test view was significant for latencies [$F(8,616) = 1.97$, $p = .048$] but not for errors [$F(8,720) = 1.19$, $p > .05$]. Finally, the three-way interaction of study view × test view × test image type was significant for errors [$F(8,720) = 2.27$, $p < .05$] but not latencies [$F(8,616) = 1.21$, $p > .05$]. Both significant three-way interactions were caused by particular test views showing idiosyncratic performance in a combination of the other independent variables.

(Manuscript received October 25, 2002;
revision accepted for publication May 28, 2004.)