Invariant recognition of natural objects in the presence of shadows

Wendy L Braje
Department of Psychology, 214 Beaumont Hall, Plattsburgh State University, Plattsburgh, NY 12901, USA; e-mail: brajewl@plattsburgh.edu
Gordon E Legge, Daniel Kersten
Department of Psychology, University of Minnesota, Minneapolis, MN 55455, USA
Received 24 June 1998, in revised form 3 October 1999

Abstract. Shadows are frequently present when we recognize natural objects, but it is unclear whether they help or hinder recognition. Shadows could improve recognition by providing information about illumination and 3-D surface shape, or impair recognition by introducing spurious contours that are confused with object boundaries. In three experiments, we explored the effect of shadows on recognition of natural objects. The stimuli were digitized photographs of fruits and vegetables displayed with or without shadows. In experiment 1, we evaluated the effects of shadows, color, and image resolution on naming latency and accuracy. Performance was not affected by the presence of shadows, even for gray-scale, blurry images, where shadows are difficult to identify. In experiment 2, we explored recognition of two-tone images of the same objects. In these images, shadow edges are difficult to distinguish from object and surface edges because all edges are defined by a luminance boundary. Shadows impaired performance, but only in the early trials. In experiment 3, we examined whether shadows have a stronger impact when exposure time is limited, allowing little time for processing shadows; no effect of shadows was found. These studies show that recognition of natural objects is highly invariant to the complex luminance patterns caused by shadows.

1 Introduction

The human visual system has a remarkable ability to recognize objects. As viewpoint and lighting change in a scene, the retinal image of an object can vary dramatically. One particular effect of changing lighting conditions is a change in the characteristics (eg shape and location) of shadows. How is object recognition affected by the presence of shadows? Image-based theories of object recognition propose that object representations retain information present in the original image, including shadows (Bülthoff et al 1995; Edelman 1995; Edelman and Bülthoff 1992; Gauthier and Tarr 1997; Poggio and Edelman 1990). Cavanagh (1991) proposed that the early processing of an image involves a crude match of the image to a memory representation, in which all image contours (including shadow contours) are used. Only once a candidate object is selected are the contours labeled as belonging to objects or shadows. If shadows are encoded as part of the object representation, they can be problematic for recognition in that they introduce spurious luminance edges that can be confused with object contours. Alternatively, encoding shadows might improve recognition by providing useful information about local object shape (eg surface orientation and curvature) and/or global scene properties (eg light-source direction). Other theories of object recognition propose that the visual system extracts invariant features, such as object edges, while discounting spurious features, such as shadows (Biederman and Ju 1988; Marr and Nishihara 1978). In this scheme, shadows should not affect recognition, provided they can be easily discounted.

We wished to examine whether shadows improve or impair object recognition. As discussed below, performance on many tasks has been shown to benefit from the presence of shadows, while other studies conclude that shadows impair performance.

However, little research has addressed the specific issue of how shadows influence performance on the task of recognizing objects.

Shadows can be classified into two types (figure 1). An *attached shadow* occurs when a surface turns away from the lighting direction, causing that region to become darker. A *cast shadow* occurs when an object is interposed between a light source and a surface, blocking the illumination from reaching the surface (Beck 1972). Cast shadows can be *extrinsic*, ie one object casts a shadow onto another; or they can be *intrinsic*, ie an object casts a shadow onto itself. All types of shadows tend to be present in real-world scenes, although intrinsic cast shadows are confined to objects with concavities.
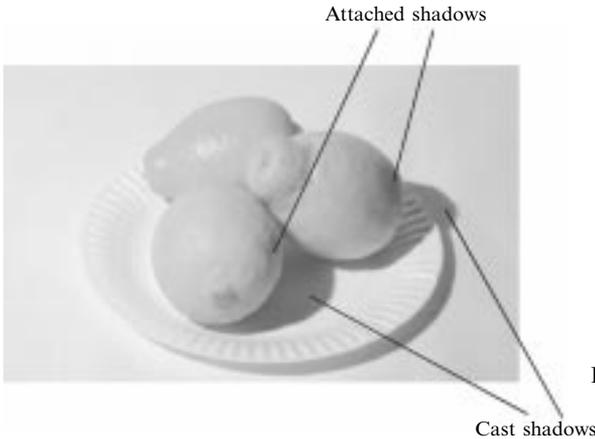

Figure 1. Cast and attached shadows.

In the present experiments, we examined the effects of shadows in general, without distinguishing between these different types. The stimuli were digitized photographs of real objects, in which the different varieties of shadows are difficult if not impossible to isolate. Real images typically contain all types of shadows, and it is nearly impossible to photograph a real object such that only one type of shadow is present. Therefore, although distinguishing among different types of shadows may be important for natural images, we have focused on the more general effects of shadows.

Shadows have been shown to be useful for performing a variety of tasks. Bülthoff et al (1994) demonstrated that shadows can cause flat objects to appear three-dimensional (3-D), and nonrigid motion to appear rigid. Cavanagh and Leclerc (1989) showed that shadows can provide information about 3-D shape. Observers can also use cast shadows to disambiguate convex from concave in shaded images (Erens et al 1993). Yonas et al (1978) showed that cast shadows can be used to infer the shape of the casting object, and the height of an object above the ground. Shadows can also influence perception of motion in depth (Kersten et al 1996), and they are sufficient to produce stereo depth perception (Puerta 1989). Waltz (1975) has used shadow geometry to computationally classify 3-D orientations of surfaces. Similarly, Shafer and Kanade (1983) used shadows to generate constraints on the orientations of surfaces in line drawings of scenes. These studies demonstrate that shadows are useful for making judgments about both local shape and more global scene properties, which may be beneficial when recognizing objects.

Although shadows can be useful, they have also been shown to be ignored or even problematic. Berbaum et al (1983) showed that, when making judgments about convexity, observers did not make use of information about direction of illumination provided by cast shadows. Shadows also have no effect on judgments of slant and tilt, and can impair judgments of illumination direction by causing the illumination to appear more oblique (Mingolla and Todd 1986). Shadows are very difficult to identify computationally because cast-shadow boundaries are generally unrelated to object boundaries

(Cavanagh 1991). Some models of object recognition argue that extraction of object edges is a crucial stage in recognition, and object edges may also be important in the determination of shape and depth. Thus, if shadows introduce edges that are confused with object edges, they should impair recognition performance.

Do these findings extend to the task of object recognition? Warrington (1982) demonstrated that patients with right posterior lesions had difficulty recognizing photographs of common objects containing shadows. Moore and Cavanagh (1998) showed that two-tone images of novel objects with shadows are difficult to recognize. These two findings are consistent with the notion that shadows introduce confusing edges into an image. On the other hand, Tarr et al (1998) demonstrated that cast shadows can improve recognition of novel geometric objects, suggesting that shadows provide useful information about shape or lighting conditions. The stimuli in the latter two studies were unfamiliar objects. Our aim here was to explore whether either of these results extends to recognition of more natural stimuli.

In the present experiments, digitized images of fruits and vegetables were used to study the impact of shadows on the recognition of natural objects. There are two reasons for using natural objects, rather than man-made or computer-generated objects. First, humans have evolved to recognize objects in the natural world, and therefore experiments with such stimuli should provide useful information about how human vision operates. Second, most studies of vision in general and object recognition in particular have used more artificial stimuli, and it is unclear how well the results of those studies apply to real-world stimuli. Common lighting models used for computer-rendered objects often ignore real-world effects such as mutual illumination and cast shadows. Although there is some evidence that computer-generated stimuli can be as effective as man-made objects for certain tasks with simple objects (Johnston and Curran 1996), such effects may not hold for more complex or natural stimuli.

Fruits and vegetables were chosen because they embody a variety of shapes and colors encountered in the natural world. It is also likely that human vision evolved in part to recognize biologically relevant stimuli such as food. Finally, such stimuli have been used in previous recognition studies (Ostergaard and Davidoff 1985; Wurm et al 1993).

In the present experiments, observers identified food images presented on a computer screen. The images were presented with or without shadows, in order to determine whether shadows help or hinder recognition performance. We also examined whether the impact of shadows on recognition is influenced by the presence of information useful for identifying shadows. Shadows can often be identified by certain characteristics, such as a penumbra or continuity of chromaticity across a shadow boundary. These characteristics were manipulated to make shadow identification easier or more difficult, which should alter the impact shadows have on recognition.

2 Experiment 1: Shadows, color, and blur

The effect of shadows on recognition may depend on other information available in the image. For example, the presence of color, which has been shown to improve object recognition (Markoff 1972; Ostergaard and Davidoff 1985; Wurm et al 1993), should alter the influence shadows have on recognition. In general, the hue and saturation of a surface patch will be very similar on both sides of a shadow boundary, allowing for easier labeling of an edge as a surface boundary than as a shadow boundary. Computational techniques for finding shadows sometimes involve finding regions of low intensity that are similar in hue to a neighboring region (Nagao et al 1979; Rubin and Richards 1982). In color images, shadows should be easier to identify. Therefore, if shadows normally impair recognition, their influence should be smaller in color than in gray-scale images.

It should be noted that color can only aid identification of shadows if the direct and ambient illuminations in the scene are similar in their spectral distributions.

In natural scenes, for example, there may be a tendency for shaded regions to contain more short-wavelength ('blue') energy, because such regions are illuminated primarily by the ambient light from the blue sky. However, Hailman (1979) demonstrated that the spectral distributions of light in sunny and shaded areas of several natural scenes (forests) are in fact very similar. Thus, the assumption of similar spectral distribution of direct and ambient light appears to be a reasonable one.

A second factor that may affect the impact shadows have on recognition is the spatial resolution of the image. Resolution can be reduced by blurring an image, which will disrupt the labeling of shadows in at least two ways. First, fine texture will be degraded or even lost. In a high-resolution image, the continuity of texture across a shadow boundary can signal the presence of a shadow, but blurring an image removes this information. Second, shadows often have a penumbra, or 'fuzzy contour'. In a high-resolution image, this blurry boundary can distinguish a shadow contour from the sharper object contours. In a blurry image, however, all contours are similar to shadow penumbrae. In fact, Hering (1874/1964) showed that occluding the penumbra of a shadow causes the contour to appear as an albedo change, rather than an illumination change. Therefore, in a blurry image, shadow edges are more easily confused with surface dges. If shadows normally impair performance, then making them more difficult to identify by blurring the image should cause even larger deficits in recognition.

The overall impact of shadows on recognition as well as these expected interactions between shadows, color, and image resolution were tested in experiment 1. Observers identified images that were displayed with or without shadows, color, and blur.

2.1 Methods

2.1.1 Observers.
Twenty-four undergraduate psychology students (ages 18 to 31 years) at the University of Minnesota participated for class credit. All were fluent English speakers, had normal or corrected-to-normal visual acuity (Snellen acuity of 20/20 or better) and normal color vision (as measured by the D-15 color test), and gave informed consent.

2.1.2 Stimuli and apparatus.
The stimuli were digitized images of 24 types of fruits and vegetables (table 1). Each food appeared in two different forms, such as slices, groups of the same item, or different varieties of the same food. These will be referred to as different poses. Care was taken to select objects and poses that were likely to produce shadows (eg groups of foods and foods with concavities).

The stimuli were photographed onto color slides (Kodak Ektachrome 100 film) with a Nikon FM2 camera fitted with a Tamron SP 90 mm $F2.5$ lens. Each item sat on a white 8-inch paper plate, against a backdrop of a white sheet. Photographs were taken from an angle of $45°$ above the stimuli. The distance from the lens to the stimulus was 3 feet. Ambient illumination was provided by three 100 W soft white bulbs and one 90 W halogen bulb. Directional illumination was provided by a Vivitar 283 electronic flash, subtending roughly 9.3 deg × 5.5 deg. Subjectively, this resulted in fairly sharp shadows (see figure 2).

The presence of shadows was manipulated by changing the direction of illumination. Each stimulus was photographed with the flash illumination from three different directions: for the no-shadow condition, the stimulus was photographed with the illuminant near the camera's lens,[1] ie on the viewing axis (figure 2, top left). For the shadow condition, the illuminant was placed $45°$ to the right or left of the viewing axis, while maintaining the same distance to the stimulus (figure 2, top right). The types of shadows present included both attached and cast, and both intrinsic and extrinsic.

[1] The illuminant was located slightly above the camera lens, rather than right at the lens. Therefore the resulting images did contain a very small amount of shadow (subjectively there were practically no shadows), and this condition could be more accurately termed the 'almost no shadow' condition. For brevity and clarity, however, it will be referred to as the 'no-shadow' condition.

Table 1. List of stimuli used, and the two poses of each.

Stimulus	Pose 1	Pose 2
Apples	Group red#	Slices red
Bananas	Bunch#	Pile (not in a bunch)
Cantaloupe	1/4 slices	1/8 slices
Carrots	Bunch##	Sticks
Celery	Bunch##	Sticks
Grapes	Bunch green	Bunch red
Kiwi	Group	Slices
Lemons	Group##	Slices
Lettuce	Green leaf#	Red leaf
Limes	Group	Slices
Mushrooms	Group	Slices
Olives*	Group black	Group green
Onions	Group green	Group yellow
Oranges	Group	Slices
Peanuts*	Group in shells	Group shelled
Peas*	Group of pods	Group not in pods
Peppers	Group green##	Group red
Pickles*	Spears	Slices
Pineapple	Chunks	Rings
Potatoes	Group brown	Group red
Radishes	Group with greens##	Group without greens
Strawberries	Group#	Slices
Tomatoes	Group#	Slices
Watermelon	1/4 Slices	1/8 Slices

*Items used in practice trials only.
#Items judged as the more prototypical pose on 100% of the trials.
##Items judged as the more prototypical pose on at least 65% of the trials.

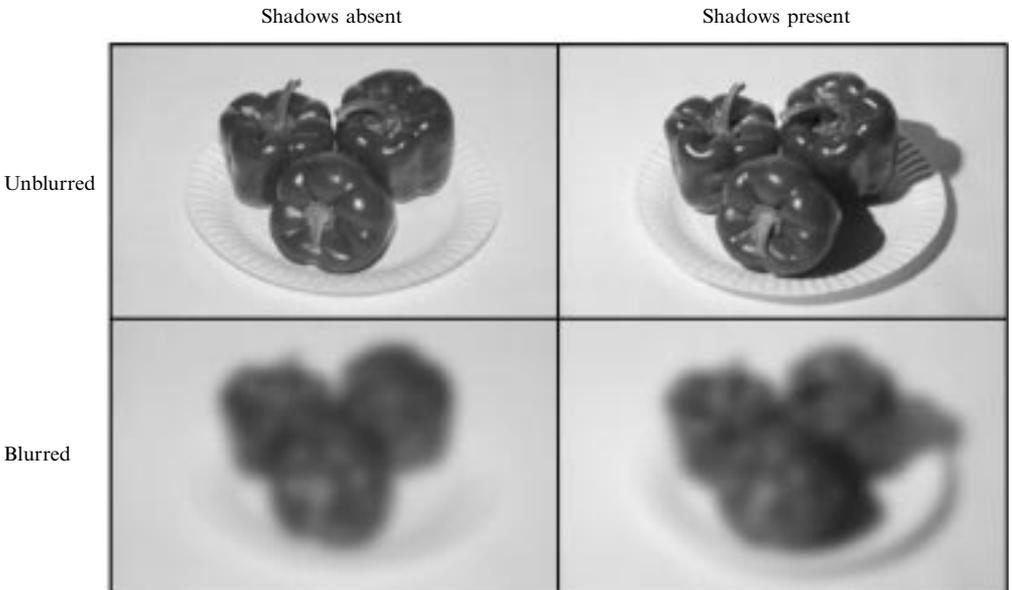| Shadows absent	Shadows present

Figure 2. Peppers rendered in gray scale with and without shadows present, and in high-resolution and blurred. Full-color versions of these four conditions were also used.

The slides were scanned into a Power Macintosh 7100 computer with a Polaroid SprintScan slide scanner. The resulting images were 300 pixels (10.6 cm) horizontally by 187 pixels (6.6 cm) vertically, equivalent to 3.03 deg by 1.89 deg at a viewing distance of 2 m. The images were presented on a white background (67 cd m^{-2}) on a 640 by 480 pixel Apple 17-inch color monitor.

Eight-bit gray-scale images were derived from the color images with the use of Equilibrium Technologies DeBabelizer software. The results of the conversion were checked photometrically to confirm that there were no differences in luminance between the color and gray-scale images. These measurements revealed only small differences, almost always less than 5%, and never greater than 15%.

The color and gray-scale images were low-pass spatial-frequency filtered with the convolve program on a Silicon Graphics Power Indigo 2 computer. A Gaussian filter with a $1/e$ bandwidth of 6 cycles deg^{-1} (20.4 cycles/object) was used. This is similar to the shape and bandwidth of the filter used by Wurm et al (1993). The blurred images were produced by the convolution of the filter and the input images; for color images, the filter was convolved separately with the red, green, and blue channels.

Examples of gray-scale images of a pepper in each shadow $\times$ blur condition are shown in figure 2. The color versions of these images contained color but were otherwise identical to the gray-scale images.

The experiment was run with RSVP software (Williams and Tarr 1998). Naming latencies were recorded by means of the CMU Button Box (accurate within 1 ms). The observers' verbal responses (the names of the foods) were compiled by the experimenter using an IBM XT personal computer.

2.1.3 *Procedure.* Observers were provided with a written list of the foods they would be asked to recognize. The list contained the 20 food category names (eg apples), but not the specific poses (eg sliced apples). They were told that the listed items were the only ones that would be presented, and that each item could appear in different forms (groups, slices, etc). Observers were given as long as they needed to review the list of foods, and they were allowed to refer to it between trials during the experiment.

Each trial in the experiment consisted of a fixation cross presented for 500 ms, followed by a brief tone, and then a food image. The purpose of the tone was to signal the beginning of a trial to the experimenter, who could not see the images being presented to the participants. The observer's task was to name the food aloud as quickly and accurately as possible. Observers were instructed to name the food category, rather than the specific pose. Each image remained on the screen until the observer responded or 10 s had passed. Reaction time was measured as the time between the onset of the stimulus and the observer's verbal response. No feedback was provided.

There were 8 blocks of 40 trials. A single shadow $\times$ color $\times$ blur condition was tested in each block. Within each block, each food was presented once in each of its two poses. For the conditions containing shadows, illumination direction was chosen randomly on each trial.

The eight shadow $\times$ color $\times$ blur conditions were presented in a different random order to each observer. The order was constrained such that each condition (shadow $\times$ color $\times$ blur) was presented in a given block (1st, 2nd, ... 8th) to at least one observer and no more than four observers throughout the experiment (eg between one and four observers ran the shadow – gray – high-resolution condition first, between one and four observers ran this condition second, etc).

Prior to the experiment, observers completed 16 practice trials. These trials contained the 4 food items that pilot subjects demonstrated to be the most difficult (took the longest time) to recognize. The practice items are tagged in table 1 with an asterisk.

2.2 Results

Median reaction times for trials in which a correct response was made are plotted in figure 3a. The main finding was that shadows did not influence performance, regardless of the color and resolution conditions. Reaction times (averaged across color and resolution conditions) were about the same whether shadows were present or absent (1037 versus 1051 ms). An ANOVA run on the log median reaction times revealed no main effect of shadows ($F_{1,23} = 0.87$, $p = 0.361$) and no interaction between shadows and color ($F_{1,23} = 0.06$, $p = 0.806$) or shadows and resolution ($F_{1,23} = 0.01$, $p = 0.508$).

Consistent with the findings of Wurm et al (1993), response times were faster when the stimuli were presented in color than in gray scale (951 versus 1137 ms) ($F_{1,23} = 98.16$, $p < 0.001$), and when they were presented in high resolution than blurred (962 versus 1126 ms) ($F_{1,23} = 88.94$, $p < 0.001$). There was no significant interaction between color and resolution ($F_{1,23} = 2.47$, $p = 0.130$).
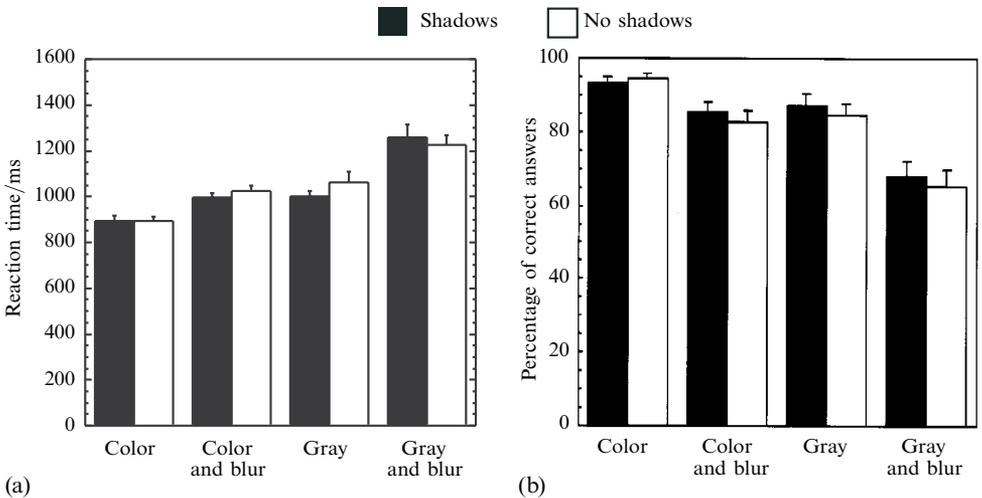

Figure 3. (a) Reaction times for experiment 1. Each reaction time is the average of the median reaction times for each observer in each condition. Standard errors are shown. (b) Accuracy for experiment 1. Each percentage is the average of the percentages for each observer in each condition. Standard errors are shown.

Accuracy is plotted in figure 3b. As with reaction times, shadows did not influence performance. Averaged across color and resolution conditions, accuracy was about equal for images with and without shadows (83% versus 82%). An ANOVA run on the arcsine of the proportion of correct responses revealed no main effect of shadows ($F_{1,23} = 1.08$, $p = 0.309$) and no interaction between shadows and color ($F_{1,23} = 1.08$, $p = 0.309$) or shadows and resolution ($F_{1,23} = 0.98$, $p = 0.332$).

Accuracy was higher when the stimuli were presented in color than in gray scale (89% versus 76% correct) ($F_{1,23} = 92.85$, $p < 0.001$), and when they were presented in high resolution than blurred (90% versus 75%) ($F_{1,23} = 85.28$, $p < 0.001$). Additionally, there was a significant interaction between image resolution and color ($F_{1,23} = 5.87$, $p < 0.030$)—blurring the images led to a greater decline in performance for gray-scale images (86% versus 67%) than for color images (94% versus 84%).

2.3 Discussion

Performance was best with color images and with high-resolution images, consistent with the findings of Wurm et al (1993). However, shadows did not influence recognition, even in images where shadows should be more difficult to identify, ie gray-scale and blurred images. We explored several possible explanations for the lack of an effect of shadows.

2.3.1 *Not enough shadows.* One possibility is that there simply were not many shadows in the images used here. A measure of the 'amount' of shadows in the images would therefore be useful. Although there is currently no satisfactory objective (computational) method for identifying shadows in an image, behavioral measures of whether the shadows are noticeable can be obtained.

An experiment was run to examine whether the shadows were perceptually salient in the images. Eight observers viewed two of the food images (from experiment 1) presented side-by-side, and indicated with a key-press which one contained more shadows. The two images were both of the same food in the same pose, and were presented in full color at high resolution. One of the images was a no-shadow image from the original experiment, and the other was the corresponding shadow image. Each pair of images remained on the screen until the observer responded, and no feedback was provided.

There were 8 blocks of 40 trials. For the shadow images, there were equal numbers of trials containing left-illuminated and right-illuminated stimuli. The shadow images were presented on the right and left sides of the screen an equal number of times.

If the shadows were not perceptually salient, then observers should perform this task rather poorly—they should not be able to tell which of the images is the one with more shadows. The results showed, however, that observers had little trouble discriminating the shadow images from the no-shadow images, achieving 96% accuracy in performing the task. This suggests that most of the shadows were perceptually salient.

2.3.2 *Prototypicality.* Another factor that may determine whether shadows influence recognition is the food pose used. Specifically, the effect of shadows on recognition might depend on the prototypicality of the food pose. For example, prototypical poses might contain a great deal of information (eg about object shape) useful for recognizing an object, and such information may be sufficient to override any confusion caused by shadows. Less prototypical poses, lacking this information, might therefore be more strongly affected by shadows.

An experiment was run to determine which poses appeared more prototypical. Five experienced psychophysical observers were presented with the two poses of each food side-by-side on the computer screen. The images were displayed in full color at high resolution with no shadows. The name of the food was displayed above the images. Observers were asked to decide which of the two poses was more prototypical for that item, and to respond with a key-press. They were given as long as needed to perform this task. Each pair of images was presented four times, with each image shown twice on the left and twice on the right. There were 4 blocks of 20 trials, and all pairs were presented twice by the end of block 2.

The results revealed no correlation ($r^2 = 0.005$) between the proportion of times an image was rated as more prototypical and normalized shadow effect (difference between shadow and no-shadow reaction times divided by the average reaction time for shadow and no-shadow conditions). This means that shadows had the same influence (or lack of influence) regardless of how prototypical an image was judged to be. There was a low correlation ($r^2 = 0.12$) between prototypicality and reaction time, indicating a slight tendency for more prototypical poses to be recognized more quickly.

The lack of correlation between prototypicality and shadow effect cannot be explained by the absence of prototypical images. Five food poses were highly prototypical—they were judged as the more prototypical pose on 100% of the trials. These items are tagged with a '#' in table 1. Five other food poses were judged as more prototypical on at least 65% of the trials. These items are tagged with a '##' in table 1. For the remaining items, neither pose was more likely to be judged as prototypical. The food images thus included both prototypical and non-prototypical poses.

2.3.3 *Priming.* Throughout the experiment, each food category (eg apple) was presented several times. It is possible that priming occurred for the category names, ie performance on the later presentations of a food item may have improved relative to its first presentation. Such priming could be strong enough to overcome small effects of shadows. However, if priming was present at all in experiment 1, it was very weak. There was no difference in accuracy for the first versus second presentation of a food, and median reaction times improved by only 8 ms between the first and second presentations. This reaction time difference is much smaller than was found by Biederman and Gerhardstein (1993) for example, who found improvements of over 100 ms for primed objects. Thus, priming cannot explain the lack of influence of shadows.

Another possibility is that priming occurred between different viewing conditions. For example, recognition in a shadow condition may have improved after completing a no-shadow condition. This possibility was also unsupported by our data: performance was unaffected by the order of presentation of shadow conditions.

2.3.4 *Difficulty of recognition.* Shadows might have an effect only on the foods that are most difficult to recognize. An ANOVA was run on the median reaction times from the trials containing the 15 food poses that were incorrectly identified the most often. The findings were the same as those obtained when all images were used: performance was best with color ($F_{1,23} = 48.9$, $p < 0.001$) and high-resolution ($F_{1,23} = 14.0$, $p < 0.001$) images, but shadows had no significant effect ($F_{1,23} = 0.97$, $p = 0.34$). Thus, difficulty of recognition cannot account for the lack of a shadow effect.

2.3.5 *Effects with specific foods.* Although shadows had no effect on recognition overall, they may have affected recognition of particular foods. A Tukey HSD test ($\alpha = 0.05$) revealed that for most foods shadows had no significant effect; however, a few foods were significantly affected by the presence of shadows, although the results reveal no clear pattern. Recognition of shadow images was slower than that of no-shadow images by more than 700 ms for red lettuce and red potatoes; recognition of shadow images was faster than that of no-shadow images for lemon slices, orange slices, groups of oranges, and radishes with greens.

Slower recognition with shadow images may be a result of confusion between albedo changes and shadow boundaries, particularly for the red lettuce. However, it is unclear why such confusions would not arise in other food images. Faster performance with shadow images of lemon slices was likely a result of enhanced contrast at the edges of the lemon slices, as the cast shadows provided a darker background for the light-colored food. It is not clear why shadows would offer a specific benefit to the recognition of the orange and radish images.

3 Experiment 2: Two-tone images

A possible explanation for the lack of a shadow effect in experiment 1 is that the images, even when degraded by blurring or removing color, still contained substantial information for performing the task—information that could override any shadow effects. Another way to degrade the images in a way that may reveal shadow effects is to threshold them, creating a 'two-tone' image. This involves setting all pixel luminances above a certain value to white, and all those below it to black. The result is a black-and-white image in which illuminated areas are white and shadowed areas are black (Mooney 1957).

In a two-tone image, all edges (object and shadow) are defined in the same way, ie by an abrupt luminance change. This can greatly impair object recognition, in that cast shadow areas are often interpreted as part of the object (Cavanagh 1991). In fact, Moore and Cavanagh (1998) demonstrated that two-tone images of objects with shadows are difficult to recognize. An effect of shadows may therefore emerge when two-tone images of the food stimuli are used.

3.1 Methods

3.1.1 *Observers.* Eight undergraduate psychology students at the University of Minnesota participated in this experiment for class credit. All had normal or corrected-to-normal visual acuity (Snellen acuity of 20/20 or better), had normal color vision (as measured by the D-15 color test), were native English speakers, and gave informed consent.

3.1.2 *Stimuli and apparatus.* The stimuli were thresholded versions of the gray-scale high-resolution images used in experiment 1. Equilibrium Technologies DeBabelizer software was used for thresholding. The threshold for each image was taken as 30% of the maximum image luminance. This threshold value preserved both the general object shape and the shadows in the images. All pixel values above the threshold level were set to white (73 cd m^{-2}), and all values below it were set to black (< 1 cd m^{-2}). A sample image is shown in figure 4.

Figure 4. A two-tone image of peppers with shadows present.

3.1.3 *Procedure.* The procedure was similar to that described in experiment 1. The task was to verbally name each food as quickly and accurately as possible. Before the experiment began, observers were presented with a 'slide show' of the full-color high-resolution food images (both shadow and no-shadow) on the computer screen. They were also given a set of 8 practice trials with two-tone images.

There were 40 no-shadow images (20 foods, 2 poses each) and 80 shadow images (20 foods, 2 poses, 2 illumination directions). Each image was presented once. The experiment contained 2 blocks of 60 trials, and the observer was allowed to rest between blocks. The trials were presented in random order, and no feedback was provided.

3.2 Results

Median reaction times for trials in which a correct response was made did not differ significantly across shadow condition or block (figure 5a). However, an ANOVA revealed a significant interaction between block and shadow condition ($F_{1,7} = 7.76$, $p < 0.03$). In the first block of trials, observers were faster with no-shadow than with shadow images (1191 versus 1334 ms) (confirmed by a Tukey HSD test, $\alpha = 0.05$). In the second block, reaction times were not significantly different for the two shadow conditions (1291 versus 1239 ms).

Similar results were obtained with the accuracy data (figure 5b): there was no main effect of block or shadow condition, but there was a significant shadow × block interaction ($F_{1,7} = 23.87$, $p < 0.01$). As with reaction times, performance was better with no-shadow than with shadow images in the first block of trials (71% versus 63%) (Tukey HSD test, $\alpha = 0.01$), but not in the second block (64% versus 69%).

3.3 Discussion

Shadows impaired performance on the early trials (ie in the first block) of the experiment. This is consistent with Moore and Cavanagh's (1998) finding that shadows impair recognition of two-tone stimuli. The impairment found here is likely due to the difficulty in distinguishing shadow contours from object contours in the impoverished
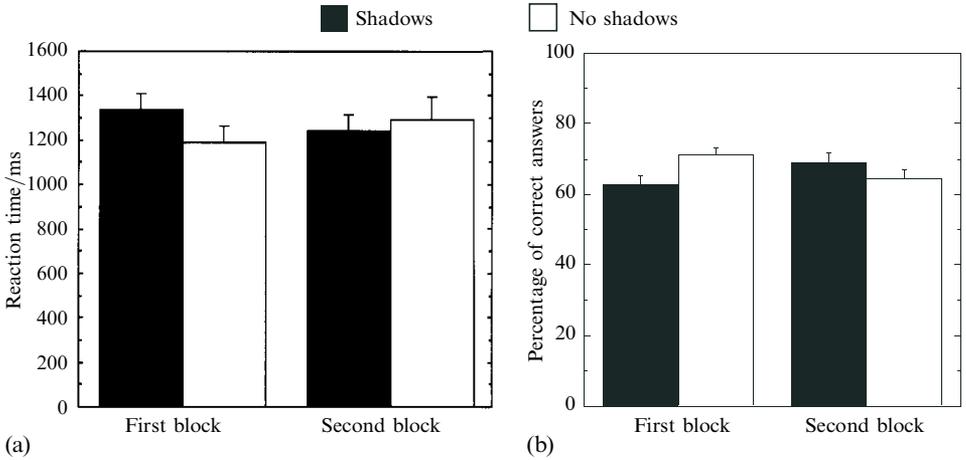
Figure 5. (a) Reaction times for experiment 2. Each reaction time is the average of the median reaction times for each observer. Standard errors are shown. (b) Accuracy for experiment 2. Each percentage is the average of the percentages for each observer. Standard errors are shown.

image, since all contours were defined in the same way. It has been suggested (Moore and Cavanagh 1998) that top–down processing is required for labeling shadows in two-tone images, and this may explain the deficit associated with shadow images in this experiment.

The effect of shadows disappeared for later trials of the experiment. It is possible that observers were able to 'learn' about shadow or illumination conditions during the early trials, and then use this information later on. This is consistent with the argument that higher-level mechanisms are used in processing shadows (Cavanagh 1991). Observers may have also obtained sufficient information about the objects themselves during the early trials (eg their shapes or sizes), such that shadows contributed relatively little noise to the later trials.

4 Experiment 3: Limited exposure time

Another reason shadows may have had no detrimental effect on recognition in experiment 1 is that observers were given plenty of time (up to 10 s) to respond. Even the 1–2 s reaction time that was typical in experiment 1 may have allowed ample time for processing shadows. Cavanagh (1991) has suggested that the labeling of shadows occurs rather late, requiring top–down processing, and that early bottom–up processes do not distinguish shadow contours from object contours. The procedure used in experiment 3 attempted to limit processing time, so that there would be little time to identify shadows. The stimuli were presented very briefly, followed by a mask. Such backward-masking paradigms have been used to control the amount of time observers spend processing a stimulus (eg Reynolds 1981). If labeling of shadows does not occur until later, then larger shadow effects should occur when a limit is imposed on the amount of time observers have for processing shadows.

4.1 Methods
4.1.1 *Observers.* Three volunteers (ages 28 to 32 years) at the University of Minnesota participated for payment, and fourteen volunteers (ages 18 to 30 years) at St Cloud State University participated for course credit. All observers had normal or corrected-to-normal visual acuity (Snellen acuity of 20/20 or better), had normal color vision (as measured by the D-15 color test), were native English speakers, and gave informed consent.

4.1.2 *Stimuli and apparatus.* The apparatus and stimuli were identical to those used in experiment 1. There was also a mask, consisting of a collage of the food images (see figure 6). The mask included all foods and poses, and both shadow and no-shadow images. The mask was presented in color for the color conditions, in gray scale for the gray-scale conditions, and blurred (with the same filter as in experiment 1) for the blur conditions. The masks were the same size as the food images.
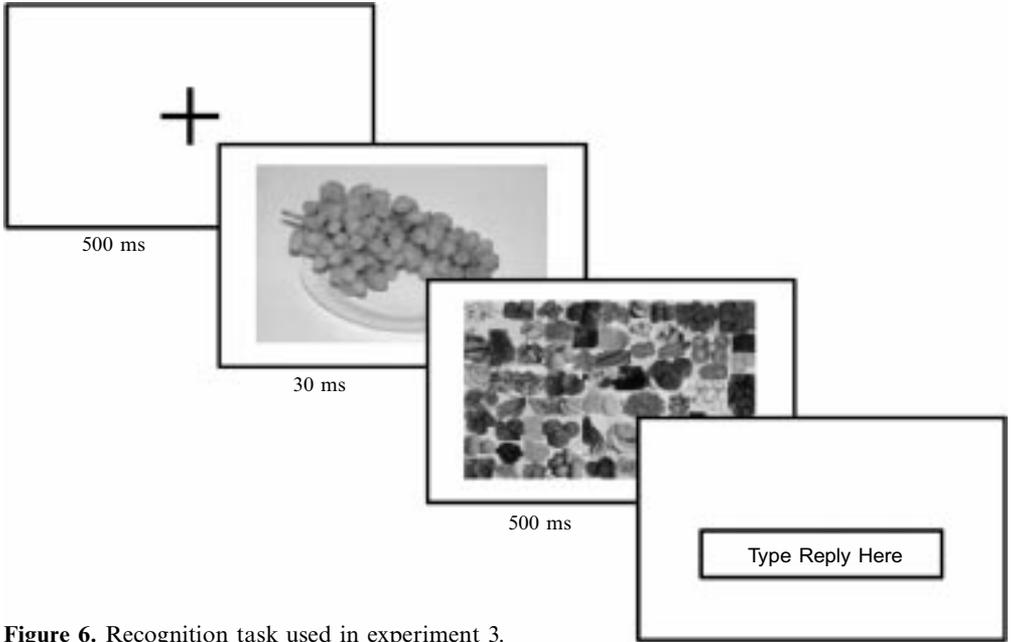

Figure 6. Recognition task used in experiment 3.

4.1.3 *Procedure.* Each trial consisted of a fixation cross for 500 ms, followed by a food image for 30 ms, the mask for 500 ms, and finally a reply box on the screen, in which the observers were instructed to type the name of the food (figure 6). No feedback was provided. As in experiment 1, observers were given a list of the foods they would be asked to recognize. They ran 8 practice trials before the experiment began.

There were 8 blocks of 40 trials. Each block contained one shadow × color × blur condition. In a block, each food was presented once in each of its two poses. The conditions were presented in a different random order to each observer.

4.2 *Results and discussion*

Shadows did not influence observers' accuracy in recognizing the foods ($F_{1,15} = 0.52$, $p = 0.48$), as shown in figure 7. Performance was significantly better with color than gray-scale images ($F_{1,15} = 137.61$, $p < 0.001$) and with high-resolution than blurred images ($F_{1,15} = 98.19$, $p < 0.001$). No interactions were found. In all conditions, accuracy was lower than in experiment 1 (by 30% on average).

Shadows did not affect recognition even when observers were given only 30 ms to visually process the objects and shadows. It is possible that the exposure durations were still long enough to allow for labeling of shadows. However, it seems that, in general, no matter how difficult the task is, shadows have little if any effect on recognition.
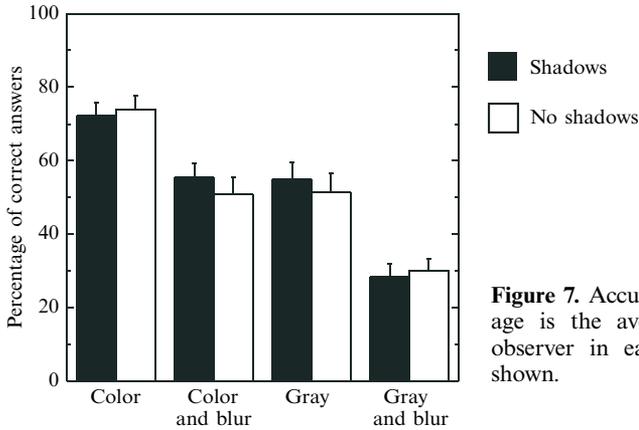
Figure 7. Accuracy for experiment 3. Each percentage is the average of the percentages for each observer in each condition. Standard errors are shown.

5 General discussion

The overall finding was that recognition of natural objects was highly invariant to the presence of shadows. Several experimental conditions were explored here, but none revealed any strong effect of shadows. One interpretation of these findings is that the visual system contains a mechanism for processing shadows quickly and early on. If this is the case, then the mechanism is very fast, completing its work within the first 30 ms of processing.

The findings argue against purely image-based approaches to recognition, in which the visual system must take time to detect and label shadows. Instead, the results are consistent with more abstract feature-based object representations, in which objects are represented by features that are invariant to spurious effects like shadows. What might these features be? In the present studies, color, texture, and boundary sharpness were explored. However, even when these cues were degraded or removed, shadows had no impact on recognition. This implies that these features were not necessary for invariant recognition, and that other shadow-invariant cues were still present.

One such 'cue' could be the global shape of the stimuli. For example, apples tend to be round, bananas curved and elongated, etc. In blurry gray-scale images, and even in two-tone images, this global shape cue is still present. Since shadows tend to be fairly local, global shape information is not disrupted by their presence, and recognition can proceed with the use of this feature.

The familiar stimuli used here apparently contained plenty of redundant information for recognition, such that shadows contributed relatively little to the recognition process. Cavanagh (1991) has suggested that recognition begins with a simple match between an image and a stored representation. Any information (such as shadows) that is not matched at this early stage is processed later, through higher-level mechanisms. When there are many redundant cues to the identity of an object, this early match might be quite successful, and shadows may provide very little extra information, or add very little extra noise. Although the degraded renderings used here (blurred, two-tone, or gray-scale images) did lead to a reduction in recognition performance, these images apparently still contained sufficient information for recognition in the presence of shadows.

It is also possible that shadows had no effect because the task required distinctions between basic-level categories (eg apple versus orange), rather than subordinate-level categories (eg Red Delicious apple versus Granny Smith apple). In subordinate-level recognition, observers need to rely on finer distinctions between texture, color, and precise contour information. Moreover, global shape information is typically less useful for making such discriminations. Shadows might impair performance under

such conditions. Braje et al (1998) have shown that, consistent with this prediction, cast shadows do impair face recognition in a same–different matching task.

Shadows did not impair recognition in the present studies, but neither did they improve recognition. This lack of improvement suggests that, at least under the stimulus conditions used here, shadows were not used for the extraction of 3-D surface shape or the location of objects in 3-D space. Tarr et al (1998) have suggested that, in order to use shadows for shape extraction, the visual system must do two things. First, it must be able to identify a region as a shadow, rather than as a change in surface orientation or albedo. This could be accomplished by using low-level information, such as contrast invariance, color invariance, or a shadow's penumbra. However, removing such low-level information did not induce shadow-dependent performance in the present experiments. This implies that either such information is not used for identifying shadows, or that any information gained about shadows is not used at higher levels involved in recognition. A second requirement proposed by Tarr et al is that the visual system must relate the shadow to the casting object. Observers appeared to have no difficulty in doing this, in that shadows usually did not impair performance. However, the fact that shadows did not improve performance suggests that the observers did not make use of shadows for extracting shape information (or that such shape information was simply not needed for these tasks). In fact, Cavanagh (1995) has suggested that shadows are not normally used for recovering shape. It may be the case that shadows are only a useful cue when novel shapes are used (as in Tarr et al's study), or when no other information is available.

It should be noted that the few shadow effects that were obtained in the present experiments could be due not to shadows but to the change in lighting direction that was used to produce the shadow images. It would therefore be useful to perform similar experiments with computer-generated objects, where shadows can be manipulated independently of other scene parameters.

Another important issue is the particular contributions made by the different types of shadows to the recognition process. For example, both attached and cast shadows might aid recognition by providing information about lighting direction or other global scene properties, thus indirectly contributing to shape recovery; or, recognition processes could make use of the precise shapes of shadows to determine local object shape. Likewise, intrinsic and extrinsic shadows may play different roles in recognition. Intrinsic shadows are always cast onto the same 'background,' ie the object itself. They may therefore be well-learned for familiar stimuli (such as used in the present experiments), and may even be helpful for recognition. Extrinsic shadows, on the other hand, might be problematic, since it would be prohibitive to learn all possible background images. However, Braje et al (1998) found that intrinsic cast shadows impaired face recognition, suggesting that they are not well learned. The present data cannot speak to these issues, since all types of shadows were present in the stimuli. Further research is necessary to examine the particular contributions of different types of shadows.

The effect of the experimental design must also be examined. In experiment 2, the shadow effect disappeared with practice, implying that observers can quickly learn about shadows, illumination conditions, or other stimulus conditions, and use this information for later recognition. This result is highly relevant to experiments 1 and 3. In these two experiments, each observer was exposed to every condition over many trials. If observers can learn about stimulus conditions in a small number of trials, any effects of shadows on performance may disappear when data are averaged across blocks and observers. It would therefore be useful to repeat these experiments using a design in which each observer is exposed to only a small number of trials of a single experimental condition.

Finally, our experiments have examined only one type of object (fruits and vegetables). Although this class of objects appears to be broadly representative of the shapes and colors of objects encountered in the natural world, further research is needed to determine whether these findings extend to other types of objects.

This study has shown that recognition of natural objects is remarkably invariant to the presence of shadows over a wide range of viewing conditions. Even with highly degraded stimuli, the visual system can still use the available information to recognize objects regardless of the presence of shadows. This is probably a consequence of strong evolutionary pressure for the development of object representations that can effortlessly cope with shadows under everyday conditions. Our findings argue against image-based models of object recognition in which shadows must somehow be detected, labeled, and discounted. Instead, our findings favor a more abstract representation of object features that are immune to the presence of shadows.

References
Beck J, 1972 *Surface Color Perception* (Ithaca, NY: Cornell University Press)
Berbaum K, Bever T, Chung C S, 1983 "Light source position in the perception of object shape" *Perception* **12** 411 – 416
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, Ju G, 1988 "Surface versus edge-based determinants of visual recognition" *Cognitive Psychology* **20** 38 – 64
Braje W L, Kersten D, Tarr M J, Troje N F, 1998 "Illumination effects in face recognition" *Psychobiology* **26** 371 – 380
Bülthoff H H, Edelman S Y, Tarr M J, 1995 "How are three-dimensional objects represented in the brain?" *Cerebral Cortex* **5** 247 – 260
Bülthoff I, Kersten D, Bülthoff H H, 1994 "General lighting can overcome accidental viewing" *Investigative Ophthalmology & Visual Science* **35**(4) 1741
Cavanagh P, 1991 "What's up in top – down processing?", in *Representations of Vision: Trends and Tacit Assumptions in Vision Research* Ed. A Gorea (Cambridge: Cambridge University Press) pp 295 – 304
Cavanagh P, 1995 "A horse of a different color: shadows have to be darker but shading does not" *Investigative Ophthalmology & Visual Science* **36**(4) S184
Cavanagh P, Leclerc Y G, 1989 "Shape from shadows" *Journal of Experimental Psychology: Human Perception and Performance* **15** 3 – 27
Edelman S, 1995 "Representation, similarity, and the chorus of prototypes" *Minds and Machines* **5** 45 – 68
Edelman S, Bülthoff H H, 1992 "Orientation dependence in the recognition of familiar and novel views of three-dimensional objects" *Vision Research* **32** 2385 – 2400
Erens R G F, Kappers A M L, Koenderink J J, 1993 "Perception of local shape from shading" *Perception & Psychophysics* **54** 145 – 156
Gauthier I, Tarr M J, 1997 "Becoming a 'greeble' expert: Exploring mechanisms for face recognition" *Vision Research* **37** 1673 – 1682
Hailman J P, 1979 "Environmental light and conspicuous color", in *The Behavioral Significance of Color* Eds J Edward, H Burtt (New York: Garland STPM Press) pp 291 – 354
Hering E, 1874/1964 *Outlines of a Theory of the Light Sense* translated from the German by L M Hurvich, D Jameson (Cambridge, MA: Harvard University Press) [originally published in 1874]
Johnston A, Curran W, 1996 "Investigating shape-from-shading illusions using solid objects" *Vision Research* **36** 2827 – 2836
Kersten D, Knill D C, Mamassian P, Bülthoff I, 1996 "Illusory motion from shadows" *Nature (London)* **379** 31
Markoff J I, 1972 "Target recognition performance with chromatic and achromatic displays" Report No SRM-148, Honeywell, Minneapolis, MN
Marr D, Nishihara H K, 1978 "Representation and recognition of the spatial organization of three-dimensional shapes" *Philosophical Transactions of the Royal Society of London, Series B* **200** 269 – 294

Mingolla E, Todd J T, 1986 "Perception of solid shape from shading" *Biological Cybernetics* **53** 137 – 151

Mooney C M, 1957 "Age in the development of closure ability in children" *Canadian Journal of Psychology* **11** 219 – 227

Moore C, Cavanagh P, 1998 "Recovery of 3D volume from 2-tone images of novel objects", *Cognition* **67** 45 – 71

Nagao M, Matsuyama T, Ikeda Y, 1979 "Region extraction and shape analysis in aerial photographs" *Computer Graphics and Image Processing* **10** 195 – 223

Ostergaard A L, Davidoff J B, 1985 "Some effects of color on naming and recognition of objects" *Journal of Experimental Psychology: Learning, Memory, and Cognition* **11** 579 – 587

Poggio T, Edelman S, 1990 "A network that learns to recognize three-dimensional objects" *Nature (London)* **343** 263 – 266

Puerta A M, 1989 "The power of shadows: shadow stereopsis" *Journal of the Optical Society of America A* **6** 309 – 311

Reynolds R I, 1981 "Perception of an illusory contour as a function of processing time" *Perception* **10** 107 – 115

Rubin J M, Richards W A, 1982 "Color vision and image intensities: when are changes material?" *Biological Cybernetics* **45** 215 – 226

Shafer S A, Kanade T, 1983 "Using shadows in finding surface orientations" *Computer Vision, Graphics, and Image Processing* **22** 145 – 176

Tarr M J, Kersten D, Bülthoff H H, 1998 "Why the visual recognition system might encode the effects of illumination" *Vision Research* **38** 2259 – 2275

Waltz D, 1975 "Understanding line drawings of scenes with shadows", in *The Psychology of Computer Vision* Ed. P H Winston (New York: McGraw-Hill) pp 19 – 91

Warrington E K, 1982 "Neuropsychological studies of object recognition" *Philosophical Transactions of the Royal Society of London, Series B* **298** 15 – 33

Williams P, Tarr M J, 1998 "RSVP: Experimental control software for MacOS [Online]", available at http://psych.umb.edu/rsvp/

Wurm L H, Legge G E, Isenberg L M, Luebker A, 1993 "Color improves object recognition in normal and low vision" *Journal of Experimental Psychology: Human Perception and Performance* **19** 899 – 911

Yonas A, Goldsmith L T, Hallstrom J L, 1978 "Development of sensitivity to information provided by cast shadows in pictures" *Perception* **7** 333 – 341