Abstract

Shadows provide valuable cues for many aspects of visual perception. This thesis discusses the definitions of cast and attached shadows, and the different types of shadow borders that exist. Eight experiments investigated whether the presence of shadows affects the speed or accuracy of human object recognition performance. Experiments 1 to 4 investigated the contributions of attached shadows to the recognition of novel objects, using a sequential-matching task modelled on that of Tarr, Kersten, and Bülthoff (1998). Their finding, of faster reaction times associated with the presence of shadows, was not replicated. Reaction times were not affected by the presence or absence of shadows. Across the four experiments, discrimination was either unaffected by shadow presence, or was at its highest when there were no shadows present. In Experiments 5 to 7, the effects of cast-shadow presence on object recognition were assessed. Visual cues about the shape of the objects were constrained by manipulating the degree of foreshortening of both the objects, and the shadows cast by the objects. Shadow presence was only of benefit to recognition in highly constrained situations: where the objects were severely foreshortened, while their cast shadows were not. Experiment 8 assessed the affect upon recognition of manipulating shape-from-shading cues independently from shadow-border cues. Shadow presence was only beneficial where shading was negligible. It is suggested that shadow presence may only provide observable benefits to object recognition when other cues to an object's identity, such as bounding contour and shape-from-shading, are minimal. It appears that shadows have the potential to facilitate object recognition, but in most situations their presence will not produce any discriminatory, or reaction time, benefit.

Acknowledgements

I would like to acknowledge the help and support of my supervisors, Dr. James McEwan, Dr. John Perrone, and Associate Professor T. Mary Foster. James, you have been my boss for several years, thank you for being good at your job! To my PhD colleagues, Simon Webber, Trudy Pocock, Eric Messick, and Karen Smith, thanks for keeping me entertained by talking about all those things other than research. Last but not least, Allan Eaddy, Ross Oliver, and Rob Bakker; in the years that I've been studying I've been trouble free where computers, networks, servers, backups, and advice are concerned, quite a feat. Cheers everyone!

Contents

List of Figures

List of Tables

Everyday, most of us use visual means to negotiate our way through the world. As we do this, we are constantly processing images of our environment. Our visual systems group these images by perceptual similarity: a categorisation task we call object recognition[1]. The ways in which we accomplish this have yet to be established. Two groups of object-recognition theories have recently led to a large amount of research: view, or feature, based theories, e.g., Edelman's *Chorus* theory (1998), and volumetric, or structural description, theories such as Biederman's (1987) *Recognition by Components*.

View-based theories suggest that the visual system detects various features present in the visual field. A hierarchy of feature detectors is then used to build an object "representation", by providing information about the features that are present in the visual field and their relative positioning (Vecera, 1998). This feature-detection mechanism constrains feature-based models. Termed the viewpoint consistency constraint, "all the features of an object are interpreted as being consistent with viewing that object from a single viewpoint" (Vecera, 1998, pg. 283). Based upon this constraint, feature-based models predict that our recognition system's shape sensitivity would be independent of changes in the size and translation of an object, but dependent upon changes in rotation in space (Vecera, 1998).

Volumetric, or structural description, models propose that the visual system identifies the 3-D structure and arrangement of segments of an object, and that recognition is based upon this. Because these models theorise that the 3-D structure of objects is represented by the visual system, they are not bound by the viewpoint consistency constraint. Thus, they predict that object recognition should be relatively (but not exclusively) insensitive to the effects of the rotation of objects through multiple viewpoints. For reviews of visual object-recognition theory see: Bülthoff, Edelman, and Tarr (1995); Hummel (2000); Logothetis and Sheinburg (1996); Tarr (2003); Ullman (1995); and Vecera (1998).

1. Rosch, Mervis, Gray, Johnson and Boyes-Braem's (1976), review categorisation with respect to object recognition. Also see Herrnstein (1990), and Zayan and Vauclair (1998), for reviews of categorisation.

Procedures Used in Object Recognition Research

There are several common procedures used to assess object recognition. When object recognition is being tested explicitly, verbal naming procedures are often used, where the latency to response, and the level of discrimination are normally the dependent variables, e.g., Braje, Legge, and Kersten (2000). In these procedures the "to be recognised" objects are presented on a computer screen, and voice activated recording of response latencies is employed. Two draw-backs are the requirement of a voice-activated latency recorder, and the inability to assess participants at the same time.

Other options, instead of verbal naming, are word/image matching tasks, e.g., DeCaro and Reeves (2002), and image/image matching, e.g., Braje, Kersten, Tarr, and Troje (1998), Braje et al. (2000), Hayward (1998), and Tarr, Kersten, and Bülthoff (1998). Both word/image, and image/image, matching usually involve the presentation of images on a computer screen, and the recording of response times and accuracy via responses to either the key board or mouse.

A word/image matching procedure is similar to a naming procedure except that the initial stimulus, S1, is a written description of an object, and the participant's task is to respond whether or not the comparison stimulus (the image), S2, matches the description, i.e., the participant must respond whether the name given matches the object seen. Both word/image, and image/image matching procedures often present masking stimuli between the initial and comparison stimuli, and after the comparison stimulus, in an attempt to prevent afterimages of the stimuli and avoid the detection of an immediate change in object bounding contour, e.g., Tarr et al. (1998). A blank in interstimulus interval has also been employed, e.g., DeCaro and Reeves (2002).

An image/image matching procedure usually involves the presentation of an image as the sample stimulus, S1, for a given period of time, followed by a masking stimulus, another image comparison stimulus, S2, and the masking stimulus again. The participant's task is to respond, via either the keyboard or the mouse, whether the comparison stimulus was the same object as the initial stimulus. Because both naming and word/image matching procedures require the participant to be able to classify an image into a verbal category, they are not suitable for testing the recognition of novel objects. When using novel objects, image/image matching is more suitable, the participant can indicate that they recognise the object as the one they just viewed, but do not need to be able to name it.

Across multiple experimental methodologies, research has demonstrated that a change in view-point sometimes results in a change in recognition performance (e.g., Tarr, Bülthoff, Zabinski, & Blanz, 1997), as predicted by view-based models, and in other cases a change in view-point does not result in a change in recognition performance (e.g., Biederman & Gerhardstein, 1993), as is predicted by volumetric models. A problem with interpreting the research, has been that consistent recognition performance across multiple views is not necessarily an indicator of an underlying viewpoint independent process (i.e., that recognition is based upon the perception of 3-D volume). This is because view-based models suggest that the greater the exposure a participant has had to different views of an object, the greater the likelihood of a fast and accurate response at any viewpoint.

To remedy the problems related to the familiarity of an object, when assessing the viewpoint consistency constraint there has been a shift towards the use of novel objects or unusual views of familiar objects. By using novel objects or unusual views, participants' exposure to different views can be controlled (e.g., Perret, Oram & Ashbridge, 1998; Tarr, Kersten & Bülthoff, 1998). The use of novel objects has seen the definition of "familiarity" come under greater scrutiny (e.g., Gauthier, Williams, Tarr & Tanaka, 1998).

Definition of Familiarity

With respect to visual object recognition, familiarity relates to an individual's history of having experienced (i.e., seen and classified) a stimulus before. For example, we say we are familiar with televisions because we have a history of classifying some objects as televisions, and when we do, our peers understand what we are talking about. However, the impossibility of conducting an empirical evaluation of a person's past experiences, means that the operational definition of familiarity has to rely on recent (and thus measurable) performances.

A procedure used by Gauthier et al. (1998) has been to train participants until they are as fast at discriminating between objects within one category, as they are at discriminating between objects of that and other categories. At this stage the observer can be described as an “expert” at recognition of individual exemplars of that category type. For example, someone familiar with American muscle cars could correctly classify a 1969 Mustang at the same speed that they could classify it as simply being a car, and not a truck, house or boat.

Research on Familiarity

The recent increase in the experimental use of novel versus familiar objects has highlighted gaps in scientific knowledge concerning the interaction of familiarity and other aspects of the object-recognition process, e.g., the visual system’s use of configural information in face recognition (Gauthier & Tarr, 1997), and the visual system’s use of shadows (Tarr et al., 1998). Tarr et al. investigated the effects of the presence of shadows on object recognition using a sequential matching procedure. The participants in Tarr et al.’s study had to determine whether two sequentially presented novel stimuli were the same or different. The presence versus the absence of shadows, and the direction of illumination, were manipulated. Tarr et al. found that there was a response-time cost when the direction of illumination of the stimuli was altered with shadows present. When shadows were not present, there was no difference in the participants’ response times taken from the two illumination directions. However, without the shadows present recognition was slower and less accurate. Tarr et al. (1998) concluded that shadows provide useful information about 3-D structure that facilitates recognition.

Tarr et al.’s (1998) result differed from the findings of Braje, Kersten, Tarr and Troje (1998). Braje et al. conducted an experiment on the effects of illumination upon face recognition; recognition performance decreased in the presence of cast shadows. Tarr et al. (1998) suggested two reasons why shadows may not aid recognition for a class such as faces. First, when discriminating between very similar complex objects, such as faces, strong shadows may obscure critical details in the images. Second, the extensive history of experience (familiarity) with the three dimensional shape of faces may result in a reduction in any potential benefit provided by cast shadows.

The initial aim of the current research was to test this hypothesis: that the effects of shadows upon object recognition are modified by the viewer's familiarity with the stimuli.

Physical Definition of Shadows

The question this raises is, what are shadows? Shaded areas and shadows are only differentiated by whether or not they are directly illuminated by light. A shadow is an area that is blocked from direct light, shading is due to the variation in reflected flux as the angle between the incident light and the surface varies (Cavanagh & Leclerc, 1989).

Shadows are usually classified as cast or attached. However, the distinction is often ambiguous in the literature on the effects of shadows on object recognition. Differentiation between attached and cast shadows dates back to at least Leonardo da Vinci (Yonas, et al., 1978). Yonas et al. (1978) provide definitions of attached and cast shadows based upon da Vinci's distinctions:

> An attached shadow occurs when the shadow of an object is visible on that same object, as in the shading on a face due to differences in the orientation of surfaces of the face relative to a light source... a cast shadow occurs when the shadow of an object is seen on another object, as when a person's shadow is cast on the ground (p. 333)

This distinction suffers from the confusion between shadowing and shading in defining an attached shadow. Otherwise, it is the same as that drawn by Cavanagh and Leclerc (1989) (Figure 1.1 illustrates cast and attached shadows). Cavanagh and Leclerc state that:

> Shadows are generally classified as cast (an object's shadow falling on another surface), or attached (an object's shadow that falls on itself - a self-shadow) (p. 6)

However, several authors have used terminology that appears to be based upon the definition of the cast shadow given by Beck (1972), e.g., Braje (2003), Braje et al. (1996), Braje et al. (1998), and Tarr et al. (1998). Beck's (1972) definition is as follows:

> Cast shadow. A darkened area on a surface from which light rays have been blocked by interposing an object between the surface and the light

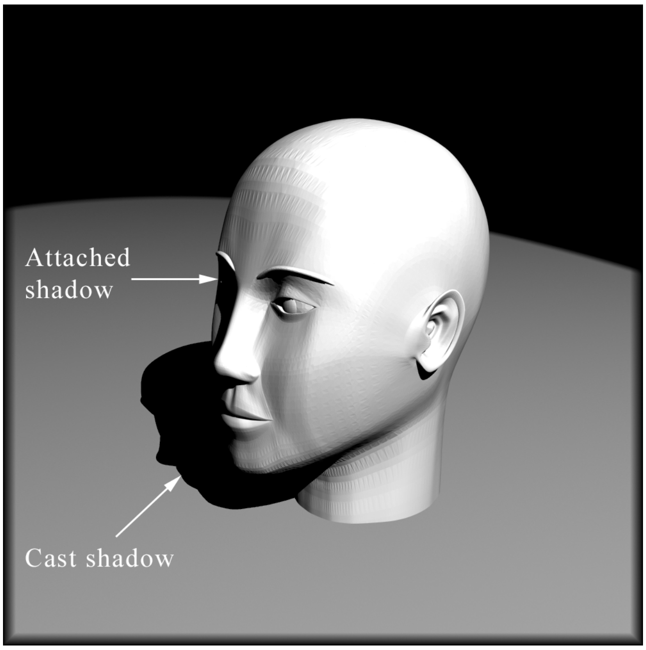

Figure 1.1. Illustration of attached versus cast shadows according to the definitions of Cavanagh and Leclerc (1989). A distinction drawn by Leonardo da Vinci (Yonas, Goldsmith, & Hallstrom, 1978). Cast shadows are those shadows that fall on another surface and attached shadows are those that fall on the same surface of the object.
(Head model: by Anto Matkovic, obtained from 3D Cafe, http://www.3dcafe.com)

> source; a cast shadow may be seen as lying on a surface and is distinguished by its fuzzy edge or penumbra. (p. 180)

Braje et al. (2000) provide definitions of attached and cast shadows, breaking cast shadows into two types, intrinsic and extrinsic:

> Shadows can be classified into two types... 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. (p.384)

Beck (1972) made no mention of attached shadows, or of intrinsic or extrinsic cast shadows, but did identify “Object Shadows”. His definition parallels that of the attached shadows discussed by Yonas (1978), again there is nothing to differentiate the shadow from shading:

> Object shadow. The shadow lies on an object and is created by the shape and spatial orientation of the object to the light source; an example of object shadows are the numerous shadows present on a crumpled towel which are not seen as shadows but as folds and creases of the towel. (p. 185, Beck, 1972)

The definitions of cast shadows are very different to those given by Yonas et al. (1978), and Cavanagh and Leclerc (1989). The definitions of intrinsic and extrinsic cast shadows correspond to the definitions of attached and cast shadows provided by Cavanagh and Leclerc. Thus, dependent upon whose definition is being used, an attached shadow (according to Cavanagh and Leclerc) could also be called a cast shadow (Braje and colleagues). The definitions provided by Braje et al. (2000) fail to provide a valid distinction between an “attached shadow” and “intrinsic cast shadow”. How is an attached shadow, where “a surface turns away from the lighting direction, causing that region to become darker”, not a subset of the intrinsic cast shadow, where “an object casts a shadow onto itself” (p.384, Braje et al.)?

Furthermore, why should an intrinsic cast shadow be confined to an object that has concavities?

The sphere illustrated in image A of Figure 1.2 (an adaptation of figures given in Cavanagh, 1991) shows how a solely convex object can cast shadows upon itself. Is this sort of shadow only an attached shadow because the object is not concave? The block illustrated in the same figure also casts a shadow onto itself. The block is neither convex nor concave, therefore, according to Braje et al. (2000) the shadows cannot be intrinsic cast shadows and must be attached shadows. Braje (2003), Braje et al. (1996), Braje et al. (1998), and Tarr et al. (1998) refer to their stimuli as containing cast shadows. This suggests that the above researchers have been inaccurate in their usage of the terms attached and cast shadow.

The definitions provided by Cavanagh and Leclerc (1989) are simple, unambiguous, follow from a long tradition of use in the arts, and refine the definitions provided by Yonas et al. (1978). Cavanagh and Leclerc also specify the difference between shadowing and shading. The definitions provided by Braje et al. (2000) contain ambiguities, and in doing so fail to provide adequate differentiation between attached and cast shadows. This detracts from the clarity of any scientific explanation. It is proposed that future investigations into how shadows affect our perceptions employ the definitions provided by Cavanagh and Leclerc.

However, the question remains as to how useful the distinction between attached and cast shadows is. Cavanagh and Leclerc (1989) and Cavanagh (1991) suggest that paying specific attention to a shadow's borders is more important than whether the shadow is attached or cast. Cavanagh and Leclerc (1989) note that shadows have both attached and cast borders. As noted earlier here, they also suggested another three types of shadow borders (shown in Figure 1.3): 1. the same-surface border of either a cast shadow, or of an attached-shadow boundary along the points of the object that are normal to the direction of the illuminant; 2. the joined-surface border where an attached contour is at a sharp discontinuity in surface orientation; and 3. the occluded-surface border where the illuminated background is occluded by the extremal contour of the object that is in shadow. Given the occluded-surface border, a fourth type of shadow border is logical, the

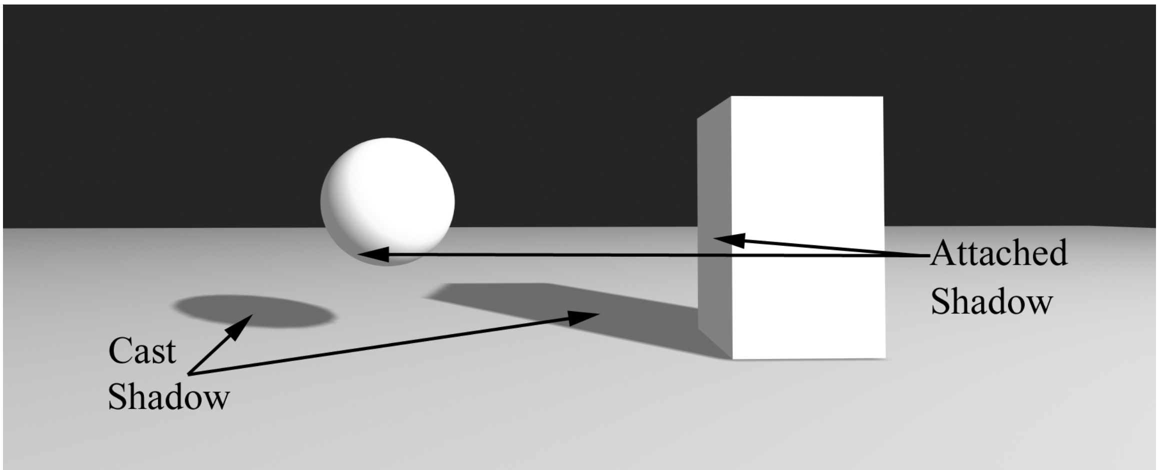

Figure 1.2. Illustration of how shadows can form upon both a solely convex object (the sphere), and an object that is neither convex nor concave (the block).

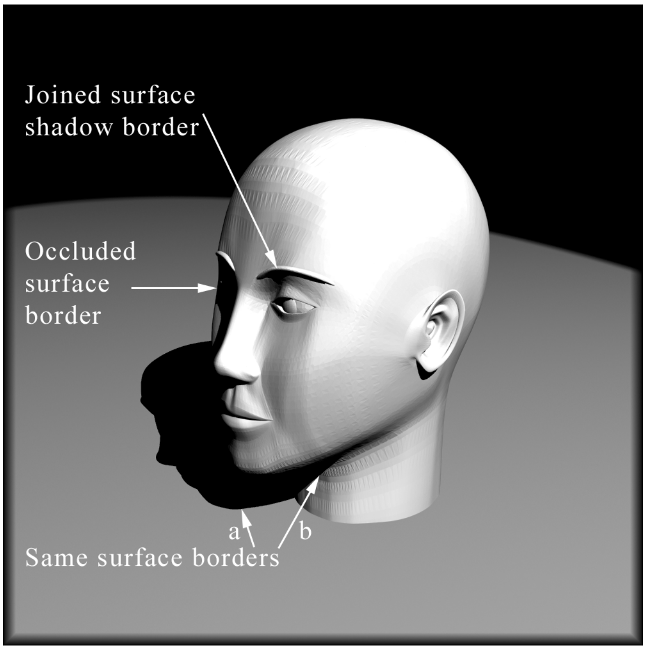

Figure 1.3. Illustration of the different border types associated with shadows. Cavanagh and Leclerc (1989) define the different border types: a joined-surface border is a terminator contour that is attached to a sharp discontinuity in surface orientation; an occluded-surface border occurs where the background surface is occluded by a part of the object that is in shadow; and a same-surface border occurs where the shadow falls across, and divides, a continuous surface, of note is that this can occur either as a cast shadow, as in *a*, or as an attached shadow, as in *b*.

occluded-background border, where an illuminated surface occludes a shadowed background.

In a later paper, Cavanagh (1991) addresses only attached borders, cast borders, and external borders, at times using the term contour in the place of border (Figures 1.4, 1.5, & 1.6 illustrate these border types). Cavanagh states that: “Shadows have two types of borders: attached borders where the direction of the illumination is perpendicular to the surface normal (the light just grazes the surface); and cast borders where the shadow cast by one surface falls on a second surface” (p. 297) also noting that “An object’s external borders are only visible where the background and the object have a (sic) different brightnesses” (p. 297).

As described above, occluded-surface borders, and occluded-background borders are shadow borders that fit the requirements of external borders, but not attached or cast borders. Thus, it is the suggestion of this thesis that there are three shadow border types: attached borders, cast borders, and a subset of external borders. To avoid confusion, from this point forward, the term external border is used with reference to an object’s border that is shadowed on one or both sides, while bounding contour will be used with reference to an object’s outer/extremal contour, irrespective of shadow presence.

Cavanagh (1991) argues that cast borders have a special status in images because they do not correspond to any discontinuity in the object, but to a discontinuity in illumination, and therefore, they are not a material border and need to be excluded from the analysis of an image. His point is illustrated by the comparison of Figure1.7 images A and B, cast contours are a product of the shapes of both the casting surface and the receiving surface. The cast border is a source of information about the bounding contour of the object, that is spatially separate from the object itself, and its shape is partially a function of the contour of the receiving surface. The attached and external borders provide information about an object’s contour at that border. Thus, dependent upon the contrast between shadowed and non-shadowed areas, they could serve to either highlight, or obscure, the contour of the object at that point.

Many shadow border types have been described: the same-surface borders, joined-surface borders, and occluded-surface borders described by

Figure 1.4. Illustration of attached borders.

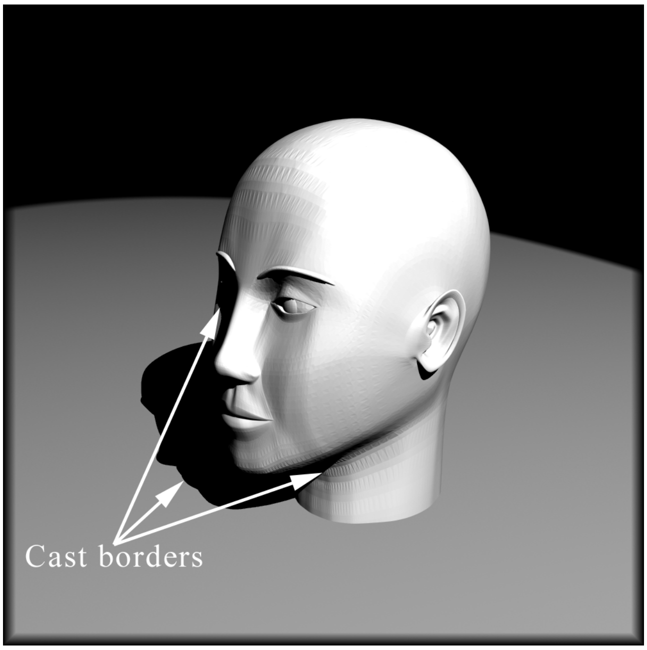

Figure 1.5. Illustration of cast borders. Cast borders may be part of either an attached shadow or a cast shadow.

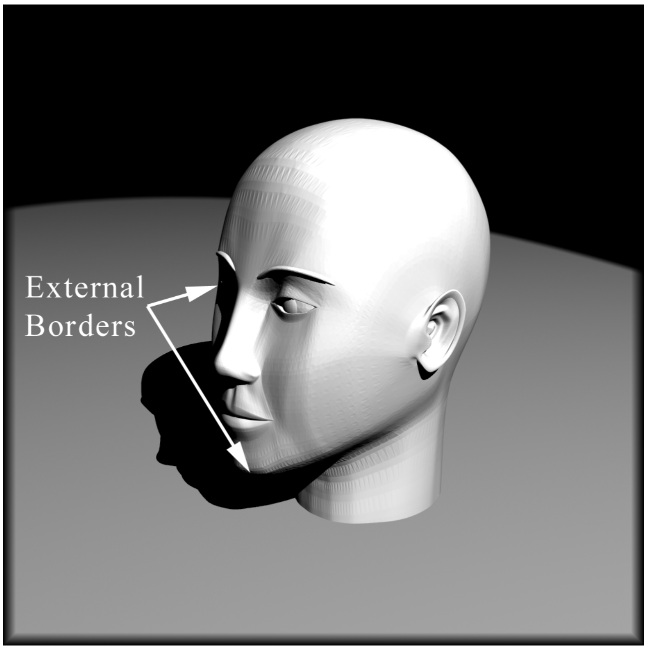

Figure 1.6. Illustration of external borders. Two external borders are illustrated, one where the shadowed object occludes the illuminated background (an occluded surface border), and a second where the illuminated object occludes the shadowed background (an occluded object border).

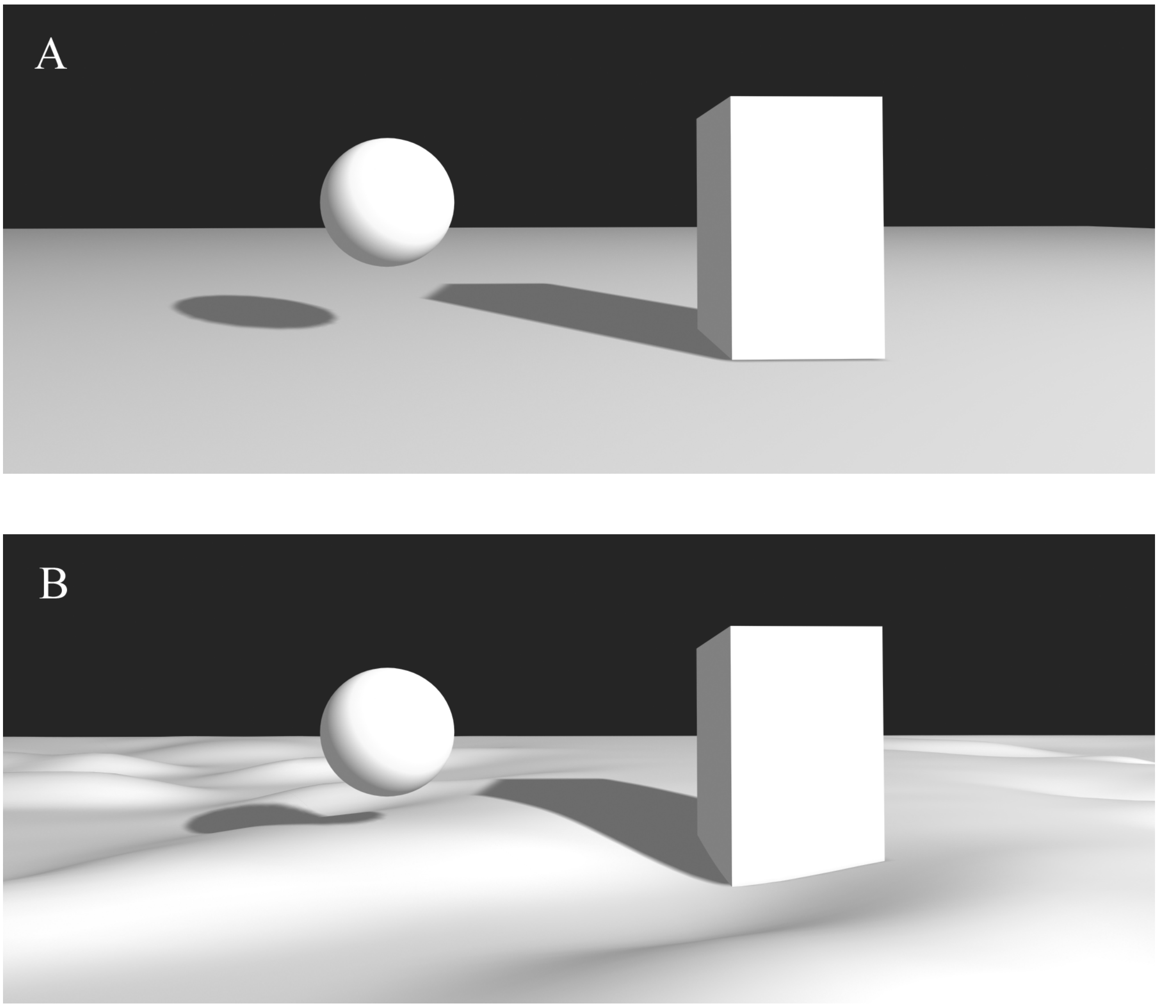

Figure 1.7. Illustration of the effect of a non-uniform receiving surface upon cast borders (compare image A to B).

Cavanagh and Leclerc (1989), plus the similar occluded-background border; and the attached, cast, and external, borders described by Cavanagh (1991) and here. The attached, cast, and external, borders seem to define the borders at a more general level than when considering same-surface, joined-surface, occluded-surface, and occluded-background borders. Thus, attached borders could be considered to be comprised of both, one, the subset of same-surface borders that are caused by the direction of illumination being perpendicular to the surface normal, and two, joined-surface borders, where shadowing is caused by a sharp discontinuity on the object surface. The distinction between these two sorts of borders is really how sharp a discontinuity is present, e.g., the edge of a sphere versus the edge of a box. Cast borders are those same surface borders that result from the shadow of one surface being cast onto another surface, the shape of the resultant boundary being a function of the shapes of both the casting and receiving surfaces. External borders are those where the object's bounding contour has shadowing either on the object side of it, an occluded-surface border, on the background side of it, an occluded-background border, or on both sides of it[1].

To conclude the discussion of various shadow types and their borders, the following suggestions are made. It is suggested that the terms cast and attached shadows should be used consistently by all researchers, and that the definitions offered by Cavanagh and Leclerc (1989) should be employed, as opposed to the ambiguous definitions provided in Braje et al. (2000). It is also suggested, with respect to the different boundaries of shadows, that considering the three major groupings of attached, cast, and external boundaries, should provide the greatest utility in determining how shadows may affect our perceptions.

Perceptual Definitions of Shadows

As well as providing definitions of the types of shadow borders, Cavanagh and Leclerc (1989) assessed what information our visual system uses

1. Note: the interactions between the physical properties of shadows, interreflections, and the object and background, means that where the bounding contour of an object is in shadow, a contrast gradient could still exist across this boundary, and that this boundary is in part due to the shadowing present.

to identify areas as shadow. They assessed the cues of luminance, colour, texture, motion, and depth due to binocular disparity, to find out which supported the perception of a shadow. Perception of shadowing was deemed to occur when an image was identifiable through the shape-from-shadow cue or cues being provided. Shape, when defined by texture, motion, binocular disparity or colour, failed to provide the perception of shadowing. The only cue to support the participants' perception of shadows was luminance. Simply, the shadowed area had to be darker than the surrounding area, along the entirety of its border. Cavanagh and Leclerc's research also demonstrated that the interior of the shadow area could be the same luminance as the no-shadow area, and still support the perception of shadowing (as illustrated in Figure 1.8, image F). Thus, a shadow border can be defined as "a consistent polarity of luminance contrast both from point to point along its length... and across scales at each point" (p. 20, Cavanagh & Leclerc, 1989).

Although it is a very comprehensive analysis into what the cues are that our visual system uses to define shadows, the research by Cavanagh and Leclerc (1989) needs to be supplemented by other experimental findings. The actual penumbra of a shadow is important for our perception of a shadow. Hering (1874) demonstrated this: he observed that the perception of a shadow can be destroyed by drawing a line around the shadow, covering the shadow's penumbra (depicted in Figure 1.9). MacLeod (1940) performed an experiment on lightness constancy, using a procedure designed to replicate the effect Hering found. Eight of MacLeod's 38 participants reported seeing the shadow enclosed by a line as a coloured surface, and the unadulterated shadow as a shadow. For these participants, the line covering the penumbra destroyed the perception of the shadow.

Kennedy and Bai (2000) conducted another experiment, with implications for the effect found by Hering (1874). They manipulated the luminance of the bounding contours of Mooney Images, which enabled them to evaluate the conditions under which the images were readily perceivable. When the bounding contour was darker than both the shadow area and the non-shadow area, perception was disrupted. When the contour was either lighter than, or of equal luminance to, the shadow area, perception of the images was not disrupted. Thus, Kennedy and Bai concluded that for shape from shadow perception to occur, a luminance gradient, from dark to light,

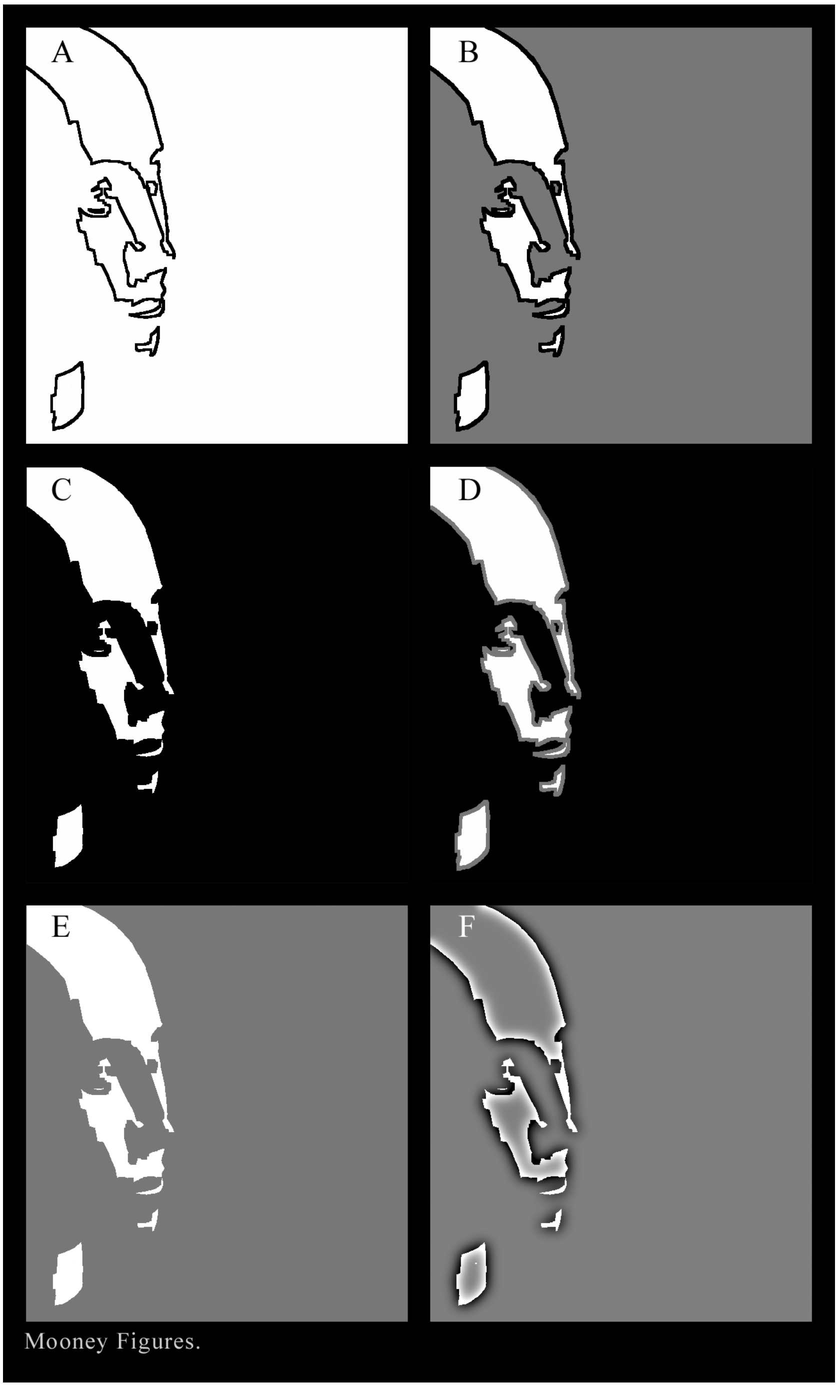

Figure 1.8. Various adaptations of Mooney (two-tone) images. Image F, is a high pass filtered image of a face, where the central regions of the shadow and non-shadow areas are of the same luminance. Cavanagh and Leclerc (1989) found that shape-from-shadow perception was provided in the case of high pass filtered images, even though they have a similar luminance gradient to that illustrated in image B. Kennedy and Bai (2000) concluded that for shape-from-shadow perception to occur, a luminance gradient, from dark to light, across the shadow to non-shadow areas was necessary: illustrated in images C, D and E. When this luminance border was reversed, perception was disrupted: as seen in images A and B.

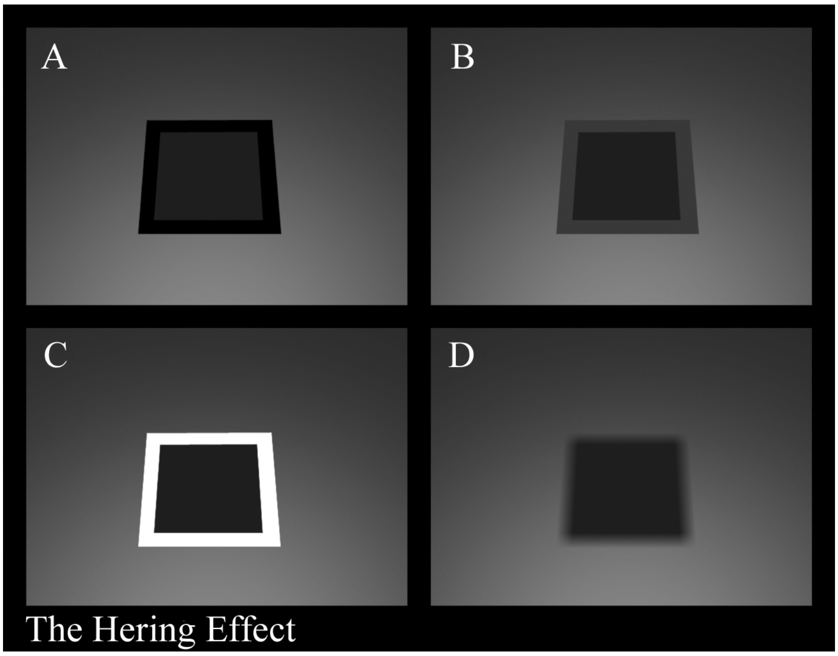

Figure 1.9. Illustration of the Hering effect. A shadow with penumbra is illustrated in image D: it looks like a shadow. Images A through to C illustrate the effect of drawing a line around the shadow, to cover the penumbra. The perception of the dark patch as a shadow is destroyed.

across the shadow to non-shadow areas was necessary: when this luminance border was reversed, perception was disrupted (illustrated in Figure 1.8). Kennedy and Bai suggested that the effect they found may explain the disruption of shadow perception identified by Hering.

These experiments suggest that a luminance border is not the only necessity for recognition of a shadow, but that the width of the border is also important. When using a wide border, as illustrated in Figure 1.8, image F, Cavanagh and Leclerc (1989) demonstrated that the border increases in luminance on the shadow side, and still supports shape from shadow, while Kennedy and Bai (2000) found that with a fine border, only a gradient of light to dark (when shifting from no-shadow to shadow areas) supported shape-from shadow perception.

While the above research relates to aspects of the shadow itself, further experimental evidence indicates that top-down scene recognition processes may also determine if a dark patch is perceived as a shadow: i.e., the perception of a shadow is controlled by the plausibility of its existence in relation to other objects in the scene. In a study on perceived illumination, Logvinenko and Menshikova (1994) demonstrated that the perception of a shadow can be destroyed by inverting the apparent depth of the shadow casting object. Participants viewed a cone that was illuminated so as to cast a shadow onto a vertical white screen that its base was attached to. When viewed through a pseudoscope[1], the cone was perceived as a conical hole in the screen, and the subsequently impossible shadow as a darkly pigmented area (Logvinenko & Menshikova). The perception of the shadow as a pigmented area did not occur immediately for Logvinenko and Menshikova's participants, but took a few seconds to stabilise. Logvinenko and Menshikova relate that during this time the shadow appeared as a shadow even though the participants were aware that the situation was physically impossible. Once the shadow was perceived as a pigmented area, the participants could not reverse the shift in perception. Thus, the perception of the shadow

1. The apparent depth of an object can be inverted by viewing it through a pseudoscope: a pseudoscope reverses apparent depth by using two dove prisms to reverse the retinal disparity of the images presented to each eye (Logvinenko & Menshikova, 1994).

was destroyed by the replacement of an appropriate shadow casting object from the scene with an inappropriate shadow casting object.

Wimmer (1994), further discusses the correspondence needed between shadow shape and object shape, before a dark area is considered a shadow. Employing different geometric shapes, he found that when multiple shadows are present, those most similar to the shadow casting object are associated with it, and that, as the proximity of the shadow and object increased, the likelihood of the association of the shadow with the object increased. Wimmer found depth perception to be enhanced when the shape of the shadow and the shadow casting object are similar, but that a high degree of correspondence was not necessary for a cast shadow to function as a depth cue.

However, Wimmer (1994) noted a dichotomy in that subjects could accurately judge depth based upon a non-corresponding shadow, and were still aware that the shadow did not match the casting object. Depth perception was also evident when the shadows were of a higher luminance than the surrounding area, contrary to Cavanagh and Leclerc's (1989) finding that a shadow region is defined by the fact that it is darker than the surrounding area all along its border. It may be the case that the participants could still judge depth from the shadow, even when it did not appear as a natural shadow to them.

These experimental results suggest that the perception of a shadow is regulated by a top down process that takes into account the physical relationship of a dark area, defined by a luminance boundary, to objects in the scene. If there are no suitable shadow-casting objects, the shadow area may be perceived as a pigmented area on the receiving surface (Logvinenko & Menshikova, 1994). When there are multiple shadow candidates, the shadow that bears the greatest similarity to the object in the scene is associated with that object, a process that is also affected by proximity (Wimmer, 1994). Lastly, some functions shadows serve in scene interpretation, such as depth perception, may remain relatively unaffected by changes in the physical properties of the shadows, e.g., changes in luminance and the correspondence between shape and shadow (Wimmer, 1994).

The Use of Shadows by the Visual System

There are numerous studies demonstrating the effects of shadows upon visual perception. We can recognise our own shadow from about 3 years of age onwards (Cameron & Gallop, 1988), and our visual system treats shadows differently to shading (Cavanagh, 1995). Furthermore, it has been demonstrated that shadows can affect our perceptions of many attributes of their casting objects: shadows are powerful cues to an object's movement (Kersten, Knill, Mamassian, & Bülthoff, 1996; Kersten, Mamassian & Knill 1997; Price, O'Toole & Dambach, 1998), an object's depth in the visual field (Allen, 1999, 2000; Puerta, 1989; Wimmer, 1994; Yonas, Farr & O'Conner, 2001), and contact between objects (e.g., Madison, Thompson, Kersten, Shirley, & Smits, 2001). Similarly, moving shadows can produce the perception of a three dimensional object (e.g., the kinetic depth effect: Day, 1989; Lucas & Taylor, 1979; Norman & Todd, 1994).

Several investigations have been conducted into the relationship between our perceptions of shape and of shadow. As described above, Wimmer (1994), and Logvinenko and Menshikova (1994) have shown that our perceptions of shadows are affected by the shape of the casting objects: the less possible a shadow is, the less likely it is to be treated as a shadow. Yonas (1978) has demonstrated the opposite effect, that shadows can control the perception of object shape, instead of object shape controlling the perception of shadows. In Yonas' experiment cast shadows disambiguated whether an ellipse viewed in a perspective drawing was an ellipse standing vertically, or a circle lying horizontally. Similarly, Berbaum, Bever, & Chung (1984), and Erens, Kappers, and Koenderink (1993), used cast shadows to disambiguate concave versus convex surface-relief and lighting direction. Bülthoff, Kersten, and Bülthoff (1994), have also illustrated that shadows can influence the perception of object shape: shadowing produced 3-D perception of solid shapes in situations where the shapes otherwise appeared flat because of being presented from accidental views.

These results show that the visual system does process the relationships between objects and shadows in a scene. However, in all the situations where shadows have affected the perception of shape, there has been a high degree of ambiguity (e.g., Yonas, 1978; Berbaum, et al.1984; Erens et al. 1993; and Bülthoff et al. 1994). Ambiguity results from of a paucity of cues available for shape perception. In most scenes there are many cues available that can contribute to object recognition, e.g., biological motion (Battelli, Cavanagh, & Thornton, 2003), colour (Naor-Raz, Tarr, & Kersten, 2003), bounding contour (Hayward, 1998; Hayward et al., 1999), the configuration of internal features (Gauthier & Tarr, 2002), shading (Liu, Collin & Chaudhuri, 2000), binocular disparity (Liu et al., 2000), and shadowing (Castiello, 2001).

With regard to object recognition, rather than the perception of shape, the benefits or disadvantages of shadow presence have been hard to identify. It is potentially possible for the visual system to use shadows to extract the surface shape of objects, to label and attach contours to objects, and to determine light source direction (Knill, Mamassian & Kersten, 1997). However, research into the extent to which the visual system can, or does, use shadows to aid object recognition has returned mixed results (see Braje et al., 1998; Braje et al., 2000; Castiello, 2001; & Tarr et al., 1998).

Changes in the illumination of an object impose severe problems upon the object-recognition system. Thus, understanding what the effects of illumination are, i.e., the effects of variations in shading and shadowing, is fundamental to the understanding of object recognition itself. For instance, altering the illumination of a person’s face can result in greater changes to that image than those produced by a change in identity (see Moses, Adini & Ullman, 1994). Thus, when viewing several images, the images of two different objects may be more similar than images of the same object under different illumination conditions.

Tarr et al. (1998) have studied the visual system’s ability to cope with such changes in illumination, when recognising novel objects. Using a sequential matching task, they investigated the effects of the presence of attached shadows and changes in illumination direction. They found that when attached shadows were present, responses were more accurate, and response times faster, than when they were not present. This result occurred irrespective of illumination direction. Furthermore, there was also a response-time cost when the direction of illumination of the sample stimulus differed to that of the comparison stimulus, but only when attached shadows were present. Tarr et al. (1998) concluded that the shadows aided recognition by providing information about the three dimensional structure of the objects.

This finding, of decreased response latencies in the presence of shadows, is in accordance with the results of a small number of experiments that have found the presence of shadows to be beneficial to object recognition. Freeburg (1966) found an interaction between shadow presence and task difficulty, in the sequential matching of textured surfaces: moderate levels of shadowing were benefited accuracy at moderate levels of discrimination difficulty. Castiello (2001) also found a reaction time benefit due to the presence of cast shadows in the recognition of familiar objects. In contrast, no benefit of shadows to recognition has been found in investigations of face recognition (Braje, Kersten, Tarr, Troje & Nikolaus, 1998; Braje, Kersten & Troje, 1996), recognition of natural objects (fruit) (Braje, et al., 2000), and in judgments of the slant and tilt of ellipsoids (Mingolla & Todd, 1986).

This apparent dichotomy in research findings could be resolved by asking "When are shadows useful for object recognition?" instead of the usual question of "Does the visual system use shadows to aid object recognition?". Tarr et al. (1998) contrasted their results, using novel stimuli, with those of Braje et al. (1998), using faces, and suggested that familiarity may play a role in whether shadows are valuable as cues for recognition. This suggestion provides a testable hypothesis: that shadows may be of use in determining the 3-D structure of novel stimulus, but that as we become increasingly familiar with a stimulus, our visual system weights other cues (e.g., configuration or bounding contours) as more salient than shadows, removing any benefit derivable from the shadows' presence. Braje et al. (2000) state a similar hypothesis "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" (p. 396).

Tarr et al. (1998) stated that the novel objects they used had completely unknown 3-D shapes, and because of this, the extra information available from the shadows, about the shape of the objects, may have been beneficial for recognition. They suggested that recognition of novel and familiar objects may be conducted in different manners, where familiarity with the 3-D shape of an object may reduce any benefit available from shadows. Gauthier, Williams, Tarr, and Tanaka (1998) have suggested that we may make greater use of configural information with increasing familiarity with previously novel objects. The increased use of configural information, instead of particulate features, could negate any benefit that shadows provide in the case of novel object recognition.

However, the results of Castiello (2001), who tested participants' recognition latencies of familiar objects (e.g., fork, mug, tennis racket), suggest that familiarity may not be a factor in the visual system's usage of the information provided by shadows. When using these familiar objects Castiello found that the time required to identify objects correctly was longer for objects without a cast shadow, than for objects with a congruent cast shadow and a congruent attached shadow. However, Castiello used objects that his participants were already expected to be familiar with. Thus, he did not assess the effect of the participants' familiarity with the stimuli. Furthermore, Castiello's methodology differed to that of Tarr et al.; Castiello's control trials, which had no cast shadows, still contained attached shadows, whereas Tarr et al. (1998) compared attached-shadow conditions with no-attached-shadow conditions.

Experiment 1

The first series of experiments in this thesis was aimed at investigating whether familiarity with objects moderates any potential benefit of the presence of shadowing. The first experiment employed similar methodology to that of Tarr et al. (1998). The rationale for conducting this experiment was to demonstrate control of the effects (upon object recognition) of the presence or absence of attached shadows when recognising novel objects.

Achieving experimental control would allow for three avenues of research. First, the investigation of the amount of attached shadowing required for an effect upon object recognition. Second, the parametric investigation of whether familiarity is a variable influencing the effect of attached shadows upon object recognition. Third, the investigation of what information shadows provide that may be of benefit to us in object recognition. Specifically, whether shadows provide useful information about the 3-D structure of objects, or whether they operate in a manner similar to 2-D pattern matching, which has been suggested to be the basis of shape-from-shading (Liu, Collin, & Chaudhuri, 2000).

Based upon Tarr et al.'s findings it was hypothesised that: 1., the Shadow Condition (described below) would produce lower reaction times than the No-Shadow Condition; and 2., that an effect of illumination direction change would be evident in the Shadow Condition, where when S1 and S2 have different illumination directions reaction times would be slower than if they were the same.

Method

Participants

Twenty-four Psychology Department undergraduates from the University of Waikato participated in the experiment for course credit. They were randomly assigned to the two experimental conditions, the No-Shadow Condition and the Shadow Condition.

Apparatus

The images used were accurate copies of the images presented in Tarr et al. (1998). Twelve 3-D objects were generated using *3D Studio Max,* on a Pentium II 400 MHz computer, from which the images were rendered. Tarr et al. gave four distinctive properties of their objects: they were novel; produced intrinsic (attached) shadows; were rendered with uniform albedo; and were illuminated from either the left or the right. There were twelve different objects in total, comprised of two versions of each of six object body types, each version sharing the same shape appendages as its pair, but in a different arrangement. The six object body types were all qualitatively different geometric volumes (Tarr et al.). Reference images of all the objects are presented in Appendix 1, when compared to the images presented in the paper by Tarr et al. (1998), they are nearly identical.

Two light sources, an ambient light, and a spot light, illuminated all the objects. The ambient lighting was set at RGB levels of 38, 38, 38 with hue and saturation at 0. The spot light was set at RGB levels of 180, 180, 180 with hue and saturation at 0, and an overall multiplier of 1.5. The light was positioned either 30° to the left or the right of the front elevation and 40° above the plane the objects sat upon, at a distance approximately 4 times the height of the objects. The objects were rendered from a camera angle of 25° above the midpoint of the objects at a distance of approximately 3 times the height of the objects. A 50 mm lens was used, and similar to Tarr et al. (1998) orthographic projection and ray tracing of shadows were employed. Each object was rendered against a white background under four lighting conditions, from the left with cast shadows, from the left without cast shadows, from the right with cast shadows, and from the right without cast shadows. Although the objects were rendered in a mid grey colour, RGB values of 200, 200, and 200 respectively, the rendered bitmaps were converted to greyscale in PhotoShop. This reduced each bitmap's file size but produced no visible changes in the images.

The images were presented on Dell Optiplex GX 1, PIII 450 MHz computers and 43 cm Trinitron screens (actual screen size 15.9" diagonal) with a horizontal scan rate of 75 hertz. The images for each trial were pre-cached in the experimental program at the start of each trial to prevent any lag in their presentation due to loading time. The visual angle of the images was 14.3° (vertical) by 18.9° (horizontal) at a 0.6 m viewing distance. The program presenting the images recorded the Participants' responses and reaction times. The exact distance the participants sat from the screen was not controlled, but was approximately 0.6 m.

Procedure

The presentation times used were the same as those employed by Tarr et al. (1998). A sequential matching to sample procedure was used (shown in Figure 1.10). Each trial was comprised of the sequential presentation of five images. The first image was a small white fixation cross on a black background, displayed for 750 ms. The second image was of the sample stimulus, displayed for 200 ms. The third image was of a masking stimulus, displayed for 750 ms. The masking stimulus was a composite of the parts of the different stimuli. The fourth image was of the comparison stimulus, presented for 100 ms. The fifth image was the masking stimulus again, presented for 500 ms. Each participant's task was to determine whether the object depicted in the comparison stimulus was the same as that in the sample stimulus. Participants responded via the keyboard pressing either "Q" to indicate "Yes" the objects were the same, or "P" to indicate "No" the objects were different. Labels were provided on the monitor to remind the participants of the appropriate key for the "Yes" and "No" responses. The participants were asked to respond as quickly as possible. The participants could respond from the onset of the presentation of the comparison stimulus, until 1500 ms after it disappeared. If they did not respond within this time, their response was not recorded on that trial. Each participant's reaction time and key press were recorded for each trial. The presentation order of the different trials was generated in a quasi-random fashion so that the same initial stimulus could not appear more than two times in a row. The presentation order was the same for all the participants.

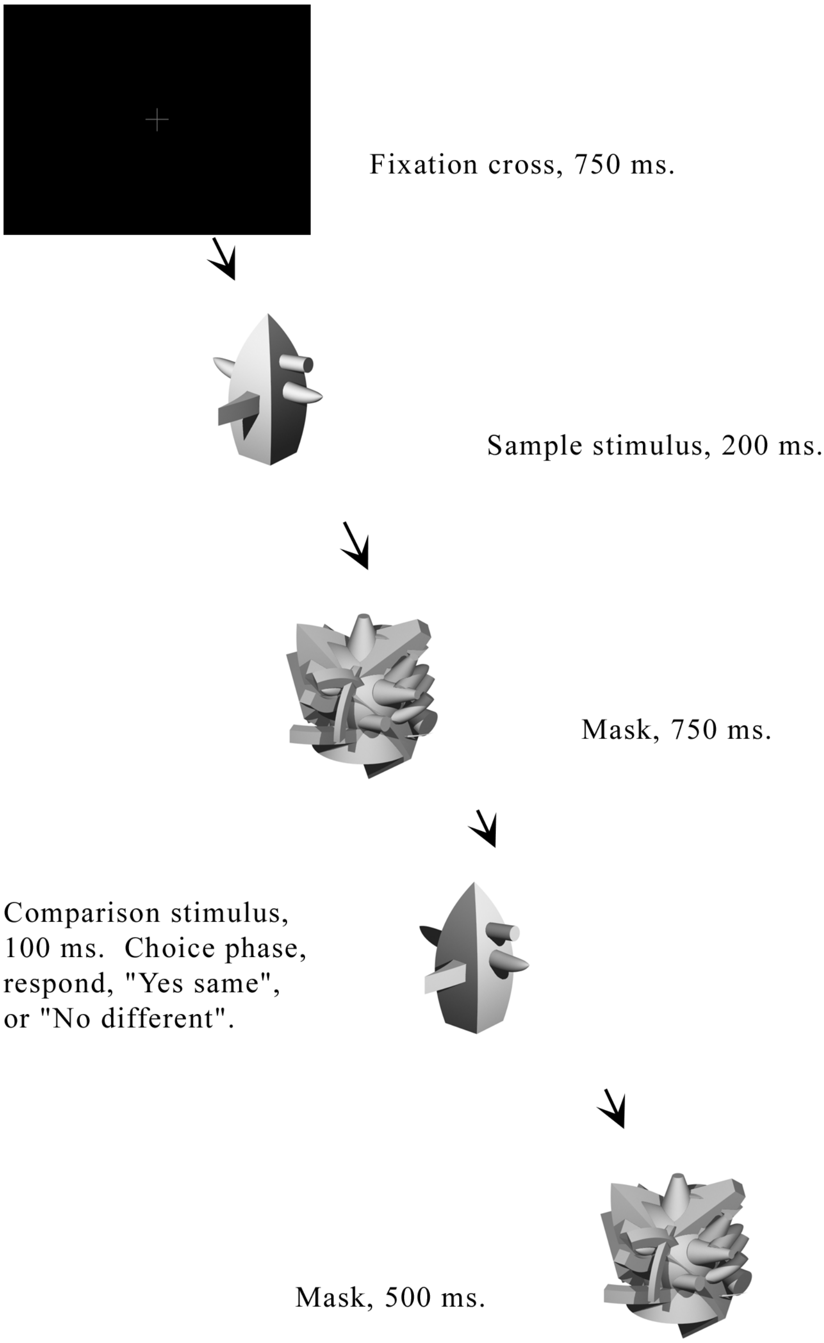

Figure 1.10. Example of the sequential matching to sample procedure used in Experiment 1.

In total there were 288 trials presented to each participant. Each of the twelve objects was presented as a sample stimulus 24 times, half of these presentations showed the object illuminated from the right and half from the left. On half of these trials the object was paired with itself, equally often in left and right illumination directions, and on the other half the object was paired with its alternate version, also equally often in left and right illumination directions. Thus, for any given trial, the sample and comparison stimuli could differ on illumination direction and/or object version.

Instructions were given to the participants verbally before the participants started the experiment and on the computer screen before they could commence the experiment. The instructions, as given on the computer, are provided in Appendix 2. The participants could advance through the trials at their own rate by starting the next trial with the press of the spacebar key. No feedback[1] was given on whether the participants had pressed the correct key for a given trial. The participants took about 25 minutes to finish the experiment.

1. Tarr et al. (1998) employed feedback, presumably to encourage learning about the stimuli to aid discrimination in the latter stages of the sessions. Feedback was not employed in this experiment. Using feedback to promote learning about the stimuli is contrary to the purpose of using novel objects. Furthermore, there were usually several participants completing the experiment at one time in the same computer laboratory, and using audible feedback beeps could have been distracting from the task.

Results

Twenty-three of the 24 participants' results were analysed. The other participant responded solely on one key, with reaction times often considerably faster than the presentation time of the second stimulus. It was concluded that the participant was not performing the sequential matching task and thus their results were excluded. As a result there were 12 valid data sets for the Shadow-Condition group and 11 valid data sets for the No-Shadow-Condition group.

Tests of the Experimental Hypotheses

Accuracy on the matching task was evaluated using the signal detection measure *d*-prime (*d*'), as was used by Tarr et al. (1998). The measure *d*' is the difference between the *z*-scores of the probability of incorrectly reporting the presence of a stimulus (a "false alarm") and the probability of correctly reporting the presence of a stimulus (a "hit", see: the signal detection matrix, Appendix 3 Figure A3.1). Thus, *d*' is a standardised score that can be compared across experiments (*z*-scores use the standard deviation as the single unit, therefore, *z*-scores and *d*' are comparable across distributions). The logic behind *d*' is that it reflects the probability of correctly responding that a stimulus is present (a hit) while taking into account the probability that the "present" response would occur even if the stimulus was absent (false alarm).

The No-Shadow group's *d'* value was 1.27, compared to the Shadow group's value of 1.30, see Table 1.1. The difference between the two groups was assessed using an independent groups *t*-test. The difference was not significant ($t(21)= -0.1074$, $p>0.05$).

Correct Responses Only

The usual analysis of response times is the analysis of latency to correct responding. The term "object recognition" infers correct categorisation, therefore, analysing the latencies to correct responding will show whether the experimental manipulation has produced effects with respect to recognition, as opposed to effects with respect to responding in general.

For the analysis of correct responses, differences in the latencies of correct responses, across the two experimental groups, No-Shadow and Shadow, were evaluated with a within-subjects ANOVA, with experimental group as a between-subjects factor (see Appendix 4, Table A4.1). Illumina-

Table 1.1.

Results of Experiments 1 to 4, Comparisons of d' in the No-Shadow Condition and Shadow Condition, Using Independent Groups t-tests.

	Experiment 1	Experiment 2	Experiment 3	Experiment 4
No Shadows *d*'	1.2673	1.4058	2.3064	1.9315
Shadows *d*'	1.3008	1.3658	1.8586	1.7415
t	*d.f.* (21) -0.1074	*d.f.* (22) 0.1479	*d.f.* (26) 2.5569*	*d.f.* (24) 0.673
η_p^2				

Note: Any differences that are significant at an alpha level of 0.05 are indicated by an asterix (*).

tion direction change between the two stimuli in each trial, and object version change between the two stimuli in each trial were the two within-subjects factors. The main effects of experimental group ($F(1,21) = 0.034$, $p>0.05$)and illumination change ($F(1,21) = 0.000$, $p>0.05$) were not significant. The main effect of change in the versions of the objects presented in a trial was significant ($F(1,21) = 14.351$, $p < 0.05$, $\eta_p^2 = 0.406$). Responses in the trials where there was no change in the object from initial stimulus to comparison stimulus were on average faster than when there was a change, 610 ms compared to 687 ms. None of the interactions were significant: group and illumination direction ($F(1,21) = 0.447$, $p > 0.05$); group and object version ($F(1,21) = 0.611$, $p > 0.05$); illumination direction and object version ($F(1,21) = 1.824$, $p > 0.05$); and illumination direction, object version and group ($F(1,21) = 0.001$, $p > 0.05$).

The analysis of the latencies to correct responding (recognition) did not show a significant effect of shadow presence, unlike the result of Tarr et al. (1998). The *d*' analysis indicated that error rates were high in this experiment, therefore, an analysis of correct responses only dealt with a moderate proportion of the data set. To see if the presence of shadows had a general effect upon speed of responding (irrespective of accuracy) the same analyses were performed on the entire data set (correct response and incorrect responses). Lastly, the latencies to incorrect responding were analysed to see if shadow presence produced an effect when the participants answered incorrectly.

All Responses (Correct and Incorrect Responses)

Differences in response times, across the two experimental groups, No-Shadow and Shadow, were evaluated with a within-subjects ANOVA, with experimental group as a between-subjects factor (see Appendix 4, Table A4.1). Illumination direction change between the two stimuli in each trial, and object version change between the two stimuli in each trial were the two within-subjects factors. The main effects of experimental group ($F(1,21) = 0.001$, $p > 0.05$) and illumination change ($F(1,21) = 1.682$, $p > 0.05$) were not significant. The main effect of change in the versions of the objects presented in a trial was significant ($F(1,21) = 8.043$, $p < 0.05$, $\eta_p^2 = 0.277$). Responses on the trials where there was no change in the object from initial stimulus to comparison stimulus were on average faster than when there was a change, 620 ms compared to 654 ms. None of the interactions were significant: group and illumination direction ($F(1,21) = 1.038$, $p > 0.05$); group and object version ($F(1,21) = 0.113$, $p > 0.05$); illumination direction and object version ($F(1,21) = 0.130$, $p > 0.05$); and illumination direction, object version and group ($F(1,21) = 0.107$, $p > 0.05$).

Incorrect Responses Only

Differences in the latencies of incorrect responses, across the two experimental groups, No-Shadow and Shadow, were evaluated with a within-subjects ANOVA, with experimental group as a between-subjects factor. Illumination direction change between the two stimuli in each trial, and object version change between the two stimuli in each trial were the two within-subjects factors. The main effects of experimental group ($F(1,21) = 0.002$, $p > 0.05$) and illumination change ($F(1,21) = 0.002$, $p > 0.05$) were not significant. The main effect of change in the versions of the objects presented in a trial was significant ($F(1,21) = 6.263$, $p < 0.05$, $\eta_p^2 = 0.230$). Responses in the trials where there was a change in the object from initial stimulus to comparison stimulus were on average faster, 641 ms compared to 696 ms, than when there was no change, i.e., responses were faster when the stimuli were different and the participants incorrectly responded that they were the same. This is in contrast to the finding for correct responses (i.e., when the stimuli were the same and the participants responded that they were the same).

Summary of Accuracy and Response Times

To summarise, the analyses of accuracy and response times indicated that there were no significant differences between the two experimental groups. Over both conditions, the participants' response latencies were faster when they responded that the two stimuli were the same, than when they responded that they were different, irrespective of whether they were correct or not. None of the interactions were significant: group and illumination direction; group and object version; illumination direction and object version; and illumination direction, object version and group.

Analysis of Bias Towards "same" Responses

Given there were differences in the participants' response times to the same object trials, and the different object trials, bias towards saying that the images were the same was measured. Bias towards saying "yes the objects were the same" was calculated using the measure log *c*[1].

The measure of bias log *c* is complementary to the discrimination measure log *d*, which is comparable to measures of discriminability derived from signal detection theory (Johnstone & Alsop, 1996). White and Wixted (1999) report that log *d* is linearly related to *d*', and that it satisfies the requirement of Macmillan and Creelman (1991), that both hits and false alarms contribute to the discrimination measure. Log *d* (discrimination) is based upon the ratio of correct to error responses following each sample, log *c* (bias) is based upon the ratio of stimulus "A" responses to stimulus "B" responses following each sample stimulus (White & Wixted). Signal detection theory is dependent upon the assumption that the response an individual makes is determined by evaluation against a decision criterion, which is set at a point along the stimulus continuum (White & Wixted). This sets up an inconsistency in the explanation of the cause of a response: the criterion value (which is response bias, see: Macmillan & Creelman, 1991) is the basis for the individual determining which response to make, but the response bias (the criterion) is determined from responses that the individual has made. White and Wixted promote the measures log *d* and log *c* as they follow an approach to discrimination and bias that is similar to signal detection theory, but which does not rely upon the theory that response selection is determined by a set decision criterion, or rule.

There is another measure of bias associated with log *c*, termed log *b*. Log *b* is a measure of response bias, and is comparable to the signal detection measures of response bias such as the criterion, *c*, and log β (Johnstone & Alsop, 1997). Log *c* is termed intrinsic bias (Davison & McCarthy, 1988), and is the portion of response bias (log *b*) that is not explained by any bias due to differences in the frequency of consequences scheduled for correct responding to the two stimuli types (Davison & McCarthy, 1988). Therefore, when specific consequences are not provided to a participant in a matching task, log *c* will equal log *b*, and both the measures are comparable to the criterion, *c*, but they do not require the assumption of a decision criterion.

In Experiment 1, the average bias towards saying the two objects were the same was log *c* = 0.42, *SD* = 0.28, for the Shadow Group and log *c*

1. See Appendix 3 for further discussion of the measure of bias log *c*, and the associated discrimination measure log *d*.

= 0.36, *SD* = 0.33 for the No-Shadow group. An extreme bias would be log *c* = 2.46. Therefore, there was a small bias present towards saying that the objects were the same, as well as faster response times for the "Yes same" response.

Within-Subject Analyses

Latencies

To assess if the between groups analysis was obscuring any systematic differences between the conditions, the participants' latencies to responding were plotted individually. This was done by combining all trial types, and for same object trials and different object trials, and shown in Figures 1.11-1.12, 1.13-1.14 and 1.15-1.16 respectively. The plots revealed that the distributions of response latencies were idiosyncratic and did not systematically differ across the two experimental conditions.

Learning and Fatigue Effects

To assess whether the participants were learning about the task, or suffering from fatigue over the course of the experimental session, each participant's cumulative error rate was plotted, as presented in Figures 1.17 and 1.18. An increase in error rate towards the end of the session would indicate fatigue, and a decrease in error rate would indicate learning. The error rates of some of the participants increased towards the end of the experiment, e.g., Participants 1, 7, 16 and 22. These plots also show that for most of the participants error rates were very high, often close to 50%, e.g., Participants 6 and 9. The error rates of some of the participants, e.g., Participants 1, 3, and 22, were lower in the second two thirds of the experiment than in the first. However, the error rates of other participants, e.g., Participants 9, 10, 18 and 21, remained constant and high throughout the course of the experiment. Overall, error rates were idiosyncratic and did not systematically differ across the two conditions.

Within-Group Analyses

Image Change Due to Shadow Presence

Having assessed differences between the two experimental conditions, shadow versus no-shadow, and found none, a within group analysis of the Shadow Condition results was conducted. To do this, the difference between the No-Shadow-Condition images and the Shadow-Condition images, due to shadowing was calculated. The greyscale values for all the pixels in an image were summed to give a total value for that image (excluding the white

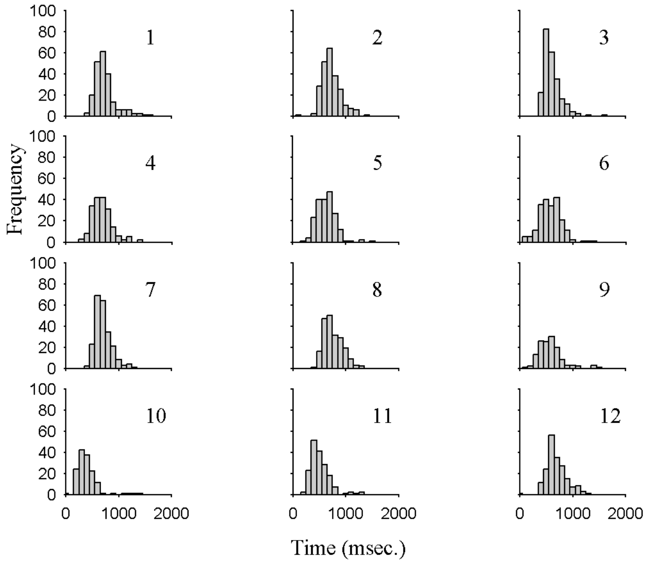

Figure 1.11. Individual histograms of reaction time for Participants 1 to 12, the Shadow Group, taken from all trial types.

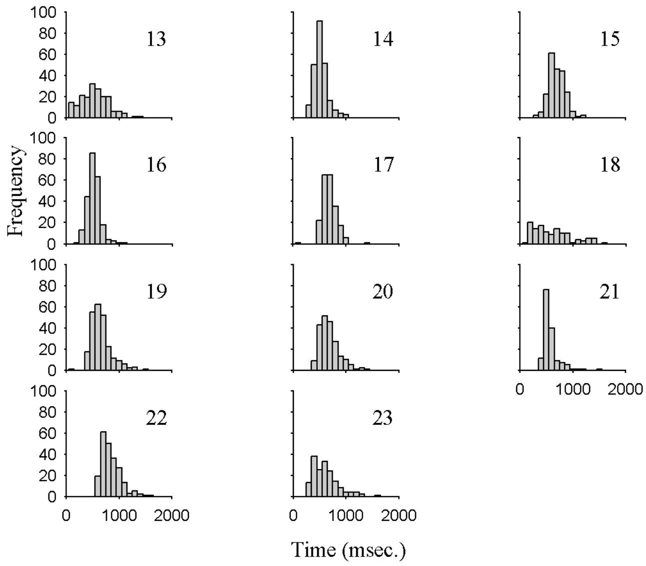

Figure 1.12. Individual histograms of reaction time for Participants 13 to 23, the No-Shadow Group, taken from all trial types.

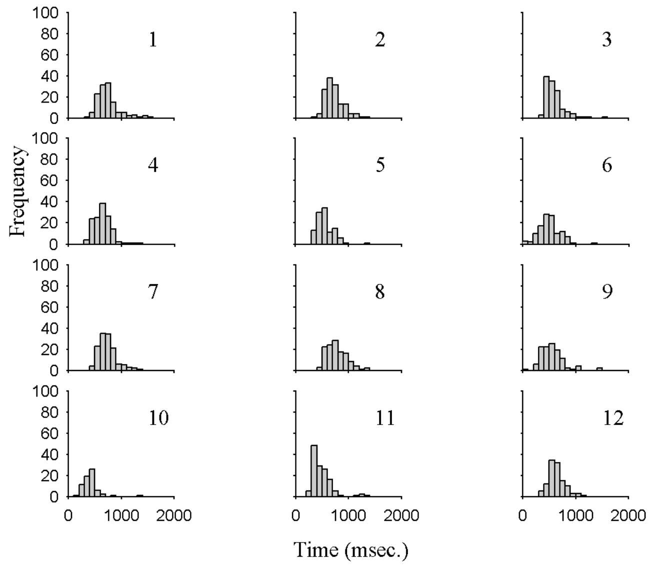

Figure 1.13. Individual histograms of reaction time for Participants 1 to 12, the Shadow Group, taken from trials in which the initial and comparison stimuli were the same.

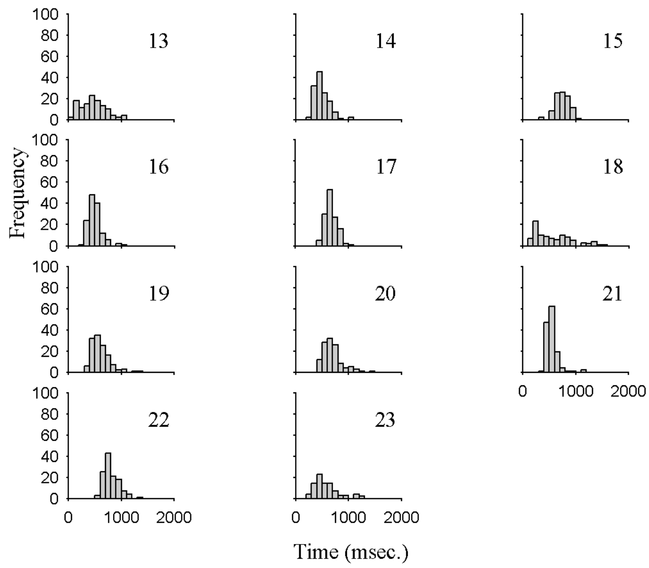

Figure 1.14. Individual histograms of reaction time for Participants 13 to 23, the No-Shadow Group, taken from trials in which the initial and comparison stimuli were the same.

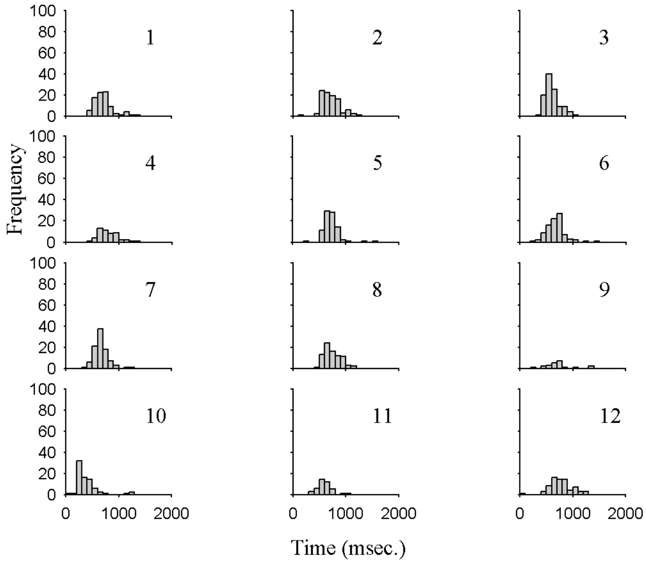

Figure 1.15. Individual histograms of reaction time for Participants 1 to 12, the Shadow Group, taken from trials in which the initial and comparison stimuli were different.

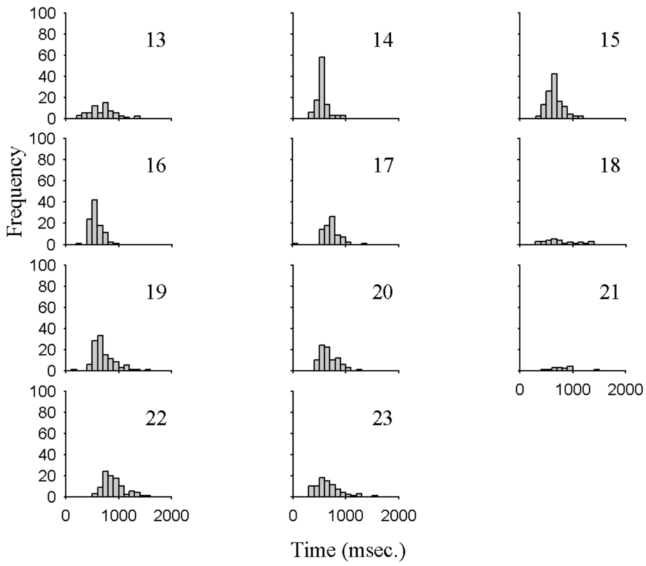

*Figure 1.16. I*ndividual histograms of reaction time for Participants 13 to 23, the No-Shadow Group, taken from trials in which the initial and comparison stimuli were different.

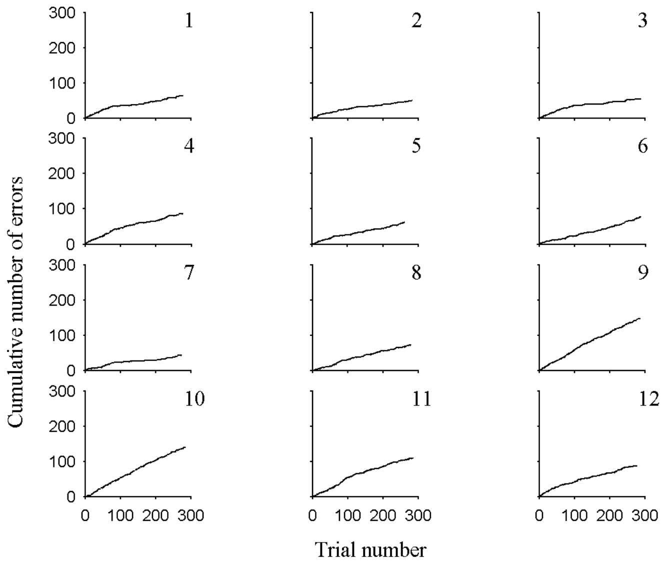

Figure 1.17. Plots of cumulative error rate for Participants 1 to 12, the Shadow Group. There were 288 trials, so a cumulative total of 144 errors would indicate overall performance was at chance levels. Learning is indicated by a reduction in error rate of as the number of trials increases, and conversely, fatigue is indicated by a rise in error rate.

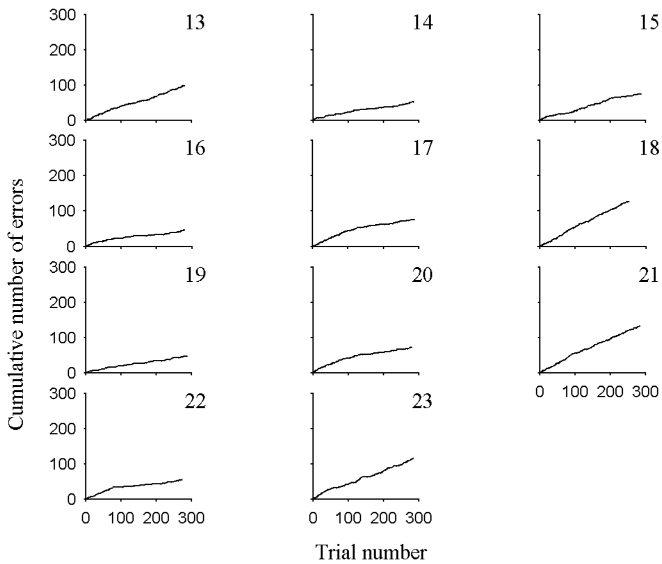

Figure 1.18. Plots of cumulative error rate for Participants 13 to 23, the No-Shadow Group. There were 288 trials, so a cumulative total of 144 errors would indicate overall performance was at chance levels. Learning is indicated by a reduction in error rate as the number of trials increases, and conversely, fatigue is indicated by a rise in error rate.

background, i.e., all values of 255, there were no white pixels in the objects themselves). As the only difference between the No-Shadow-Condition images, and the Shadow-Condition images was the presence of the cast shadows, the difference between the totals for the corresponding images in the two conditions gave a measure of the degree of shadowing in each Shadow-Condition image (see Table 1.2).

Correlation of Amount-of-Shadow With Response Latencies and Proportion Correct

After calculating the amount of shadow contained in each of the Shadow-Condition images, see Table 1.2, the proportion of the image that was in shadow was correlated with the number of correct responses and the average latency for each image. There was no correlation between the amount of shadow in the Shadow-Condition images and the number of correct responses for those images ($r_s(24) = -0.151$, $p > 0.05$), or between the amount of shadow and response latencies ($r_s(24) = -0.164$, $p > 0.05$).

Differential Task Difficulty by Object Type

Responses of some of the participants during debriefing indicated that some of the object pairs may have been more difficult to discriminate between than others. To assess this, the number of incorrect responses for each of the 24 object pairings were compared. A repeated measures ANOVA was conducted on the number of incorrect responses, with object pair as a within-subjects factor, and Shadow Presence as a between-subjects factor. There was a significant difference in number of corrects by object pair ($F(23,483) = 25.219$, $p < 0.05$, $\eta_p^2=0.546$), but no effect of Shadow Presence ($F(1,21) = 0.013$, $p > 0.05$), or interaction between Shadow Presence and object pair ($F(23,483) = 0.606$, $p > 0.05$). Figure 1.19 illustrates the differences.

The pairings of Objects 51 and 52, and 52 and 51, showed the highest number of incorrect responses, with some participants getting all these trials incorrect. The average number of incorrect responses was approximately 9 out of the possible 12. These pairs can be contrasted with the pairings of 61 and 62, and 62 and 61, where the average number of incorrect responses was approximately 3 out of the 12. The other trend observable in Figure 1.19 is the difference in number of incorrect response between the same version trials and the different version trials. More incorrect responses were recorded on the different version trials, paralleling the participants' bias towards saying the versions were the same.

Table 1.2

Amount of Shadow in Each Image, and Grouped Into Four Levels of Shadowing

Object version and illumination direction*	Shadow image total as a percentage of No-Shadow image total	How many percent darker the shadow image is when compared to the No-Shadow image	Shadow level by quartiles of "percent darker" column.
111	94.21%	5.79%	4
112	97.76%	2.24%	2
121	96.53%	3.47%	3
122	96.80%	3.20%	3
211	98.62%	1.38%	1
212	97.05%	2.95%	3
221	97.62%	2.38%	2
222	98.24%	1.76%	2
311	98.27%	1.73%	1
312	96.47%	3.53%	3
321	97.64%	2.36%	2
322	96.22%	3.78%	4
411	97.92%	2.08%	2
412	97.27%	2.73%	3
421	95.65%	4.35%	4
422	97.41%	2.59%	3
511	91.44%	8.56%	4
512	99.82%	0.18%	1
521	93.74%	6.26%	4
522	98.78%	1.22%	1
611	98.48%	1.52%	1
612	98.26%	1.74%	2
621	99.38%	0.62%	1
622	93.65%	6.35%	4

**The first value indicates the object (1-6), the second value indicates the version (1 or 2), and the third value indicates the illumination direction (1 left or 2 right).*

Note 1: For any pixel, black takes a value of 0 and white takes a value of 255 (although the white background was excluded). Thus, the higher the total value, the lighter the image, and the No-Shadow Condition image will return a higher total value.

Note 2: The shadow levels were based upon the interquartile range of the difference between the No-Shadow Condition and the Shadow Condition image summed pixel values.

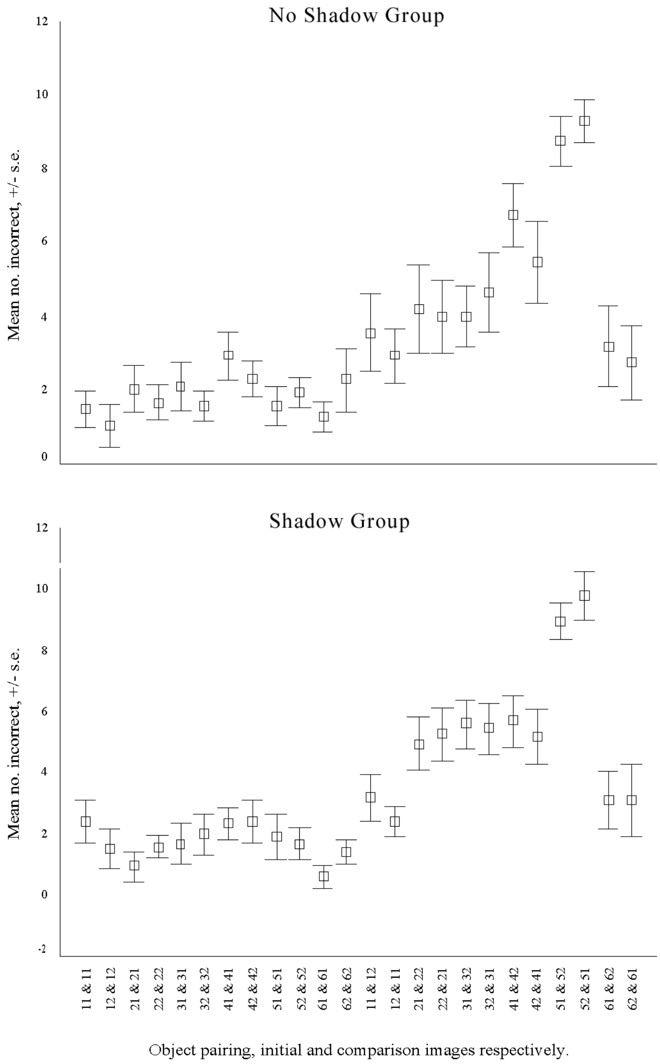

Figure 1.19. Mean number of incorrect responses by object pairing, averaged across all participants (bars represent +/- standard error). A repeated-measures ANOVA was conducted on the number of incorrect responses, with object pair as a within-subjects factor, and shadow presence as a between-subjects factor.

Results Summary

The major finding was that there was no difference between either the response latencies, or discrimination (as measured by *d*'), for the Shadow and No-Shadow groups. Response latencies differed between trials where the same versions of an object were presented, and those where different versions of the object were presented. This was a result of the participants responding “Yes, Same” faster then “No, Different” irrespective of whether it was a same image trail or a different image trial. Trials that presented the same version of the objects were, on average, completed faster. There was no interaction between this main effect and the experimental group the participants were in.

Post-hoc analyses of the shadow images revealed that there were large differences within the image set in terms of the amount of shadow present in the images. The amount of shadow present in the images was not found to be correlated with the number of correct responses, or response latency.

Discussion

The hypothesis of this experiment was that the Shadow and No-Shadow groups would differ in terms of average latency to responding. Based on Tarr et al.'s (1998) findings, it was predicted that the Shadow group participants would respond faster. Contrary to this expectation, the average latencies were the same for each group (633 ms) when rounded to zero decimal places. These latencies were considerably faster than those reported by Tarr et al. (1998) which ranged from 827 ms in their *no illumination change, different object, with shadows, condition*, to 967 ms in their *illumination change, same object, no shadows, condition*.

The latencies reported for Experiment 1 are only slightly longer than those reported by Miller and Low (2001), who conducted a study on reaction times using a standard computer keyboard, and, as in this experiment, the participants were students at a New Zealand University. Miller and Low tested a simple one stimulus task, a go/no-go task, and a choice task.

The choice task used by Miller and Low (2001) is reasonably analogous to that used here: the participant was required to select the hand to respond with at the onset of the second stimulus. The results of Experiment 1 fall in the upper range of reaction times found by Miller and Low. Miller and Low cued an initial response, and then provided another stimulus to indicate whether the participant would have to respond with their other hand. On trials where the initially cued response did not change, reaction times averaged 441 ms at 97.8 percent correct. When the initially cued response did change, reaction times averaged 566 ms at 92 percent correct. Miller and Low suggest that it is reasonable to suppose that the fastest responses in an uncued-choice task would be intermediate in value to those found for their validly-, and invalidly-, cued choice trials.

The participants' error rates in Experiment 1 were very high, often near chance (d'(Shadow, S) = 1.30 (2 d.p.) & d'(No Shadow, NS) = 1.27 (2 d.p.)). In comparison, Tarr et al. (1998) obtained high levels of discrimination (d' ranging from 2.35 to 2.64). In sequential matching tasks, the presentation durations of the stimuli are used to maintain performance below a ceiling level of accuracy. The presentation times used here were identical to those reported by Tarr et al., S1 = 200 ms and S2 = 100 ms. However, given the differences in computer hardware and software, it is possible that the presentation times were not effectively the same in both experiments. Thus, the differences in accuracy between the results reported by Tarr et al., and those reported here may have been due to differing effective presentation durations. Modifying presentation durations could also provide a means of increasing the accuracy of participants in further experiments in this series.

The images were presented for 200 ms and 100 ms for the initial stimulus (S1) and the comparison stimulus (S2) respectively. A review of similar object-recognition procedures indicated that these presentation times were relatively short, but were in the range usually used, e.g., presentation times range from 50 ms display times of lowpass and highpass faces (Schyns & Oliva, 1999) to 450 ms display times of faces for MRI scans (Kanwisher, Tong, & Nakayama, 1998). However, when viewing rendered stimuli that are very similar to each other, longer durations may be required if the participants are to discriminate between them. Biederman and Bar (1999) state that in their pilot testing, it was clear that exposure durations of 200 ms for S1 and 100 ms for S2 (that were sufficient for line drawings) were insufficient for clear perception of rendered stimuli. They employed durations of 400 ms for S1 and 300 ms for S2.

In a priming procedure, Stankiewicz, Hummel, and Cooper (1998) note that they chose their prime stimuli to last less than 200 ms so that their participants could not perform a saccade. If a duration of 200 ms or less does not allow for a saccade to be made, the short duration times used in Experiment 1 would have made it impossible for the participants to search the images visually and to use this information to base their discriminations upon.

Given that no saccades were possible, the size of the image could also have affected discriminability. At any given distance as an image is increased in size the amount of the image that is viewed with the fovea is decreased. Tarr et al. (1998) used an image size of 5.7° by 5.7° at approximately 0.6 m (about 60 mm by 60 mm on the screen). The images used in Experiment 1 were 14.3° (vertical) by 18.9° (horizontal) at a 0.6 m viewing distance, about 150 mm by 200 mm on the screen. Thus, the images used in Experiment 1 were approximately 8.3 times larger (in overall area) than those used by Tarr et al.

Hecht (1998) reports the dimensions of the anatomy of the human eye, from which it can be calculated that foveal vision covers approximately 1° of visual angle. Thus, in Tarr et al.'s (1998) experiment, the size of the image on the retina would have been about 5.7 times the size of the foveal visual region, whereas in this experiment the image would have been about 14.3 times the size of the foveal visual region. As a result, much less of the image would have been viewable through the foveal region without making a saccade.

In retrospect, it is possible that either the length of presentation of the stimuli, or the size of the stimuli, may have resulted in the discrepancy between the results obtained from this experiment and those obtained by Tarr et al. (1998). These factors will be addressed in the two following experiments. In Experiment 2, the effect of reducing the size of the stimuli on latency to responding was assessed. The size of the stimuli was set to that used by Tarr et al., and the results are compared to those obtained in Experiment 1. In Experiment 3, the effect of increasing the stimuli's presentation time was assessed. The presentation times used were 400 ms (S1) and 300 ms (S2), the times given by Biederman and Bar (1999) as appropriate for rendered 3-D objects. Again the results are compared with those obtained in Experiment 1.

Experiment 2 Reduced Stimulus Size

The first experiment in this series, modelled after Tarr et al. (1998), was based on a between-groups sequential-matching procedure. The participants in the experimental group were shown images of novel objects with attached shadows, while those in the control group were shown images of the same objects, sans the attached shadows. The two groups did not differ in terms of the participants' response latencies or their accuracy in performing the matching task.

These results contrast with those found by Tarr et al. (1998). They obtained faster response latencies overall for the Cast-Shadow Group than the No-Shadow Group, and also found a response latency cost, for the Cast-Shadow Group, when the direction of illumination was different for the initial and comparison stimuli. The participants' levels of accuracy on the matching task were also much higher in Tarr et al.'s experiment than in Experiment 1 of this series.

It was hypothesised that the discrepancy between the two sets of results may have been a result of the different sized images presented to the participants. The images used in Experiment 1 were about 2.5 times larger (taller) than those used by Tarr et al. (1998). To counter this, the image size in this experiment was reduced so that its height was the same as that used by Tarr et al. (i.e., 5.7°). It was expected that the reduction in the size of the images would result in an increase in *d'* values for both the Shadow and No-Shadow groups, and that response latencies would be faster for the Shadow group.

Method

Participants

Twenty-four Psychology Department undergraduates from the University of Waikato participated in the experiment for course credit. They were randomly assigned to the two experimental conditions, No-shadows and Attached-shadows.

Apparatus

The apparatus was the same as that used in Experiment 1, except the images were 5.7° (vertical) by 7.6° (horizontal) at a 0.6 m viewing distance. On a 43 cm monitor, this was approximately 60 mm by 80 mm on the screen. The exact distance the participants sat from the screen was not controlled.

Procedure

A sequential matching procedure was used. It was the same as the procedure used in Experiment 1, except for the change in image size noted above.

Results

Tests of the experimental hypotheses

Discrimination

The participants' discriminatory performance was evaluated using the measure *d*-prime (*d*'). The No-Shadow group's *d*' value was 1.41, compared to the Shadow group's value of 1.37. The difference between the two groups was assessed using an independent groups *t*-test. The difference was not significant ($t(22) = 0.1479$, $p > 0.05$) (see Table 2.1).

Correct responses

Differences in response times of the correct responses across the two experimental groups, No-Shadow and Shadow, were evaluated with a within-subjects ANOVA, with experimental group as a between-subjects factor (see Appendix 4, Table A4.1). Illumination direction change between the two stimuli in each trial, and object version change between the two stimuli in each trial were the two within-subjects factors.

The main effects of experimental group ($F(1,22) = 1.295$, $p > 0.05$) and illumination change ($F(1,22) = 0.185$, $p > 0.05$). The main effect of change in the versions of the objects presented in a trial was significant ($F(1,22) = 18.307$, $p < 0.05$, $\eta_p^2 = 0.454$). Responses in trials where there was no change in the object from initial stimulus to comparison stimulus were on average faster, 692 ms compared to 762 ms, than when there was a change. None of the interactions were significant: group and illumination direction, ($F(1,22) = 0.035$, $p > 0.05$); group and object version ($F(1,22) = 3.283$, $p > 0.05$); illumination direction and object version ($F(1,22) = 0.020$, $p > 0.05$); and illumination direction, object version and group ($F(1,22) = 0.003$, $p > 0.05$).

All Responses (Correct and Incorrect Responses)

Differences in response times across the two experimental groups, No-Shadow and Shadow, were evaluated with a within-subjects ANOVA (see Appendix 4, Table A4.1), illumination direction change between the two stimuli in each trial, and object version change between the two stimuli in each trial were the two within-subjects factors.

The main effect of experimental group was not significant ($F(1,22) = 0.582$, $p > 0.05$), i.e., there was no difference in average response latency between the Shadow and No-Shadow groups. While the main effects of illumination change ($F(1,22) = 7.652$ $p < 0.05$, $\eta_p^2 = 0.258$), and change in the versions of the objects presented in a trial ($F(1,22) = 10.541$, $p < 0.05$,

Table 2.1.

Results of Experiments 1 to 4, Comparisons of d' in the No-Shadow and Shadow Conditions, Using Independent Groups t-tests.

	Experiment 1	Experiment 2	Experiment 3	Experiment 4
No Shadows *d'*	1.2673	1.4058	2.3064	1.9315
Shadows *d'*	1.3008	1.3658	1.8586	1.7415
t	*d.f.* (21) -0.1074	*d.f.* (22) 0.1479	*d.f.* (26) 2.5569*	*d.f.* (24) 0.673
η_p^2				

Note: Any differences that are significant at an alpha level of 0.05 are indicated by an asterix (*).

$\eta_p^2 = 0.277$), were significant. The response latencies were faster when the illumination direction was not changed within a trial, 717 ms compared to 769 ms, and response latencies were faster in trials where there was no change in the object from initial to comparison stimulus, 705 ms compared to 752 ms. None of the interactions were significant: group and illumination direction ($F(1,22) = 0.035$, $p > 0.05$); group and object version ($F(1,22) = 3.283$, $p > 0.05$); illumination direction and object version ($F(1,22) = 0.027$, $p > 0.05$); and illumination direction, object version and group ($F(1,22) = 0.143$, $p > 0.05$).

Incorrect Responses

Differences in response times of the incorrect response were assessed across the two experimental groups, No-Shadow and Shadow, with a within-subjects ANOVA (see Appendix 4, Table A4.1). Experimental group was a between-subjects factor. Illumination direction change between the two stimuli in each trial, and object version change between the two stimuli in each trial were the two within-subjects factors.

The main effects of experimental group ($F(1,22) = 0.123$, $p > 0.05$) and illumination change ($F(1,22) = 0.764$, $p > 0.05$) were not significant. The main effect of change in the versions of the objects presented in a trial was significant ($F(1,22) = 9.305$, $p < 0.05$, $\eta_p^2 = 0.318$). Responses in trials where there was a change in the object from initial stimulus to comparison stimulus were on average faster, 803 ms compared to 748 ms, i.e., responses were faster when the stimuli were different and the participants responded that they were the same. This is in contrast to the latencies found for the correct responses, where responses were faster when there was no change in the object version, i.e., when the stimuli were the same and the participants responded that they were the same. Thus, irrespective of whether they were correct or not, the participants were faster in responding when they responded that the two stimuli were the same, than when they responded that they were different. None of the interactions were significant: group and illumination direction, ($F(1,22) = 0.398$, $p > 0.05$); group and object version ($F(1,22) = 0.645$, $p > 0.05$); illumination direction and object version ($F(1,22) = 0.145$, $p > 0.05$); and illumination direction, object version and group ($F(1,22) = 0.627$, $p > 0.05$).

Comparison of Experiments 1 and 2

When comparing the results of Experiments 1 and 2, the participants in Experiment 1 were faster in their responses than those in Experiment 2 (as shown in Figure 2.1). A One-Way ANOVA was performed on latency of all responses, using the different experiments as the independent variable (Note: this included Experiments 3 and 4). The ANOVA indicated that, the differences in average latency to responding were not significant ($F(3, 97) = 2.130$, $p > 0.05$).

Discrimination, *d'*, was also compared across Experiments 1 and 2 (as shown in Figure 2.2) and a One-Way ANOVA (Note: this included Experiments 3 and 4) returned a significant result ($F(3,97) = 8.524$, $p < 0.05$, $\eta_p^2 = 0.209$). A Scheffe's post hoc test indicated that there was no significant difference between the average *d*' values for Experiments 1 and 2 (mean *d*' values are presented in Table 2.2). Individual plots of response latency distributions, as presented in Experiment 1 are not included here, as they did not differ to those obtained in Experiment 1.

The comparison of the error rates found in Experiment 1, to those found in Experiment 2, i.e., the comparison of Figures 1.17 and 1.18 with Figures 2.3 and 2.4, shows that error rates were similar in the two conditions. Overall, the results did not differ to those from Experiment 1. As in Experiment 1, there were no differences between the two groups, Shadow and No-Shadow on either of the dependent variables, discrimination, or response latency.

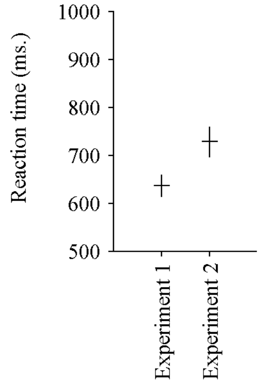

Figure 2.1. Average response latencies across Experiments 1 and 2. All trials (i.e., both correct and incorrect responses) were considered. The horizontal bars indicate the means, and the vertical bars indicate the standard errors of the means.

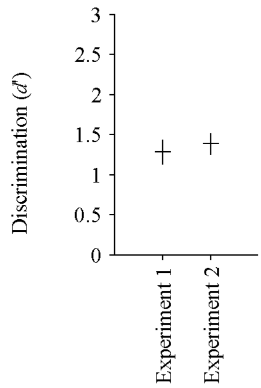

Figure 2.2. Average values of *d*' across Experiments 1 through to 4. The horizontal bars indicate the means, and the vertical bars indicate the standard errors of the means. There was no significant difference between the *d*' values of Experiments 1 and 2.

Table 2.2.

Scheffe's Post-hoc Test on the ANOVA of d' by Experiment. The Homogeneous Subsets of d' Values are Shown. The Average d' Values are Displayed For Each Experiment.

Scheffe[a,b,c]

		Subset		
Experiment	N	1	2	3
1.00	23	1.2848		
2.00	24	1.3858	1.3858	
4.00	26		1.8365	1.8365
3.00	28			2.0825
Sig.		.959	.118	.618

a. Uses Harmonic Mean Sample Size = 25.107.

b. The group sizes are unequal. The harmonic mean of the group sizes is used. Type I error levels are not guaranteed.

c. Alpha = .05.

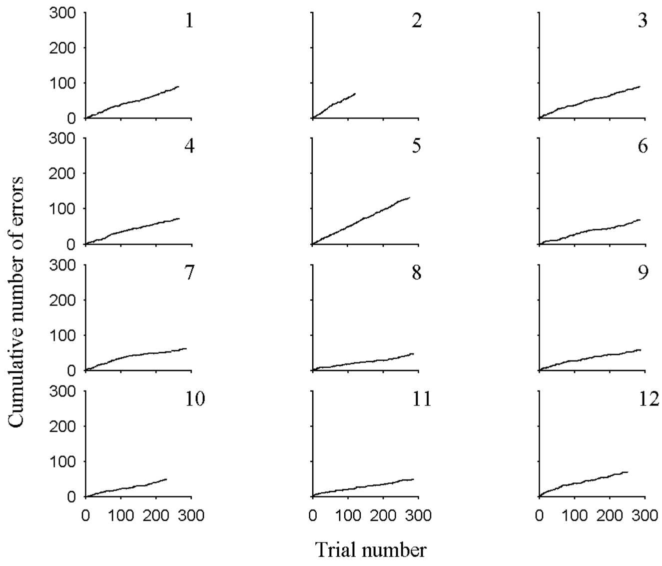

Figure 2.3. Plots of cumulative error rate for Participants 1 to 12, the Shadow Group. There were 288 trials, so a cumulative total of 144 errors would indicate overall performance was at chance levels. Learning is indicated by a reduction in error rate as the number of trials increases, and conversely, fatigue is indicated by a rise in error rate

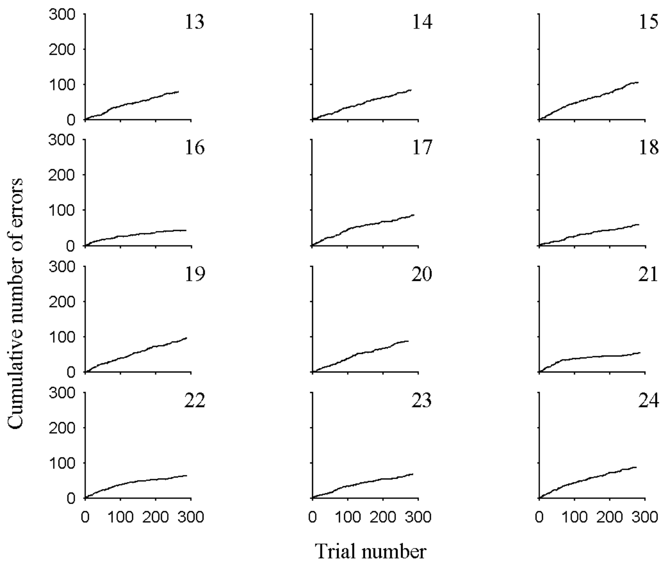

Figure 2.4. Plots of cumulative error rate for Participants 13 to 24, the No-Shadow Group. There were 288 trials, so a cumulative total of 144 errors would indicate overall performance was at chance levels. Learning is indicated by a reduction in error rate as the number of trials increases, and conversely, fatigue is indicated by a rise in error rate.

Discussion

In reviewing Experiment 1 it was hypothesised that a reduction in the size of the stimuli used in the procedure might allow the participants to view more of each of the stimuli in the single fixation that the presenting durations allowed. The decrease in absolute size of the stimuli would mean that the amount of each image to fall on the foveal region of each eye would be increased. It was also hypothesised that when the participants could see more of each image, overall discrimination would be improved and the differences in response latencies between the two experimental groups, Shadow and No-Shadow found by Tarr et al. (1998) may be reproduced.

When this change in procedure was implemented in Experiment 2, the participants' response latencies were, as in Experiment 1, equivalent across the two conditions. This indicated that the information available from the attached shadows in the Shadow Condition was not being used to assist the participants in the discrimination task. At 690 ms (No-Shadows) and 763 ms (Shadows), the response latencies were slightly longer than those obtained in Experiment 1, 653 ms (No-Shadows) and 645 ms (Shadows). The difference between the Shadow Condition and the No-Shadow Condition was not statistically significant, and still faster than those found by Tarr et al. (1998).

Discrimination, as measured by *d'*, remained at similar levels to those obtained in Experiment 1, i.e., the participants' discrimination levels were equivalent across the two experimental groups, and indicated a low level of accuracy. One of the possibilities raised in the discussion of Experiment 1 was that the duration of presentation of the stimuli may have been too short for participants to perform the discrimination well, even though they were matched to the times used by Tarr et al. (1998).

In Experiment 2, the reduction in the presentation size of the stimuli was not associated with any detrimental effects on the measure of discrimination, *d*', so it was concluded that the reduced image size should be retained for the next experiment. Given the possibility that the presentation times used in the first two experiments were too fast to enable accurate discrimination (see Biederman & Barr, 1999), it was proposed that the next experiment should assess the effects of longer presentation times, while employing the smaller stimuli used in Experiment 2.

Experiment 3 Extended Presentation of the Stimuli: 1

In the first two experiments of this series, discrimination (as measured by *d'*) was poorer than that found by Tarr et al. (1998). For some individuals, the *d'* values indicated near chance performance. Thus, some of the participants were not discriminating between the stimuli used. In Experiment 2 the image size was reduced in comparison to that used in Experiment 1, so that a considerably larger amount of each image would be presented to the participants' foveal regions without the need for a saccade. This failed to improve discrimination in the matching task.

As suggested in the discussion of Experiment 1, another possible way to improve discriminability would be to present the stimuli for longer periods. Biederman and Bar (1999) suggested that appropriate presentation durations for rendered objects were 400 ms (S1) and 300 ms (S2). Given the failure of Experiment 2 to show an increase in the participants' discrimination, it was proposed that the presentation times given by Biederman and Barr be used in this experiment. The extended presentation lengths would allow for a saccade to occur.

It was hypothesised that the participants' ability to discriminate between the different version of the objects would increase and be reflected in higher values of *d*'. It was also predicted that with the longer presentation durations the participants would have more time to use the cues available from the attached shadows, and thus differences in either response latency or discrimination would be evident between the two experimental groups, No-Shadow and Shadow. Given the results of Tarr et al. (1998) it was hypothesised that the performance of the participants in the Shadow Group would be faster and more accurate than that of the participants the No-Shadow group.

Method

<u>Participants</u>

Twenty-eight Psychology Department undergraduates from the University of Waikato participated in the experiment for course credit. They were randomly assigned to the two experimental conditions, No-shadows and Cast-shadows.

<u>Apparatus</u>

The apparatus was the same as that used in Experiment 2.

<u>Procedure</u>

A sequential matching procedure was used. It was the same as the procedure used in Experiment 2, except the presentation times were changed. The presentation times were 400 (S1) and 300 ms (S2). Thus, the procedure was: a small white fixation cross on a black background, displayed for 750 ms; the sample stimulus, displayed for 400 ms; a masking stimulus, displayed for 750 ms; the comparison stimulus, displayed for 300 ms; and another masking stimulus, for 500 ms. The participants could respond from the onset of the comparison stimulus until 1500 ms had elapsed after the comparison stimulus disappeared.

Results

Tests of the experimental hypotheses

Discrimination

The participants' discriminatory performance was evaluated using the measure *d*'. The No-Shadow group demonstrated a higher level of discrimination between the objects than the Shadow Group. The No-Shadow group's *d'* value was 2.31, compared to the Shadow group's value of 1.86. The difference between the two groups was assessed using an independent groups *t*-test. The difference was significant ($t(26) = 2.5569$, $p < 0.05$).

Correct Responses

Differences in response times of correct responses, across the two experimental groups, No-Shadow and Shadow, were evaluated with a within-subjects ANOVA, with experimental group as a between-subjects factor. Illumination direction change between the two stimuli in each trial, and object version change between the two stimuli in each trial were the two within-subjects factors.

The main effects of experimental group ($F(1,26) = 0.319$, $p > 0.05$) and illumination change ($F(1,26) = 0.092$, $p > 0.05$) were not significant. The main effect of change in the versions of the objects presented in a trial was significant ($F(1,26) = 12.854$, $p < 0.05$, $\eta_p^2 = 0.331$). Responses in the trials where there was no change in the object from initial stimulus to comparison stimulus were on average faster, 692 ms compared to 762 ms, than when there was a change. None of the interactions were significant: group and illumination direction ($F(1,26) = 2.804$, $p > 0.05$); group and object version ($F(1,26) = 3.076$, $p > 0.05$); illumination direction and object version ($F(1,26) = 4.157$, $p > 0.05$); and illumination direction, object version and group ($F(1,26) = 1.354$, $p > 0.05$).

All Responses

Differences in response times, across the two experimental groups, No-Shadow and Shadow, were evaluated with a within-subjects ANOVA, with experimental group as a between-subjects factor. Illumination direction change between the two stimuli in each trial, and object version change between the two stimuli in each trial were the two within-subjects factors.

The main effects of experimental group ($F(1,26) = 0.401$, $p > 0.05$) and illumination change ($F(1,26) = 0.880$, $p > 0.05$) were not significant. The main effect of change in the versions of the objects presented in a trial ($F(1,26) = 23.607$, $p < 0.05$, $\eta_p^2 = 0.476$) was significant. The response latencies were faster in trials where there was no change in the object from initial to comparison stimulus, 709 ms compared to 773 ms, than when there was a change. The interaction between group and object version was significant ($F(1,26) = 7.355$, $p < 0.05$, $\eta_p^2 = 0.221$) (as illustrated in Figure 3.1). When the same versions of an object were presented in a trial, reaction times were the same for both the Shadow Condition and No-Shadow Condition, 709 ms, to zero decimal places, but the participants in the No-Shadow Condition took longer to respond when the versions were different from those in the Shadow Condition, 808 ms compared to 737 ms. The other interactions were not significant: group and illumination direction ($F(1,26) = 2.956$, $p > 0.05$); illumination direction and object version ($F(1,26) = 1.205$, $p > 0.05$); and illumination direction, object version and group ($F(1,26) = 2.597$, $p > 0.05$).

Incorrect Responses

Differences in response times of incorrect responses, across the two experimental groups, No-Shadow and Shadow, were evaluated with a within-subjects ANOVA, with experimental group as a between-subjects factor. Illumination direction change between the two stimuli in each trial, and object version change between the two stimuli in each trial were the two within-subjects factors.

The main effects of experimental group, illumination change and object version were not significant, ($F(1,26) = 0.000$, $F(1,26) = 0.040$, and $F(1,26) = 0.275$, $p > 0.05$ respectively). None of the interactions were significant: group and illumination direction ($F(1,26) = 0.137$, $p > 0.05$); group and object version ($F(1,26) = 2.229$, $p > 0.05$); illumination direction and object version ($F(1,26) = 2.996$, $p > 0.05$); and illumination direction, object version and group ($F(1,26) = 0.000$, $p > 0.05$).

Comparison with Experiments 1 and 2

The accuracy and response time data obtained using the longer presentation times were compared to the data from Experiments 1 and 2. There was no significant effect of presentation duration on latency to responding ($F(3, 97) = 2.130$, $p > 0.05$) (as illustrated in Figure 3.2, Note: ANOVA included Experiment 4). However, the longer presentation times produced higher discrimination than that found in both Experiments 1 and 2

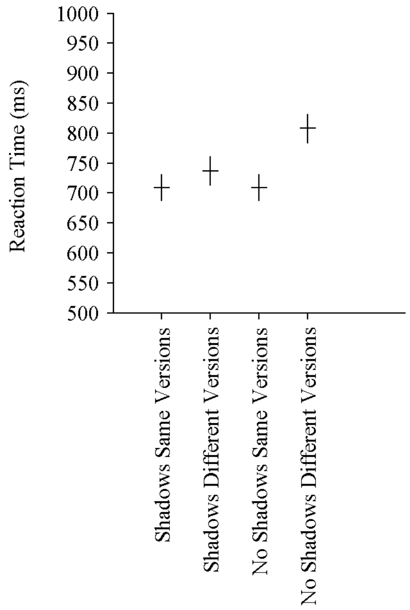

Figure 3.1. Mean reaction times for the Shadow and No-Shadow groups, split by trial type: trials in which the object versions changed from initial to comparison stimuli, and trials in which the object versions did not change. The horizontal bars indicate the means, and the vertical bars indicate the standard errors of the means. The interaction between experimental group and object change was significant ($F(1,26)=7.355$, $p<0.05$, $\eta_p^2 = 0.221$).

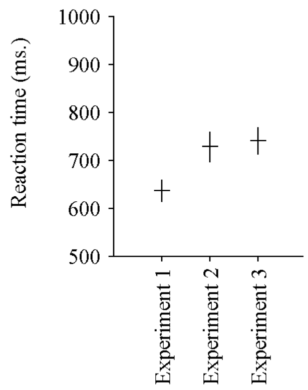

Figure 3.2. Average response latencies across Experiments 1, 2, and 3. All trials (i.e., both correct and incorrect responses) were considered. The horizontal bars indicate the means, and the vertical bars indicate the standard errors of the means.

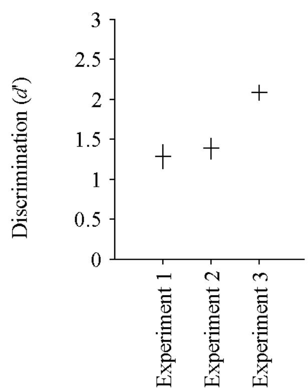

Figure 3.3. Average values of *d*' across Experiments 1, 2, and 3. The horizontal bars indicate the means, and the vertical bars indicate the standard errors of the means. Experiment 3 returned significantly higher *d*' values than Experiments 1 and 2.

($F(3,97) = 8.524$, $p < 0.05$, $\eta_p^2 = 0.209$) (as shown in Figure 3.3, and Table 2.1, note: ANOVA included Experiment 4).

Results summary

A significant difference between the two experimental groups was observed in terms of discriminability. The No-Shadow Group showed greater discrimination between the different versions of the objects. As in the previous two experiments, there was no difference between the two experimental groups in terms of latency to responding. However, a significant interaction between experimental group and object version was observed in the analysis of the data obtained from the combined correct and incorrect responses. When the same versions of an object were presented in a trial, reaction times were the same for both the Shadow Condition and No-Shadow Condition, 709 ms, but when the versions were different, the participants in the No-Shadow Condition took longer to respond than those in the Shadow Condition, 808 ms compared to 737 ms. This interaction was not significant when either the correct or the incorrect responses were assessed individually.

Discussion

The hypothesis of this experiment was that the Shadow Group, when compared to the No-Shadow Group, would benefit in terms of either response latency or level of discrimination, from the presence of attached shadows in the images. In contrast to this hypothesis, there was generally no difference between the groups in terms of latency to responding, and the No-Shadow Group displayed a higher level of discrimination, *d'* = 2.31, than the Shadow Group, *d'* = 1.86.

Based upon the experiment of Tarr et al. (1998) it was predicted that there would be a reaction time benefit for the shadow group over the no-shadow group. The response times found here were still shorter than those recorded by Tarr et al. (approximately 750 ms versus 850 - 950 ms). Tarr et al. reported faster reaction times for their shadow condition in their analysis of same trial response times for correct trials: if a similar effect was present for this experiment we would expect to see a group*object interaction on correct trials (see Appendix 4, Table A4.2). There was a significant interaction between experimental group and Object Type (same or different object in a trial), but this only occurred for the analysis of all the trials (correct and incorrect responses) ($F(1,26) = 7.355$, $p < 0.05$, $\eta_p^2 = 0.221$). When the same versions of an object were presented in a trial, reaction times were the same for both the Shadow Condition and No-Shadow Condition, 709 ms, but the participants in the No-Shadow Condition took longer to respond when the versions of an object in a trial were different, than the participants in the Shadow Condition, 808 ms compared to 737 ms. This effect occurred regardless of whether or not there was a change in illumination direction. The same pattern was evident for the correct trials analysis, but, as the interaction was not significant, it cannot be concluded that the presence of the attached shadows aided the correct discrimination of a difference between the two stimuli in a trial.

Other research into object recognition that has assessed response latencies has returned inconsistent results. Braje et al. (2000) studied the effect of shadow presence on recognition of natural objects (fruit and vegetables). When assessed across colour, greyscale, and blurred images, there were no general effects of the presence of shadows upon reaction times. However, there were significant effects for different food types.

> 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 (p. 391, Braje et al., 2000).

It was suggested that slower recognition in the presence of shadows may have been due to “a result of confusion between albedo changes and shadow boundaries” (p. 391, Braje et al., 2000) and that improved performance in the presence of shadows may have been due to enhanced contrast at the food’s bounding contour (Braje et al., 2000).

Braje et al. (1998) used a sequential face matching procedure, with masking of the initial and comparison stimuli, and the same timing as used in Experiment 1 of this series. They found a significant increase in response time (127 ms) when cast shadows were present, compared to when they were absent. However, when using a naming procedure with voice activated response time recording, Braje et al. (1998) obtained no differences in response time. Conversely, Castiello (2001) reported longer reaction times (36 ms) for the naming of computer presented images of familiar objects when cast shadows were absent, in comparison to when cast shadows (that were congruent with regards to shape of the object and the lighting direction) were present. Thus, with the exception of some individual types of food in the Braje et al. (2000) study, the Tarr et al. study is the only one to find a reaction time benefit due to the presence of attached shadows. This benefit was not replicated here, and conversely, discrimination was found to be poorer in the presence of attached shadows.

In this experiment, the sensitivity measure of the No-Shadow group was similar to that found by Tarr et al. (1998) who obtained *d*’ values of 2.35 to 2.64. However, the reduced level of discrimination for the shadow group is in contrast to the findings of Tarr et al. (with novel objects), Braje et al. (1998) (with faces), and Braje et al. (2000) (with natural objects: except in the case of two-tone images), who all found no significant differences in sensitivity due to the presence of shadows. The exception being the Braje et al. (2000) study, which used two-tone images of natural objects (derived from photos of fruit and vegetables). Braje et al. (2000) found a sensitivity cost when shadows (both cast and attached) were present in the first of two blocks of trials. The difference disappeared in the second block of trials.

An analysis of the procedures used in the different experiments does not indicate why these differences in findings may have occurred. The current experiment employed a methodology modelled after Tarr et al. (1998) and used a similar sequential matching procedure, with masking of the initial and comparison stimuli, to Braje et al. (1998). Where a different procedure was used by Braje et al. (2000), verbal naming of images presented for up to 10s, the results (for the two-tone condition) are the most similar to those of this experiment and indicate that shadows may hinder recognition under some conditions. Although the sample objects used by Braje et al. (2000), were not novel (they were fruit and vegetables), the two-tone images they used are not commonly encountered. The results of Moore and Cavanagh (1998), and Moore and Engel (2001) indicate that shadow areas in two tone images are often mistaken for parts of the object itself. If this occurred it would be expected to produce a reduction in sensitivity as found by Braje et al. (2000). The results of Braje et al. (2000), where the effect of shadow presence was only found in the first of two conditions, also indicated that familiarity with a discrimination may moderate the effect of the presence of shadows: that is, as a discrimination becomes more familiar, any costs associated with the presence of shadows may be nullified.

The difference in sensitivity found between the Shadow and No-Shadow groups in Experiment 3 implies that shadows are not an aid to recognition in tasks such as this one, and may in fact be a hindrance. Tarr et al. (1998) suggested that shadows could help to resolve the 3-D structure of the objects being viewed. The result here suggests that if there was any potential benefit from the attached shadows, in terms of 3-D resolution, this must have been negated by the shadows either masking other cues, or providing spurious information. It is not clear why the results of this experiment differed to those found by Tarr et al. (1998).

The lengthened presentation times for S1 and S2 in this experiment resulted in improvements in *d'* for both the Shadow Group and No-Shadow Group, and for the No-Shadow Group to a value similar to the lower end of the values found by Tarr et al. (1998). It was therefore possible that a further lengthening of presentation times would again increase sensitivity. This was tested in the next experiment. Reproducing the same difference in sensitivity found in this experiment, again using longer presentation times would also demonstrate the reliability of the results, and the claim that attached shadows may hinder the object recognition of novel objects.

Experiment 4 Extended Presentation of the Stimuli: 2

In the previous experiments, accuracy, and latencies to responding, increased after the size of the stimuli was reduced, and the presentation times of the stimuli were increased. In Experiment 3, the latencies and *d'* values both increased in the direction of those found by Tarr et al. (1998) although still being lower than Tarr et al.'s.

In order to establish whether accuracy and response latencies could be further increased, it was proposed that presentation durations would be increased again. Increases of the same magnitude (200 ms) as between Experiment 2 and Experiment 3, were used between Experiment 3 and Experiment 4. Thus, the presentation times were increased from 400 ms (S1) in Experiment 3 and 300 ms (S2) in Experiment 3, to 600 ms (S1) and 500 ms (S2).

Based on the results of Experiment 3, it was predicted that the increased presentation duration would result in another improvement in accuracy, as the increased time available to view the image would enable more information about the image to be processed. A corresponding increase in response latency was also predicted, due to response latency being calculated from the onset of the comparison stimulus, and the time available to view the comparison stimulus would be increased. Another possibility was that enough time was already available for the participants in Experiment 3 to perform the task with relative accuracy, and thus, the additional lengthening of presentation times would not produce any changes in accuracy or response latency.

Method

Participants

Twenty-six Psychology Department undergraduates from the University of Waikato participated in the experiment for course credit. They were randomly assigned to the two experimental conditions, No-shadows and Cast-shadows.

Apparatus

The apparatus was the same as used in Experiment 3.

Procedure

A sequential matching procedure was used. It was the same as the procedure used in Experiment 3, using the reduced image size in comparison to Experiment 1, except the presentation times were changed. The presentation times were 600 ms (S1) and 500 ms (S2). Thus, the procedure was: a small white fixation cross on a black background, displayed for 750 ms; the sample stimulus, displayed for 600 ms; a masking stimulus, displayed for 750 ms; the comparison stimulus, displayed for 500 ms; and another masking stimulus, for 500 ms. The participants could respond from the appearance of the comparison stimulus until 1500 ms had elapsed after the comparison stimulus disappeared.

Results

Tests of the experimental hypotheses

The participants' discriminatory performance was evaluated using the measure *d*'. The No-Shadow Group's *d'* value was 1.93, compared to the Shadow Group's value of 1.74. The difference between the two groups was assessed using an independent groups *t*-test (see Table 1.1.). The difference was not significant ($t(24) = 0.673$, $p > 0.05$).

Correct Responses

Differences in the response times of correct responses, across the two experimental groups, No-Shadow and Shadow, were evaluated with a within-subjects ANOVA, with experimental group as a between-subjects factor. Illumination direction change between the two stimuli in each trial, and object version change between the two stimuli in each trial were the two within-subjects factors.

The main effect of experimental group was not significant ($F(1,24) = 0.238$, $p > 0.05$). While the main effect of illumination change ($F(1,24) = 22.997$, $p < 0.05$, $\eta_p^2 = 0.489$) was significant. The response latencies were faster when the illumination direction was not changed within a trial, 709 ms, compared to when it was changed, 762 ms. There was no effect of change in the versions of the objects presented in a trial ($F(1,24) = 0.6871$, $p > 0.05$). None of the interactions were significant: group and illumination change ($F(1,24) = 0.112$, $p > 0.05$); group and object version ($F(1,24) = 0.221$, $p > 0.05$); illumination direction and object version ($F(1,24) = 3.091$, $p > 0.05$); and illumination direction, object version, and group ($F(1,24) = 0.010$, $p > 0.05$).

All Responses

Differences in response times, across the two experimental groups, No-Shadow and Shadow, were evaluated with a within-subjects ANOVA, with experimental group as a between-subjects factor (see Appendix 4, Table A4.2.). Illumination direction change between the two stimuli in each trial, and object version change between the two stimuli in each trial were the two within-subjects factors.

The main effect of experimental group was not significant, ($F(1,24) = 0.174$, $p > 0.05$). While the main effect of illumination change ($F(1,24) = 19.516$, $p < 0.05$, $\eta_p^2 = 0.448$), was significant. The response latencies were faster when the illumination direction was not changed within a trial, 720 ms, compared to when it was changed, 757 ms. There was no effect of change in the versions of the objects presented in a trial ($F(1,24) = 3.831$, $p > 0.05$). None of the interactions were significant: group and illumination change ($F(1,24)=0.114$, $p > 0.05$); group and object version ($F(1,24) = 1.074$, $p > 0.05$); illumination direction and object version ($F(1,24) = 2.435$, $p > 0.05$); and illumination direction, object version, and group ($F(1,24) = 0.647$, $p > 0.05$).

Incorrect Responses

Differences in the response times of the incorrect responses, across the two experimental groups, No-Shadow and Shadow, were evaluated with a within-subjects ANOVA, with experimental group as a between-subjects factor. Illumination direction change between the two stimuli in each trial, and object version change between the two stimuli in each trial were the two within-subjects factors.

The main effect of experimental group was not significant ($F(1,24) = 0.015$, $p > 0.05$). While the main effect of object version ($F(1,24) = 5.201$, $p < 0.05$, $\eta_p^2 = 0.178$) was significant. The response latencies were faster when the same versions of an object were presented in a trial, 766 ms, compared to when different versions were presented 801 ms. There was no effect of change in illumination direction ($F(1,24) = 0.405$, $p > 0.05$). None of the interactions were significant: group and illumination change, ($F(1,24) = 2.855$, $p > 0.05$); group and object version, ($F(1,24) = 9.066$, $p > 0.05$); illumination direction and object version ($F(1,24) = 0.965$, $p > 0.05$); and illumination direction, object version and group ($F(1,24) = 2.923$, $p > 0.05$).

Comparison with Experiments 1 to 3

The accuracy and response time data obtained using the presentation times of 600 ms (S1) and 500 ms (S2), were compared to the data from Experiments 1 to 3 (as shown in Figure 4.1). There was no significant effect of presentation duration on latency to responding ($F(3, 97) = 2.130$, $p > 0.05$). However, Figure 4.2 and Table 2.1 show that the longer presentation times produced higher discrimination than that found in Experiment 1 but no difference in comparison to Experiment 3 or 4 ($F(3,97) = 8.524$, $p < 0.05$, $\eta_p^2 = 0.209$).

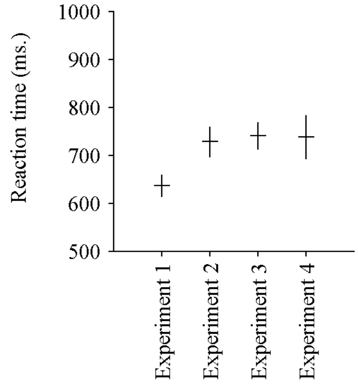

Figure 4.1. Average response latencies across Experiments 1 to 4. All trials (i.e., both correct and incorrect responses) were considered. The horizontal bars indicate the means, and the vertical bars indicate the standard errors of the means.

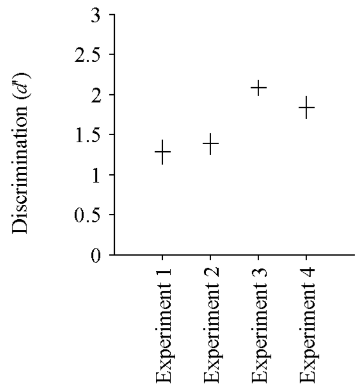

Figure 4.2. Average values of *d*' across Experiments 1 to 4. The horizontal bars indicate the means, and the vertical bars indicate the standard errors of the means. Experiment 4 returned significantly higher *d*' values than Experiment 1.

Results summary

There was no difference in accuracy of matching between the two experimental groups, shadow and no-shadow. With respect to response times, there was no main effect of experimental group, shadow versus no-shadow, either when all the data were analysed, or when correct responses and incorrect responses were analysed separately. There was a main effect of illumination change in both the analysis of all responses and correct responses: response latencies were faster when there was no change between the objects. There was a main effect of object version only for incorrect responses, where responses were faster when the same versions of an object were presented in a trial.

Experiments 1 to 4 Results and Discussion

The four previous experiments were conducted to assess whether the presence of attached shadows would aid in an object-recognition task using novel objects. Based upon the results of Tarr et al. (1998) it was hypothesised that when attached shadows were present, recognition would be faster than when they were absent. The analyses conducted in Experiments 1 to 4 failed to provide any evidence of a differentiation in reaction time associated with the presence or absence of attached shadows.

The first four experiments manipulated the size of the initial and comparison images and the duration of their presentation. Experiment 2 decreased the size of the stimuli from 14.3° height by 18.9° width, to 5.7° height by 7.6° width, and Experiments 3 and 4 lengthened the presentation time of S1 and S2, from the 200 ms and 100 ms used in Experiment 1 and by Tarr et al. (1998), to 600 ms and 500 ms in E4 (while still employing the reduced image size). Response latencies in all the four experiments were faster than those recorded by Tarr et al. (1998). Response latencies increased from a minimum in Experiment 1 (637/636 ms, No-Shadow “NS”/Shadow “S”), peaking in Experiment 4 (719/758 ms, NS/S) (as shown in Figure 2.1, and Appendix 4, Tables A4.1 & A4.2). However, the difference across experiments was not significant ($F(3,97) = 2.130$, $p > 0.05$). Dependent upon condition, Tarr et al. reported latencies in the range of 807 to 967 ms.

In a footnote to their article, Tarr et al. (1998) describe conducting another trial using shadows rendered with penumbra, for which they report reaction times of 787 (no illumination change) to 793 ms (illumination change). These latencies, averaging 790 ms, are similar to those found here in Experiment 4. Also note, that they are a substantial 164 ms faster than the average of the times reported for their No-Shadow Conditions (a very large effect), and 57 ms faster than the average of the times reported for their Shadow Conditions. In the condition using shadows with penumbra, Tarr et al. also recorded sensitivity values, *d*’, of 2.19 (no illumination-change) and 2.00 (illumination-change), similar to what was found in Experiments 3 and 4. Tarr et al. could not account for the differences with respect to their other conditions.

Across Experiments 2 to 4, *d’* changed as the stimulus presentation times were manipulated (Braje et al. (2000) found a similar result). Accuracy, as measured by *d'* increased from Experiment 1 (*d'*=1.26/1.30, NS/S) through to Experiment 3 (*d'*=2.31/1.86, NS/S), before dropping slightly in Experiment 4 (*d'*=1.93/1.74, NS/S). Experiment 4 still had the second highest *d'* values. An ANOVA confirmed a significant difference between the groups ($F(3,97)=10.459$, $p<0.05$, $\eta_p^2=0.209$). Figure 2.2 indicates that Experiments 3 and 4 produced higher *d'* values than Experiments 1 and 2. As noted earlier, the changes in image size and presentation duration had the effect that by Experiment 3, the sensitivity values obtained here were similar to the range obtained by Tarr et al. (*d'*=2.00 to 2.64, including the condition with penumbra).

Feedback

This series of experiments did differ from that of Tarr et al. (1998) in the use of feedback. Tarr et al. employed an audible beep as feedback indicating an incorrect response. The use of feedback is based upon the supposition that it improves discrimination through learning. Learning about the stimuli is not a desired effect in an investigation specifically employing novel objects. The *d'* values obtained in Experiment 3 were only slightly lower than those obtained by Tarr et al. (1998), an indication that lack of discriminability, as a result of lack of feedback, was not a reason for the difference between the results of the two experiments. Therefore, the difference between the two experiments, in terms of the use of feedback, should have no bearing upon the results.

Effect of Shadow Presence on Response Latencies

For Experiment 4, as for the previous three experiments (with the exception of the significant Object by Experimental Group interaction in Experiment 3), there was no differential effect of the presence of attached shadow with respect to response latencies. This finding (of no differential effect) is similar to that of Braje et al. (2000) when testing the recognition of colour, greyscale, and blurred images of fruit and vegetables, and that of Braje et al. (1998) using a face naming procedure. However, it differs to some of the results of Tarr et al.'s (1998), Braje et al. (1998), and Braje et al. (2000). Tarr et al. found a response time reduction due to the presence of attached shadows. Braje et al. (1998), found response time increases due to the presence of shadows in a face matching procedure, and Braje et al. (2000), found a response time increase due to the presence of shadows (undistinguished attached and cast) in the first of two phases using two-tone images of fruit and vegetables.

In the face-matching task, Braje et al. (1998) found that the presence of shadows impaired response times. When the faces had attached shadows (called cast by Braje et al.) participants were slower (127 ms) but discriminability was the same. Braje et al. suggested that shadows may introduce spurious contours that people may confuse with surface contours. Analyses of shadow boundaries by Cavanagh and Leclerc's (1989) and Cavanagh (1991) support this suggestion. They suggest that cast contours would be spurious while attached (terminator) contours would relate information about surface contour (as depicted in Figure 4.3). Yet, as noted above, when participants had to name, rather than match faces, Braje et al. (1998) found no effect of the presence of attached shadows on either reaction times or sensitivity.

Thus, there is very little consensus amongst the findings of experiments that have investigated the effect of shadowing on reaction time. But overall, there is little evidence of a reliable response time benefit for recognition due to the presence of attached shadows.

Effect of Illumination Direction on Response Latency

Although it was not the primary measure of interest, reaction time differences due to changes in illumination were recorded. In the analyses of all the data for Experiments 2 and 4, and correct responses for Experiment 4, a significant difference in reaction time was found to be related to a change in illumination direction (refer to Appendix 4, Tables A4.1 & A4.2): a change in illumination direction produced a cost in recognition latency (53 ms for correct responses in Experiment 4). That is, the participants responded faster when the initial and comparison stimuli were illuminated from the same direction. This main effect occurred irrespective of whether or not there was a change in the object being presented (there was no Illumination-Direction by Object-Change interaction), and irrespective of experimental condition (no Illumination-Direction by Shadow-Presence interaction).

Braje et al. (1998), when investigating face matching, also found a significant response time cost (44 ms) due to change in illumination direction, irrespective of shadow condition (in same-face trials only). Tarr et al. (1998) reported a similar result. A 25 ms response time cost due to a change in illumination direction, but found it only in their shadow condition.

The results of the current experiments, in combination with those of Braje et al. (1998) and Tarr et al. (1998), indicate that changes in the illumination direction of images in object-recognition tasks may lengthen reaction

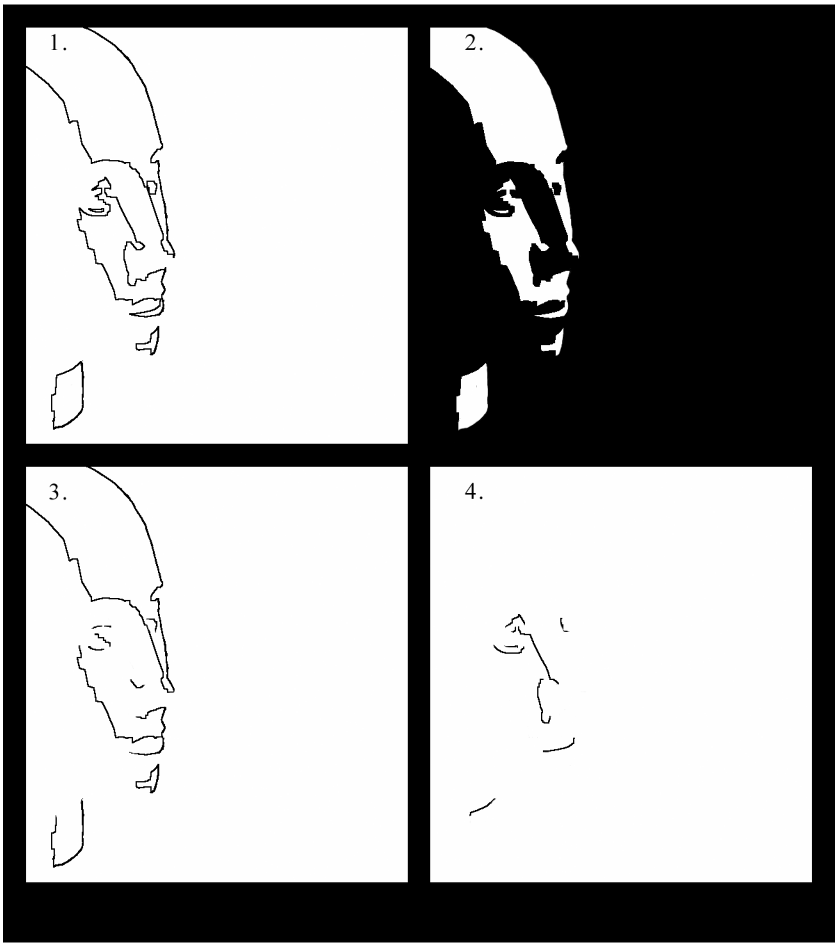

Figure 4.3. Illustration of the information available for recognition from the external, attached and cast contours of attached shadows; as suggested by Cavanagh (1991). In the first quadrant the full contour of the face is presented. This can be compared to the second image, where the full contour, plus appropriate shadowing is presented. Cavanagh (1991) suggests that external and attached contours, as given in the third quadrant, may provide cues to recognition through matching to a prototype, whereas the cast contours, as shown in the fourth quadrant, are spurious.

times. The presence of attached shadows on an object may produce greater differences between two images, when illumination direction changes, than the same situation without attached shadows. This could compound any effect of a change in illumination direction, and may be why Tarr et al. only found an effect of change in illumination direction in their shadow condition.

Effect of Object Change on Response Latencies

The effect of a change in object (between the initial and comparison stimuli) on reaction time was assessed. There was a consistent main effect across Experiments 1, 2, and 3 when the objects in the initial and comparison stimuli differed. In Experiments 1, 2 and 3, analyses of all the data, and analyses of correct responses only, demonstrated a reliable increase in reaction time when different objects were presented in a trial. Otherwise, there were no significant differences in reaction time for correct responses or for all the data in Experiment 4.

Analyses of incorrect responses in Experiments 1 and 2 demonstrated a reliable increase in reaction time when the same objects were presented in a trial. This effect was not present in Experiment 3, and in Experiment 4 the opposite effect was found. Thus, the analysis of incorrect responses fails to provide any clear picture of a reliable pattern of responding.

The Effect of Shadow Presence on Discrimination

With respect to discrimination, the results of Experiment 3 (using the small stimulus size, and 300 and 400 ms presentation times for S1 and S2) showed that discrimination was significantly better for the No-Shadow group than the Shadow group (d'=2.31/1.86, NS/S, $t(26)=2.5908$, $p<0.05$). The results of Experiment 3 fit with the hypothesis that, either the addition of shadows masks cues used for recognition, making discrimination more difficult, or act as distracters by adding spurious cues to an image. This finding was not replicated in Experiment 4 (d'=1.93/1.74, NS/S), a systematic replication of Experiment 3 increasing the presentation durations. It is concluded that any effect present is not very reliable.

Other experimental research indicating a change in sensitivity due to the presence of cast shadows is limited. Braje et al. (2000) reported a sensitivity cost due to the presence of attached and cast shadows. However, this was only obtained in the first of two phases using two-tone images of fruit and vegetables, and a 10-s maximum stimulus exposure time. In the same series of experiments, Braje et al. (2000) also assessed recognition using colour, greyscale, and blurred images, across two conditions: one with a 10-s maximum exposure time of the stimulus, and a second with a 30-ms exposure of the stimulus. They found that although accuracy was reduced by 30% in the 30-ms condition, there was no differential effect of the presence or absence of shadows in either case. Their results do not indicate that manipulating the difficulty of the discrimination by controlling stimulus exposure time had any differential effect upon recognition with respect to the presence or absence of shadows. This supports the suggestion here that the effect found in Experiment 3 is unreliable; as opposed to being genuine but restricted to presentation times between those used in Experiments 2 and 4. While Tarr et al. (1998) found a response latency benefit due to shadow presence, they also found no significant differences in *d*' between the No-Shadow and Shadow groups (*d*' figures between 2.35 and 2.62).

Neither the current experiments, nor any of those discussed here, show a benefit in terms of discrimination due to the presence of attached shadows. There is a suggestion that discrimination may, at times, be hindered by the presence of attached shadows (see: Experiment 3 and Braje et al. 2000), although this finding is unreliable. Braje et al.'s (2000) results indicate that the effect that they found with two-tone images may be moderated by familiarity with the task: since the reduction in discrimination due to the presence of shadows disappeared in the second of two blocks of trials.

Failure to Find an Effect of the Presence of Shadows

Two possibilities were raised by the fact that there were no major differential effects found due to the presence or absence of cast shadows. It could be suggested that in this sort of task, either 1, the presence or absence of attached shadows does not actually have any effect upon recognition latency or discrimination, or 2, the experimental methodology was deficient in one or more areas for testing the experimental hypotheses.

Experiments 1 to 4 may not have used a sufficient number of participants to make detection of an effect likely. A power analysis of the results of Tarr et al. (1998), shows that the sample sizes used should have been more than adequate. Tarr et al. report reaction times of 848, 827, 873, and 838 ms for the four subgroups of the Shadow Condition (no-change in illumination-direction, with same/different objects, and change in illumination-direction, with same/different objects), averaging 847 ms. For the No-Shadow Condition the means of the same subgroups are 952, 945, 967, and 950 ms, averaging 954 ms. The Shadow group had 32 participants, and the No-Shadow group 44. Standard deviations or standard errors are not stated, but standard errors can be estimated off the graphs provided: a liberal estimate is 10 ms. Standard deviations of the two groups can therefore be estimated at 57 ms (Shadow) and 66 ms (No-Shadow). Based upon a two-sample bi-directional test, with an alpha level of 0.05, power can be estimated to be approximately 100%. Using the same two distributions, but 12 participants per group, as was typical for Experiments 1 to 4, power is 98%, indicating that if the effect was present it should have been detected.

However, several factors related to the experimental methodology indicate that it would be premature to assume that shadowing has no effect upon recognition. Although the experiments compared two conditions, one with attached shadows present, and one without, within the Shadow Condition there was no control over the proportion of shadowing present in an image, or how this differed with changes in illumination direction. Thus, some of the images in the Shadow Condition were barely different from their counterparts in the No-Shadow Condition (see Table 1.2). For example, the extremes of difference between the No-Shadow Condition and the Shadow Condition can be seen with the two illumination directions of Object 51: when illuminated from the left, the image in the Shadow Condition differed, in terms of in total pixel value, by 8.56% of the No-Shadow Condition image, while when illuminated from the right the difference was only 0.18%.

In a natural situation there would also be no control over the amount of shadowing, but if a robust effect existed it would be most likely to have been evident when there were moderate levels of shadowing present (see: Freeburg, 1966). This being the case an effect may have been present, but limited to those conditions in which a certain amount of shadow information was present. The average number of correct responses to each image, and the average response latencies, were re-analysed to assess this possibility.

First, the percentage differences between the total pixel values of the Shadow and No-Shadow images (see Table 1.2), were correlated with both the Shadow Group's average latencies to responding and average number of correct responses. For Experiments 1 to 4, no relationships between amount of shadow and either correct response or response latencies were evident (see Table 4.1). A second analysis investigated the hypothesis that only a certain amount of shadow may produce an effect upon either accuracy or response

Table 4.1

The Percent Darker the Shadow Image is in Comparison to the No-Shadow Image, Correlated With the Average Number of Correct Responses, and the Latency to Responding. Analysis for Same Versions of an Object in a Trial Where Illumination Does Not Change.*

	Experiment 1	Experiment 2	Experiment 3	Experiment 4
Number of correct responses*:	$r_s(24)=-0.159$	$r_s(24)=0.284$	$r_s(24)=0.236$	$r_s(24)=0.377$
Latency to responding*:	$r_s(24)=-0.182$	$r_s(24)=-0.277$	$r_s(24)=-0.115$	$r_s(24)=-0.081$

*See Table 1.2.
None of the correlations are significant at an alpha level of 0.05.

*Average number of correct responses not normally distributed for Experiments 1 and 2, average latency normally distributed for Experiments 1 to 4 (Shapiro-Wilkes test of normality).

latency (e.g., Freeburg, 1966), as opposed to an increasing effect with increasing level of shadow. The percentage difference between the Shadow and No-Shadow images (see Table 1.2) was used as a crude measure of amount of shadow. This measure was used to divide the images into four levels defined by the quartiles of the data set (see Table 1.2).

The accuracy of responding, and latency to responding, were assessed across these four shadow levels using two-way ANOVA, with “Shadow Level” and “Experiment” as the independent variables. The Shadow Levels were nominally called Level 1, 2, 3 and 4, representing the least (Level 1) to greatest (Level 4) difference between the Shadow image and the No-Shadow image. The ANOVA returned a significant result with respect to Shadow Level ($F(3,80) = 4.919$, $p<0.05$, $\eta_p^2=0.156$). A Tukey’s post hoc test showed that the grouping of Shadow images that differed the least from their corresponding No-Shadow images, produced (on average) less correct responses (82%) than the other three levels with greater differences (87% to 89%). There was no interaction between the different Experiments and the Shadow Levels. The ANOVA on response latencies showed a similar pattern ($F(3,80) = 6.733$, $p < 0.05$, $\eta_p^2 = 0.202$): post hoc testing revealing that times were significantly longer for Level 1 (720 ms) in comparison to Levels 2 (656 ms) and 4 (673 ms), but approximately equal to Level 3 (689 ms). There was no significant effect upon response latencies of the interaction between the Experiment and the Shadow Level.

The amount of shadow in an image is determined by the particulate features of the object, and their position with respect to the light source. Therefore, it is possible that another factor (related to the features of the objects in each Shadow Level quartile) may co-vary with the Shadow Level, and actually cause an effect that could mistakenly be attributed to the amount of shadow present. The within-Shadow-Level analysis above had returned a significant result, so to control for the possibility that another factor was influencing this finding, the data from the No-Shadow group was analysed in the same manner as above, i.e., the data from the No-Shadow group were analysed as if that group had been presented the images containing shadow. If the No-Shadow group’s data varied according to these “levels of shadow” (as the Shadow group’s data did), this would demonstrate that another factor aside from shadowing was causing the effect (as the No-Shadow group did not experience the different "levels of shadow").

When the data from the No-Shadow group were analysed according to the groupings determined by the shadow images, similar results were obtained as for the Shadow group's data. The average number of correct responses varied according to the levels determined by the images containing shadow ($F(3,80) = 5.559$, $p < 0.05$, $\eta_p^2 = 0.172$), even though the No-Shadow group viewed the images without shadows. The average number correct was lower in Level 1 (81%) than Level 2 (89%) and 3 (86%), while Level 4 returned an intermediate value (86%). Therefore, it can be concluded that the covariation of number correct, with the four shadow levels, must be due to a factor common to both the Shadow and No-Shadow groups, i.e., the effect could not have been due to differences in amount of shadowing which was restricted to only the shadow group.

The effect of Shadow Level upon response latencies was also significant ($F(3,80) = 5.310$, $p < 0.05$, $\eta_p^2 = 0.166$), with response latencies slowest in Level 1 (711 ms), and fastest in Level 2 (645 ms). Level 3 (669 ms) and 4 (667 ms) did not significantly differ to either Levels 1 or 2.

Across the two analyses, correlational and ANOVA, the major finding is that there was no effect of the amount of shadow present in the images, upon either the number of correct responses, or the latency to responding. The significant effect of Shadow Level found in the ANOVA on average number of correct responses, and latency to responding, appears in both the No-Shadow-, and Shadow-, Groups' data. This means that something common to both groups (and both conditions) must co-vary with the differences in the levels of shadowing present in the four quartiles of the Shadow Level variable.

The main finding in the results so far is the failure to find any evidence that differences in the amount of shadowing in the images affected accuracy or latencies to responding. So, upon what grounds did the participants make their same/different decision?

In discussing the results of their investigation of the effects of shadowing upon the recognition of natural objects (fruit and vegetables) Braje et al. (2000) compare the lack of differentiation they found between their shadow and no-shadow groups, in terms of either reaction times, or sensitivity, with the findings of Tarr et al. (1998). They state that “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” (p. 396). In the case of the current experiment, even though the shapes are novel, there is still plenty of shape information available for the participant to use in the matching task. The participants’ responses during debriefing provide indicators of what parts of the images they were attending to during the experiments.

Debriefing

Comments from some of the participants indicated that they used the presence or absence of a particulate feature as a basis for the matching/not-matching response. A change in a particulate feature would produce a sizeable change in the overall bounding contour of an object. Others indicated that they tried to view the object as a whole rather than looking for a change in parts of the objects. A quick visual inspection of the object versions suggests that the object discriminations the participants performed poorly on may have been those that involved the least obvious change in particulate features, e.g., Figure 4.4 presents a comparison of the relatively similar objects 51 and 52 (sail-shaped object, high error rate), with the more dissimilar objects 61 and 62, (sphere-shaped object, low error rate).

Use of Bounding Contour as a Mediator for Recognition

Thus, the task was very simple, and could be performed on the basis of a few rules about the location of particulate features of the objects. Furthermore, some of the participants indicated that they did use a very simple “feature present/feature absent” choice criterion while performing the task. Given the variability in the level of shadowing present in the images (see Table 1.2), the participants in the Shadow Condition were often required to make a discrimination with very minimal shadow cues present. Cavanagh and LeClerc (1989) also suggest the extraction of shadow information is a “high-level” process whereby a cast shadow is one of the last choices the visual system considers when attempting to identify a dark area. Therefore, in these experiments, an optimal decision criterion would be one that ignored shadows, and relied on a more salient cue, such as an object’s bounding contour, than those provided by shadows.

The bounding contour of novel figures has been shown to be able to mediate recognition, while features that do not contribute to bounding contour are not recognised immediately after viewing (Rock, Halper, & Clayton, 1972). Rock et al. (1972) relate:

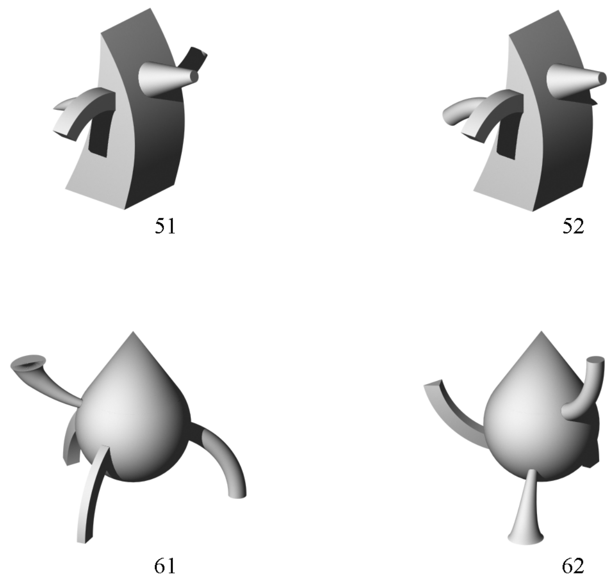

Figure 4.4. Participants performed very poorly on the discrimination between Objects 51 and 52, and performed well on the discrimination between Objects 61 and 62 (as shown in Figure 1.19). A subjective visual assessment indicates that the differences in the arrangement of the features of Objects 51 and 52 are fewer and less obvious than those between Objects 61 and 62.

> Nuances of a complex figure are generally not adequately apprehended during a single exposure and therefore fail to establish viable memory traces. Whichever components of complex figure are immaterial with regard to its global shape - such as a configuration inside an outer closed contour or minor fluctuations in the contour itself - are nuances which will suffer this fate. (p. 672)

The point that Rock et al. (1972) make, is that interior configurations or nuances cannot be recognised after a single exposure, while bounding contour can. However, Gauthier, et al. (1998) raise the possibility that people may make greater use of configural information for recognition as they become increasingly familiar with objects, and that in a verification task, participants may initially rely on the presence of particulate features to recognise objects. In an experiment such as the one just conducted, outer contour is all that needs to be attended to, in order to do the task, and when viewing novel images the internal contours may not be attended to.

Hayward (1998) has also demonstrated that bounding contour appears sufficient for the recognition of objects, and that it supports a level only slightly inferior to that of the recognition of shaded images. Hayward used a same/different sequential matching task, similar in procedure to Experiments 1 to 4, and presented objects such as geometric solids similar to those in Experiments 1 to 4, and familiar objects taken from entry level recognition categories: such as a horse, different birds, and planes. The images were either presented with shading, or in silhouette.

The results of Hayward (1998) demonstrate that bounding contour alone can support high levels of recognition, and the findings of Rock et al. (1972) indicate that, when viewing novel objects, people may not be able to report on the internal detail of those objects. Combined with the reports of the participants during debriefing, these two results suggest that it is likely the participants were using the objects' bounding contours to perform the matching task. This is a testable hypothesis. If bounding contour was used to perform the matching task, it would be expected that an analysis of trials in which the initial and comparison stimulus differ in outline, would return a specific result: that the greater the difference in the bounding contours of two objects, the greater the accuracy, and the faster the latency to responding. To test this, a measure of difference in bounding contour is needed.

In these experiments, the initial and comparison stimuli appeared in the same position on the screen, separated in time and by a masking stimulus. Ignoring the masking stimulus, a difference between the two images can be viewed as differences between the specific pixels on the screen where the images are presented. As the analysis is concerned with the global shape of the images, and not the interior, the two images can be represented by their silhouettes. Where the two silhouettes do not overlap, you have a difference in the global shape, or bounding contour of the objects (as illustrated in Figure 4.5). This difference can be measured. The silhouettes of each of the two versions of an object base were made by adjusting the contrast of the images, and the areas that they had in common were deleted (as illustrated in Figure 4.6) leaving the areas that did not overlap. The number of pixels in these areas were then summed, giving a value to the difference between the silhouettes of the two objects, with large differences between the silhouettes of the objects corresponding to large differences in bounding contour.

To test the hypothesis that change in bounding contour would increase accuracy and reduce latencies to responding, the values calculated were correlated with the number of incorrect responses for trials in which the initial and comparison stimuli differed (as presented for Experiment 1 in Figure 1.19). The pixel difference measure, and the number of incorrect responses by object pair, are presented in Table 4.2, along with the correlations between them. For Experiments 1 to 4 very strong negative correlations were found between the number of incorrect responses and the difference in pixels ($r_{s\ Experiment\ 1}$= -0.820, $r_{s\ Experiment\ 2}$= -0.876, $r_{s\ Experiment\ 3}$= -0.864, $r_{s\ Experiment\ 4}$= -0.919; *d.f.*=(12), $p<0.05$).

The correlation between latency to responding and difference in bounding contour was also assessed. Unlike the analysis of the number of correct responses, latency to responding was not correlated with the pixel difference between the object versions (see Table 4.2). The strength of the correlations between number correct and the pixel difference between versions, suggests that similarity in global shape/bounding contour was the key cue the participants used in determining whether the objects were the same. As on different object trials, when the objects were similar, the participants were more likely to respond incorrectly, that is, that the objects were the same. However, the speed the participants performed the task at was not sig-

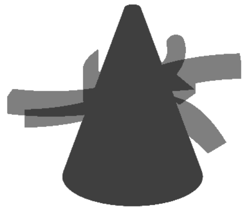

Objects 21 and 22:
14151 pixels do not overlap

Figure 4.5. This graphic shows how the silhouettes of two versions of the same object overlap. The darker area (the area the two have in common) can be removed, leaving the areas that the versions do not have in common with each other (see below, Figure 4.6). The size of this remaining area can then be used as a measure of the difference between the silhouettes of each object. Size of the area was calculated by counting the pixels in this area. If the objects were exactly the same the size of the area would equal zero.

Objects 11 and 12: 6406 pixels

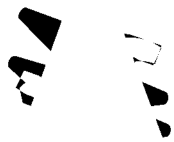

Objects 41 and 42: 5328 pixels

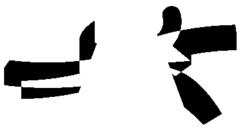

Objects 21 and 22: 14151 pixels

Objects 51 and 52: 2419 pixels

Objects 31 and 32: 2899 pixels

Objects 61 and 62: 19865 pixels

Figure 4.6. Subtraction of the silhouettes of the two versions of each object, giving the count of the pixels that the silhouettes do not have in common. Strong negative correlations between the size of the difference between the silhouettes and the number of incorrect responses on different-version trials in Experiments 1 to 4 were found (see Table 4.2).

Table 4.2.

Mean Number of Incorrect Responses on Trials Presenting Different Objects, for Experiments 1 to 4.

Object Pairing and Presentation Order	Mean Number of Incorrect Responses E1	Mean Number of Incorrect Responses E2	Mean Number of Incorrect Responses E3	Mean Number of Incorrect Responses E4	Pixel Difference
11 then 12	3.35	2.38	0.79	1.19	6406
12 then 11	2.65	2.04	1.11	1.35	6406
21 then 22	4.57	3.50	1.50	2.12	14151
22 then 21	4.65	3.67	1.39	1.81	14151
31 then 32	4.83	4.92	2.82	2.81	2899
32 then 31	5.04	4.54	3.04	2.54	2899
41 then 42	6.17	6.17	3.57	2.35	5328
42 then 41	5.30	6.08	3.25	2.69	5328
51 then 52	8.83	8.00	4.07	4.73	2419
52 then 51	9.52	8.63	6.61	6.19	2419
61 then 62	3.13	2.00	0.79	0.88	19865
62 then 61	2.91	1.00	0.68	1.12	19865
Correlation between number correct and pixel difference:	r_s=-0.820* p<0.05.	r_s=-0.876* p<0.05.	r_s=-0.864* p<0.05.	r_s=-0.919* p<0.05.	
Correlation between latency to responding and pixel difference:	r_s=0.057 p>0.05.	r_s=0.014 p>0.05.	r_s=-0.410 p>0.05.	r_s=-0.311 p>0.05.	

Note: Pixel difference is not normally distributed (Kolmogorov-Smirnov statistic=0.291, p<0.05, Shapiro-Wilkes statistic=0.813, p<0.05), therefore the non-parametric Spearman's correlations were conducted. Correlations significant at alpha=0.05 are asterixed.

nificantly influenced by the differences between the initial and comparison stimuli.

Fatigue Effects

Future experimental methodologies should aim to address the deficiencies of the first four experiments. A problem with the experiments was the number of trials performed by the participants. Completion of 288 trials proved to be an arduous task for the participants: error rates for many individuals increased towards the end of the session (e.g., Figures 1.12, 1.13, 2.3, and 2.4 show that having demonstrated a reduction in error rate through the middle of the experiment, Participants 1.1, 1.3, 1.7, 1.13, 1.14, 1.16, 1.22, then show a rise in error rate near the end of the experiment, indicating possible fatigue effects). Many of the participants also commented on the length and tedium of the task. The simplest way to alleviate these problems would be to reduce the number of trials each participant must complete.

Variation in Amount of Shadow Across Images

As discussed above, there was a large variation in the amount of shadow present in the Shadow-Condition images, when compared to their corresponding No-Shadow-Condition images (see Table 1.2). Some images contained such a small amount of shadow, that they could easily have been incorporated in the No-Shadow Condition. The amount of shadowing in each of the illumination directions was also analysed (see Table 4.3). This was calculated by subtracting the No-Shadow Condition grey-scale totals from the Shadow Condition totals for the images illuminated from the left, and dividing this by the corresponding difference for the images illuminated from the right. Thus, if the amount of shadowing was approximately equal with regard to each illumination direction, the resulting value would be close to one. The values obtained indicate that there was considerable variability in the amount of shadowing cast upon each object, dependent upon illumination direction, e.g., when Object 5 Version 1 was illuminated from the left, it had over 40 times the amount of shadowing as when it was illuminated from the right, whereas for Object 1 Version 2, the amount of shadowing was approximately equal across the two illumination directions.

In future experimentation, the amount of shadow in different images should be controlled and manipulated. For a single object, a possible way to do this would be to change the angle of illumination of the object. This would change the amount of shadow cast by the object, in a similar way to

Table 4.3.

Ratio of the Amount of Shadow in the Left Illumination Image Over the Amount of Shadow in the Right Illumination Image.

Object Type.	Ratio of amount of shadow when object is illuminated from the left, compared to when illuminated from the right.
11	2.61: 1
12	1.10: 1
21	0.48: 1
22	1.32: 1
31	0.49: 1
32	0.64: 1
41	0.77: 1
42	1.68: 1
51	41.18: 1
52	4.57: 1
61	0.82: 1
62	0.10: 1

Note: The differences were calculated by subtracting the No-Shadow Condition grey scale totals from the Shadow Condition grey scale totals for the left illumination images, and dividing this by the corresponding difference for the right illumination images. Thus, if the amount of shadowing was approximately equal under the two illumination conditions, the resulting ratio would be close to 1:1.

changes in the amount of shadow due to the sun shifting over the course of a day.

Difficulty of the Matching Task

The difficulty of the matching task was not controlled in the previous experiments, with the number of errors for some object pairs being over six times higher than for other object pairs (see Table 4.2, 61 then 62 versus 52 then 51). This lack of control could decrease the likelihood of finding any effect of the presence of cast shadows. The interaction of task difficulty with level of shadowing was assessed by Freeburg (1966). Freeburg (1966) conducted a sequential matching to sample task: matching photos of a model lunar landscape with views of illuminated sections of the model itself. Freeburg manipulated the amount of shadow (through the angle of illumination), and the difficulty of the task through the presence of particulate features. Based upon his results, Freeburg (1966) suggested that the effects of shadows are apparent in an interaction between amount of shadow and task difficulty: where moderate levels of shadowing may be useful in moderately difficult discriminations. He concluded that shadows are subordinate cues for recognition at the extremes of recognisability, where there are either highly dominant shape cues, or these cues are absent, but, that the addition of shadowing may aid recognition in the area between these extremes. It follows that in this series of experiments the utility of having shadows present could differ according to which object was in a trial (as the difficulty of the task varied by object), and again by the varying amount of shadow present across the object types and illumination directions. The lack of control across these two continuums means that the lack of any effect of shadow presence is not conclusive evidence that shadows do not play a part in object recognition.

Stimuli Presentation Durations

Experiment 1 employed the same stimulus presentation durations as Tarr et al. (1998), S1 = 200 ms and S2 = 200 ms. The stimuli presentation durations were manipulated across Experiments 1 to 4, to try to improve the initial low discriminability found in Experiment 1, and to see if varying presentation times would have any impact upon the use of shadow cues for recognition.

The results from Experiments 1 to 4 suggest that presentation times of 200 and 100 ms (for the initial and comparison stimuli respectively) are probably too short to produce reasonable levels of discrimination. If a sufficient level of discrimination is not present, the possibility of being able to manipulate the dependant variables of accuracy, and latency to responding, is minimal. The highest accuracy was achieved using presentation times of 400 and 300 ms for the initial and comparison stimuli respectively. The lengthening of presentation times above these did not improve accuracy. As these display times were also proposed by Biederman and Bar (1999) to be sufficient for the recognition of rendered images, it would be appropriate to retain them in future research. When using the 400 and 300 ms presentation durations, a differential effect of shadow presence was found: a cost in sensitivity for the group with cast shadows present (d'_{NS}=2.31, d'_{S}=1.86) (see Table 1.1). This cost could not be replicated when the presentation times were extended to 600 and 500 ms.

The possible effects of task difficulty on people's use of shadows as cues for recognition have just been discussed. The question arises whether extending the presentation times made the task too easy, removing the effect found in Experiment 3. A comparison of the *d'* values for Experiments 3 and 4 does not indicate that Experiment 4 was easier than Experiment 3: *d'* values in Experiment 4 were very similar to those for the Shadow Condition in Experiment 3 ($d'_{S\,E3}$=1.86; $d'_{S\,E4}$=1.93; $d'_{NS\,E4}$=1.74) (see Tables 1.1 and 2.1).

Sequencing Effects:

It was suggested to the author that the difference between the results of Experiments 1 to 4 (no difference in response time between the Shadow and No-Shadow groups) and the results of Tarr et al. (1998) (lower response times for the Shadow group) could be due to a particular arrangement of the trial order sequence. That is, the particular trial order used in Experiments 1 to 4 may have accidentally shifted the average response time of one group (Shadow or No Shadow) in comparison to the other. This would mask any difference in mean response times that was present in the underlying population, and because the same random order was employed for all participants, this masking would occur across all the participants.

This argument assumes that speed of response on any given trial is affected by the type of trial or trials that preceeded it. Chance arrangements of particular trials would affect the speed of response to the trials following them, and in this manner change the mean response time across the session. In order to mask a difference in average response time in the underlying populations this effect of sequencing would need to: one, operate differentially across the two conditions (Shadow and No Shadow); and two, be of a magnitude sufficient to make the difference between the underlying populations undetectable (but not to reverse the difference). Given that both conditions used the same sequence, this differential effect would have to be caused by the trial sequence in conjunction with the physical differences between the stimuli in the Shadow and No Shadow conditions. That is, a particular random sequence would have to interact with the manipulation of shadow presence to produce the opposite effect to that which is hypothesised to be present in the underlying populations (based upon Tarr et al.'s (1998) result).

The author does not know of any theoretical model that would predict, or even suggest, an interaction between a randomly distributed 288-trial sequence and shadow presence that would result in an effect the same size as that obtained by Tarr et al. (1998), but in the opposite direction. This aside, the distribution of trial types in the sequence used in Experiments 1 to 4 was examined for any irregularities. The distribution of the different initial stimuli was assessed using a Kruskal-Wallis non-parametric ANOVA upon the rank positions. The test found no significant difference in the mean rank presentation-order of the initial stimuli ($\chi^2 = 17.9$, $p > 0.05$), and this was confirmed by graphical inspection (as depicted in Figure 4.7). The likelihood of each stimulus following every other stimulus was assessed graphically for evidence of abnormalities in clustering and no unusual clusterings of trials were evident. The results of these two analyses indicated that the particular sequence of trials employed did not have any abnormal concentration of trials types, or groups of trial types, in parts of its sequence, nor any predictable pattern of stimulus presentation. It was concluded that the sequence used should not have produced shifts in the mean response times of either the shadow or no-shadow groups in comparison to each other.

A second way that the effect observed by Tarr et al. (1998) could have been masked is suggested by relations between trial order and response time that was noted in the data from Experiments 1 to 4: the reaction times for some individuals tended to reduce across the session. To quantify this, the degree of relation (between reaction time and trial order) was assessed for each participant using Pearson's *r*. As Tarr et al. only assessed response time for correct trials, only correct trials were assessed here. The correlations were significant for a proportion of the participants, typically describing a

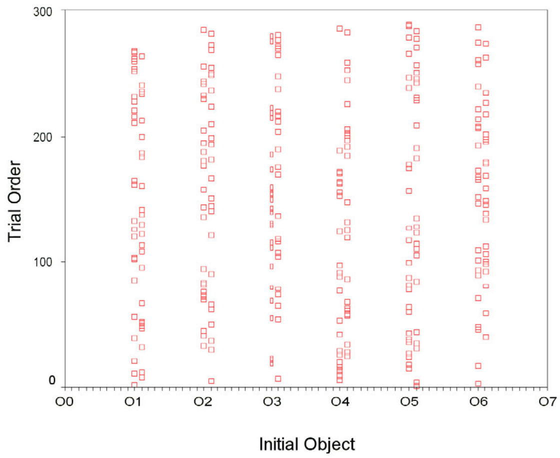

Figure 4.7. Graphical illustration of the distribution of initial stimuli across the trial order. The difference between mean rank order is not significant.

reduction in response time across the session. The average correlation in the Shadow and No-shadow conditions ranged from -0.15 to -0.32 (the average correlation for each condition is shown in Table 4.4). This raised the question that, if there was an order-effect present (i.e., if reaction time reduced with increasing trial number), could this have masked the effect demonstrated by Tarr et al. (1998)? Tarr et al. did not report any trial by trial analysis, so their result can be summarised as a significant difference between the average response time of their two groups (Shadow and No-Shadow), with this difference being equal over trials (depicted in Figure 4.8) (although note that the lack of report of a within-session effect does not imply that there was not one). The possibility that the reduction in response time across trials could mask a difference in mean response time between the groups was then assessed. Adding the same degree of reduction in response time over trials to the data from both experimental conditions would not mask the difference in mean reaction time between the two conditions. It would simply result in a change in the slopes of the regression equations fitted to the two groups of data, but no change in the difference between the means of the groups (illustrated in Figure 4.9).

In order to mask a difference in mean response times between the two groups, the reduction in response times over trials would have to be different for the two conditions. Given the assumption that there was no reduction in response times over trials in Tarr et al.'s (1998) two data sets, the slopes of the regression lines fitted to these two groups must be zero (flat). To mask the difference in means (i.e., to shift one mean relative to the other, so that they are equal), the degree the response times reduce over trials must differ in the two resulting data sets. That is, the regression lines of the two data sets must have different slopes. Such a situation is depicted in Figure 4.10, where the mean response times in one group has changed in comparison to Figure 4.8 but the mean response times in the other group has not. Note that, if the means are equal (as they are in Experiments 1 to 4, and a necessity to mask the hypothesised effect of shadow presence), then the slopes must be different (a necessity to be able to equate the means through response times reducing over trials), and the intercepts must also differ. To check that such a situation had not occurred, response times (for correct trials[1]) were regressed on trial-order for each participant. The slopes and intercepts of the participants' data were averaged over each condition, and are shown in Table 4.4. When the slopes and intercepts were compared across the Shadow and No-Shadow conditions, across all the experiments no differences were significant (assessed by *t*-tests, in all cases $p > 0.05$, shown in Table 4.4). Thus, any reduction in response time across trial order can be taken to be the same in both conditions, and therefore, could not mask a difference in means of the underlying populations.

The finding of a general reduction in response time over trial order raises one last possibility: that the difference in mean response time observed by Tarr et al. (1998) could be masked by the data sets containing a greater degree of variability, making it difficult to detect a real difference between the means of the underlying populations. Accepting this argument, it would appear that if the variation in the data due to the reduction in response times over trials could be removed from the data sets, then the data could be re-analysed and the effect found by Tarr et al. might be present. The variance in each individual's data set that is due to his/her reduction in response time over trials can be calculated using a regression procedure, and once removed leave a residual data set. The residual values represent the amount of unexplained variation around the regression line, and as such, they have a mean of zero. Without any reference point for each data set, individual's mean response times cannot be compared. However, it could be argued that each set of residuals should be referenced against the original mean of the data set, or that each set of residuals should be referenced against the intercept of the regression line. This depends upon whether there is any a-priori reason to presume that the reduction in response time over trials is due to: one, an effect of both increased response times for early trials and decreased response times for latter trials; or two, an effect of only a reduction in response time across trials from the predicted value at the first trial.

In either case, the actual amount of variation present in each individual's data set (whether the total variation, or the variation amongst the residuals) becomes irrelevant when a between groups analysis is performed. For example, the between groups *t*-test only uses one value from each individual:

1. Tarr et al. (1998) analysed response times with respect to only correct responses. The regressions described were also conducted upon the entire data sets for each of Experiments 1 to 4, but the results did not differ to those for correct trials only.

Table 4.4.

Average Correlation, Slope, and Intercept, for the Regression of Response Time on Trial Order, for Each Condition in Experiments 1 to 4. Also shown is Average Response Time of Correct Trials. Only the Response Times of Correct Trials Were Considered.

Experiment and Group	Pearson's r	Slope	Intercept (ms)	Average R.T. (ms)
Experiment 1				
No Shadow Group	-0.19	-0.58	704	618
Shadow Group	-0.23	-0.49	696	622
t(21):	*t*=0.46	*t*=0.33	*t*=-0.15	*t*=0.10
Experiment 2				
No Shadow Group	-0.15	-0.36	744	687
Shadow Group	-0.23	-0.49	814	742
t(22):	*t*=-0.75	*t*=-0.38	*t*=0.77	*t*=0.89
Experiment 3				
No Shadow Group	-0.32	-1.00	872	722
Shadow Group	-0.25	-0.72	810	704
t(26):	*t*=0.86	*t*=0.88	*t*=-0.86	*t*=-0.39
Experiment 4				
No Shadow Group	-0.30	-0.93	831	697
Shadow Group	-0.30	-0.95	878	738
t(24):	*t*=0.01	*t*=-0.13	*t*=0.49	*t*=0.49

Note: Across correlation, slope, intercept, and average response times, none of the differences between the mean of the Shadow group and the No Shadow group were significant (*t*-test: $p > 0.05$).

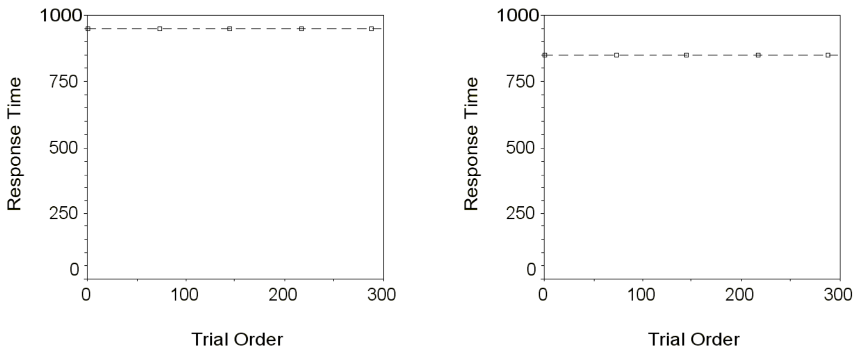

Figure 4.8. Schematic representation of the result of Tarr et al. (1998). The left graph shows a higher mean response time than the right graph.

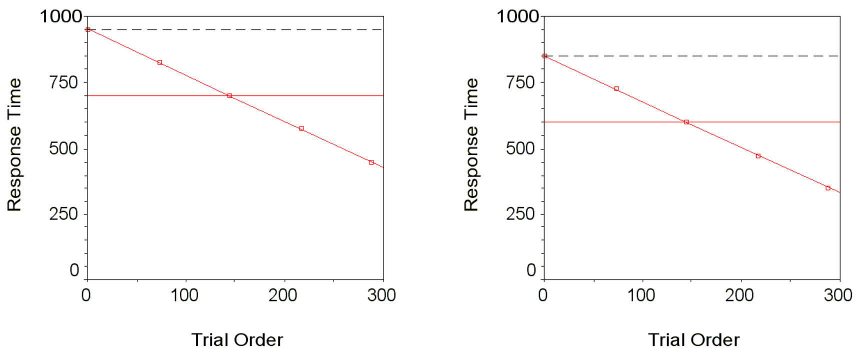

Figure 4.9. Schematic representation of the result of Tarr et al. (1998) after the same reduction in response times across trials is added to both plots. The difference between the means is unchanged. The relation between trial order and response time is represented by the diagonal line, the mean of the data set is represented by the solid horizontal line, and the original mean by the dashed line.

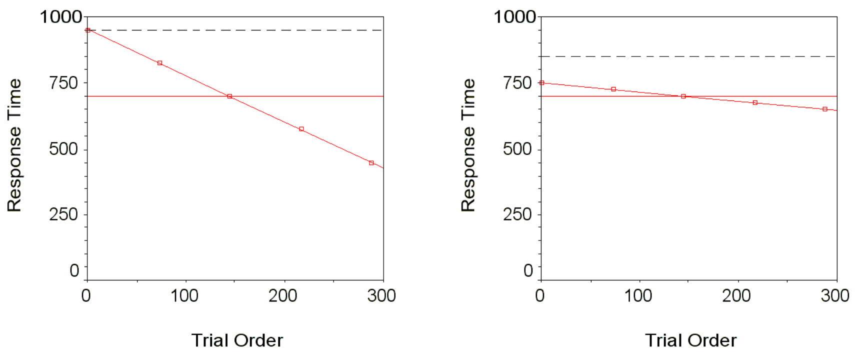

Figure 4.10. Schematic representation of the result of Tarr et al. (1998) after different reductions in response time across trials have been added to each data set. To mask difference in means of the underlying population, the slopes of the regression lines must be different. Note that when the means are equal and the slopes are different, and the intercepts must differ. The relation between trial order and response time is represented by the diagonal line, the mean of the data set is represented by the solid horizontal line, and the original mean by the dashed line.

that individual’s mean response time. As explained, removing the variance due to the reduction in response time over trials does not change the mean response times, and will not change the outcome of the between groups test: there will still be no significant difference between the mean response times of the Shadow and No Shadow groups. If it is argued that the presence of the reduction in response time over trials makes it logical to use the predicted response time at trial number one as the value that best represents each individual's data set, then the *t*-test should be performed using these values. This analysis was performed, and none of the differences between intercepts (predicted response times at trial number one) were significant (the results are displayed in Table 4.4).

Summary

The overall finding of Experiments 1 to 4 was that the presence of shadows in an object-recognition task did not have an effect upon recognition speed or accuracy. Two aspects of the previous experiments’ methodologies meant that it was impossible to rule out the possibility that this failure, to replicate the results of Tarr et al. (1998), was due the procedure itself, rather than the inability of the visual system to use the extra information provided by the attached shadows.

First, the experiments employed images created to mimick those used by Tarr et al. (1998) as accurately as possible. Across these images, the amount of shadowing was inconsistent to a point where some of the images from the Shadow Condition were not readily distinguishable from their equivalents in the No-Shadow Condition (see Table 1.2). Thus, the participants may not have used the information available from the shadows, given that it was inconsistent across the different objects, and often provided very little information in comparison to the physical changes between the objects themselves.

Second, debriefing of the participants indicated that the presence, or absence, of individual parts of the comparison objects was a highly salient cue for matching the sample to the comparison stimulus, i.e., the procedure may have been more akin to a signal detection task, of “feature present / feature absent”, than an actual object-recognition task. The statements of the participants were supported by the analysis of the relationship between bounding-contour change within a trial and the number of incorrect responses. The number of correct responses was highly correlated with the amount of bounding-contour change within a trial: the greater the change the fewer incorrect responses. Furthermore, the task did not require the use of any 3-D shape information, and performance was very highly correlated with the single measure of amount of change in bounding contour between the initial and comparison stimuli. Even if shadows are of benefit in 3-D shape resolution, it seems likely that they would be redundant in a task such as this one. It is not clear why the participants of Tarr et al.'s experiments would not have also used a similar strategy.

Although Experiments 1 to 4 did not replicate the result of Tarr et al. (1998), they did raise some practical issues. When using a Shadow Condition versus No-Shadow Condition manipulation, simply rendering images with and without shadows does not demonstrate a high level of experimental control over the amount of shadow present. A method was employed here where the matching Shadow and No-Shadow images were effectively subtracted from each other, the difference between them returning a quantitative value based upon shadow presence. Current image processing technology (e.g., *Matlab*) allows for the quick and efficient processing of images on a pixel-by-pixel basis. A simple technique was employed here, comparing the summed pixel values of the Shadow images with their matching No-Shadow images, yet the measure successfully quantifies how much darker the Shadow image is in comparison to its counterpart, a value attributable entirely to shadow presence. Although a simple measure, the technique does have some intrinsic finesse. The greater the amount of shadowing, the larger the final value will be. However, each pixel's contribution to the final value is determined by the difference between its value in the No-Shadow image and its value in the Shadow image. If a shadow falls over a dark surface in the No-Shadow image, it will have little effect upon the final value, if it falls over a light surface, it will have a large effect.

A similar technique was employed to investigate whether responding may be based upon gross display changes between S1 and S2 (shown in Figure 4.6). Specifically whether responding was based upon changes in the bounding contour of the objects presented. The technique again used the difference between two images, but in this case, the images were the silhouettes of non-matching S1 and S2 pairs. Each pixel in the silhouettes of S1 and S2 can only take on a value of 0 (black) or 255 (white). Where the images overlap the difference between pixel “x” in S1 and the same pixel “x” in S2 will be zero, where the images do not overlap, the difference between the two pixels will be 255. The number of non-overlapping pixels (those pairs with a difference of 255) can then be counted.

A technique similar to this would be valuable for quantifying changes between different silhouettes either when conducting experiments specifically into bounding contour (e.g., Hayward, 1998), or when trying to ascertain the effect of bounding contour while investigating other aspects of visual perception. It may also be useful in evaluating any differences in the results of procedures using position changes between S1 and S2 compared to procedures that do not change the position of the stimuli (also see: Biederman & Bar, 1999; Braje, 2003; Nederhouser & Mangini, 2001).

Experiment 5 Rotation of the angle of illumination.

It has been argued here that Experiments 1 to 4 had two methodological problems. One, there was no control over the difficulty of the matching task for the different objects being used, and two, there was no control over the amount of shadow present in the images during the Shadow Condition. The analyses of the data from Experiments 1 to 4 indicated that changes in the global shape of the objects were likely to be the major cues used by the participants in performing the matching task.

Whether shadows aid the recognition process is an unanswered question. Research using attached shadows has produced inconclusive evidence (e.g., the results of Freeburg, 1966, and Tarr et al., 1998, versus Braje et al., 1998). Considering cast shadows may help to provide some answers. When an object is illuminated so that shadows are cast on a flat ground plane, the shadows provide a deformation of the object's bounding contour. That is, a deformation of the object's silhouette. It makes sense that we can recognise objects from their silhouettes. We often encounter scenes where cues other than an object's bounding contour are degraded or missing, as in Figure 5.1. Situations such as this occur frequently when the sun is low in the sky, at dawn, dusk, or when levels of ambient lighting are low, such as at night.

Research into the recognition of silhouettes has demonstrated that people are quite adept at object recognition from only the external contour of an object. Lloyd-Jones and Luckhurst (2002) have demonstrated that bounding contour information alone can successfully mediate object recognition and naming, although decisions based upon silhouettes are slightly slower and less accurate than those based upon shaded objects (also see: Hayward, 1998; Hayward, Tarr, & Corderoy, 1999). Furthermore, 3-D shape can be inferred from rotating silhouettes (Norman, Dawson, Raines, & Shane, 2000), suggesting that if 3-D shape is important for object recognition (a debate not central to this thesis), moving silhouettes or shadows can provide this information.

It is mathematically possible to extract information from the cast contours of a shadow about the shape of an object's casting surface, e.g., Knill, Mamassian, and Kersten (1997), and Shafer and Kanade (1983), but the degree to which the visual system can do this has not been resolved. The shape of a shadow has benefited children's judgements of the shape of an

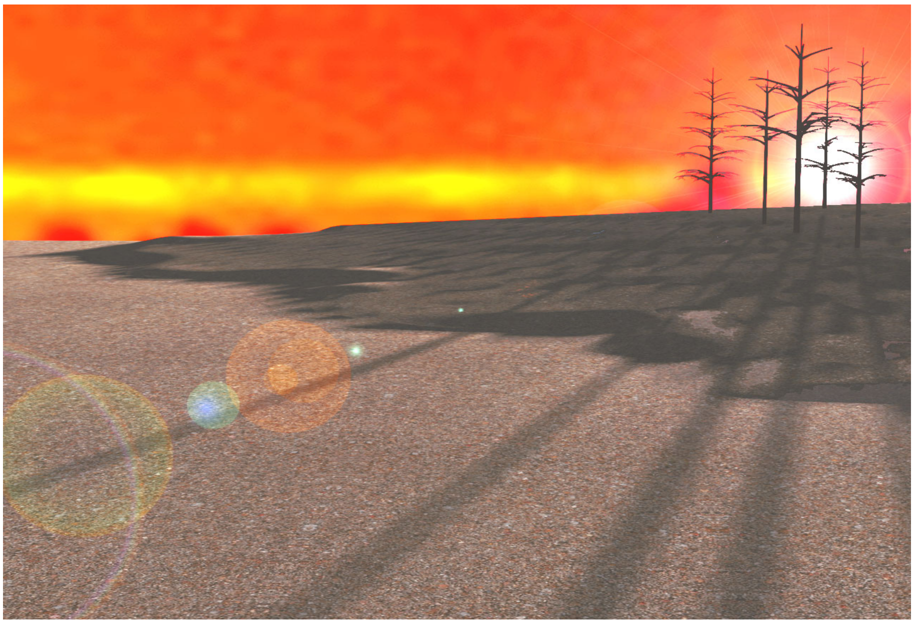

Figure 5.1. Illustration of recognition of objects from their silhouettes alone in a natural setting. We have no difficulty recognising the trees in the image, even though all we can see are their bounding contours.

object viewed in perspective (a sphere versus an ellipse) (Yonas, Goldsmith, & Hallstrom, 1978), and cast shadows have produced reaction time benefits in the recognition of familiar objects (Castiello, 2001). In Castiello's study, reaction times were faster when familiar objects were presented with a cast shadow (and the shadow's shape was congruent with the object's shape and illumination direction) than without.

In contrast, research using more naturalistic shadowing situations, such as attached shadows on faces (Braje et al. 1998), and cast and attached shadows on fruit and vegetables (Braje et al. 2000), has in general failed to find any benefit from the presence of shadows. Even in single experiments, the results are sometimes contradictory. In testing the recognition of digitised images of fruit and vegetables in a verbal naming task, Braje et al. (2000) did find a reaction time benefit of shadows for some specific objects. They suggested that faster performance in the presence of shadows was likely to be a result of enhanced contrast at the edges of the objects, where the object's cast shadows produced a darker background against the lighter coloured foods (Braje et al. 2000). This suggestion parallels that of Cavanagh (1991), in which he proposed that an external, or attached, shadow contour could provide cues for recognition (as illustrated in Figure 4.1).

Given that people can readily recognise objects from their silhouette alone, it is possible that any object recognition benefit derived from the presence of shadows may be maximal when three conditions are met. One, when both cast shadows and attached shadows are present, e.g., Castiello's (2001) research found quicker recognition of familiar objects when cast shadows were present (these shadows were congruent in shape to the object and to the illumination direction), and that these reaction times were faster than the attached shadow only condition (called the no-shadow). Two, when the cast shadows are cast upon a flat ground-plain. Although Cavanagh (1991) relates that cast shadow borders are generally unrelated to object contours, he is focusing his attention on the cast border of an attached shadow, which falls upon an uneven surface, resulting in the shape of the cast contour being a combination of the casting contour and the receiving plane. The cast contour of a cast shadow on a flat ground plane does not suffer from the same distortion. Three, shadows may be useful when there are few other cues available for recognition, as suggested by Braje et al. (2000).

There were several constraints considered for the design of the next experiments, given the methodological problems associated with the first experiments, and the robust effects upon recognition found with silhouettes. Focusing future experimentation on cast shadows looks promising. First, because a cast shadow is a deformation of an object's silhouette, and as such, presents similar cues to a silhouette. Second, because differences in cast shadows will be reflected in the global outline of the combined stimulus plus its shadow, whereas changes in attached shadows will only be reflected in the interior contours of the object itself. Interior contour changes may not be particularly relevant for recognition, at least when viewing novel objects (see Rock et al., 1972).

The research mentioned has covered two extremes of context, from one, where silhouettes have provided the only cues for recognition (e.g., Hayward, 1998, Lloyd-Jones & Luckhurst, 2002), to where relatively naturalistic images were used providing attached or cast and attached shadow cues (e.g., Castiello, 2001). The aim of the next series of experiments was to investigate the continuum between these two extremes, to assess when shadow cues are used by the visual system to aid recognition, and when they are not. To do this, it was first necessary to demonstrate experimental control over either response times or accuracy/discrimination in an object-recognition task, through the manipulation of the amount of shadow present in an image. If control could be exerted over the performance of participants in an object-recognition task, the conditions under which shadows may provide valuable cues for recognition could then be investigated.

In light of the findings of Freeburg (1966), the difficulty of the task needed to be controlled, as any affect of shadows may be regulated by the difficulty of the discrimination. The amount of shadow in the shadow images also needed to be controlled, so that at least an ordinal manipulation of amount of shadow could be conducted.

One way to approach the problem of task difficulty is to use only two different objects. This would result in a constant level of difficulty when discriminating between the objects. Further manipulation of task difficulty could later be performed by manipulating the number of cues available from shading or colour. That is, is it the case that shadowing is only useful in the absence of other cues, as suggested by Braje et al. (2000), or does the use of shadows interact with task difficulty in a manner similar to that described by Freeburg (1966)?

To control the amount of shadow in the images, either a continuum of intensity of shadow, or a continuum of cues available from the shadows is necessary. The latter can be achieved by rotating a light source from the foreshortened axis of an object, to an axis presenting a more informative viewpoint. That is, when an object is illuminated from behind a foreshortened view, the shadow cast is of that foreshortened view, and it presents no more shape cues to the viewer than the foreshortened view does itself. When the object presents the foreshortened view, but the light source is rotated around the object, the shadow cast by the object will present increasing amounts of information about the view orthogonal to the foreshortened view. To minimise the likelihood a ceiling effect in discrimination in a design such as this, the two objects would have to be very similar in their foreshortened views, while still sufficiently different that the task could be completed without the use of any shadowing information.

Of the 3-D models available, rifles fitted these requirements. Two rifles can have very similar ventral silhouettes, while their side profiles present greater variability. The two objects chosen as stimuli were a Generic Rifle with a scope, and an Assault Rifle, with pistol grip and a large magazine. Figure 5.2 presents the objects under four rotations, illustrating the differences in the particulate features and external contour of the different views.

While viewing the ventral surface of the rifles, the direction of illumination can be adjusted so that the shadows cast could provide either essentially the same amount of information as the bounding contour of the ventral view, or an increasing amount of information about the side profile of the rifles. This rotation of the illuminant was performed to produce a procedure that manipulated the amount of information available to the participant provided by the shadows. As the rifles were selected to be very similar in their ventral view, it was expected that when the shadows were projections of the ventral view then discrimination between them would be poor. It was predicted that as the illuminant was rotated, to produce shadows that contained more information about the profile of the rifles, discrimination would improve and reaction times would decrease.

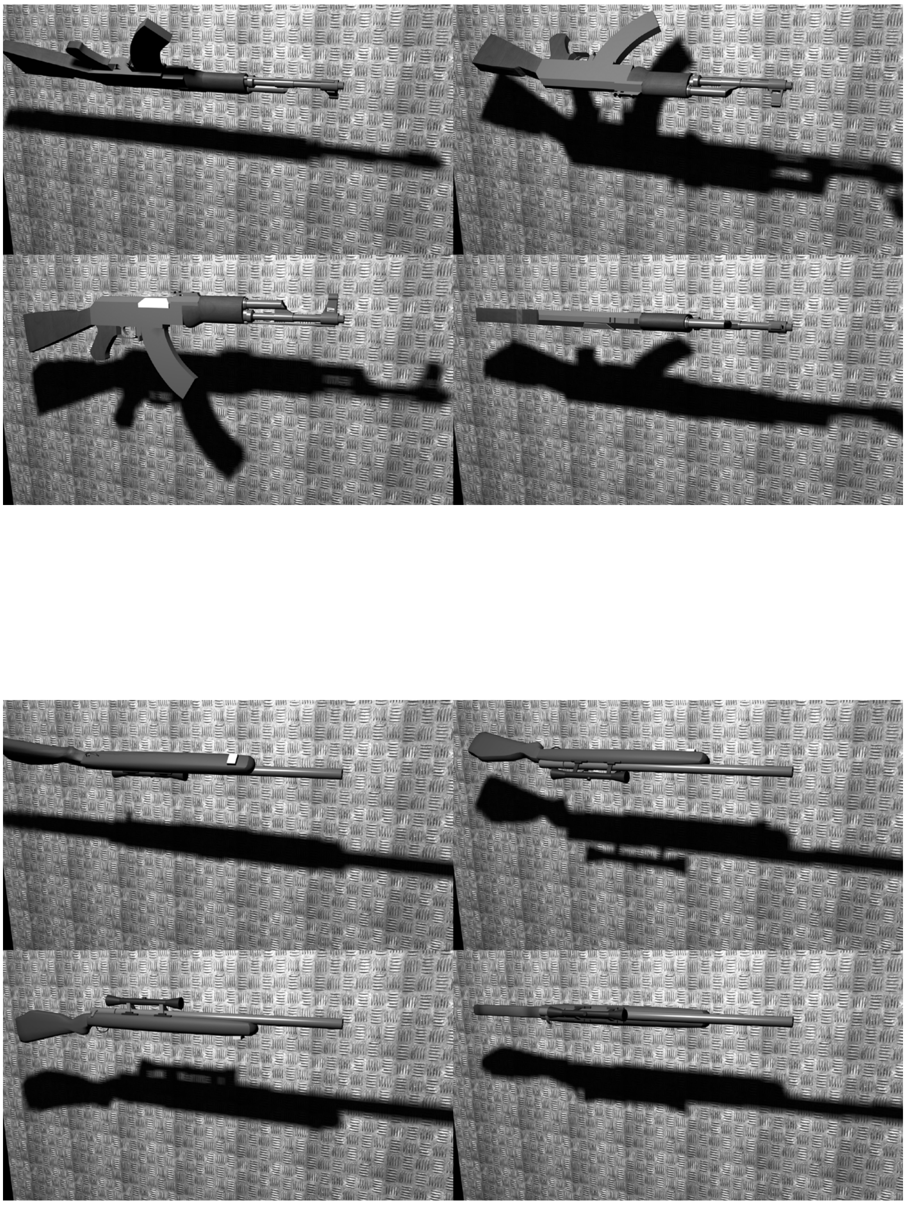

Figure 5.2. Four rotations of the Assault Rifle, top panel, and four rotations of the Generic Rifle, bottom panel. These rifles are similar in bounding contour when viewed from the top or bottom, see the bottom right quadrants, but very different when viewed from the side, see the bottom left quadrants.

Method

Participants

Eight students at the University of Waikato participated in the experiment. Those who were undergraduates participated for course credit.

Apparatus

The experimental sessions were conducted on a Dell Pentium II 400 MHz computer with a 43 cm Trinitron screen with a 75 hertz refresh rate, and on a Dell Pentium III 1.1 GigHz computer also with a 43 cm Trinitron screen with a 75 hertz refresh rate.

Images of a Generic Rifle and an Assault Rifle were rendered using *3D Studio Max*. The images were rendered from models freely available from 3DCafe (http://www.3dcafe.com/asp/meshes.asp (the Generic Rifle, listed as a sniper's rifle, was modelled by Adam Snow; the Assault Rifle, listed as an AK47, was modelled by PMW: foofighterA320@hotmail.com)).

The images were rendered with the guns in an upright position, barrels pointing upward, with the underside of the gun in the direct line of sight. Thus, the side profiles of the rifles were foreshortened to such an extent that they were not visible. The rendering environment in *3D Studio Max* used a global lighting level of 1.5, and the ambient lighting value was set at 230. Each image was rendered with a white background on a white base. The guns themselves were all grey, using a standard material with RGB values of 65, 65, 65, and ambient and diffuse colour locked. Blinn shading was used with shininess = 25, shininess strength = 5, self illumination = 0, and opacity = 100. The objects were positioned at the XYZ co-ordinates of (0, 0, 0) and were approximately 200 units high. The light source was an omni light with a 1.5 multiplier, RGB values of 180, 180, 180, diffuse and specular checked, and no decay. When shadows were required, they were rendered using shadow maps to produce a penumbra and a natural looking shadow. The shadow maps had a map bias value of 2.0, size of 1024, sample range of 5, raytrace bias of 0.2, and absolute map bias checked. The high levels of ambient lighting produced mid-grey, rather than less natural looking black shadows. The images were rendered from a virtual camera positioned at the XYZ co-ordinates of (0, -500, 250) with a target at (0, 200, -100). The camera used a 35-mm lens and a 54.432° field of view.

Seven views were rendered of each rifle (59.9 mm by 59.9 mm at 150 dpi), all from the same camera position. The no-shadow images were rendered with the light source elevated in the line of sight at (0, 550, 600), but without shadows. The zero-degree of rotation images were rendered with the light source and camera in the same positions and with shadows. The remaining images with shadows were then rendered shifting the light source 5°, 10°, 15°, 20° and 30° around the X, Y, Z, co-ordinates (0, 0, 0), in a clockwise rotation with a radius of 550 units, illustrated in Figures 5.3 and 5.4.

The images were reduced to a size of 60 mm by 60 mm at 72dpi in Adobe Photoshop. It was intended that using a screen resolution of 1024 by 768 this would produce an on-screen size of approximately 60 mm by 60 mm, however, the image was approximately 51 mm by 51 mm on the screen, displacing a visual angle of 4.9° at a distance of 0.6 m.

Procedure

Each trial commenced with the presentation of a fixation cross in the middle of the screen for 750 ms. This was followed by the presentation of one rifle stimulus for 400 ms, before the stimulus was covered by a masking stimulus for 750 ms. Two A4 renderings of the rifles (shown in Figure 5.5) were placed between the keyboard and monitor. The picture of the Generic Rifle was placed to the left of the screen and the picture of the Assault Rifle was placed to the right of the screen. When an image was presented on the computer screen, the participants were required to respond by pressing the "Q" key if they thought that the image was that of the Generic Rifle (the left comparison), or "P" if they thought the image was of the Assault Rifle (the right comparison). The participants were asked to respond as fast as they could.

Each session was comprised of 140 trials. There were 14 different images used, seven each of the two rifles. Each was presented 10 times, for a total of 140 trials. The No-Shadow Condition images were presented in trials 1 to 10 and in trials 131 to 140. This was to enable post hoc testing of learning across the session with regard to the objects sans shadows. The images with cast shadows were block randomised in ten blocks of twelve trials, so that each of the twelve different images containing shadows were presented before any one could be presented again. These twelve trial types

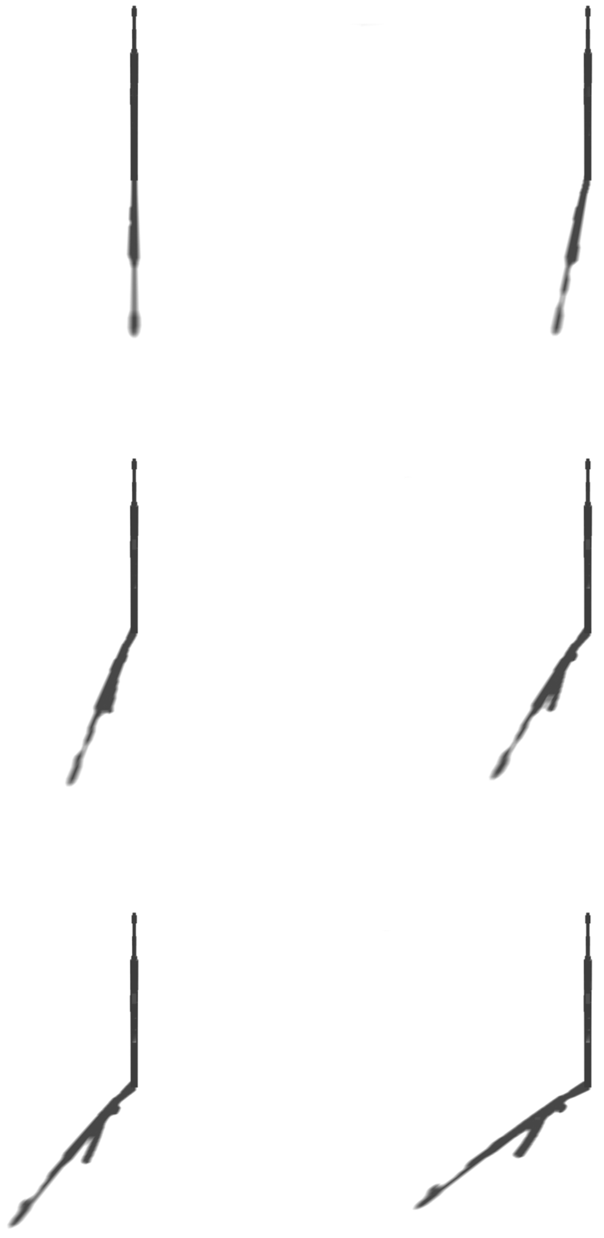

Figure 5.3. Images of the Assault Rifle, as used in Experiment 5. The rotations of the illuminant, clockwise from top-left are 0°, 5°, 15°, 30°, 20° and 10°. The amount of information, from the cast shadow, regarding the profile of the rifle increases as the illuminant is rotated. The No-Shadow Condition is not shown.

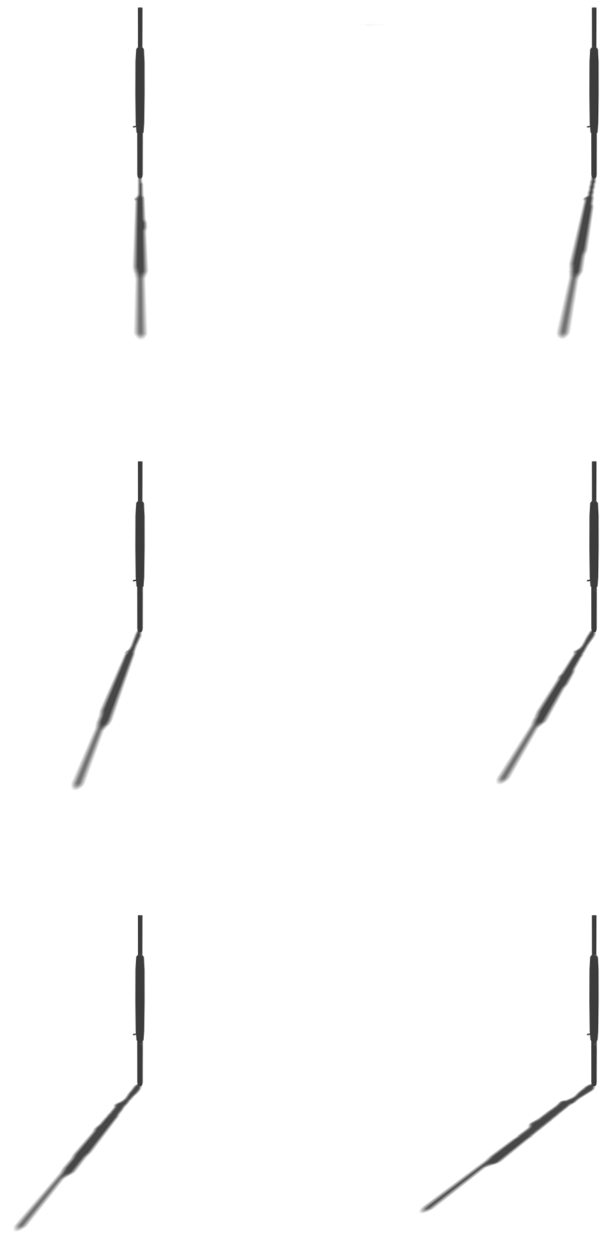

Figure 5.4. Images of the Generic Rifle, as used in Experiment 5. The rotations of the illuminant, clockwise from top-left are 0°, 5°, 15°, 30°, 20° and 10°. The amount of information, from the cast shadow, regarding the profile of the rifle increases as the illuminant is rotated. The No-Shadow Condition is not shown.

Figure 5.5. Images of the Assault Rifle, top, and Generic Rifle, bottom, as used as references by the participants in Experiment 5. The reference plates the participants were given were A4 in size.

were used as the other six conditions for each of the two rifles, and are referred to as: the 0° Condition; 5° Condition; 10° Condition; 15° Condition; 20° Condition; and 30° Condition. Thus, trials 11 to 130 were comprised of ten blocks of 12 randomised trials, each block containing the six conditions for each rifle.

Instructions were given to the participants verbally before the experiment started, and also on the computer screen before they commenced the experiment (see Appendix 5). For each trial, reaction time (in ms) and response ("Q" or "P") were recorded. The participants could respond from when the stimulus was presented until 1500 ms after the masking stimulus appeared. The sessions were self-paced, in that the participants started each new trial themselves by pressing the spacebar key on the keyboard.

Results

The data were analysed with respect to latency to responding and discrimination. In Experiments 1 to 4, discrimination was assessed using *d'*; for ease of comparison to Tarr et al. (1998). In Experiment 5, discrimination was assessed using the measure log *d*[1], and the participants' biases towards the "Generic Rifle" response, or the "Assault Rifle" response, were evaluated using the measure log *c*. The use of log *d* and log *c* as measures of discrimination and bias is discussed in the results section of Experiment 1.

Reaction Time (Correct Trials Only)

On average, the effect of rotation of the illuminant on reaction times was restricted to the trials in which the Assault Rifle was presented (shown in Figure 5.6), where reaction times tended to decrease as the rotation of the illuminant increased. However, Participant 1's (and possibly Participant 7's) reaction times decreased for both trial types, as the amount of shadow increased. A repeated measures ANOVA was conducted on reaction times by amount of rotation of the illuminant. There were two missing values in the data for the Assault Rifle. These occurred where a participant made no correct responses for a particular Illumination Condition. This happened once in the No-Shadow Condition and once in the 0° Condition. Missing values were replaced with the condition mean, so that the ANOVA could be performed without having to remove the rest of the data corresponding to the two participants with a missing value. For the Assault Rifle, longer reaction times were found for the 0° and 5° of Conditions than for the No-Shadow, 10°, 15°, 20°, and 30° Conditions ($F(6,42) = 3.594$, $\eta_p^2 = 0.339$, $p < 0.05$); e.g., reaction time at 0° = 956 ms and at 30° = 657 ms, with an average of 753 ms (± 138 ms). With respect to the Generic Rifle there was no significant effect of rotation of the illuminant ($F(6,42) = 0.807$, $\eta_p^2 = 0.103$, $p > 0.05$); the average reaction time being 703 ms (± 104 ms).

Discriminability

With regard to discriminability, as the angle of illumination was increased discriminability improved (log *d* values for each Illumination Condition are presented in Table 5.1). A repeated measures ANOVA was conducted on discriminability across changes in rotation of the illuminant. Increased rotation produced significant ($F(6,42) = 14.265$, $\eta_p^2 = 0.671$, $p < 0.05$) increases in discrimination up to 10° rotation (the trend is illus-

1. See Appendix 3 for a further review of the derivation of log *d* and the log *c*.

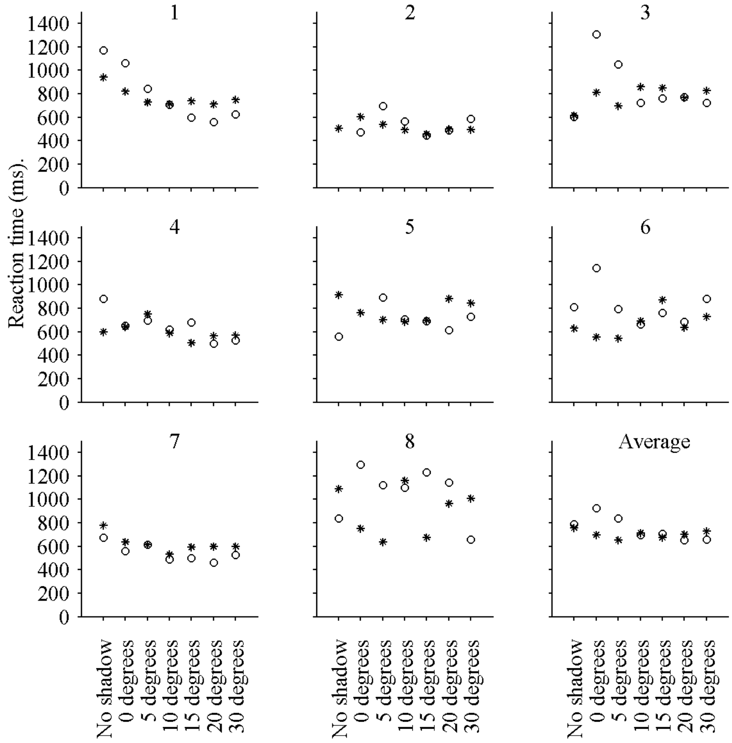

Figure 5.6. Reaction times (ms.) of each participant, and the group averages. Reaction times are given for each angle of illumination and for the no-cast shadow trials. The asterixes represent trials where the Generic Rifle was presented, the circles represent trials where the Assault Rifle was presented. The reaction times are calculated from correct responses only. Participants are identified by the number at the top of each graph.

Table 5.1

Mean Values of Log d for Each Condition

Condition	Mean	Standard Deviation
No Shadow	-0.05	0.18
0 degrees	0.05	0.22
5 degrees	0.32	0.31
10 degrees	0.81	0.40
15 degrees	0.76	0.40
20 degrees	0.86	0.35
30 degrees	0.77	0.49

trated in Figure 5.7). Pairwise comparison's of the means of the individual conditions shows that the No-Shadow and 0° conditions produced the lowest discrimination (chance levels); the 5° Condition produced slightly greater discrimination (log *d* = 0.315), and the 10° to 30° Conditions produced the highest levels of discrimination (averaging log *d* = 0.80 (2dp), comparisons of all means are presented in Table 5.2). To provide a perspective for these log *d* values, a conversion between percent correct (with no biases), and log *d,* is given in Table 5.3, (calculated with 10 trials for each initial stimulus and using the Hautus correction, e.g.: a log *d* value of 0.80 corresponds to a bias free 90% correct, while a log *d* value of 0.33 corresponds to 70% correct).

Each individual's, and the average, level of discrimination (log *d*) by Illumination Condition is plotted in Figure 5.7. A trend of increasing levels of discrimination with increasing rotation of the illuminant, is evident across seven of the eight participants (Participant 8 failed to show any evidence of discrimination during the experiment). Figure 5.7 also indicates that the relationship between discriminability and rotation is not a linear one; there was no improvement in discriminability above the 10° rotation level.

Bias

An ANOVA was performed to assess the size and significance of any change in bias with illuminant rotation ($F(2.7,18.6) = 3.563$, $\eta_p^2 = 0.337$, $p < 0.05$; as the assumption of sphericity was violated, the Greenhouse-Geisser adjustment of degrees of freedom was used). Across the subjects, bias (log *c*) towards saying that the image was of the Generic Rifle decreased as the level of rotation increased. Figure 5.8 depicts this trend: the average log *c* values decreased from 0.33 and 0.50 in the No-Shadow and 0° rotation conditions to an average of 0.06 across the rotation levels 10° to 30°.

Learning

Learning across the course of the experiment was evaluated with respect to the No-Shadow Condition. A within-groups ANOVA with trial position (first 10 trials versus last 10 trials) and rifle type as the two independent variables, and percentage of correct responses as the dependent variable, was performed. The main effect of trial position was not significant ($F(1,7)=3.115$, $\eta_p^2 = 0.308$, $p>0.05$), indicating that there was no improvement in accuracy across the course of the experiment; while the main effect of rifle type was significant ($F(1,7)=5.618$, $\eta_p^2 = 0.445$, $p<0.05$), the Generic Rifle was correctly identified 62.5% of the time compared to the Assault

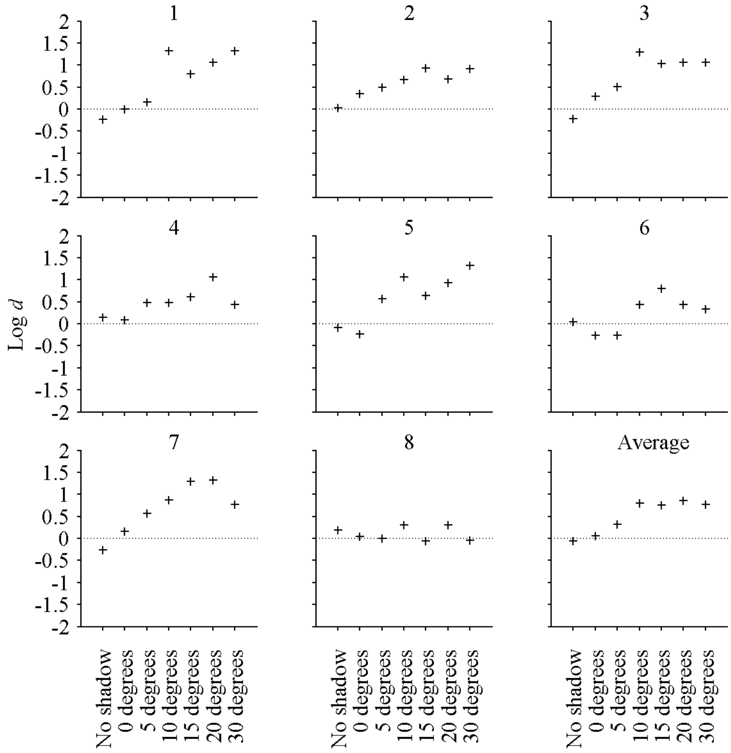

Figure 5.7. Accuracy of each participant, as measure by log *d*, and the average across all participants. A Hautus correction was employed in the calculation of log *d*. The horizontal lines indicate chance discrimination. Participants are identified by the number at the top of each graph.

Table 5.2

Experiment 5: Pairwise Comparisons of Average Log d For Each Condition.

(I) SHADOW	(J) SHADOW	Mean Difference (I-J)	Std. Error	Sig.[a]	95% Confidence Interval for Difference[a] Lower Bound	Upper Bound
No Shadow	No Shadow					
	0 degrees	-.105	.109	.365	-.363	.152
	5 degrees	-.367*	.147	.041	-.714	-1.934E-02
	10 degrees	-.858*	.197	.003	-1.324	-.392
	15 degrees	-.811*	.192	.004	-1.264	-.357
	20 degrees	-.909*	.175	.001	-1.323	-.495
	30 degrees	-.818*	.224	.008	-1.347	-.289
0 degrees	No Shadow	.105	.109	.365	-.152	.363
	0 degrees					
	5 degrees	-.261*	.096	.029	-.488	-3.467E-02
	10 degrees	-.752*	.149	.001	-1.105	-.400
	15 degrees	-.705*	.137	.001	-1.029	-.382
	20 degrees	-.804*	.125	.000	-1.098	-.509
	30 degrees	-.713*	.184	.006	-1.148	-.277
5 degrees	No Shadow	.367*	.147	.041	1.934E-02	.714
	0 degrees	.261*	.096	.029	3.467E-02	.488
	5 degrees					
	10 degrees	-.491*	.132	.007	-.803	-.179
	15 degrees	-.444*	.133	.013	-.759	-.129
	20 degrees	-.542*	.086	.000	-.745	-.340
	30 degrees	-.451*	.144	.017	-.792	-.111
10 degrees	No Shadow	.858*	.197	.003	.392	1.324
	0 degrees	.752*	.149	.001	.400	1.105
	5 degrees	.491*	.132	.007	.179	.803
	10 degrees					
	15 degrees	4.709E-02	.136	.739	-.274	.368
	20 degrees	-5.149E-02	.108	.649	-.308	.205
	30 degrees	3.965E-02	.075	.614	-.138	.217
15 degrees	No Shadow	.811*	.192	.004	.357	1.264
	0 degrees	.705*	.137	.001	.382	1.029
	5 degrees	.444*	.133	.013	.129	.759
	10 degrees	-4.709E-02	.136	.739	-.368	.274
	15 degrees					
	20 degrees	-9.857E-02	.104	.374	-.344	.147
	30 degrees	-7.437E-03	.150	.962	-.362	.347
20 degrees	No Shadow	.909*	.175	.001	.495	1.323
	0 degrees	.804*	.125	.000	.509	1.098
	5 degrees	.542*	.086	.000	.340	.745
	10 degrees	5.149E-02	.108	.649	-.205	.308
	15 degrees	9.857E-02	.104	.374	-.147	.344
	20 degrees					
	30 degrees	9.114E-02	.136	.524	-.230	.412
30 degrees	No Shadow	.818*	.224	.008	.289	1.347
	0 degrees	.713*	.184	.006	.277	1.148
	5 degrees	.451*	.144	.017	.111	.792
	10 degrees	-3.965E-02	.075	.614	-.217	.138
	15 degrees	7.437E-03	.150	.962	-.347	.362
	20 degrees	-9.114E-02	.136	.524	-.412	.230
	30 degrees					

Based on estimated marginal means

*. The mean difference is significant at the .05 level.

a. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

Table 5.3.
The Correspondence Between Percentage Correct and Log d, When There is Zero Bias Present.

Percent Correct	log *d*
100.0%	1.3222
97.5%	1.1357
95.0%	1.0000
92.5%	0.8921
90.0%	0.8016
85.0%	0.6532
80.0%	0.5315
75.0%	0.4260
70.0%	0.3310
65.0%	0.2430
60.0%	0.1597
55.0%	0.0792
50.0%	0.0000

Note: Log *d* values calculated using a Hautus correction (adding a constant 0.5 to each cell of the signal detection matrix), and using a total of 20 trials (10 of each initial stimulus type).

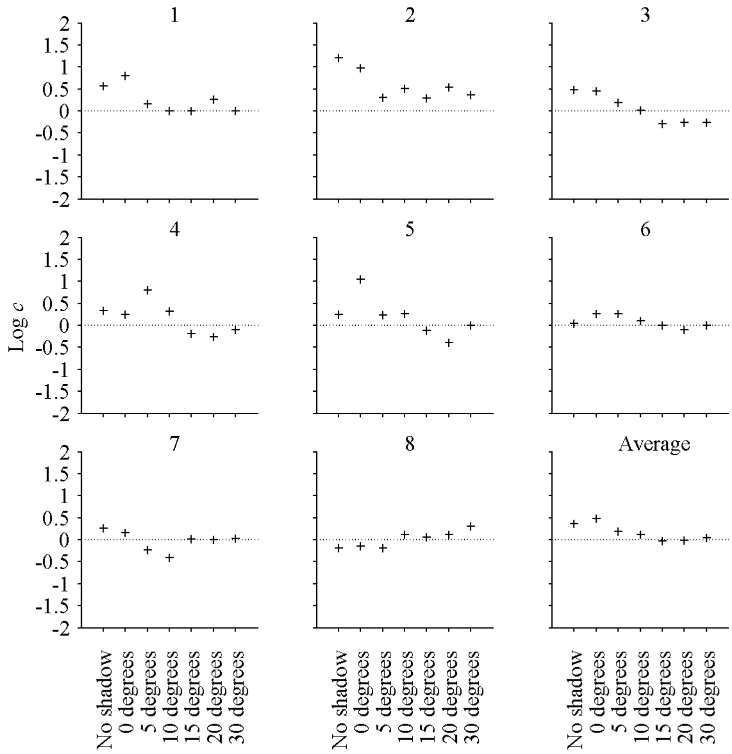

Figure 5.8. Bias for each participant, as measure by log *c*, and the average across all participants. The horizontal lines indicate zero bias, data points above the line indicate a bias towards responding that the image was of the Generic Rifle. A Hautus correction was employed in the calculation of log *c*. Participants are identified by the number at the top of each graph. The averages indicate that there was a general trend across participants to responding that the object was the Generic Rifle at lower levels of rotation.

Rifle 28.75%. The interaction between trial position and rifle type was not significant (F(1,7)=0.054, $\eta_p^2 = 0.008$, p>0.05).

Summary

In summary, latency to responding was associated with rotation for the Assault Rifle: responses were faster when the rotation of the illuminant was above 5°. Discriminability (log *d*) was strongly associated with rotation of the illuminant, where discriminability increased with increasing rotation up to 10°. For at least five of the eight participants, in the conditions with lower rotation of the illuminant, there was a bias towards responding that the stimulus was the Generic Rifle. Analysis of the percentage of correct trials in the No-Shadow Condition indicated that no learning took place in the No-Shadow Condition across the course of the experiment.

Discussion

Experiments 1 to 4 lacked experimental control over the difficulty of the discrimination between the pairs of objects employed, and the amount of shadowing presented by the objects. In controlling these factors, Experiment 5 employed a very constrained procedure, using only the foreshortened views of two objects, and varying the information available from shadowing systematically. Thus, Experiment 5 was a simple experiment, with a simple aim: it was designed to demonstrate control of the dependent variables, discrimination and reaction time, through the manipulation of the cues to object shape provided by cast shadows.

The results indicate that for discrimination (log *d*) this aim was achieved. Discrimination improved as the information from the shadows regarding the profile of the objects was increased. For correct responses, reaction times decreased for the Assault Rifle as the angle of illumination was increased, whereas, reaction times were consistent across illumination rotations for the Generic Rifle. A small bias, towards the Generic Rifle response, was found at low levels of rotation.

The general bias towards the Generic Rifle, present when the shadows were foreshortened, suggests that the participants may have been using the presence of cues only associated with the Assault-Rifle shadows in order to make their decisions. There were two distinct differences between the profiles of the rifles, the presence of the magazine, and the pistol grip stock. Both of these features were present only on the Assault Rifle. If a decision criterion was based upon the presence of these features, then when the features were partially or totally hidden, due to foreshortening of the stimulus, a bias would be expected towards responding that the object was the Generic Rifle. This bias was present; the result indicates control of responding by the cast shadows, rather than the rifles themselves.

The cast shadows also controlled discrimination, and to lesser degree, response latency. The shadow manipulation accounted for 67.1% of the variance in discrimination ($\eta_p^2 = 0.671$), and 33.9% of the variance in reaction time for the Assault Rifle ($\eta_p^2 = 0.339$). The presence of the objects themselves appears incidental to the results: the shadow manipulation accounted for a high percentage of variance in discrimination, there was very poor discrimination in the 0° trials (as shown in Figures 5.7 and 5.8), and at the end of the experiment, the No-Shadow trials were being performed with chance accuracy (47.5%). Thus, the aim of the experiment was achieved, cast shadows controlled discrimination, but the task no longer appeared to be one of object recognition. The task now appeared like shadow recognition; the results offered little indication that the task involved discrimination of the rifles themselves.

The aim of this experiment was to investigate whether shadows contribute to object recognition. That is, to demonstrate the benefit of cast shadows in an object-recognition task. Although shadows mediated responding, it cannot be asserted that they contributed to the recognition of the objects present. The next experiment would need to show that object recognition did occur in the first place, to be able to ask if cast shadows contribute to object recognition.

In Experiment 5 the rifles were presented in a foreshortened view, so that the discrimination would not be too simple. The results of Freeburg (1966) had indicated that an effect of cast-shadow presence would be unlikely given an easy discrimination. So is foreshortening the problem? Poor discrimination is not necessarily a result of foreshortening per se. For instance, the author can perform the No-Shadow task in Experiment 5 with very high accuracy. The lack of familiarity with the foreshortened view is a problem: there are cues present in the stimuli to enable recognition in the foreshortened view, the problem arises in whether these cues are attended to or not. For example, picture a screw, nail, or bolt. When viewing the head of each we see a very foreshortened view. Yet we are familiar with these views, and can discriminate between these objects very rapidly. Consider other cylindrical objects, a pen and biro. We are very familiar with these objects too, but we are not familiar with seeing them end on, and may fail to recognise one in such a foreshortened view. If the participants are familiarised with the foreshortened views of the rifles, the question of whether cast shadows contribute to object recognition can then be asked.

Therefore, Experiment 6 investigated whether there would be any effect of the presence of cast shadows if the participants could already perform the task to a high level of accuracy without them. Would the participants use the extra shape cues provided by cast shadows if they could do the task well without them, or would the extra cues be redundant?

Experiment 6 Familiarity with the stimuli

In Experiment 5 the stimuli, rifles, were each presented in an extremely foreshortened view, so that their profiles were not visible. The results indicated that the degree of foreshortening of the shadows exerted a high degree of control over discrimination. The effect of foreshortening the rifles was to make any non-foreshortened shadow a more salient cue to recognition than any cues available from the rifles themselves.

At the conclusion of Experiment 5, it was reasoned that foreshortening itself may not be problematic, but that people's level of familiarity with a view, whether foreshortened or not, may control that view's utility. In Experiment 5 there was no indication of control of responding by the rifles themselves, therefore, in the literal sense it did not seem that "object" recognition taking place. Experiment 6 aimed to generate control of responding in the absence of cast shadows, to produce an object-recognition task, rather than a shadow recognition task. The contribution of cast shadows could then be investigated.

A familiarisation period was used so that the participants could learn to discriminate between the foreshortened views of the rifles, in the absence of cast shadows. Thus, when cast shadows were introduced to the experiment, the participants were already familiar with the discrimination without them. This parallels what occurs in natural viewing situations (where we can readily recognise objects without shadows), and allow the testing of the contribution of shadowing to an object-recognition task.

Tarr et al. (1998) raised the possibility that familiarity with objects might decrease the visual system's reliance upon the cues provided by shadows. This proposal was based upon the reaction time benefit they found using novel objects with shadows, contrasted with the reaction time cost due to shadow presence reported by Braje et al. (1998, also see Braje at al.,1996) in a face recognition experiment.

If familiarity does reduce any benefit of the cues available from shadowing, then adding shadows to an already familiar task should not improve discrimination or reduce reaction times. Thus, based upon Tarr et al.'s (1998) postulation, discrimination and response latencies in Experiment 6 should be stable across all levels of shadowing. Stable discrimination across could be due to participants only using the information available from the views of the rifles, or from both the rifles and their shadows. If stable discrimination was found, then future experimentation would be required to differentiate whether it was based solely upon the objects present, or was contributed to by both the object and shadow.

An alternate possibility is that the additional presence of shadows could result in a decline in performance, as found by Braje et al. (1996) and Braje et al. (1998). Braje et al. (1998) suggest that shadows could interfere in the recognition process: by masking informative features; by requiring time to be identified as shadows, and discounted as a source of information about shape; and by adding spurious contours that may be confused with object contours.

One other outcome of this experiment is possible. The addition of shadows could improve performance above the level obtained in their absence (any ceiling effect would hide this effect in terms of discriminability, or a floor effect in terms of reaction time). Contrary to the prediction of Tarr et al. (1998), Castiello (2001) reported a response time benefit due to the presence of congruent cast shadows in a recognition task using familiar objects.

It was predicted that familiarity with the discrimination in the absence of shadows would improve performance in the more foreshortened conditions, to a level similar to that found for the non-foreshortened conditions in Experiment 5. That is, it was predicted that familiarity with the discrimination would result in performance that was stable across the variations in illumination direction, and shadow presence or absence.

Method

Participants

Eight students at the University of Waikato participated in the experiment. Those who were undergraduates participated for course credit.

Apparatus

The experimental sessions were conducted on a Dell Pentium II 400 MHz computer with a 43 cm Trinitron screen with a 75 hertz refresh rate, and on a Dell Pentium III 1.1 GigHz computer also with a 43 cm Trinitron screen with a 75 hertz refresh rate. The images were the same as those used in Experiment 5.

Procedure

In general, the procedure was the same as that used in Experiment 5. However, in Experiment 6 each session was comprised of 200 trials. The first 60 where familiarisation trials, where the two images were presented without shadows. These 60 trials were not used in the data analysis. The next 20 trials were also conducted without shadows being present, and for the purposes of the data analyses these were the No-Shadow Condition trials. The next 120 trials all presented images containing shadow, and corresponded exactly to trials 11 to 130 in Experiment 5 (in Experiment 5, of the total 140 trials, the first and last ten were No-Shadow Condition trials).

Results

Latency to responding and discriminative performance are the measures presented here. The discriminability measure log *d* was calculated. Any biases the Participants had to either the "Generic Rifle" response, or the "Assault Rifle" response, were assessed using log *c*. The first 60 trials were treated as training trials and excluded from the analysis. Twenty trials were available for analysis per Illumination Condition (the No-Shadow, 0°, 5°, 10°, 15°, 20°, and 30° Conditions), as in Experiment 5. Participant 5 discriminated well between the two stimuli, but demonstrated reversed discrimination in the No-Shadow Condition (as shown in Figure 6.1).

Reaction Times (Correct Trials Only)

Reaction times for the two rifle types did not show any changes across Illumination Condition. Figure 6.2 presents the reaction times for each participant and the average across all the participants, there are no trends evident in either the individual plots, or the plot of the means of each condition. The average reaction time for the Generic Rifle trials was 546 ms (± 80 ms), and for the Assault Rifle 512 ms (± 97 ms). Repeated measures ANOVAs confirmed that there was no effect of Illumination Condition upon reaction times ($F(6,42)_{Generic\ Rifle} = 0.634$, $\eta_p^2 = 0.083$, $p > 0.05$; $F(6,42)_{Assault\ Rifle} = 0.749$, $\eta_p^2 = 0.076$, $p > 0.05$).

Discrimination

Across the last 20 trials of the No-Shadow Condition the participants discriminated well between the two images without any shadow present (log *d* = 1.01) (see the No-Shadow Condition in Figure 6.1). However, this level of performance was not maintained across the rest of the experiment (as illustrated in Figure 6.1 and Table 6.1). Once shadows were introduced to the experiment, the discriminative performance of five of the eight participants decreased. On average, this decrease was countered by the higher levels of rotation of the illuminant. The graphs of log *d* (in Figure 6.1) illustrate that discrimination, when shadows were present, was on average poorer than at the end of the training period. For some of the participants there was a trend for discrimination to increase as the angle of rotation increased. This trend is visible in the results of Participants 1, 2, 4, and 5 (if ignoring the reversed discrimination during training). However, the results of Participant 6 shows a disruption of discrimination when shadows were added to the task, but no evidence of a clear trend in discrimination across the changes in illu-

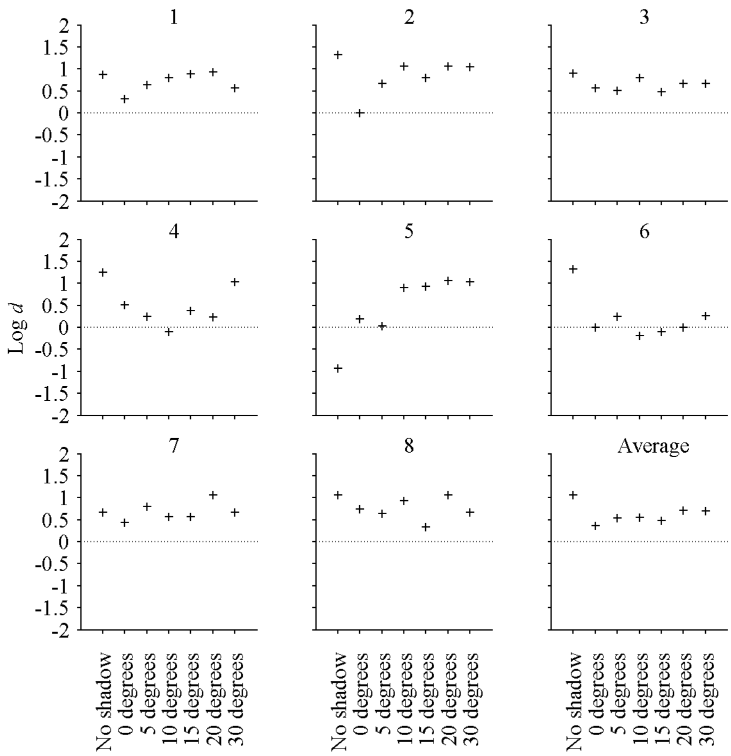

Figure 6.1. Discrimination in Experiment 6, as measured by log *d*. Participant number is given at the top of each graph. Participant 5 produced reversed discrimination during the initial training period. For the majority of the participants, discrimination during the No-Shadow Condition was at least as good as discrimination in any of the other conditions; this is reflected in the averages. Because of his/her reversed discrimination in the No-Shadow Condition, Participant 5 was excluded from the calculation of the averages. A Hautus correction was employed in the calculation of log *d*.

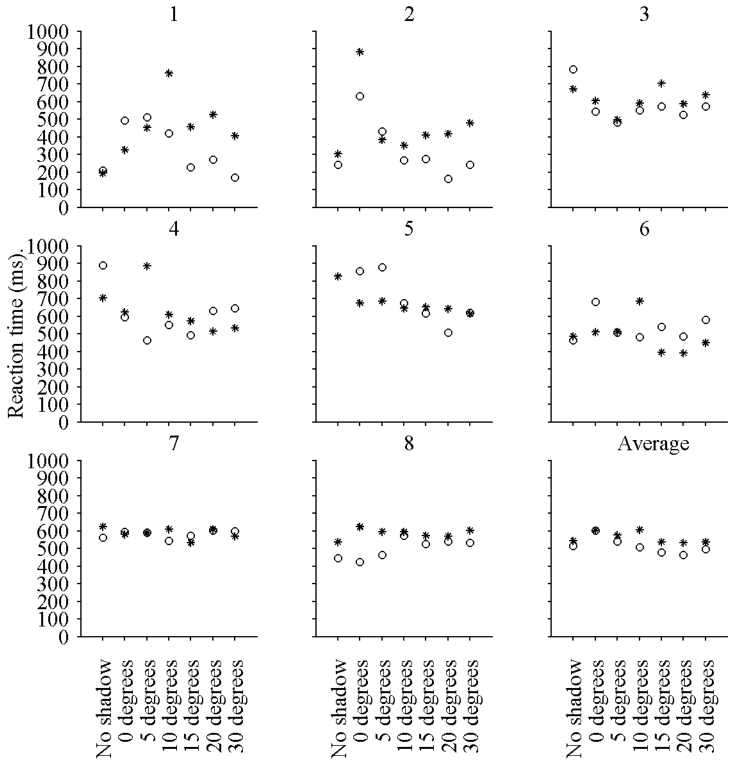

Figure 6.2. Latency to responding by Object Type for correct trials only: Assault Rifle = o, Generic Rifle = *. Participant 5's data are included in the average, as these are correct trials only, thus removing virtually all of his/her responses in the No-Shadow Condition.

Table 6.1

Pairwise Comparisons of Log d for Each Condition.

Condition	Significantly Different Conditions (α=0.05)	log d
No Shadow	0°, 5°, 15°	1.012
0°	No Shadow, 30°	0.364
5°	No Shadow, 20°	0.483
10°	None	0.545
15°	No Shadow, 20°	0.454
20°	5°, 15°	0.730
30°	0°	0.670

minant rotation. Participants 3, 7, and 8, show relatively stable discrimination across the course of the experiment.

When viewed on average, the results look similar to those obtained in Experiment 5. There is a trend of increasing discrimination as rotation of the illuminant increases (log *d* = 0.36 in the 0° Condition, 0.67 in the 30° Condition), but discrimination was highest in the No-Shadow Condition (log *d* = 1.01 in the No-Shadow Condition) (see Table 6.1). The differences between the means were tested using a repeated measures ANOVA assessing log *d* by Condition. Participant 5 was excluded from the analysis because of his/her reversed discrimination in the training period. There was a large and significant effect of Illumination Condition on average discrimination ($F(6,36) = 4.039$, $\eta_p^2 = 0.402$, $p < 0.05$), where the manipulation accounted for 40.2% of the variance in discrimination. Discrimination in the presence of any of the levels of shadowing did not exceed that obtained at the end of the training period where there were no shadows present (log *d* = 1.01 in the No-Shadow Condition). Pairwise comparisons revealed that the No-Shadow Condition had a higher level of discrimination than the 0°, 5°, and 15° Conditions (all comparisons are presented in Table 6.1), while not significantly differing from the 20° and 30° Conditions. Among only those conditions in which shadows were present, the results are less clear, but discrimination was generally lower in the lower rotation conditions: discrimination across the 0°, 5°, 10°, and 15° Conditions did not differ, while discrimination in the 20° and 30° Conditions was sometimes higher than in the lower rotation conditions (see Table 6.1). This effect is not as marked as that seen in Experiment 5 (compare to Table 5.2), where discrimination got progressively better across the 0° to 10° rotation conditions.

Bias

As was found in Experiment 5, bias (log *c*) was associated with shadow foreshortening. A relationship (depicted in Figure 6.3) is evident the averages, where bias shifts from towards the Generic Rifle response at low rotations of the illuminant (log *c* = 0.28 in the 0° Condition), to a bias towards the Assault-Rifle response at the larger rotations (log *c* = - 0.28 in the 30° Condition) (further log *c* values are presented in Table 6.2). The results of Participants 1, 4, 5, 6, and 8 all illustrate this trend, while those of Participants 2, 3, and 7 show no trend in bias across the rotation levels. A repeated measures ANOVA assessing log *c* by Illumination Condition con-

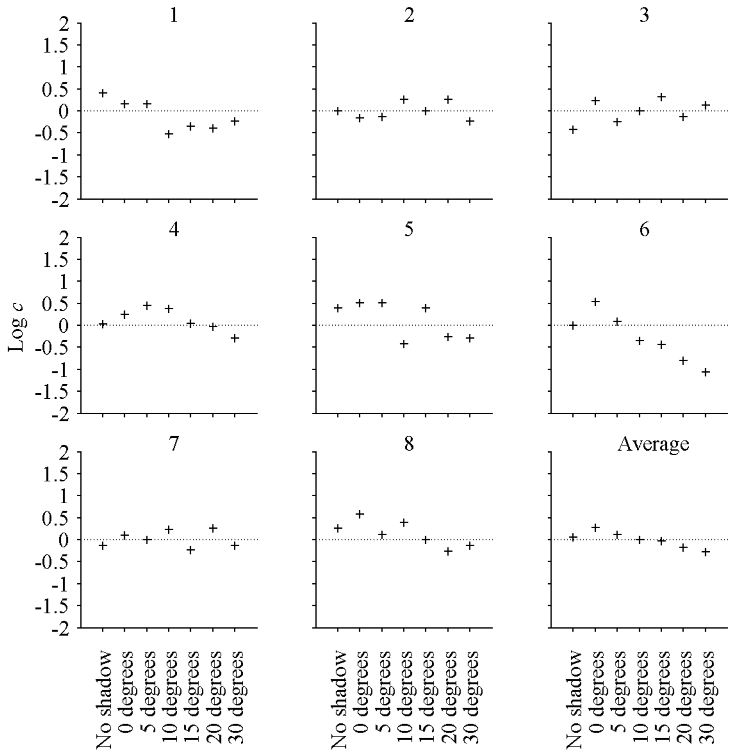

Figure 6.3. Bias, as measured by log *c*, for each participant. Participant number is given at the top of each graph. The horizontal lines indicates zero bias, data points above the line indicate a bias towards responding that the image was of the Generic Rifle. A Hautus correction was employed in the calculation of log *c*, and Participant 5's data are included in the analysis.

Table 6.2

Pairwise Comparisons of Log c for Each Condition.

Condition	Significantly Different Conditions (α=0.05)	log c
No Shadow	30°	0.065
0°	None	0.276
5°	None	0.117
10°	None	-0.004
15°	None	-0.032
20°	None	-0.171
30°	No Shadow	-0.280

Note: Positive values indicate a bias towards the Generic-Rifle response, and negative values indicate a bias towards the Assault-Rifle response.

firmed that, bias changed as Illumination Condition changed, and that the manipulation of Illumination Condition accounted for a reasonable proportion of the variance in bias ($F(6,42) = 3.010$, $\eta_p^2 = 0.301$, $p < 0.05$). Post-hoc pairwise comparisons indicated that bias in the 30° Condition was towards the Assault Rifle compared to no bias in the 15° Condition, and a bias towards the Generic Rifle in the 0° Condition.

Summary

The participants were given the opportunity to learn to discriminate between the two foreshortened views of the Generic Rifle and Assault Rifle, prior to the introduction of cast shadows. When the cast shadows were introduced, discrimination dropped for five of the participants in the more foreshortened shadow containing conditions. For four of those five participants, discrimination in the 20° and 30° Conditions remained at levels similar the No-Shadow Condition. For the other three participants, there was no reliable change in discrimination after the shadow trials were introduced. As in Experiment 5, there was a bias towards responding that the sample stimulus was the Generic Rifle in the low illuminant rotation conditions, and a bias towards responding that the sample stimulus was the Assault Rifle in the high illuminant rotation conditions.

Discussion

To enable the participants to discriminate between the rifles, familiarity with their foreshortened views was trained in Experiment 6. It was hoped that this would enable the assessment of any contribution to object recognition through the presence of cast shadows. Training familiarity was successful: in the last twenty trials of the No-Shadow familiarisation condition discrimination was very high (log d = 1.01). Figure 6.4 shows pattern of discrimination after the introduction of cast shadows. In contrast to the stable performance expected across Illumination Conditions, or a benefit of shadow presence being evident, the pattern was similar to that reported in Experiment 5.

In Experiment 5, shadow foreshortening produced increases in the response times for the Assault-Rifle trials. This effect was not apparent in Experiment 6, where the participants were familiar with the foreshortened views. There was no differentiation of response times resulting from the Illumination Conditions, although response times in Experiment 6 were on average faster than those obtained in Experiment 5. In Experiment 6, trials in which the Generic Rifle was the initial stimulus produced response times of 546 ± 80 ms, versus 703 ± 104 ms in Experiment 5. In Experiment 6, response times for the Assault Rifle were 512 ± 97 ms, versus 753 ± 138 ms in Experiment 5. The increased speed, in comparison to Experiment 5, does not appear to have come at the cost of accuracy: Figure 6.4 shows that discrimination is comparable across Experiments 5 and 6.

While the response time data suggest that familiarisation made the shadow cues redundant, the discrimination results do not. Contrary to the response time results, there was a tendency for shadow presence to disrupt discrimination. In the post-familiarisation conditions, discrimination was related to shadow foreshortening: discrimination was lowest when the cast shadow was of the foreshortened view of the stimulus (log d = 0.36), and it peaked in the 20° and 30° Conditions (log d = 0.73 and 0.63 respectively). The greater amounts of shadow that were interpreted as beneficial in Experiment 5, were associated with a cost in discrimination in Experiment 6; where the No-Shadow Condition now produced the highest level of discrimination (log d = 1.01). In light of the results of Experiment 6, the levels of discrimi-

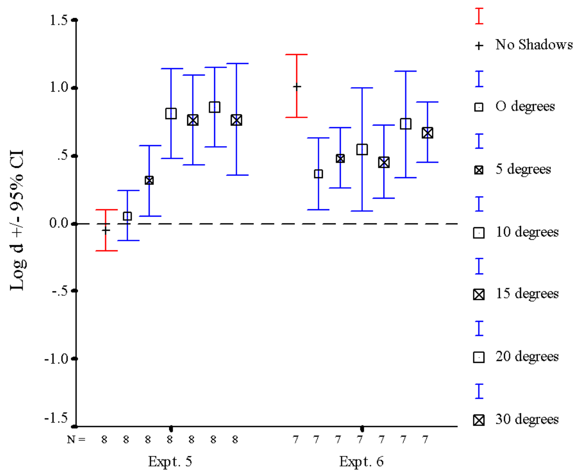

Figure 6.4. Comparisons of average log *d* values ± 95% confidence intervals, across illumination conditions for Experiments 5 and 6. The highest level of discrimination was achieved in the training (No-Shadow) condition of Experiment 6.

Note: In Experiment 6, Participant 5's data was excluded from the log *d* analysis because he/she exhibited reversed discrimination in the training phase.

nation enabled by the cast shadows in Experiment 5 are now interpretable as sub-optimal.

There are three possible reasons why the addition of shadow cues had no effect upon discrimination for Participants 3, 7, and 8: one, the participants only attended to the rifles themselves, making the extra cues redundant; two, the participants used either one of the rifle or shadow cues on different trials, but all to the same effect; or three, the participants only attended to the shadow cues alone, as these were sufficient once discrimination had been learnt for the foreshortened views. The important finding is that there was no benefit of shadow presence evident. This indicates that there is no benefit of the combination of shadow plus object, suggesting that discrimination was a product of only one of the cues, object shape, or shadow shape.

Half of the participants displayed the same pattern of responding (shown in Figure 6.1) as that reported in Experiment 5. This implies that, as in Experiment 5, the shadows, rather than the rifles themselves were mediating recognition. Participant 6 showed no discrimination after the introduction of the shadows, suggesting that the shadows acted as distracters, hindering performance. Therefore, the familiarisation procedure did not generate any evidence of control of responding by the rifles in over half of the participants.

Thus, the significant result for Experiment 6 is that for over half the participants, recognition performance in the presence of shadowing was worse than in its absence. The results suggest that the participants are attending only to the shadows, and not the rifles themselves, as in Experiment 5. However, even if this is occurring, the most foreshortened views of the shadows are very similar to the foreshortened images used for the familiarisation task, and that task could be performed well. These participants had previously demonstrated efficient discrimination between the rifles, so why would their performance now approximate that of the participants in Experiment 5?

As mentioned, decrements in performance due to shadow presence have been reported before. Braje et al. (1998) found a reaction time cost due to attached shadow presence in a face recognition task. They suggested three reasons why this may occur: one, because shadows mask informative features; two, because time is required to identify shadows, and then discount them as a source of information about shape; and three, because people may confuse spurious shadow contours with surface contours.

None of these reasons appear applicable in this context. In Braje et al.'s (1998) task there was the potential for attached shadows to mask information relevant to the face recognition process. However, as illustrated in Figure 6.5, in this experiment, there were minimal cues available for recognition aside from global shape, and global shape could not be masked by any attached or cast shadows.

On average, response times were equally fast across all conditions, countering the possibility that time is required to identify shadows and then discount them as a source of information about shape. Examination of Figure 6.2 also shows that where there could be individual trends in response times, they are more likely to be decreasing as the amount of shadowing increases (as there is less foreshortening of the shadow). Furthermore, as opposed to discounting the shape cues available from shadows, to a large degree discrimination was based upon the shape cues provided by the shadows: manipulating the Illumination Conditions controlled 40.2% of the variance in log *d* across the group.

Spurious shadow contours are unlikely to be confused with surface contours in this experiment, as the cast shadows used in Experiment 6 all portrayed the outer contour of the stimuli. In addition, in some cases they provided more information about the 3-D shape of the stimuli than that present from the stimuli alone. In this sense, the shadow contours are not spurious: unlike the cast contours of attached shadows on faces (see Cavanagh, 1991, who also reasons that cast contours of attached shadows on faces are spurious).

While it is unlikely that the participants misinterpreted shadow contours as contours of the actual rifles, the results show that to a small degree responding, whether correct or incorrect, was controlled by the Illumination Condition. As was the case in Experiment 5, log *c* varied according to the Illumination Conditions. For five of the eight participants there was an orderly swapping of bias as the rotation of the illuminant increased, at the lower rotation conditions, they exhibited a small tendency to respond that the image was of the Generic Rifle, and at higher rotations they exhibited a ten-

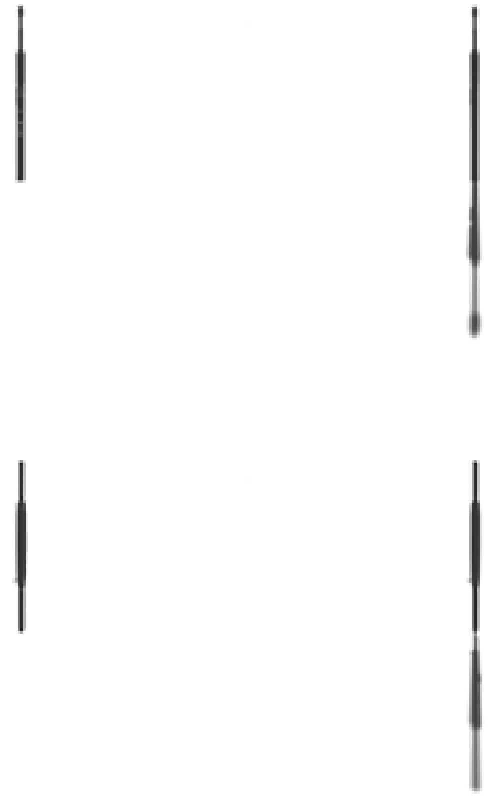

Figure 6.5 Images used in Experiment 6. There are very few cues for attached shadows to mask. Top left is the Assault Rifle in the No-Shadow Condition, top right is the Assault Rifle in the 0° Condition. Bottom left and right are the Generic Rifle in the No-Shadow and 0° Conditions respectively.

dency to respond that the image was of the Assault Rifle. Both rifles did have features apparent in their profiles, but the Assault Rifle had very prominent features. The biases indicate that responding may have been partially based upon a criterion of if there is are any features evident (higher rotations of both rifles) pick the Assault Rifle, if there are no features present (foreshortened views of both rifles) pick the Generic Rifle.

Braje et al.'s (1998), suggestion that shadows may produce spurious information, receives some support from the fact that the variation in shadowing produced a changing bias in the responses. Thus, the shadows could be considered somewhat spurious. However, in this case, any spurious information presented by the shadows has not affected discrimination. The drops in discrimination seen in the conditions with shadow cannot be attributed to the biases present. The presence of a bias does necessitate accuracy at levels below 100% correct, but the discrimination measure log *d* represents accuracy once bias has been accounted for. Therefore, the results show that at high degrees of rotation of the illuminant, shadowing contributed to both high levels of discrimination, and towards a bias (in five of eight participants) to the Assault-Rifle response. It is suggested that this bias at higher rotations is due to the Assault Rifle having more prominent features in profile, and the higher rotations being associated with more features across both the objects. Thus, if a feature is present (likely in a high rotation condition for either object) the participant is likely to be biased to respond that the stimulus was the Assault Rifle (as it has more prominent features than the Generic Rifle at any non-foreshortened view).

Several possible causes of reduced performance in the conditions with shadows can now be ruled out: the masking of informative contours by attached shadows; the extra processing required to identify and discount shadow contours; and being unable to disambiguate shadow contours from object contours. The association of discrimination with the amount of shadow, indicates that reduced accuracy is likely to be due to the participants' reliance on shadow shape as a major cue for recognition, even though it only provides more information than the objects themselves when the illuminant is rotated.

It is suggested that half of the participants used the non-foreshortened shadows shapes as their primary, or only, basis for discrimination: even though they could perform the task accurately without them, and even though their ability to discriminate between the stimuli fell. Perhaps this drop in performance was because once the shadow cues were present, the participants were looking for the very salient shadow cues only available in the higher rotation conditions. This would have resulted in a reduction in performance in the conditions where the shadows were no more informative than the rifles themselves.

The question that was initially posed was whether shadows can contribute to an object-recognition task. There was no observable benefit from the addition of shadows to a familiar object-recognition task for three of the eight participants in this experiment. For the remaining participants, there was little evidence that they were discriminating between the rifles once the shadows were introduced. Four showed a pattern of discrimination suggesting that the shadows present, rather than the rifles themselves, controlled responding, and the other participant did not discriminate between the stimuli.

The results of Experiment 6 suggests that when familiar with the task without shadows, shadow presence may act as a distracter from more reliable (but in this case it seems less salient) object shape cues, reducing performance to those levels obtainable from the shadow information. Thus, shadows may provide discrimination benefits in an unfamiliar task such as Experiment 5, where discrimination in their absence is negligible, but there is no indication that these benefits produce better performance than that achievable from the object shape cues evident when the views of the objects are familiar.

Experiment 7 Foreshortening of the objects

Experiment 6 attempted to investigate the effects of shadow presence on object recognition. The participants were trained so that they were familiar with the foreshortened views of the stimuli, without shadows present. It was anticipated that this would produce a task that did involve recognition of the objects (rifles), as opposed to recognition of their shadows. In five out of eight cases, there was little indication that discrimination was affected by the views of the actual rifles; instead, discrimination appeared to be controlled by the shadows. Consequently, Experiment 6 was not successful in demonstrating that the majority of the participants were performing an object-recognition task (as opposed to a shadow-recognition task).

Salient shape cues were present in the cast shadows of the 10° to 30° Conditions of both Experiments 5 and 6. Similar cues were not available from the rifles. It is advanced that this difference resulted in the majority of the participants in Experiment 6 directing their attention to the cast shadows in the image, rather than viewing the image as a whole, or observing the rifles. The results of Experiment 6 indicate that discrimination was generally controlled by the shadows present, rather than by the rifles themselves. This occurred even though the participants could perform the discrimination more adequately in the absence of the shadows.

For Experiment 7, the aim was to design the task so that the focus was on recognition of the rifles, and how this may be assisted by the presence of shadow cues, as opposed to simply recognition of the shadows that are present. Experiment 6 attempted to shift the participants' attention to the rifles by making the participants familiar with the foreshortened views of them. However, this was only partially successful. In Experiment 7, a different method was employed. The extra shape information that is sometimes available from the cast shadows, was sometimes available from the rifles too. There was a No-Shadow Condition, but not a No-Rifle Condition, so, on any one trial there was a slightly greater probability that the rifle was present more side-profile information than the shadow. It was proposed that this would nullify any advantage the cast shadows may have over the rifles in directing attention to themselves (the results of Experiment 6 suggest the cast shadows controlled discrimination for more of the participants than the rifles did).

The effect of familiarisation with the foreshortened views is difficult to quantify in Experiment 6: three of the participants showed relatively stable discrimination across the Illumination Conditions, and it is assumed that this was a result of the familiarity training. Thus, the procedure used in Experiment 6 did not establish a reliable effect of familiarisation that could be employed in Experiment 7. For instance, if familiarisation training had an effect for half of the participants in Experiment 7, interpreting any effect of rotating the view of the rifles would be difficult. Thus, Experiment 7 did not use the familiarisation procedure.

The combined results of Experiments 5 and 6 give rise to some expectations for Experiment 7. It was predicted that when both the rifle and its shadow were foreshortened, discrimination would be poor, and when one, or both, were not foreshortened discrimination would be high. With respect to response times, it was predicted that the Assault-Rifle times would be fastest when one, or both, of the rifle and shadow were not foreshortened.

Method

Participants

Eight students at the University of Waikato participated in the experiment. Those who were undergraduates participated for course credit.

Apparatus

The experimental sessions were conducted on a Dell Pentium II 400 MHz computer with a 43 cm Trinitron screen with a 75 hertz refresh rate, and on a Dell Pentium III 1.1 GigHz computer also with a 43 cm Trinitron screen with a 75 hertz refresh rate. A new image set was used in Experiment 7, although the images remained the same size as in Experiment 5. The initial stimuli were rendered as in Experiment 5, but in this experiment there were three rotations of the rifles used, 0°, 5°, and 30°, and four levels of shadowing, No-Shadow, 0°, 5°, and 30°. These three levels of illuminant rotation were chosen on the basis of the results of Experiment 5. In Experiment 5, log *d* values shifted from 0.05 at 0°, to 0.32 at 5°, to 0.81 at 10°, and then remained static, being 0.77 at 30°. Thus, of the rotations used the 5° Condition was close to the middle of the range in terms of accuracy, and the 0° and 30° Conditions were at each end of the spectrum. In creating each image, the degree of shadow was calculated from the foreshortened view of the rifle, for instance, the 30° Shadow Condition always presented a shadow cast by a light source rotated 30° from the foreshortened axis of the rifle (this is illustrated in Figures 7.1 and 7.2).

The comparison stimuli were also changed in this experiment. Instead of a single view of each rifle being presented as in Experiments 5 and 6 (as in Figure 5.5), four views were presented on each reference plate (as in Figure 5.2). This change was made in response to the reversed discrimination exhibited by Participant 5 in the No-Shadow Condition in Experiment 6. It was anticipated that providing more information about the shape of the rifles in the comparison images would reduce the likelihood of the recurrence of reversed discrimination.

Procedure

In general, the procedure was the same as that used in Experiment 5. However, for Experiment 7 there were 240 trials, comprised of ten blocks of 24 randomised trials that correspond to the 24 different initial images (presented in Figures 7.1 and 7.2). Unlike Experiment 5, the No-Shadow Condi-

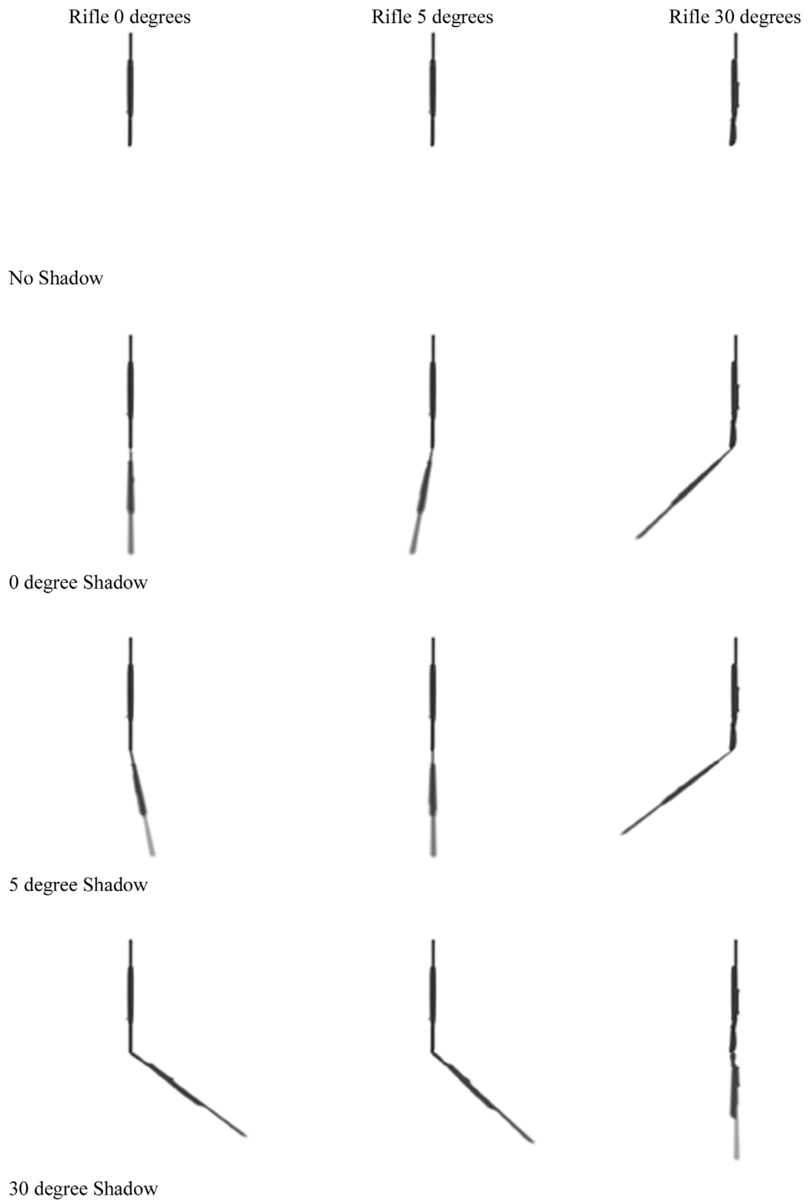

Figure 7.1. Images of the Generic Rifle as used in Experiment 7 (note that the on-screen presentation size was slightly larger than the size of the images presented here). Three rotations of the rifle were used, 0°, 5°, and 30°, in combination with four levels of shadowing: No-Shadow, 0°, 5°, and 30°.

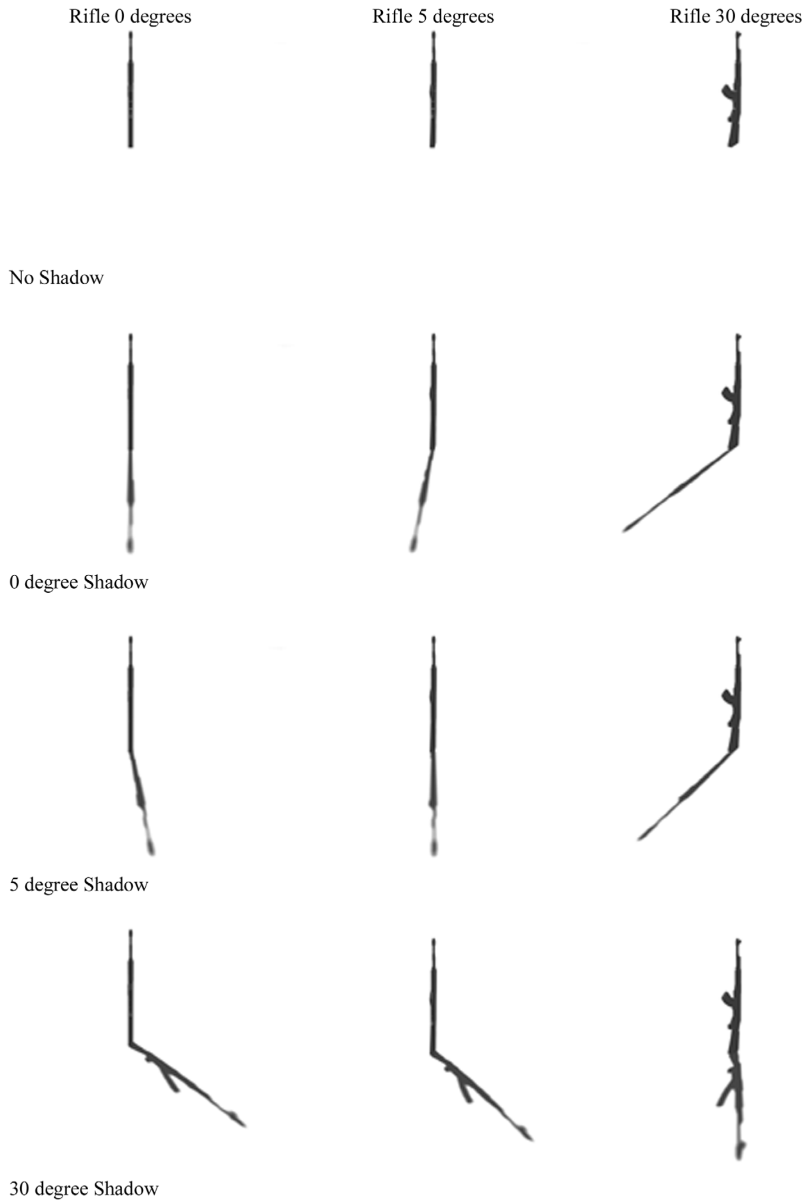

Figure 7.2. Images of the Assault Rifle as used in Experiment 7 (note that the on-screen presentation size was slightly larger than the size of the images presented here). Three rotations of the rifle were used, 0°, 5°, and 30°, in combination with four levels of shadowing: No-Shadow, 0°, 5°, and 30°.

tion trials were randomised among the rest of the trials. It was anticipated that this might assist in directing attention to the entire image, as on one trial in 24 the shadow would be absent.

Results

The participants' responses were analysed with respect to reaction time, discrimination (log *d*), and bias (log *c*). Table 7.1 presents a summary of the results.

Response Latencies (Correct Trials Only)

A two-way repeated measures ANOVA was conducted on latencies to responding for each of the two rifle types, with Illumination Condition, and Rifle Rotation, as the two independent variables. For the Generic Rifle, there was no significant effect upon latencies to responding due to rotation away from the foreshortened views of either the shadow ($F(3,21) = 1.326$, $\eta_p^2 = 0.159$, $p > 0.05$) or the rifle ($F(2,14) = 2.639$, $\eta_p^2 = 0.274$, $p > 0.05$).

The participant's response latencies for correct Generic-Rifle trials are plotted by Illumination Condition in Figure 7.3, and the average latencies are plotted in Figure 7.4. Participants 2, 3, 5, and 8 exhibited no real trends in their response latencies as the Illumination Condition changed from the No-Shadow to the 30° Condition. Of the remaining 4 participants, Participant 7's reaction times decreased as the rotation of the illuminant increased, while the response latencies of Participants 1, 4, and 6 increased, peaking in the 30° Condition. On average there was no significant trend in response latencies across the conditions with shadows present, the average times ranging from 627 ± 95 ms in the 5° Condition to 674 ± 93 ms in the 30° Condition.

As shown in Figure 7.5 rifle rotation increased in the Generic-Rifle trials, the participants generally displayed a trend of increasing response latencies, although Participant 2 showed the reverse trend, and Participant 1 showed no changes in reaction time across rotation levels. Figure 7.6 illustrates the trend of increasing response latencies with increasing rotation in the average response latencies, and in the plots of Participants 3, 5, 7, and 8. However, under statistical testing the means were not significantly different from each other (ANOVA above). As for the Illumination Conditions, rotation of the Generic Rifle produced the longest reaction times in the 30° Condition (688 ± 105 ms), while the 0° Condition produced the shortest (601 ± 93 ms). Thus, there is a trend evident, in the data from about half of the participants, of increasing reaction times as the Generic Rifle or its shadow are rotated, but when assessed across the whole groups' data, this trend is not significant.

Table 7.1

Summary of Results of Experiment 7.

Variable	Measure	ANOVA Result	Effect Size η_p^2	Description
Illumination Condition (Generic Rifle)	Reaction Time	Not Significant	0.159	No general trend.
Rifle Rotation Condition (Generic Rifle)	Reaction Time	Significant	0.274	Decreasing RT as rotation of rifles increased.
Interaction Illum*Rifle (Generic Rifle)	Reaction Time	Not Significant	0.120	No general trend.
Illumination Condition (Assault Rifle)	Reaction Time	Significant	0.329	Increasing RT over NS, 0°, and 5° Conditions, decreased again in 30° Condition.
Rifle Rotation Condition (Assault Rifle)	Reaction Time	Not Significant	0.487	Increasing RT as rotation of rifles increased.
Interaction Illum*Rifle (Assault Rifle)	Reaction Time	Not Significant	0.054	No general trend.
Illumination Condition	Discrimination log *d*	Significant	0.370	Increasing log *d* over NS through to 30° Conditions.
Rifle Rotation Condition	Discrimination log *d*	Not Significant	0.242	Increasing log *d* over NS through to 30° Conditions.
Interaction Illum*Rifle	Discrimination log *d*	Significant	0.413	In the 30° Illumination Condition Rifle Rotation had no effect upon log *d*.
Illumination Condition	Bias log *c*	Significant	0.724	Bias to Generic Rifle response in the NS to 5° Conditions, but bias to Assault Rifle response in the 30° Condition.
Rifle Rotation Condition	Bias log *c*	Not Significant	0.342	Four of the participants showed a bias to the Generic Rifle response in the NS to 5° Conditions, but bias to Assault Rifle response in the 30° Condition.
Interaction Illum*Rifle	Bias log *c*	Significant	0.390	The main effect of Illumination Condition on log c was not apparent in the 30° Rifle Rotation Condition.

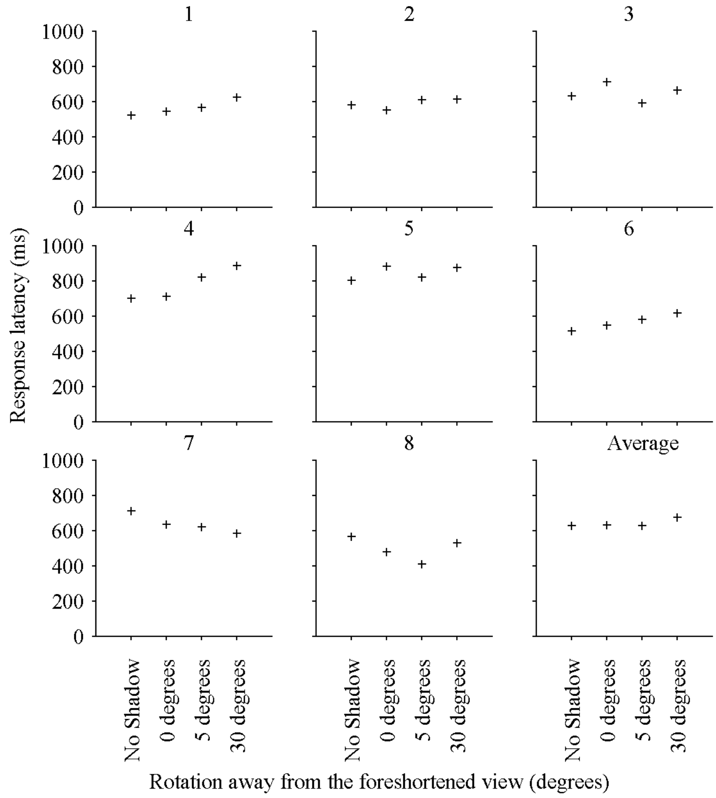

Figure 7.3. Mean response latency for each Illumination Condition and each participant, for the Generic Rifle. Participants are numbered at the top of each graph.

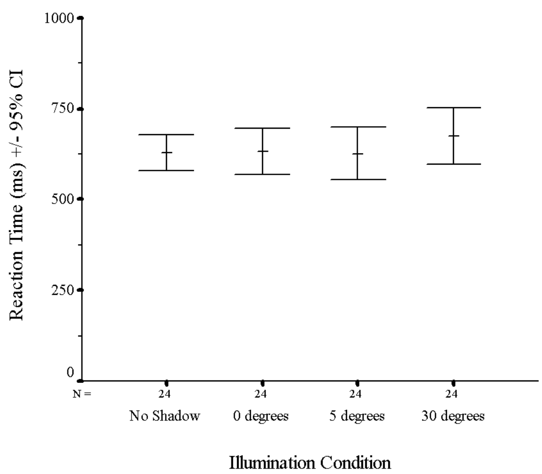

Figure 7.4. Mean response latency for each Illumination Condition for the Generic Rifle.

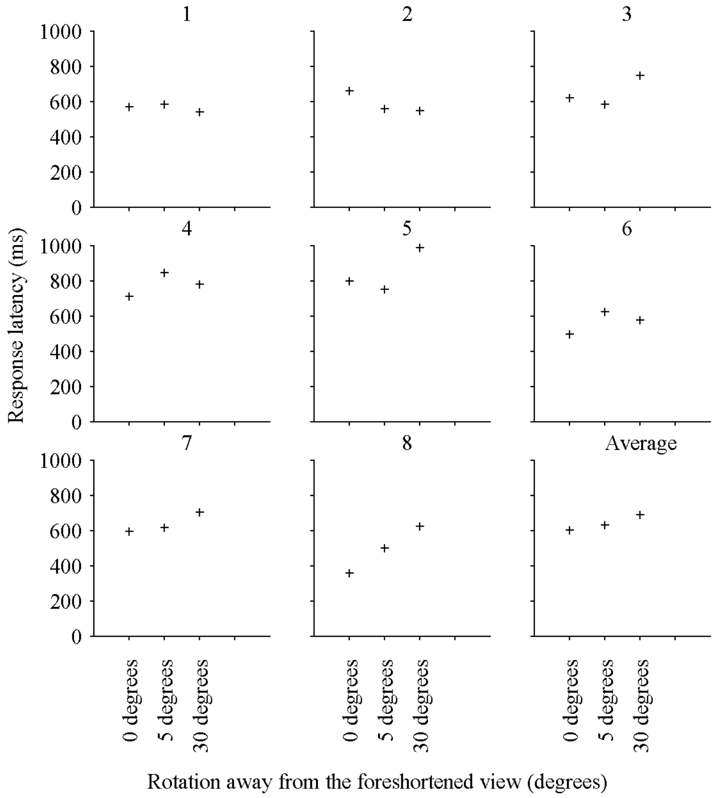

Figure 7.5. Mean response latency for each Rifle Rotation and for each participant, for the Generic Rifle.

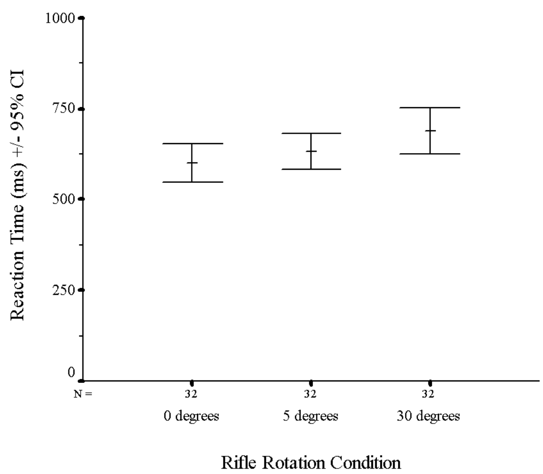

Figure 7.6. Mean response latency for each Rifle Rotation for the Generic Rifle.

For the Assault Rifle, the effect of Illumination Condition upon reaction time was significant ($F(3,21) = 3.426$, $\eta_p^2 = 0.329$, $p < 0.05$), and this trend is illustrated in Figures 7.7 and 7.8. There were three missing values in the data set upon which the ANOVA was conducted, these occurred where the participants did not return any correct responses for a particular trial type. Participants 2 and 5 made no correct responses for the No-Shadow Condition, and Participant 5 also made no correct responses in the 0° Condition. The repeated-measures ANOVA procedure uses equal group sizes in the analysis. Therefore, if a participant has a missing data point in one subgroup (e.g., the 0° Condition), as a remedy, that participant's data are removed from all subgroups of the analysis. This has the effect of removing a considerable amount of valuable data. A second option is to replace the missing value, and therefore retain the participant's data in the analysis. The second option was employed here; the missing values were replaced with the mean for that condition (given that the mean is the most representative value for that condition).

Post-hoc testing using pairwise comparisons of the means of each condition indicated that the 30° Condition had a lower mean (546 ± 95 ms) than the 5° Condition (696 ± 135 ms), but was not significantly different to either the No-Shadow (587 ± 96 ms), or 0° Conditions (653 ± 119 ms) (as illustrated in Figure 7.8). The ANOVA shows a trend of increasing latency to responding over the No-Shadow, 0°, and 5° Conditions, before a reduction again in the 30° Condition. Individually, this trend is evident (in Figure 7.7) for half of the participants, Participants 3, 4, 5, and 6; Participant 8's fastest response latencies were also in the 30° Condition; while Participants 1 and 2, show no trends across the conditions; and the response latencies of Participant 7 increased as foreshortening of the shadows decreased.

When the main effect of Rifle Rotation was assessed for the Assault Rifle, the reverse trend to that suggested for the Generic Rifle was found (as shown in Figures 7.9 and 7.10) ($F(2,14) = 6.656$, $\eta_p^2 = 0.487$, $p < 0.05$; missing values replaced with condition means for Participants 2 and 5 at 0°, and Participant 5 at 5°): the 30° Condition produced the fastest response latencies (532 ± 80 ms), faster than those of the 0° (649 ± 96 ms) and 5° Conditions (679 ± 125 ms). This same trend can be seen in the individual results

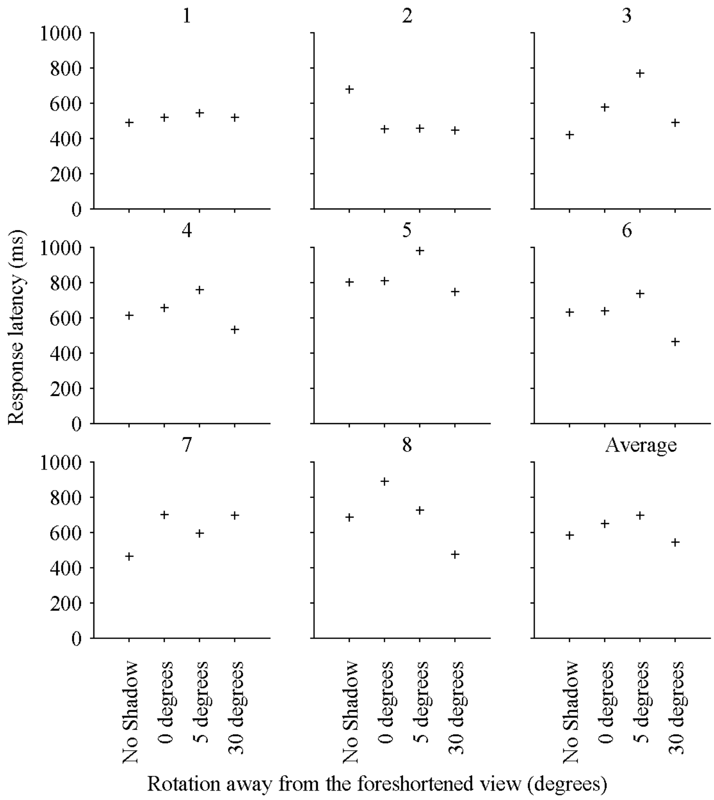

Figure 7.7. Mean response latency for each Illumination Condition and for each Participant, for the Assault Rifle.

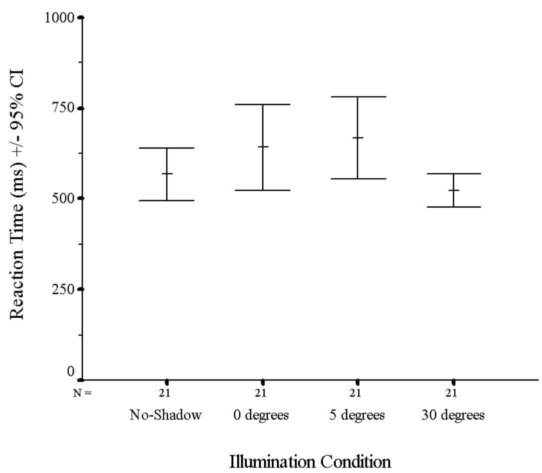

Figure 7.8. Mean response latency for each Illumination Condition for the Assault Rifle.

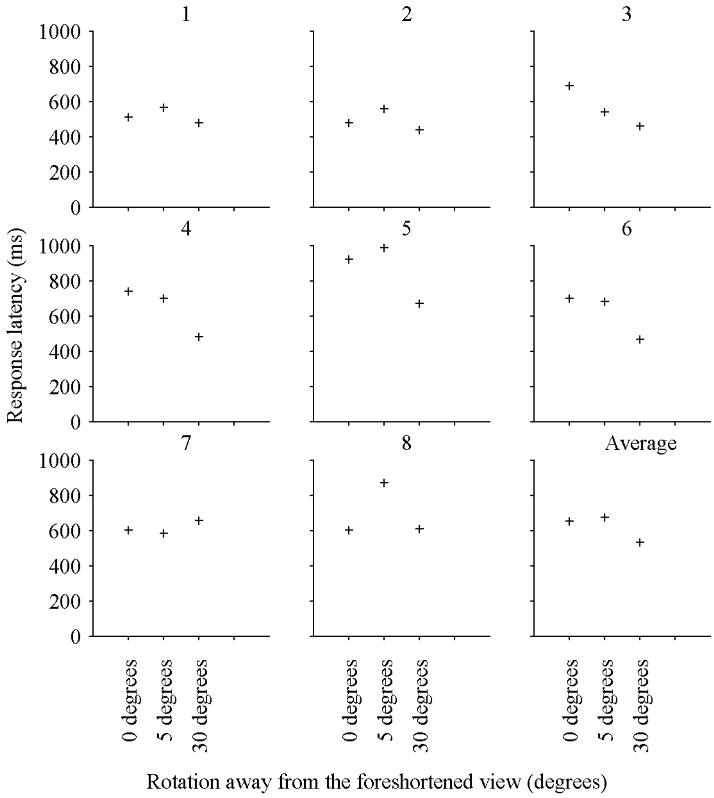

Figure 7.9. Mean response latency for each Rifle Rotation and for each Participant, for the Assault Rifle.

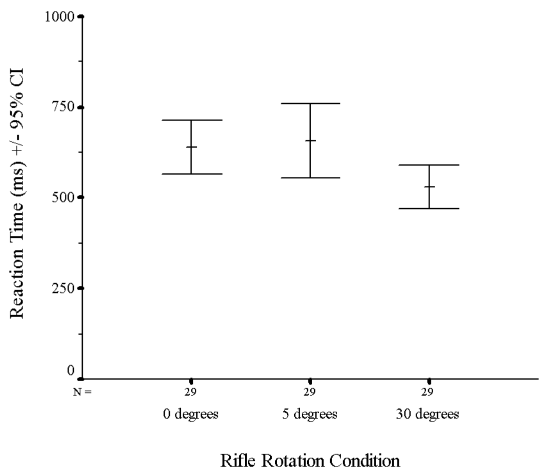

Figure 7.10. Mean response latency for each Rifle Rotation for the Assault Rifle.

of Participants 3, 4, 5, and 6, while for Participants 1, 2, 7, and 8, there is no trend in response latencies across the rifle rotation conditions.

The response latencies indicated that there were no significant interactions between Illumination Condition and Rifle Rotation ($F_{Generic\ Rifle}(6,42) = 0.956$, $\eta_p^2 = 0.120$, $p > 0.05$; $F_{Assault\ Rifle}(2.586,18.102)^* = 0.729$, $\eta_p^2 = 0.054$, $p > 0.05$, *Greenhouse-Geisser adjustment used as the assumption of sphericity was violated). Figures 7.11 and 7.12 present response latencies for each participant, and on average, for both Illumination and Rifle Rotation Conditions, and for the Generic Rifle and Assault Rifle respectively: there are no notable trends across the participants.

Discrimination

Log *d* was assessed using a two-way repeated measures ANOVA, with Illumination Condition, and Rifle Rotation as the independent variables. The best discrimination occurred in the 30° Conditions for both Illumination Condition and Rifle Rotation. The effect of Illumination Condition is depicted in Figures 7.13 and 7.14; log *d* significantly increased ($F(3,21) = 4.111$, $\eta_p^2 = 0.370$, $p < 0.05$), as the Illumination Condition changed from the No-Shadow Condition (log $d = 0.154$) through to the 30° Condition (log $d = 0.457$). Pairwise comparisons revealed that only the No-Shadow Condition and 30° Condition had significantly different mean values. Figure 7.13 shows that Participants 4 and 8 failed to discriminate between the two initial stimuli, the Generic Rifle and the Assault Rifle. Of the remaining participants, Participants 1 and 2 showed little indication of any differentiation across the Illumination Conditions, while Participants 3, 5, (possibly) 6, and 7, displayed a general trend of increasing discrimination as the shadows were less foreshortened.

Comparison of Figures 7.15 and 7.16 with Figures 7.13 and 7.14 shows that the effect of Rifle Rotation upon discrimination is similar to the effect of Illumination Condition: discrimination increased from the 0° Condition (log $d = 0.232$) to the 30° Condition (log $d = 0.390$). As illustrated in Figure 7.15, half of the participants showed increasing levels of discrimination with increasing rotation of the rifles; of the others, Participants 4 and 8 again demonstrated no discrimination across any of the levels, and Participants 1 and 3 showed stable and comparatively high levels of discrimination across the rifle rotation conditions. However, Figure 7.16 shows that, when assessed across the group, the major trend visible in Figure 7.15 was

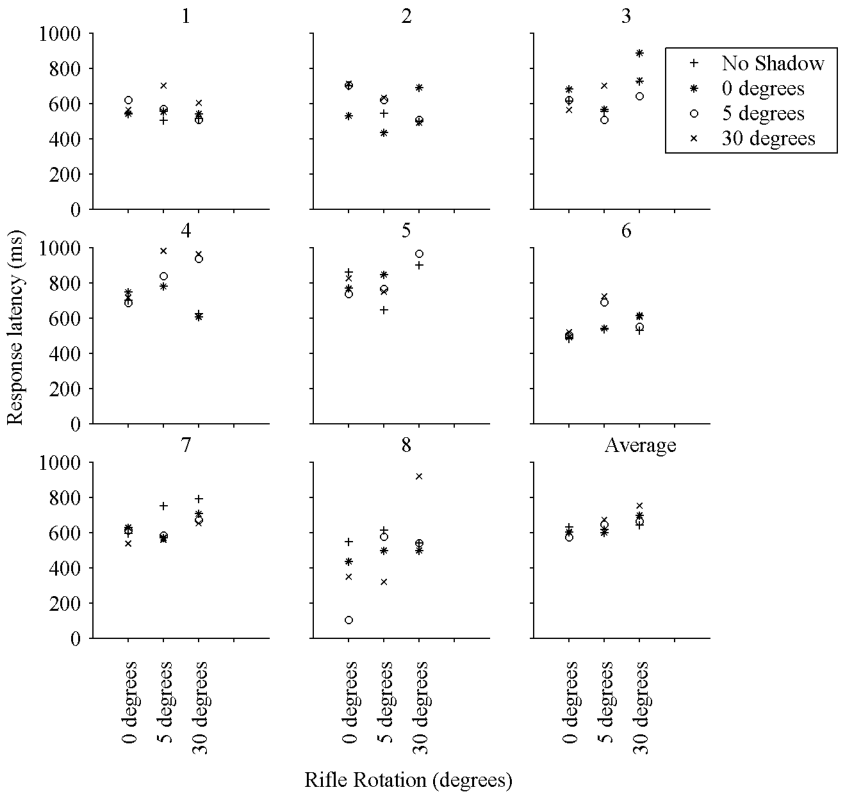

Figure 7.11. Mean response latency for each Illumination Condition and Rifle Rotation, for each participant for the Generic-Rifle Trials.

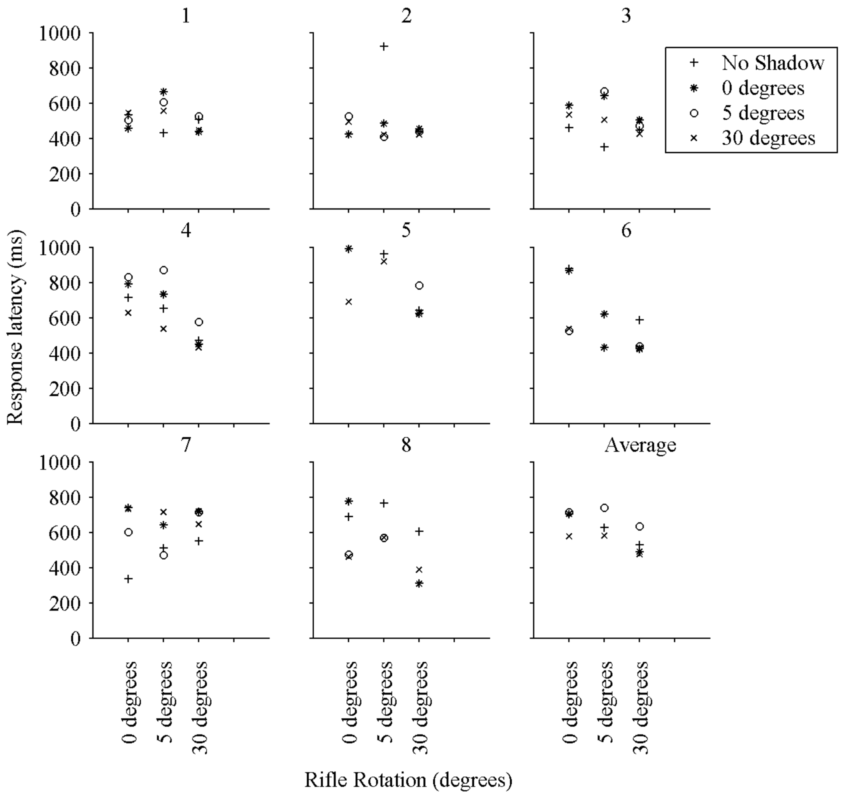

Figure 7.12. Mean response latency for each Illumination Condition and Rifle Rotation Condition, for each participant for the Assault-Rifle Trials.

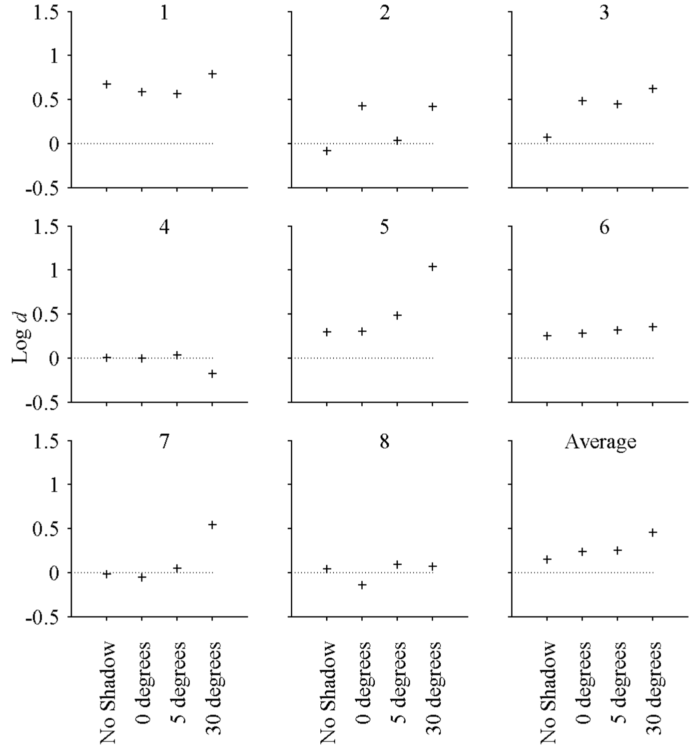

Figure 7.13. Mean log *d* for each Illumination Condition for each participant.

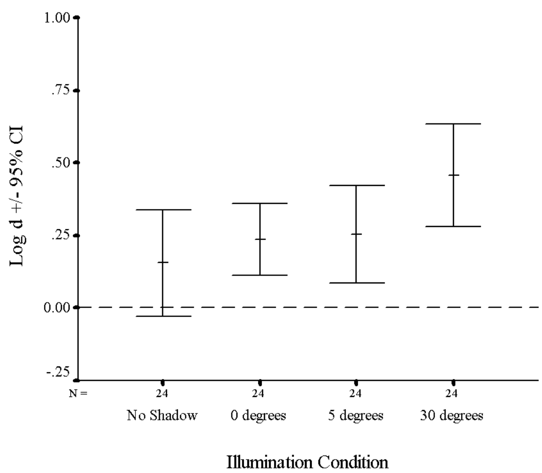

Figure 7.14. Mean log *d* for each Illumination Condition, taken across all participants.

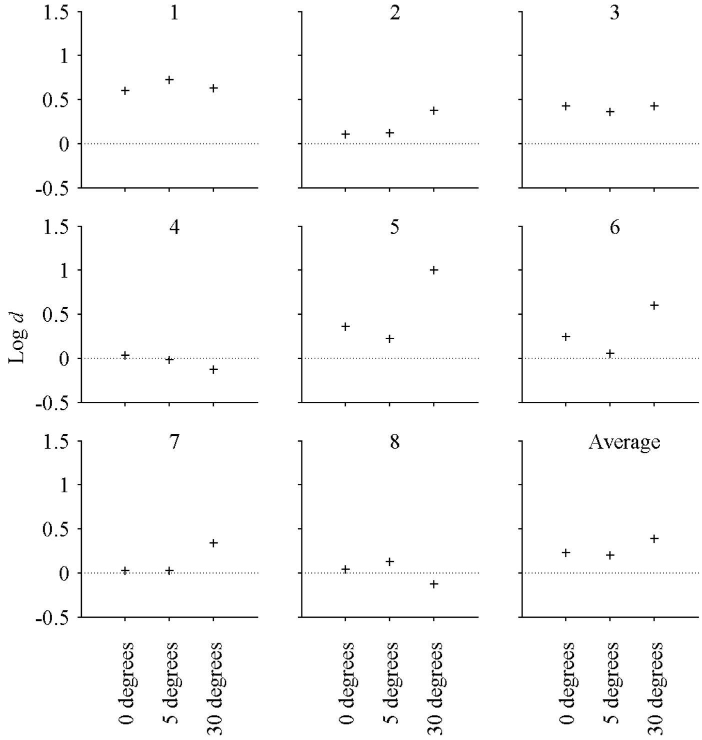

Figure 7.15. Mean log *d* for each Rifle Rotation for each participant.

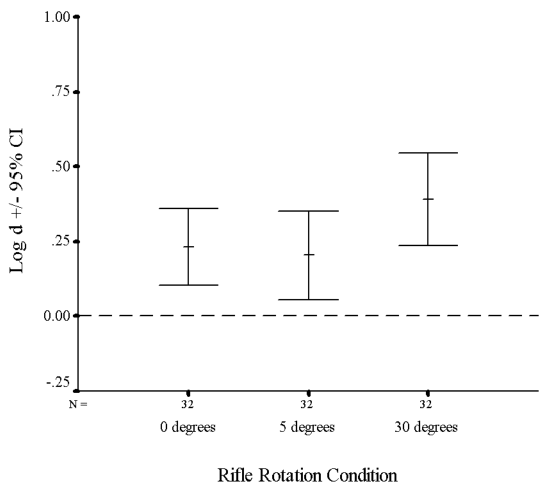

Figure 7.16 Mean log *d* for each Rifle Rotation taken across all participants.

not significant ($F(1.084,7.585)* = 2.235$, $\eta_p^2 = 0.242$, $p > 0.05$, *Greenhouse-Geisser adjustment used as the assumption of sphericity was violated).

The interaction between Illumination Condition and Rifle Rotation is illustrated in Figures 7.17 and 7.18. The significant ($F(6,42) = 4.918$, $\eta_p^2 = 0.413$, $p < 0.05$) interaction between Illumination Condition and Rifle Rotation is evident where there is no overlap between the confidence intervals around the means of the 5° Rifle Rotation/30° Illumination Condition subgroup, and the 5° Rifle Rotation/0° Illumination Condition subgroup. Figure 7.18 illustrates the tendency of any stimuli that contains a 30° Condition, whether in the Rifle Rotation or Illumination Conditions, to enable the participants to achieve a high level of discrimination.

Bias

As in the previous experiments, bias was assessed using the log *c* measure. In this case, when a positive value is calculated, this indicates a bias towards responding that the image was of the Generic Rifle, and a negative value indicates that the bias was towards the Assault Rifle. There was a general trend across the participants' data (with the exception of those of Participant 4), where bias shifted from being towards the Generic-Rifle response in the No-Shadow, 0°, and 5° Conditions, to being towards the Assault-Rifle response in the 30° Condition (as shown in Figure 7.19,). This tendency is depicted in Figure 7.20, and reflected in the large and significant main effect of Illumination Condition upon bias ($F(3,21) = 18.407$, $\eta_p^2 = 0.724$, $p < 0.05$), where the means of the biases present in the No-Shadow, 0°, and 5° Conditions differed from the 30° Condition.

There was no main effect of Rifle Rotation upon log *c* detected ($F(1.084,7.595)* = 3.643$, $\eta_p^2 = 0.342$, $p > 0.05$, *Greenhouse-Geisser adjustment used as the assumption of sphericity was violated). Although, as illustrated in Figure 7.21, five of the participants, Participants 1, 3, 5, 6, and 7, demonstrated the same pattern of changing bias as seen for the Illumination Condition: bias shifting from being towards the Generic-Rifle response in the 0° and 5° Conditions, to being towards the Assault-Rifle response in the 30° Condition. Figure 7.22 illustrates this trend was evident when the data were averaged across all participants. Participant 8 showed the reverse trend, while Participants 2 and 4 showed no bias.

The effect of the interaction between Rifle Rotation and Illumination Condition upon log *c* was significant ($F(6,42) = 4.466$, $\eta_p^2 = 0.390$,

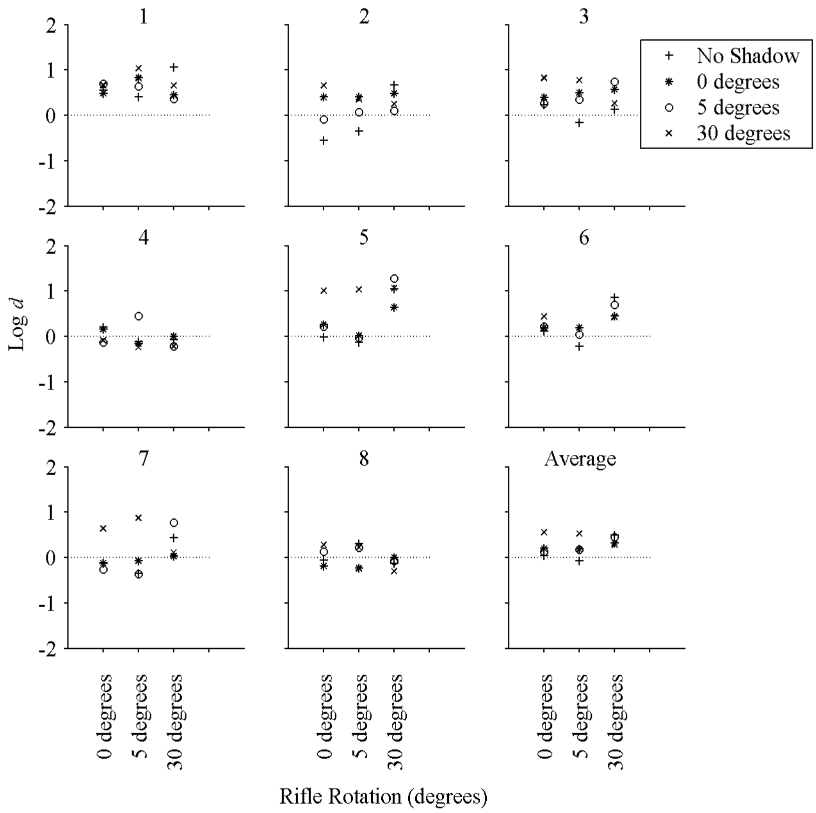

Figure 7.17. Mean log *d* for each Shadow Rotation and Rifle Rotation, for each participant, and averaged across all participants.

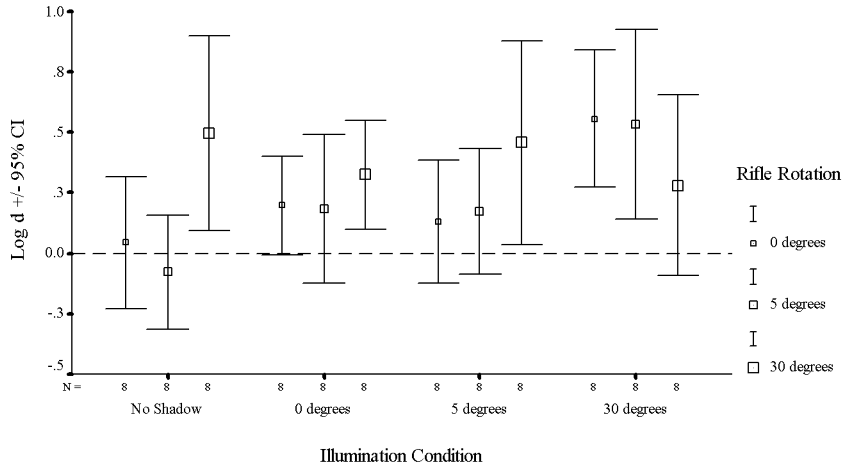

Figure 7.18. Mean log *d* for each Rifle Rotation and Illumination Condition, taken across all participants.

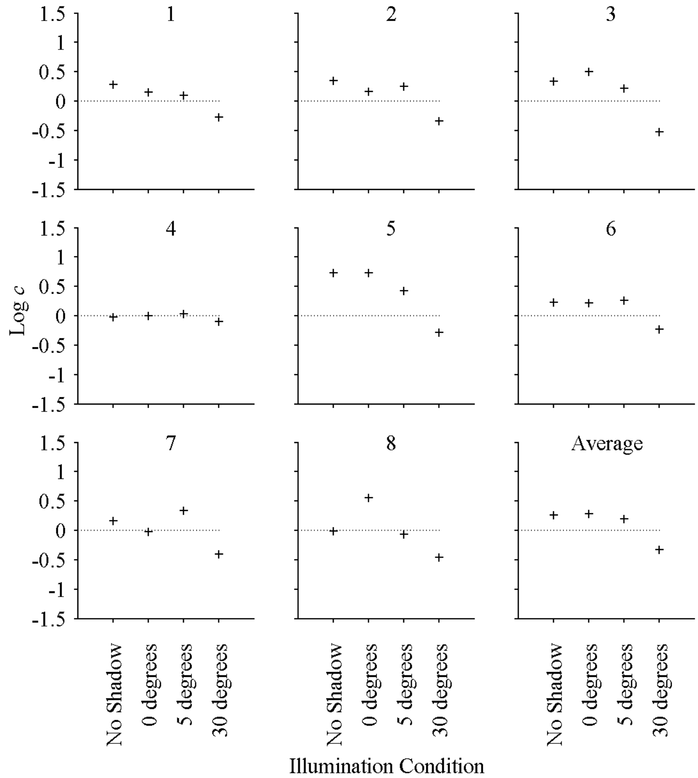

Figure 7.19. Average biases (log *c*) for each Illumination Condition, and for each participant. Positive values of log *c* indicate a bias towards the Generic Rifle.

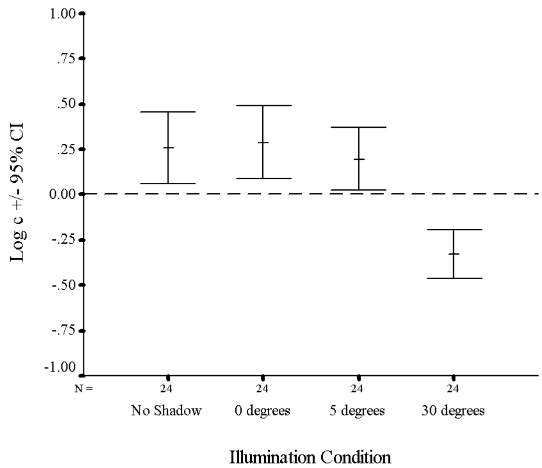

Figure 7.20. Average biases (log *c*) for each Illumination Condition, taken across all participants. Positive values of log *c* indicate a bias towards the Generic Rifle.

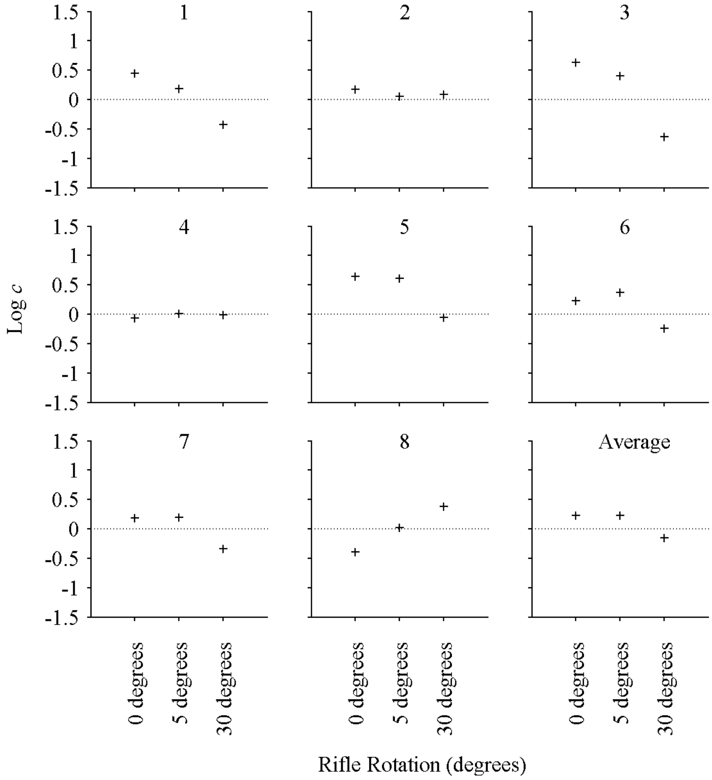

Figure 7.21. Average biases (log *c*) for each Rifle Rotation and for each Participant. Positive values of log *c* indicate a bias towards the Generic Rifle.

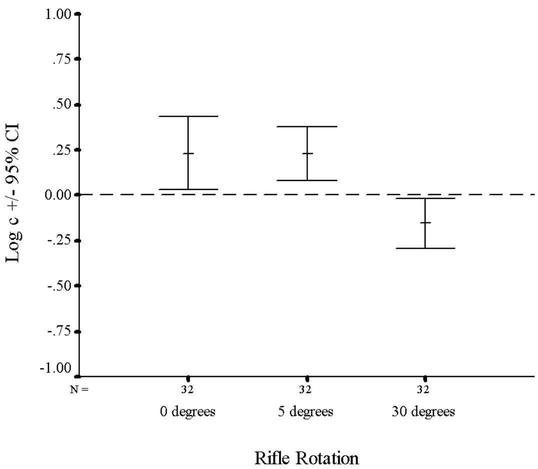

Figure 7.22. Average biases (log *c*) for each Rifle Rotation, taken across all participants. Positive values of log *c* indicate a bias towards the Generic Rifle.

p < 0.05). Figures 7.23 and 7.24, illustrate the interaction effect: differentiation in log *c* due to Illumination Condition, is present in the 0° and 5° Rifle Rotation conditions (where log *c* is positive in the No-Shadow, 0°, and 5° Illumination Conditions, and negative in the 30° Illumination Condition) but in the 30° Rifle Rotation, the differentiation is considerably negated.

Summary

High levels of recognition performance are evident in the figures and statistical analyses, either when the views of the rifles, their shadows, or both, are not foreshortened. This trend is evident in the plots of reaction time for the Assault Rifle, and in the plots of discrimination. Consistent biases to the Generic-Rifle response were found in the No-Shadow and foreshortened conditions, with a change in bias, towards the Assault-Rifle response in the 30° Conditions. The results are summarised in Table 7.1.

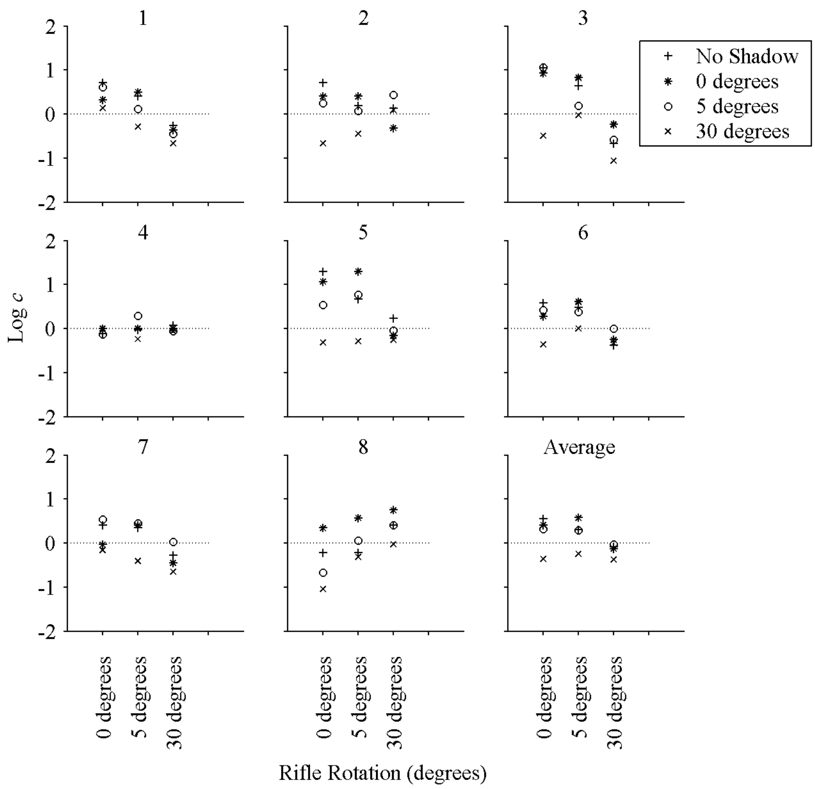

Figure 7.23. Average biases (log *c*) for each Rifle Rotation and Illumination Condition, for each participant.

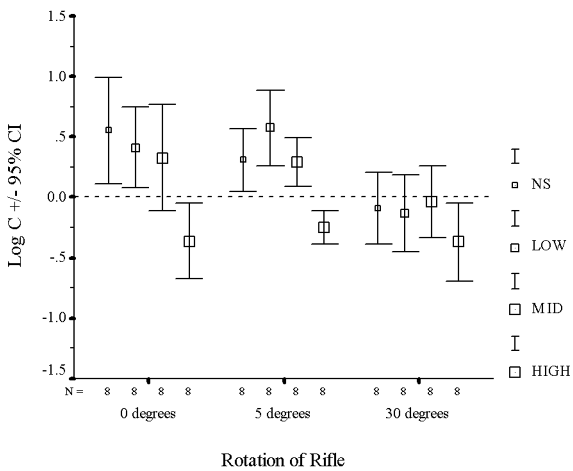

Figure 7.24. Average biases (log *c*) for each Rifle Rotation and Illumination Condition, taken across all participants.

Discussion

Experiment 7 was designed to be an object-recognition experiment that contained multiple levels of shadowing. In most cases, the responding of the Participants in Experiments 5 and 6 appeared to be controlled by the shadows present, with no influence from the rifles themselves; it was therefore arguable whether those experiments actually constituted tests of object recognition. In Experiment 7, both the shadow-, and rifle-, portions of the stimuli were associated with increased levels of discrimination when they were not foreshortened. This indicated that (at least on a proportion of the trials) the rifle portions of the stimuli had produced a degree of control over responding.

The reaction time results show relatively little differential effect of the rotation of either the rifles, or their shadows and it is therefore difficult to draw any conclusions regarding how the shadows, or rifle views, may have been controlling responding. There were no significant trends in reaction time for the Generic Rifle in either the Rifle Rotation or Illumination Conditions. For the Assault Rifle, the rotating the rifle 30° produced the lowest reaction times across the Rifle Rotation variable, and rotating illumination 30° produced the equal lowest reaction times across the Illumination Condi-variable. However, this was equal lowest with the No-Shadow and 0° Conditions.

As mentioned above, the cast shadows did produce a significant effect in discrimination across the Illumination Conditions, whereby discrimination was higher in the 30° Condition than the No-Shadow Condition, and there was no significant effect of Rifle Rotation upon discrimination. However, using multiple views of the rifles was partially successful in generating a level of control over discrimination by the rifles (as opposed to control only by the cast shadows). There was an interaction present where, if the Rifle Rotation was 30°, there was no differentiation by Illumination Condition (as shown in Figure 7.18). Thus, the trend in discrimination was that if either the rifles, or their shadows, were rotated 30°, then discrimination was at its highest. The interaction shows that the presence of the rifles did affect responding, i.e., responding was not solely based upon the cast shadows. The plot of the interaction between Rifle Rotation and Illumination Condition (Figure 7.18) also supports the hypotheses that discrimination would be low if both the rifle and its shadow were foreshortened, but high if at least one of them was not. A last point is that there is no evidence that the presence of shadows increased discriminatory performance above the levels obtained when the rifles are rotated away from the foreshortened view.

The association between discrimination and Illumination Condition and Rifle Rotation indicates that either the shadow or the rifle shape can contribute to recognition. Which one does so in a given case, appears dependent upon how salient the shape cues of each are (i.e., whichever cue, shadow shape or rifle shape, is less foreshortened, seems to determine the level of discrimination). If both cue types provide what would be useful information on their own, there is no additional benefit apparent.

The analysis of bias, log c, can also shed light upon the mechanisms the participants may have been using to make their discriminations. Across the Illumination Condition, when the shadows were foreshortened the participants were more likely to respond that the stimulus was the Generic Rifle, and when there was little foreshortening, in the 30° Condition, the participants were more likely to respond that the stimulus was the Assault Rifle. A similar bias was evident in half of the participants' data when the analysis was by Rifle Rotation (although there was no significant effect across the group). The results support the contention that participants employed the same cues from either the cast shadows, or the rifles, while they performed the discriminations. This is because a general bias across participants implies that the bias is due to the stimuli, rather than being idiosyncratic to the individuals. Where similar biases occurred across both the Rifle Rotation and Illumination Conditions (as occurred for Participants 1, 3, 5, 6, and 7), this implies that these participants employed the same sort of decision criterion in both conditions.

A possible criterion for discrimination can be formulated by considering the task and the stimuli. The task provides one initial stimulus, and two comparisons. It was designed so that the stimuli were very similar in the foreshortened view, but differed considerably in profile. In profile the Assault Rifle has large protruding features, while the Generic Rifle's features protrude less (as illustrated in Figures 7.1 and 7.2). Given that the stimuli are only visible for a short period, what would be an effective criterion for discrimination? A search for large protruding features would result in good discrimination when a rifle profile is evident, from either object or shadow, but poor discrimination when the stimuli are foreshortened. This is what occurred in Experiments 5, 6, and 7, but this does not confirm that such a criterion was in effect. However, the log c analysis does lend support to the suggestion that the majority of the participants employed a relatively crude criterion such as this.

The “large protruding features” mentioned in the criterion are larger or smaller depending upon the degree of foreshortening in any one given trial. Consider each of the two rifles across the condition types: as the Generic Rifle is rotated away from the foreshortened view, the extent to which its features protrude will increase, but the features will still be relatively small, even in the 30° Condition. As the Assault Rifle is rotated away from the foreshortened view, the extent to which its features protrude will also increase, and will do so to a considerable degree in the 30° Condition. The probability of the stimulus being the Generic Rifle is high if the extent the features protrude is nil to small; this range covers all of the Generic-Rifle trials, but only a portion of the Assault-Rifle trials. Whereas, it is more likely that the stimulus is the Assault Rifle if the features protrude by a small to large extent, this range covers most of the Assault-Rifle trials, but only the few Generic-Rifle trials where there is no degree of foreshortening. The biases seen in responding parallel these probabilities. When the stimuli are foreshortened, there is a bias to the Generic-Rifle response, and in the 30° Conditions, there is a bias towards the Assault-Rifle response. Thus, a very crude discrimination criterion based upon a search for protruding features, and their magnitude, would produce discrimination following the pattern seen in Experiment 5 to 7, as well as biases matching those found.

The reaction time results also offer an indication that the participants may be using a discrimination criterion based upon a search for large features. Reaction times for the Generic Rifle were relatively consistent as foreshortening decreased (as seen in Figures 7.3 to 7.6), while reaction times for the Assault Rifle were generally fast in the 30° Conditions (shown in Figures 7.7 to 7.10). That is, the fastest reaction times correspond to trials with the large features of the Assault Rifle evident in the 30° Conditions (illustrated in Table 7.2).

Table 7.2.
Reaction Times for the Generic and Assault Rifles in Experiment 7.

Rifle	Manipulation	Condition	Reaction Time (ms)
Generic Rifle	Illumination	No-Shadow	629
		0°	633
		5°	627
		30°	674
	Rifle Rotation	0°	601
		5°	632
		30°	688
Assault Rifle	Illumination	No-Shadow	587
		0°	653
		5°	696
		30°	546
	Rifle Rotation	0°	649
		5°	679
		30°	533

How do the results compare to other work with similar methodology? The experiment is somewhat similar in design to that of Freeburg (1966) mentioned earlier in the thesis: both used multiple levels of shadowing, and stimuli with different amounts of salient features. However, Freeburg used attached shadows cast on a textured surface; his participants had to sequentially match photographs of 3D textured surfaces resembling lunar landscapes, with views of the actual surfaces themselves. The surfaces were lit from various angles, so that low, moderate, or high levels of shadowing were present. Freeburg classed the discriminations as hard, moderate, or easy, dependent upon the presence of zero, one, or two volcano-like features in the landscapes. The conclusions Freeburg (1966) made regarding his study are quite succinct:

> At the extremes of recognizability, where other highly dominant cues - or their absence - are evident in the stimulus-material, shadow becomes subordinate as a determiner of recognition. But for some "middle ground" (the limits of which remain to be defined experimentally) there tends to be an enhancement of recognition by the addition of shadow. An exception occurs when the variable stimulus pattern being judged is relatively shadowless but the standard against which it is judged contains shadow. For such a case, the addition of shadow apparently serves as an extraneous cue and degrades recognition. (p.255)

Freeburg (1966) found that shadows were of greatest benefit to his participants when they were present in a moderate level, in combination with moderately difficult discriminations. In his high level shadow conditions, the large amount of attached shadow cast would have obscured the features of the surface, reducing discrimination as cues are neither available from the shadowing or the obscured surface itself. In the low level shadow conditions, the small amount of shadowing would do little to highlight the relief of the surface. These lighting situations would interact with the amount of features available for discrimination, so that if the discrimination was relatively easy then shadows would provide no additional benefit, and if it was very hard they would also provide little benefit. However, at an intermediate stage, moderate levels of shadowing (sufficient to highlight contours, but insufficient to obscure them) would produce discrimination at comparable levels to the “easier” conditions.

Freeburg’s (1966) results are similar to those here, in that any benefit of shadow presence upon discrimination did not exceed the performance attained when high levels of other shape information was available. The results also suggested that the benefit of shadowing may be restricted to when non-shadow cues are limited. In Experiment 7, shadow shape cues were valuable when the object was foreshortened (and object shape cues were valuable when the shadow was foreshortened), but if the object was not foreshortened there was no extra benefit attributable to the cast shadows.

Cast Versus Attached Shadows, and Shadow Borders

Experiments 5 through to 7 employed cast shadows, according the definition of Cavanagh and Leclerc (1989). Cast shadows may have an attached border, where they touch the casting object, and by definition, they will always have a cast border where the shadow falls across another surface. Both these sorts of borders have the potential to provide information to the viewer. However, in Experiments 5 to 7 the attached border of the shadow was physically negligible, it occurred where the shadow touched the butt of the rifle, and, given the lighting conditions, there was little contrast in brightness across this border. In comparison, the experiments were designed so that the cast border could provide a significant amount of information about the object’s shape.

There is an important distinction between the two sorts of shadow borders. The attached border is a border that is shared by the object and the shadow. As such, this border directly presents the object’s bounding contour. In comparison, the cast contour presents information about the object’s bounding contour that is physically separate from the object. It is a second, and spatially separate, source of information.

The results of Experiments 5 to 7 highlight this point. In Experiments 5 and 6, the stimuli usually provided these two separate sources of bounding contour information. Yet, the participants appeared to attend to only one of these two sources of information, the source provided by the cast shadow borders. In Experiment 7 the results indicated that the participants attended to either of the sources, dependent upon which provided the least foreshortened view. It looks as if discrimination was based upon one information source or the other, the bounding contour of the object, or the cast border of the cast shadow. Attached borders can occur at the bounding contour of an object, and at internal contours. Thus, they have the potential to enhance the information that is already present, as opposed to provide an alternative source of information.

If shadows can benefit recognition in more normal settings, they would need to enhance the information about the object's shape that is available in a scene. Experiments 5 to 7 demonstrated that cast borders can do this, but the benefit is restricted to when the cast shadow is more informative than the object's own bounding contour. As suggested by Cavanagh (1991) attached borders may be able to provide enhancement of an object's contours. This would be likely to be of benefit in a normal viewing situation, whereas the cast border of a cast shadow provides a separate source of information, which in most natural scenes would be less reliable than information available from the object itself. It is suggested that future experimentation address the potential information provided by the attached border, as well as the cast borders of cast shadows. Note that the cast borders of attached shadows also offer information about the shape of the object's surface, but that this is combined with the shape of the casting contour, and affected by lighting position.

Summary and Conclusions

Experiment 7 employed multiple views of the stimuli, and multiple levels of foreshortening of the cast shadows. Reaction times tended to be fastest in the 30° Conditions, when the stimulus was the Assault Rifle (the stimulus with the more salient shape cues present in the non-foreshortened view). This result indicates that fast times could result from either of the sources of shape information, the shape cues available directly from the object, or those available from the cast shadows. There was no differentiation in reaction times across either Illumination Condition or Rifle Rotation the Generic Rifle.

The presence of cast shadows produced a significant benefit to discrimination, when they were not foreshortened. A discrimination benefit due to rotating the objects away from the foreshortened view was also evident in the interaction between Illumination Condition and Rifle Rotation: the main effect of illumination condition was not apparent in the 30° Rifle Rotation condition. However, main effect of rifle rotation was not significant under testing. Overall, there was a common pattern of superior performance when either the objects or cast shadows were presented in the 30° Conditions. Therefore, there is no evidence of a special status of shadows per say, in improving recognition in this experiment. Rather, the results suggests that the visual system will use those cues that are relatively salient, whether they be shadow cues, or the object's shape itself, and that the combination of these cues will not return any additional benefit.

Experiments 1 to 7

Seven experiments were conducted into the effects of shadows on object recognition. The first four experiments used a design modelled after Tarr et al. (1998), employing novel objects to investigate whether attached shadows may improve recognition performance (through either faster reaction times, or improved accuracy). Across these experiments the size of the stimuli and the presentation durations were changed. Improvements in discrimination were generated by these changes, but at no stage did the presence of the attached shadows produce a benefit. In contrast to the results of Tarr et al. (1998), when the Shadow and No-Shadow Conditions did significantly differ, in Experiment 3, the Shadow Condition returned a lower level of discrimination than the No-Shadow Condition. When presentation times were again extended in Experiment 4, the significant difference between the conditions was not replicated. The conclusion drawn from the first four experiments is that, in experiments of this design, the attached shadows on novel objects do not benefit their recognition, and if they affect it in any way, then this would be to reduce performance. Further analyses of the data from Experiments 1 to 4 suggested that discrimination was strongly correlated with the degree of any change between the bounding contours of the initial and comparison stimuli.

In Experiments 5 and 6, the objects used (rifles) were always foreshortened, so that they presented few cues to the shape differences between them, while in Experiments 5 through to 7, the stimuli included prominent cast shadows. All three procedures constrained the number of objects to two, and controlled the viewpoints of the objects and the shape information provided by the shadows. The experiments were constrained to counter the methodological problems encountered in Experiments 1 to 4, e.g., lack of control the degree of shadowing present, and differences in the difficulty of discrimination between various pairs of objects. The procedure employed in Experiments 5 and 6 was successful in demonstrating control over responding by shadow presence, but failed to demonstrate that the majority of the participants were discriminating between the rifles, as opposed to discriminating between the shadows cast off them. Thus, the two tasks could not be genuinely classified as object-recognition tasks: it appeared that the objects were to a large extent immaterial.

In Experiment 7 both the rifles, and the cast shadows, presented in views from foreshortened to rotated 30°. When this occurred, both shadow cues and rifle cues contributed to discrimination. When shape cues from the rifles were prominent, in the 30° Rifle Rotation Condition, the results did not show any improvement in discrimination due to shadow presence. This suggests that the different cues, from shadow or from object, do not benefit recognition in an additive manner. That is, the level of discrimination is determined by cue that provides the most salient information, and if both provide information, there is no additional benefit.

In light of the previous seven experiments, some hypotheses regarding the contribution of shadowing to object recognition can be made. If we have little more than a cast shadow to view, then we should be able to identify the object via that cast shadow (at least if it is cast on a flat ground plane, and the shadow is not of an unfamiliar view, as in Experiments 5 and 6). This situation is little different to the recognition of silhouettes, which people are very adept at (Hayward, 1998). However, if the object and shadow present similar shape cues, then there is no indication of an additional benefit provided by the presence of both object and shadow cues (as in Experiment 7).

A question that arises is how often objects and shadows present similar information? If the view is taken that (in the vast majority of situations) objects provide much more salient shape cues than shadowing does, it could be argued that shadows would virtually always be redundant cues for recognition. Whereas, if the view is taken that objects will typically provide more salient shape cues than shadowing, but shadows may be able to enhance one, or some, of these cues, then there is still the possibility that shadowing could benefit object recognition. Are there situations where shadows enhance the shape cues provided by an object? Freeburg (1966) has indicated that shadows can enhance surface relief, and Braje et al. (2000), have suggested that shadows may enhance contrast at an object's bounding contour.

Experiment 7 investigated the effects of shadowing by using shadows that provided a second source of shape information very similar to that available from the objects themselves. The results suggested that people could utilise either source of information, dependent upon which provided the most shape information. Experiment 8 was conducted to investigate how shadows may contribute to object recognition through the enhancement of shape cues available from the objects themselves.

Experiment 8 The effects of shadows upon object recognition

In Experiment 7 it was suggested that shadows could provide a benefit to object recognition when they enhance those cues present in an object, rather than by providing an alternate source of information about the object's bounding contour. Shadowing may provide a possible benefit to recognition when it highlights a discontinuity in the object (e.g., highlighting relief, see Freeburg, 1966), or highlighting a discontinuity between the object and the background (see Braje et al., 2000). In order to investigate these possibilities, the procedure used in Experiments 5 to 7 needed to be reviewed.

Problems Encountered in Experiments 5 to 7

Using the Cast Borders of Cast Shadows

The cast shadows used in Experiments 5 to 7 provided information about object shape that was spatially separated from the objects. This was problematic, because correct matching could be performed on the basis of either of two spatially-separated cues; the shape of the rifle, or the shape of the cast-border of the rifle's cast shadow. In Experiment 7 it was demonstrated that the participants could use either cue, dependent upon which provided the greatest amount of information regarding bounding contour. While the cast-borders of the cast-shadows employed were sufficient to enable recognition, it was suggested that other shadow borders could provide a benefit to object recognition by enhancing the cues to shape available from the object itself.

To extract shape from the shadows in an image, Cavanagh (1991) suggests that the attached border is of importance, whether of an attached or a cast shadow, and that the cast border should be ignored (for an illustration, see Figure 4.1). Cavanagh (1991) discusses how a shadow's cast borders have a special status in images: they correspond to a discontinuity in illumination, rather than a discontinuity in the object. Therefore, he sees the cast border as irrelevant to the perception of object shape from shadows. However, other border types could serve to identify object discontinuities (as shown in Figure 1.3), such as joined-surface and occluded-surface shadow borders (Cavanagh & Leclerc, 1989), or the border between a lit surface and a shadowed background. To recount the definitions of these borders, a joined-surface border is where an attached contour is at a sharp discontinuity in surface orientation, and occluded-surface borders are where the illuminated background is occluded by the extremal contour of the object that is in shadow (Cavanagh & Leclerc, 1989). In a similar vein, the border where an illuminated surface occludes a shadowed background could be termed an occluded-background border.

Discontinuities between parts of an object, or an object and background will be more or less obvious, dependent upon the contrast across any of these borders. In a desaturated image, any discontinuity, whether in surface illumination, or in the object itself, is represented by a change in brightness. Cavanagh (1991) highlights this point, saying that an object's external borders are only visible when the object and background have different brightness. Therefore, the perception of any discontinuity in an object will be affected by illumination conditions.

As noted above, an object's bounding contour can be a primary source of information for object recognition (Hayward, 1998; Hayward, et al. 1999; Lloyd-Jones & Luckhurst, 2002; Norman, Phillips, & Ross, 2001). At an object-background border, the object's external contour either can be highlighted by the contrast in brightness over this border, or alternatively, it can be obscured. If the object is light in intensity and viewed against a light background, then any bounding contour would be highlighted where the background was in shadow. If the object is dark against a light background, then any extremal border would be obscured where the background is in shadow.

The situations described show how different shadow borders could potentially aid recognition. As previously noted, the difficulty of a discrimination may influence the utility of shadow information (Freeburg, 1966). The results of Experiments 5 and 6 also suggest this, the discrimination between the two rifles was difficult, and shadow cues were used in lieu of object cues. However, in a post-hoc analysis of reaction times by level of task difficulty, Braje et al. (2000) did not find any effect of shadow presence.

Problems With a Small Feature Set

In Experiments 5 and 6, there were only two objects used. This was to control the difficulty of the task, with both objects being presented in very foreshortened views. In Experiment 7, there were still only two objects, but there were three views of each. The reliance upon only two objects is problematic. It produces a task in which discrimination can be based upon a very small feature set (the prominent grip and magazine of the Assault Rifle), and possibly a feature absent/present criterion: in Experiment 7 it was postulated that a single decision criterion was employed by the participants, based upon the presence of only one or two features of the Assault Rifle. There is the possibility that a considerable degree of discrimination of the Generic Rifle was based upon the absence of the features of the Assault Rifle, as opposed to the presence of a feature of the Generic Rifle. Using only two objects also produced a task that bears little resemblance to everyday object recognition, or to other experimental object-recognition procedures. Therefore, for Experiment 8 it was proposed that multiple objects should be employed.

Implications for Methodology

Experiments 5 to 7 highlighted the difficulty of designing a procedure that appears to be a genuine test of object recognition, while concurrently demonstrating a high level of control over both shadow cues, and the difficulty of the task. Two procedural factors have been discussed; the focus upon the cast-borders of cast shadows, and the use of a limited number of stimuli. The design of Experiment 8 needed to address these points, while exerting control over task difficulty and the amount of shadowing present.

Object-Recognition Task Versus Shadow-Recognition Task

As discussed in Experiment 7, the participants' responses in Experiments 5 and 6 appeared to be controlled by the shadows present, with little influence from the rifles themselves. The non-foreshortened views of the rifles used in Experiment 7 produced a greater degree of control over responding by the rifles themselves, but the procedure can still be described as only quasi-object-recognition. On any given trial, it appeared that recognition could be based upon one or other of the stimuli available, the shadow or object. Thus, on any given trial, the task might have been "shadow recognition" or it might have been "object recognition". Given that the participant could still respond accurately while ignoring the object's present, this methodology did not lend itself to testing how shadowing may assist in the recognition of objects. Therefore, the next procedure needed to be a valid object-recognition task, where the question of whether shadows enhance the recognition of the object itself can be tested. To do this, the experiment needed to utilise shadows cast so that they may enhance the bounding, or joined-surface, contours of their associated objects.

Number of Stimuli

Using multiple stimuli necessitates changes to the methodology employed in Experiments 5 to 7, where the two comparison images were present all the time. Multiple stimuli were employed in a sequential matching task in Experiments 1 to 4, but this task suffered from the very high degree of control that was exerted over the participants' "same/different" responses by changes in the silhouettes of the initial and comparison images. One solution to this problem is to employ multiple views of the same object, so that the bounding contours of matching objects do not necessarily match. From a practical standpoint, this can be problematic, in that employing several views of a single object increases the number of trials significantly. However, the potential for one of several views is similar to natural recognition.

Another solution is to use a sequential word-picture verification task (e.g., Decaro & Reeves, 2002). Because the initial stimulus is a word, the word-picture task prevents problems with discrimination being performed solely on the basis of any change in the bounding contours between the initial and comparison images. As a result, good performance in a word-picture verification task is likely to indicate genuine object recognition, while good performance in a picture-picture matching task does not imply recognition of the pictures, just matching. As previously mentioned, one of the reasons that picture-picture matching is often employed is that novel objects can be used. Novel objects are of use when experimenters wish to control the previous experience of the participants, such as when evaluating whether recognition is dependent or independent upon viewpoint (recognition at novel views of objects allows the researcher to make inferences regarding whether recognition is based upon views/features or structural descriptions). When novel objects are not required, the word-picture task can offer the advantages discussed.

Task Difficulty

To maximise the likelihood of finding any contribution of shadow presence to object recognition, a procedure was required that would assess the effects of shadow presence across several levels of task difficulty. The difficulty of a visual object-recognition task is necessarily related to the availability of cues about object shape: if there are no cues to shape, then the task cannot be performed through visual means. Experiment 7 indicated that shadows were more likely to be beneficial to recognition when the objects presented few other cues to shape. In Experiment 7, various amounts of foreshortening were employed to control the amount of information available from the objects' bounding contours. As well as bounding contour, shading is a major cue to shape (see Ramachandran, 1988).

Task Difficulty and Shape-From-Shading

In a desaturated image, the only information presented to the visual system is differences in levels of grey. From the greyscale information in an image we can discriminate between different objects. These different levels of grey are determined by the illumination of the scene. Classical shading, shadowing, and interreflections between surfaces, are the three broad divisions of illumination effect (Langer, 1999). Shading is due to the variation in reflected flux as the angle between the incident light and the surface varies, while shadows are areas blocked from direct illumination (Cavanagh & Leclerc, 1989). Interreflections are the result of light bouncing between multiple surfaces (Madison, et al., 2001). Interreflections may contribute to perception of object contact (Madison, et al., 2001), but their contribution to object recognition is not documented. However, shading cues have a well documented contribution to object recognition (see: Berbaum et al., 1984; Mingolla & Todd, 1986; Erens et al., 1993; Ramachandran, 1988) and importantly, they can be controlled independently of shadow cues. With respect to a shadow, changing the ambient light level has the effect of adding or subtracting a constant to the darkness of the shadow, but there is no change to the contours of the shadow (except the possibility of a small change across the penumbra). With respect to shading, changing the ambient light level has the effect of increasing or decreasing the amount of shading in the image.

Therefore, changing shading in an image can have an effect upon contrast at discontinuities between an object and the background. For example, without shadows there is little contrast between a white object and a white background: where the background is in shadow, a white object is contrasted with a grey background. However, in the two situations, no-shadow and shadow, changing ambient lighting has different effects. In both, the amount of shading in an image is reduced by increasing the ambient lighting levels, but, the degree to which an object contrasts with its background is only reduced in the absence of an object/shadow border (the no-shadow situation). An increase in ambient lighting does not change the contrast across the border between object and background when the background is in shadow; it adds the same amount of brightness to both the object and shadow (with the exception of when the object's brightness is at a ceiling level already). When there is no shadowing present, an increase in ambient lighting is very likely to change the contrast across the border between an object and a white background, as a white background is at, or very close to, a brightness ceiling, while the object is not. In this situation, increasing the amount of ambient light decreases the contrast between the object and the background, because the object gets brighter, but the background is at a ceiling.

Viewing a white object on a white background is not a normal recognition situation. We normally have an abundance of texture and colour cues to aid us. Therefore, a more natural image condition should also be used to provide a baseline measure of performance across the Shadow and No-Shadow Conditions. Such a condition should employ techniques like colour rendering, and texture mapping, to create more realistic objects with an abundance of the cues typically available for recognition. The potential of colour to affect response times and accuracy in recognition and/or categorisation tasks has been widely investigated (e.g., Braje et al., 2000; Delorme, Richard, & Fabre-Thorpe, 1999, 2000; Gegenfurtner & Rieger, 2000; Wurm, Legge, Isenberg, & Luebker, 1993), with recent research indicating that any benefit appears constrained by how diagnostic colour is in the categorisation of individual objects (Nagai & Yokosawa, 2003, also see: Naor-Raz, Tarr, & Kersten, 2003). Therefore, the addition of extra cues such as colour, would, if anything, make shadows more likely to be redundant as cues for recognition.

Procedural Requirements

As mentioned above, the procedure used in Experiment 8 needed to redress the focus of Experiments 5 to 7 which was upon the cast borders of cast shadows, and their use of a limited number of stimuli. Experiment 8 also needed to control task difficulty and the amount of shadowing present. Several objects were selected for use in the experiment, with each object being presented from two different viewpoints. Multiple objects were used to reduce the potential for the participants to select their response on the basis of the presence or absence of only one or two simple features. To disambiguate object recognition from image matching, a sequential word-picture verification task was chosen as the basis for the experimental procedure (as opposed to picture-to-picture matching). The amount of shadow in the images was controlled by measuring the contribution shadows made to each image, in a similar manner to that used in the post-hoc testing of the images from Experiments 1 to 4. Thus, a set of images was selected in which the amount of shadowing was relatively constant. Task difficulty was controlled by manipulating illumination levels, which had the effect of adjusting the amount of shading present in each image. The shading manipulation would also be reasonably independent of the difference in brightness across attached-, and external-, shadow borders.

Predictions

A general hypothesis was formed from the results of Braje et al. (2000), Freeburg (1966), and Experiments 5 to 7: that the presence of shadows will aid recognition when other salient cues are diminished or absent. In this experiment, the level of ambient lighting would be manipulated in a series of images rendered with and without shadows present. By increasing the levels of ambient lighting, contrast across the object to background borders would be significantly reduced in the No-Shadow Conditions, and remain comparatively high in the Shadow Conditions. Thus, the difficulty of the task, as dependent upon shading cues, would change across the spectrum of ambient lighting, without an equivalent decrease in the cues from shading. It was hypothesised that shadows would be of benefit to object recognition when cues to shape from shading were reduced (high levels of ambient lighting), but that there would be no differences evident between the Shadow Conditions and No-Shadow Conditions under normal viewing conditions.

Method

<u>Participants</u>

Forty-two undergraduate students at the University of Waikato participated in the experiment. Each received a 1% credit towards a Level 1 Psychology paper.

<u>Apparatus and Stimuli</u>

The experimental sessions were conducted on Dell Pentium II 400-MHz, and Dell Pentium III 1.1 GigHz, computers with 43 cm Trinitron screens and 75 Hz refresh rates. Images of 14 different objects were rendered using 3D Studio Max. The images were of an: Ant; Bee; Beetle; Can; Cellphone; Cross; Cup; Fork; Handycam; Knife; Lighter; Snail; Tennis racquet; and Vase. The 3D models were obtained from those available at 3DCafe http://www.3dcafe.com/asp/meshes.asp. The images (excluding background) displaced approximately 5° of visual angle at a 600 mm viewing distance. Each entire image displaced approximately 14.3° (height) by 18.9° (width) of visual angle, at a 600 mm viewing distance, as was used in Experiment 1.

Each object was rendered in two views, each view was rendered in five lighting conditions, and each lighting condition was rendered with and without shadows. The five lighting conditions used in rendering the images were: 1. Colour object, texture mapping, ambient lighting level = 50 (called the Colour Condition); 2. White object, ambient lighting level = 50 (the Ambient 50 Condition); 3. White object, ambient lighting level = 100 (the Ambient 100 Condition); 4. White object, ambient lighting level = 150 (the Ambient 150 Condition); 5. White object, ambient lighting level = 200 (the Ambient 200 Condition) (ambient lighting scale = 0 to 255: blackout to whiteout conditions).

Shadows were rendered using shadow mapping (as opposed to raytracing) to produce a penumbra. The maps used a map bias of 2.0, size of 1024, and a sample range of 5.0. The illuminant was an omni-directional light at a height of 75 units, and a distance from the object of 100 units, 15° to the left of the line of sight. The lighting environment used a global lighting level of 2, with ambient lighting levels of either 50, 100, 150, or 200 (range 0 to 255). The objects were less than 23 units in height, with their width/length being determined by the height to width ratio of the individual objects. In the no-colour conditions the objects were rendered using a white material: Blinn shading, RGB values of 225, diffuse and ambient colours locked, shininess value of 25 (range 0 to 100), shininess strength 5 (range 0 to 100), and opacity 100 (range 0 to 100). The camera was positioned at a distance of 100 units from the object, at a height of 100 units, and used a 45° field of view.

To control for the amount of shadowing present in the images, from the original 14 objects, six were selected for use in the experiment. For selection of the objects, an image analysis was conducted similar to that in Experiment 1. A total pixel value was calculated for each image (background included). The images had been rendered in RGB format and there were 640 (width) x 480 (height) x 3 (the RGB layers) datum per image, with each pixel taking three values of 0 to 255. For the purposes of object selection, images without colour, in the Ambient 50 Condition, were analysed. Therefore, the three values for each pixel were the same. The values were inverted with respect to normal; so that maximum darkness equalled 255 and maximum brightness equalled 0. Thus, a totally white image would have a value of 640 x 480 x 3 x 0 = 0, and a totally black image would have a value of 640 x 480 x 3 x 255 = 235,008,000.

Given that the darker Shadow Condition image has a larger value than the No-Shadow Condition image, the difference between the two is a measure of the contribution of shadow to the image. The standardised score (z-scores) of each of these differences was calculated, along with the average difference. Using these standardised scores six objects were selected. Each object had two views and thus two standardised scores. For each object, the smaller standardised score was ignored, and from the larger standardised scores the six closest to zero were selected. This restricted the range of difference between the images with, and without, shadows, and centred it around the average difference across the 14 objects and two views. The largest differences in pixel value was 2.44 times the smallest difference (compared to 26.28 times the smallest difference across all 14 objects), and on average the change due to shadowing accounted for 6.75% of the shadow image total pixel values (ranging from 3.78% to 8.80%, see: Table 8.1). Both views of each of the objects used are presented in Appendix 6, across all shadow and lighting conditions.

Table 8.1.

Image Analysis of the Objects Employed in Experiment 8.

Object	View	Standardised (z-score) difference in pixel value:	Pixel difference as a percentage of the Shadow Image total:
Ant	**1**	**0.08**	**7.28%**
Ant	**2**	**-0.27**	**5.57%**
Bee	**1**	**0.40**	**8.80%**
Bee	**2**	**-0.03**	**6.69%**
Beetle	**1**	**0.26**	**7.99%**
Beetle	**2**	**0.24**	**7.92%**
Can	**1**	**0.20**	**7.81%**
Can	**2**	**-0.61**	**3.78%**
Cellphone	1	-0.41	4.74%
Cellphone	2	-0.90	2.21%
Cross	1	-0.09	6.44%
Cross	2	-1.16	0.86%
Cup	1	2.07	16.00%
Cup	2	0.54	9.05%
Fork	1	-0.67	3.53%
Fork	2	-1.13	0.98%
Handycam	1	2.73	17.82%
Handycam	2	2.61	17.31%
Knife	1	-0.68	3.40%
Knife	2	-1.08	1.26%
Lighter	**1**	**-0.55**	**4.05%**
Lighter	**2**	**-0.13**	**6.28%**
Snail	1	-0.84	2.57%
Snail	2	0.06	7.11%
Tennis Racquet	1	-0.15	6.13%
Tennis Racquet	2	-0.76	2.99%
Vase	**1**	**0.08**	**7.15%**
Vase	**2**	**0.19**	**7.72%**
Range across all objects		3.89	16.96%
Range objects employed		**1.13**	**5.02%**
Mean across all objects		0.00	6.55%
Mean across objects employed		**0.05**	**6.75%**

Objects that were selected for the final experiment are displayed in bold.
Note: The images used for analysis were from the Ambient 50 Condition. They were rendered without colour, but in RGB format, image size was 640 x 480 pixels x 3 layers (Red, Green, Blue). Standard pixel values, where maximum darkness = 0 and maximum lightness = 255, were reversed. This makes the difference between the Shadow and No-Shadow images an additive factor, rather than subtractive. The difference between the two conditions can then be expressed as a percentage of the Shadow Condition image, i.e., the percentage contribution to the Shadow image total of the shadows present. Maximum pixel value would therefore be 235,008,000 if the whole image was black, and 0 if the whole image was white.

The initial stimuli were black words on a white background, written in times new roman. The words were approximately 14 mm high when displayed on the screen. The stimuli used were: 1. Ant; 2. Bee / Fly; 3. Beetle; 4. Can; 5. Lighter; and 6. Vase. The initial stimulus chosen to be associated with the Bee images was "Bee / Fly" because when the images were rendered without texture mapping the image could have been interpreted as a fly (given that the stripes typical of a Bee were not evident without texture mapping). From this point forward this object is simply referred to as the Bee.

Procedure

A sequential word-picture verification task was employed. For each trial the stimuli were presented on a computer screen, and the participant responded, via a keyboard, whether he/she thought the image matched the previously presented word. The participants could respond from the onset of the comparison stimulus until 750 ms had elapsed after the masking stimulus had disappeared (a 1650-ms window). If they did not respond in this time, their response was not recorded. An example of a trial is presented in Figure 8.13, showing the presentation sequence of the stimuli, and the timing used in the experiment. A program written specifically for this experiment recorded the participants' responses and response latencies, and controlled the presentation of the images.

There were six objects, each presented in two views, and for each view there were five illumination conditions, each with and without shadows present. This resulted in 120 different images, each was presented as the initial stimulus twice. In half of the trials the initial and comparison stimuli matched. Selection of non-matching comparison stimuli was randomised, but controlled to produce equal employment of the objects as comparisons. This resulted in 240 trials, the presentation order of which was randomised. The instructions given to the participants are presented in Appendix 7.

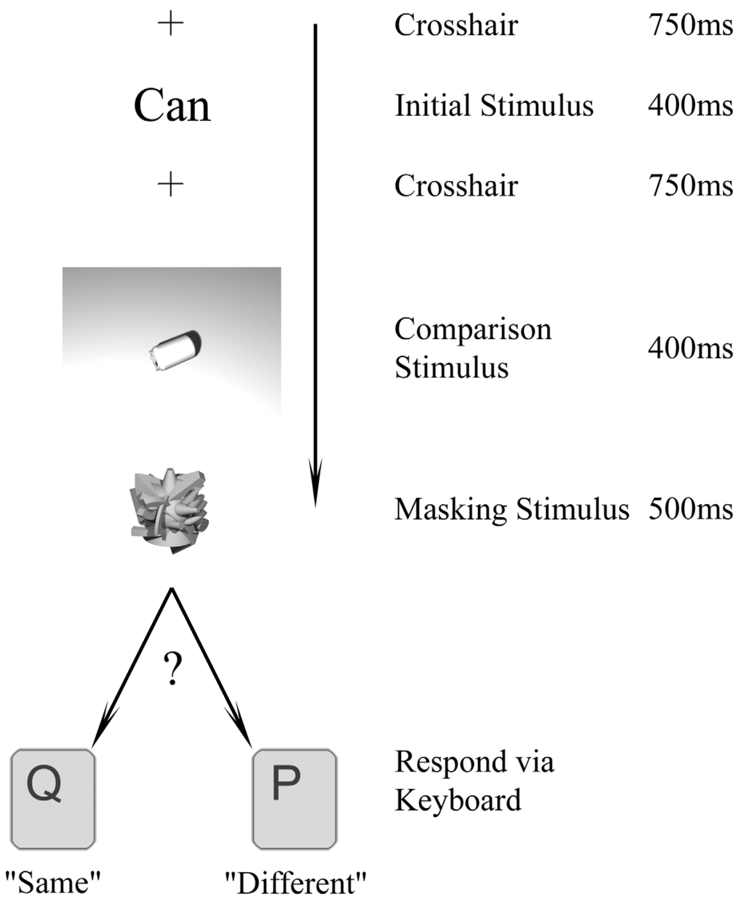

Figure 8.1. Example of a trial where the initial and comparison stimuli match. The participants could respond from the onset of the comparison stimulus until 750ms had elapsed after the masking stimulus had disappeared (a 1650-ms window).

Results

Repeated-measures ANOVA were conducted on both log *d* and response latency, across the independent variables of Shadow Presence, Illumination Condition, and Object Type. The interactions between these variables were also analysed.

Log *d*

As shown in Figure 8.2, the Shadow Condition produced slightly higher levels of discrimination than the No-Shadow Condition ($F(1,41) = 26.244$, $\eta_p^2 = 0.390$, $p < 0.05$: log $d = 0.557/0.516$ S/NS). However, this difference between the means of the Shadow and No-Shadow Conditions, was not reliable across all objects or conditions. When the effect of Shadow Presence was assessed by Object Type, the effect of Object Type was significant ($F(5,205) = 19.774$, $\eta_p^2 = 0.325$, $p < 0.05$). Superior performance in the Shadow Condition was evident for the trials in which the image of the Bee was presented (as illustrated in Figure 8.3) (assessed using the 95% confidence intervals of the means). For trials employing the image of the Can, the reverse effect was found, images of the Can without shadow produced significantly higher levels of discrimination than those with shadow. Thus, for four of the six objects, there was no significant effect of Shadow Presence (irrespective of the Illumination Condition), and for the other two, one showed a benefit and one showed a cost due to the presence of shadows.

The effect of the presence of shadows was also assessed across the Illumination Conditions, and a differential effect of Shadow Presence was found according to the Illumination Condition ($F(4,164) = 25.770$, $\eta_p^2 = 0.386$, $p < 0.05$). Discrimination in the Shadow Condition was unaffected by changes in illumination, log *d* was 0.56 in the Colour Condition and 0.56 in the Ambient 200 Condition, with no significant difference found between any of the means when their 95% confidence intervals were compared, as can be seen in Figure 8.4 (the Percent Correct scores are given in Table 8.2 for comparison with log *d*). The No-Shadow Condition produced a different pattern of discrimination: discrimination was constant across the Colour, Ambient 50, and Ambient 100 Conditions, and decreased markedly in the Ambient 150 Condition, and again in the Ambient 200 Condition.

Comparing Illumination Conditions across the two Shadow Conditions, it can be seen that for the Colour and Ambient 100 Conditions, there

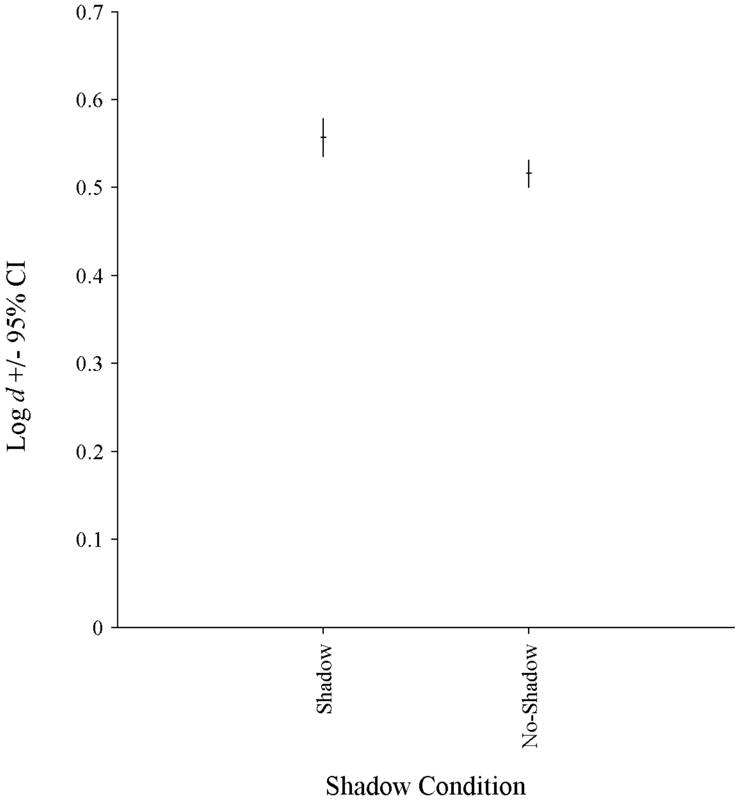

Figure 8.2. The mean value of log *d* for the Shadow, and No-Shadow, Conditions. of the main effect of Shadow Condition upon discrimination (log *d*).

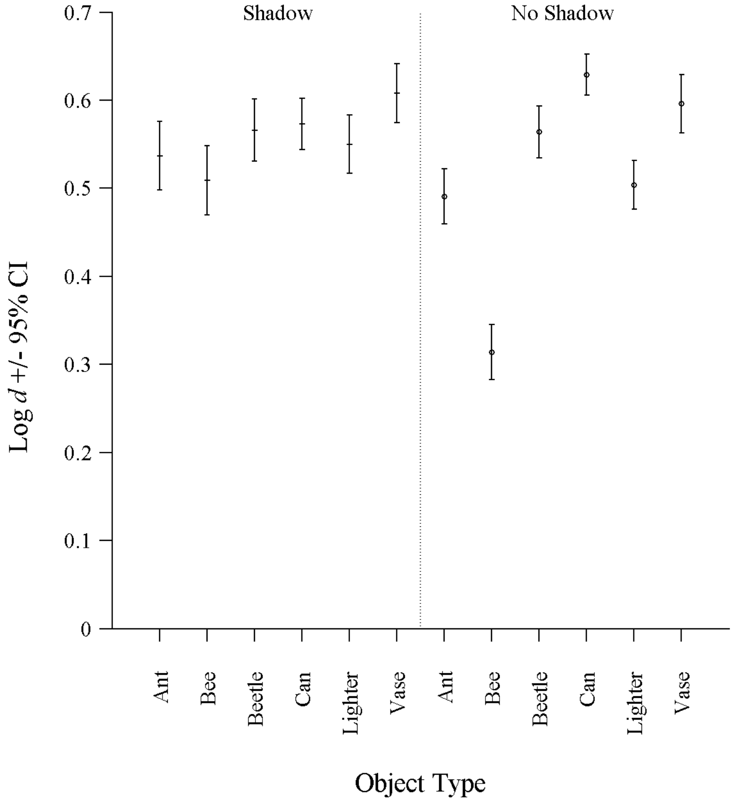

Figure 8.3. The mean value of log *d* for each object in both the Shadow, and No-Shadow, Conditions.

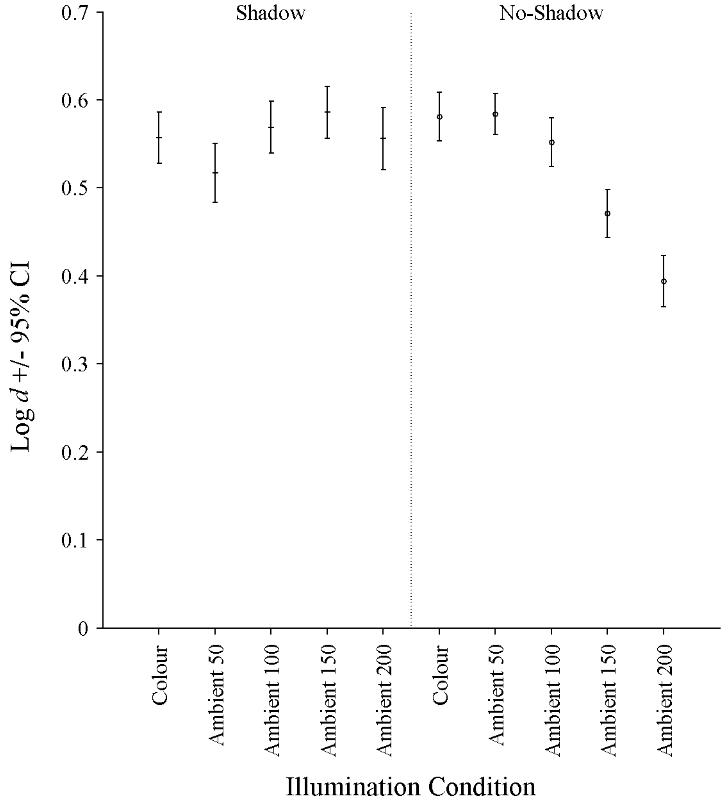

Figure 8.4. The mean value of log *d* for each Illumination Condition in both the Shadow, and No-Shadow, Conditions.

Table 8.2.
Mean Percent Correct For Each Illumination, and Both Shadow, Conditions.

Shadow Condition	Illumination Condition	Percent Correct
Shadow	Colour + Ambient 50	90.278
	Ambient 50	87.698
	Ambient 100	91.270
	Ambient 150	92.692
	Ambient 200	90.542
No Shadow	Colour + Ambient 50	92.229
	Ambient 50	92.791
	Ambient 100	90.264
	Ambient 150	84.722
	Ambient 200	78.737

was no effect of Shadow Presence upon discrimination. For the Ambient 50 Condition the presence of shadows produced a discrimination disadvantage. In the Ambient 150 and 200 Conditions, the presence of shadows produced discrimination advantages in comparison to the same No-Shadow Illumination Conditions.

The interaction between Shadow Presence and Object Type showed that any benefit of Shadow Presence, when assessed irrespective of Illumination Condition, was restricted to the Bee images. When the changes in Illumination Condition were also included in the analysis (the interaction between Shadow Presence, Object Type, and Illumination Condition), it was apparent that the effect of the presence of shadows was quite restricted. For the Beetle, Can, Lighter, and Vase images, Figure 8.5 shows that there were no notable differences between the data for the Shadow and No-Shadow Conditions across the changes in illumination. However, the ANOVA showed that interaction was significant ($F(11.670,482.156)^* = 12.857$, $\eta_p^2 = 0.239$, $p < 0.05$, *Greenhouse-Geisser adjustment used as the assumption of sphericity was violated) with differences in the effects of Shadow Presence occurring across the Illumination Conditions for both the Ant and the Bee images. For the Ant, the presence of shadows was detrimental to recognition in the colour condition. There was no difference according to Shadow Presence across the Ambient 50 and Ambient 100 Conditions, but in the Ambient 150 and Ambient 200 Conditions discrimination of the Ant was poor in the No-Shadow Condition, while remaining high in the Shadow Condition (as illustrated in Figure 8.5). Discrimination of the Bee images was unaffected by the presence of shadows in the Colour Condition, but was reduced in the presence of shadows under the Ambient 50 Condition, before returning to high levels again across the Ambient 100 to Ambient 200 Conditions. When shadows were absent, discrimination was good in the Colour and Ambient 50 Conditions, but was significantly reduced in the Ambient 100 and Ambient 150 Conditions, and below chance in the Ambient 200 Condition.

The major result to be taken from the discrimination measure is that discrimination was disrupted due to extreme levels of illumination, for the Ant and Bee images in the Ambient 100 (Bee only), Ambient 150, and Ambient 200 Conditions, but that this disruption only occurred in the

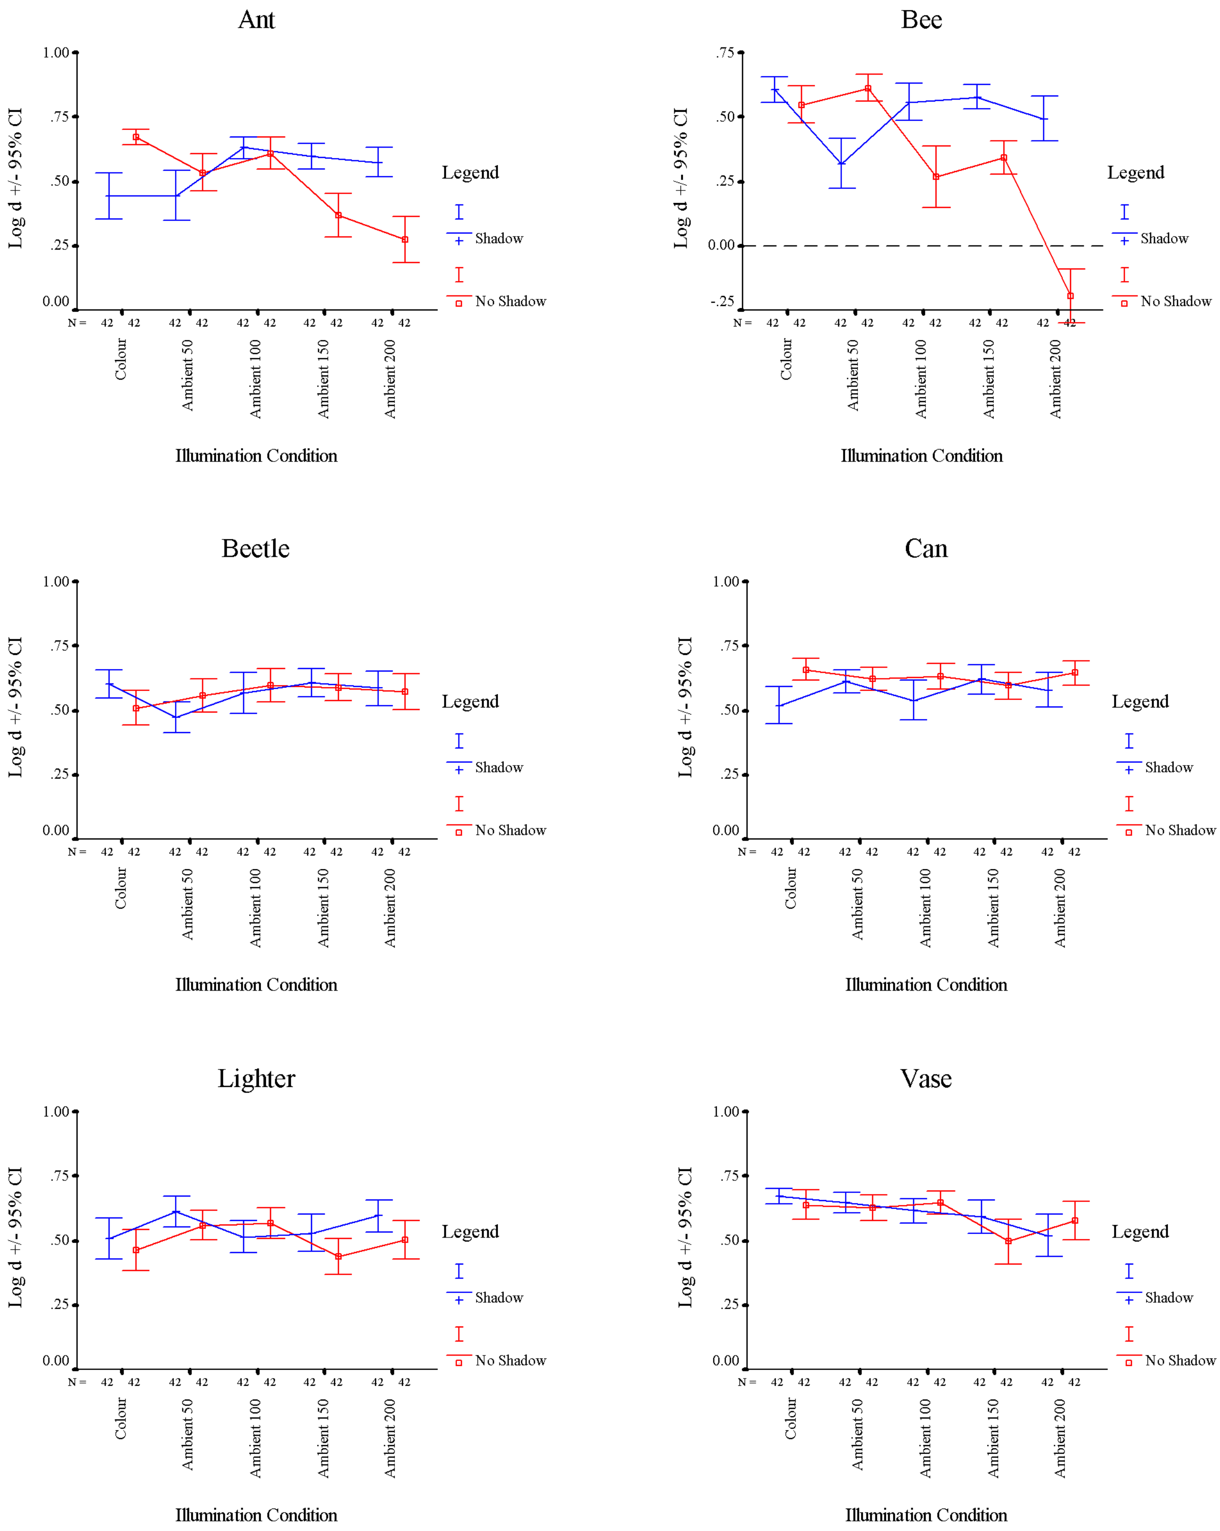

Figure 8.5. The mean log *d* for each Illumination Condition, for each object, and for both Shadow Conditions. There was a significant interaction, where the manipulation of shadow presence had an effect across the Illumination Conditions for the Ant and Bee images only. The means are joined as an aid in illustrating the trends across the Illumination Conditions.

absence of shadows. When shadows were present, discrimination under extreme levels of illumination was no different to that in the more normal Illumination Conditions.

To complete the analysis of discrimination, the main effects of Object Type and Illumination Condition were assessed, along with their interaction. As these effects were considered without accounting for Shadow Presence, they were not the major interest of the study. Discrimination was found to vary by Object Type ($F(4.051,166.105)^* = 42.296$, $\eta_p^2 = 0.508$, $p < 0.05$, *Greenhouse-Geisser adjustment used as the assumption of sphericity was violated). Figure 8.6 shows that the Bee produced the lowest levels of discrimination (log $d = 0.411$), followed by the Ant (log $d = 0.514$) and Lighter (log $d = 0.527$). Comparisons between all the objects are presented in Table 8.3.

Irrespective of Shadow Presence, there was a significant effect ($F(4,164) = 19.573$, $\eta_p^2 = 0.323$, $p < 0.05$) of Illumination Condition, illustrated in Figure 8.7, where the Colour Condition (log d =0.569) and Ambient 100 Condition produced higher levels of discrimination than the Ambient 150 Condition. The Colour Condition, and the Ambient 50, 100, and 150 conditions also all produced higher discrimination than the Ambient 200 condition (log d =0.475) (see Table 8.3).

The interaction between Illumination Condition and Object Type was significant ($F(20,820) = 11.316$, $\eta_p^2 = 0.216$, $p < 0.05$), and is presented in Figure 8.8. This figure shows that the reduction in discrimination seen in the highest lighting levels is attributable to the reduced discrimination at these levels for the Ant, Bee, and Vase. These were the only objects to show a significant reduction in performance in the most extreme Illumination Conditions. As evident in Figure 8.5, for the Ant and Bee these reduced levels of discrimination only occurred when there were no shadows present.

Response Latency

The second dependent variable was response latency, or reaction time. As with discrimination, the benefit of Shadow Presence was evident in the response latency measure: response latencies when there were shadows present (660 ms) were 24 ms significantly faster ($F(1,41) = 29.115$, $\eta_p^2 = 0.415$, $p < 0.05$) than when there were no shadows present (684 ms) (as shown in Figure 8.9). Response latencies across the Illumination Conditions showed little differentiation by Shadow Presence (as shown in Figure

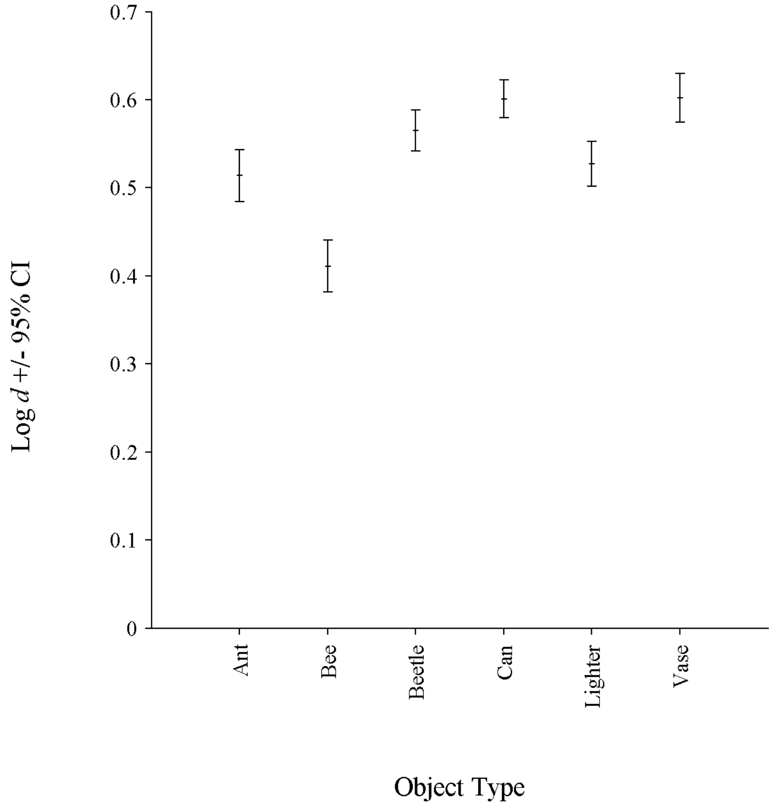

Figure 8.6. Mean log *d* for each Object Type.

Table 8.3.
Mean Log d for Each Illumination Level, Shadow Condition, and Object Type.

Condition	Level	Significantly Different Levels ($\alpha = 0.05$)	log *d*
Illumination	Colour A 50	Ambient 150, Ambient 200	0.569
	Ambient 50	Ambient 200	0.550
	Ambient 100	Ambient 150, Ambient 200	0.561
	Ambient 150	Colour A 50, Ambient 100, Ambient 200	0.528
	Ambient 200	Colour A 50, Ambient 50, Ambient 100, Ambient 150	0.475
Shadow	Shadow	No Shadow	0.557
	No Shadow	Shadow	0.516
Object Type	Ant	Bee, Beetle, Can, Vase	0.514
	Bee	Ant, Beetle, Can, Lighter, Vase	0.411
	Beetle	Ant, Bee, Can, Lighter, Vase	0.565
	Can	Ant, Bee, Beetle, Lighter	0.601
	Lighter	Bee, Beetle, Can, Vase	0.527
	Vase	Ant, Bee, Beetle, Lighter	0.602

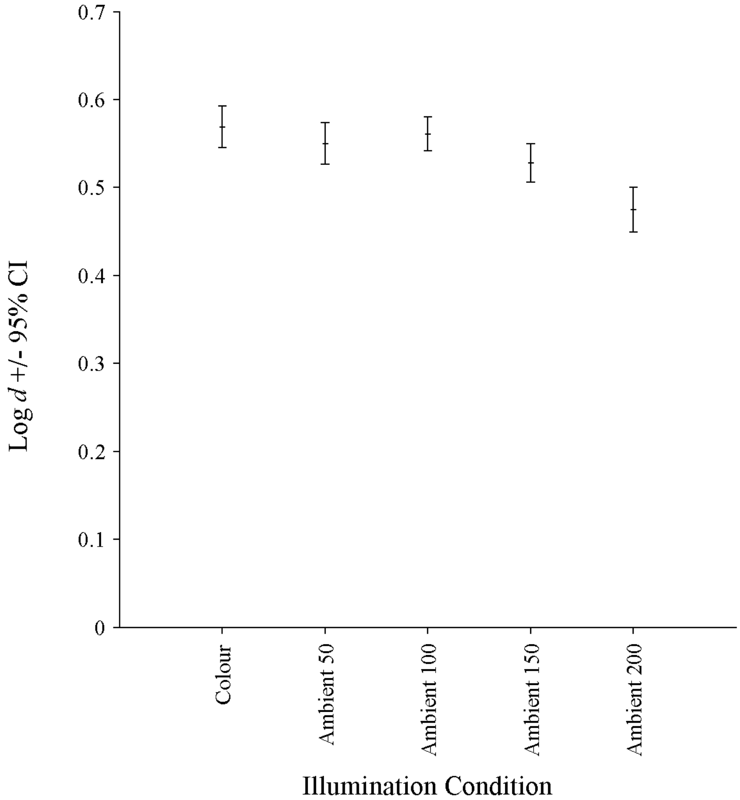

Figure 8.7. Mean log *d* for each Illumination Condition.

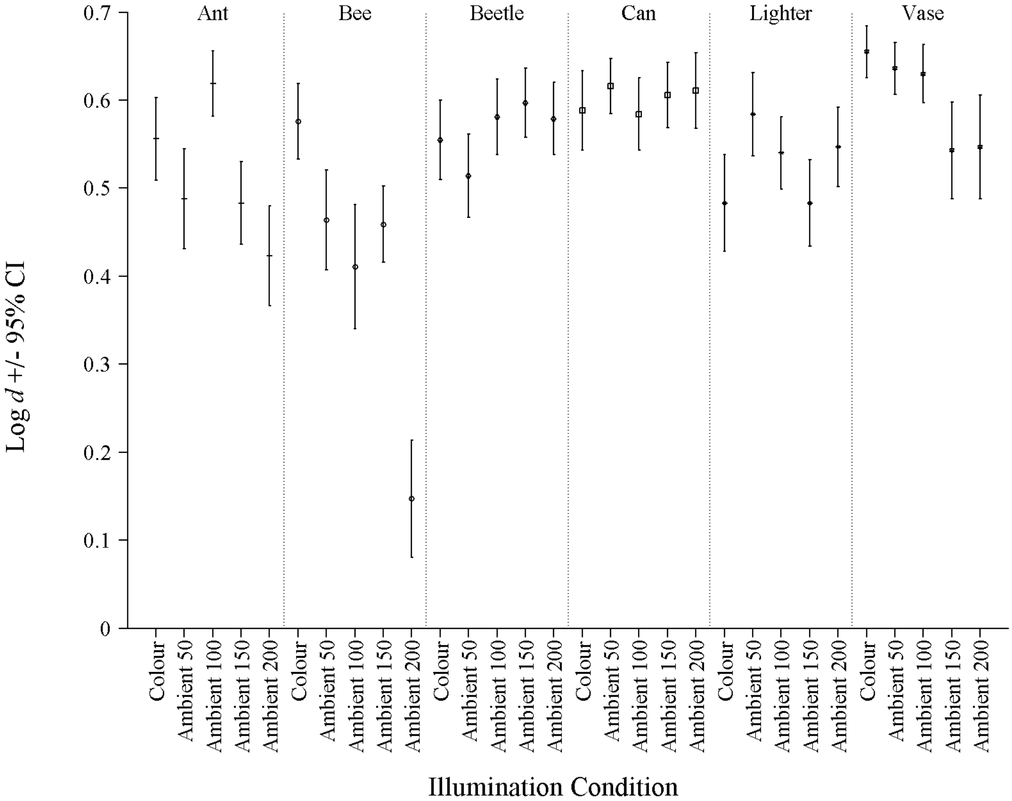

Figure 8.8. Mean log *d* for each Illumination Condition by each Object Type.

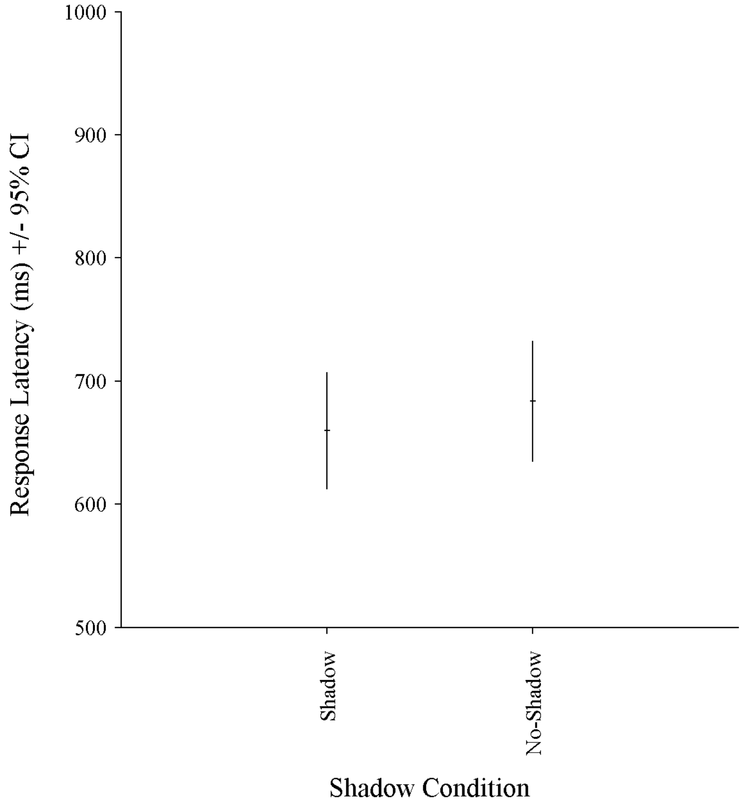

Figure 8.9. The mean response latency for both Shadow, and No-Shadow, Conditions.

8.10): the interaction was significant ($F(4,164) = 18.271$, $\eta_p^2 = 0.308$, $p < 0.05$). Figure 8.11 illustrates a trend of increasing reaction time in the No-Shadow Condition as illumination levels increased, but the No-Shadow Condition with the slowest response latencies, the Ambient 200 Condition, was only significantly different to the Shadow Condition with the fastest response latencies, the Ambient 100 Condition. Unlike the results for discrimination, across the objects the differences in response latency due to the presence of shadows was not significant ($F(3.676,150.722)^* = 1.747$, $\eta_p^2 = 0.041$, $p > 0.05$, *Greenhouse-Geisser adjustment used as the assumption of sphericity was violated).

Object Type, in conjunction with Shadow Presence and Illumination Condition, had a degree of control over discrimination. For response latencies the results were not as clear, although a significant interaction ($F(20,820) = 9.640$, $\eta_p^2 = 0.190$, $p < 0.05$) was found, and is depicted in Figure 8.12. Figure 8.13 illustrates the interaction between the illumination changes and shadow presence or absence: responses to the Ant, Bee, Beetle, and Lighter images were slower in the No-Shadow Ambient 200 Condition, than the Shadow Ambient 200 Condition, however, the confidence intervals of these means still overlap. The significant differences between the Shadow Condition and No-Shadow Condition, according to object, match differences in discrimination: the No-Shadow and Colour Condition combination for the Ant image produced faster times than the same trials with Shadow (562 vs 690 ms). This corresponds to the higher accuracy also seen for the No-Shadow Condition combined with the Colour Condition in Figure 8.5. For the Bee images, the combination of Shadow Condition trials combined with the Ambient 50 Condition produced slower response latencies than the same images without shadows (862 vs 669 ms). Response latencies for the Ant and Bee trials also increased from the Colour Condition to the Ambient 200 Condition (Ant = 562 vs 821 ms; Bee = 672 vs 880 ms) but only in the No-Shadow Condition, the times remained static in the Shadow Condition. For the other objects response latencies did not trend from the Colour Condition to the Ambient 200 Condition for either the Shadow Condition or the No-Shadow Condition.

The effect of Object Type upon response latencies was assessed, with significant differences ($F(3.845,157.634)^* = 48.047$, $\eta_p^2 = 0.540$, $p < 0.05$,

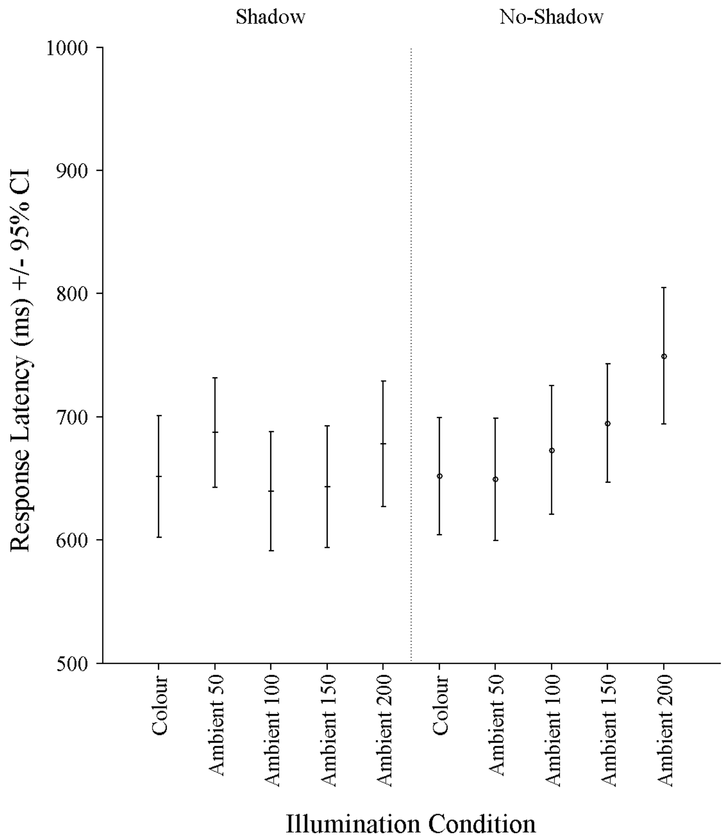

Figure 8.10. Mean response latencies for each Illumination Condition, and for both the Shadow, and the No-Shadow, Conditions.

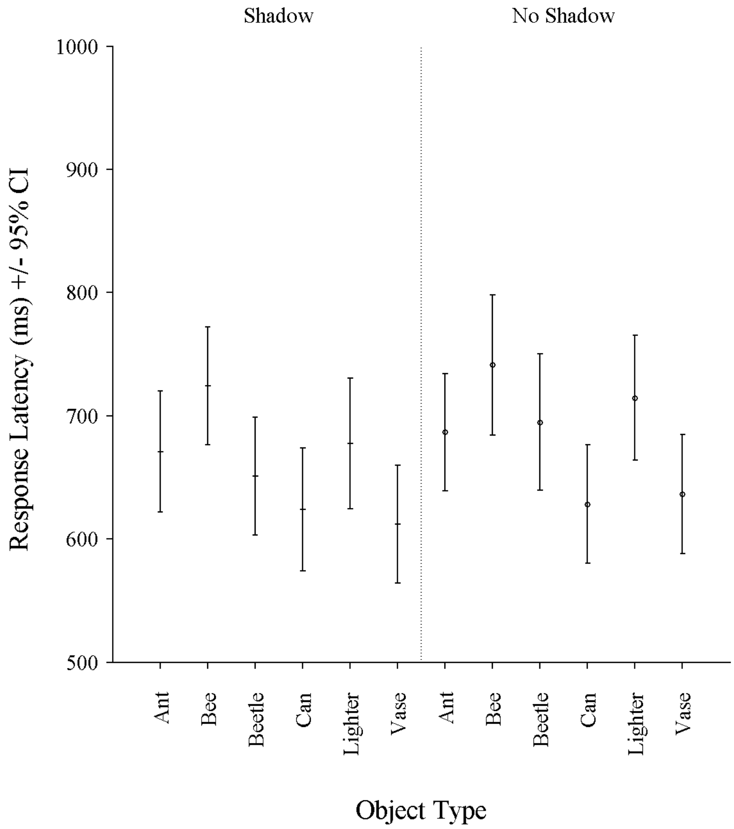

Figure 8.11. Mean response latency for each Object Type, differentiated by Shadow Condition.

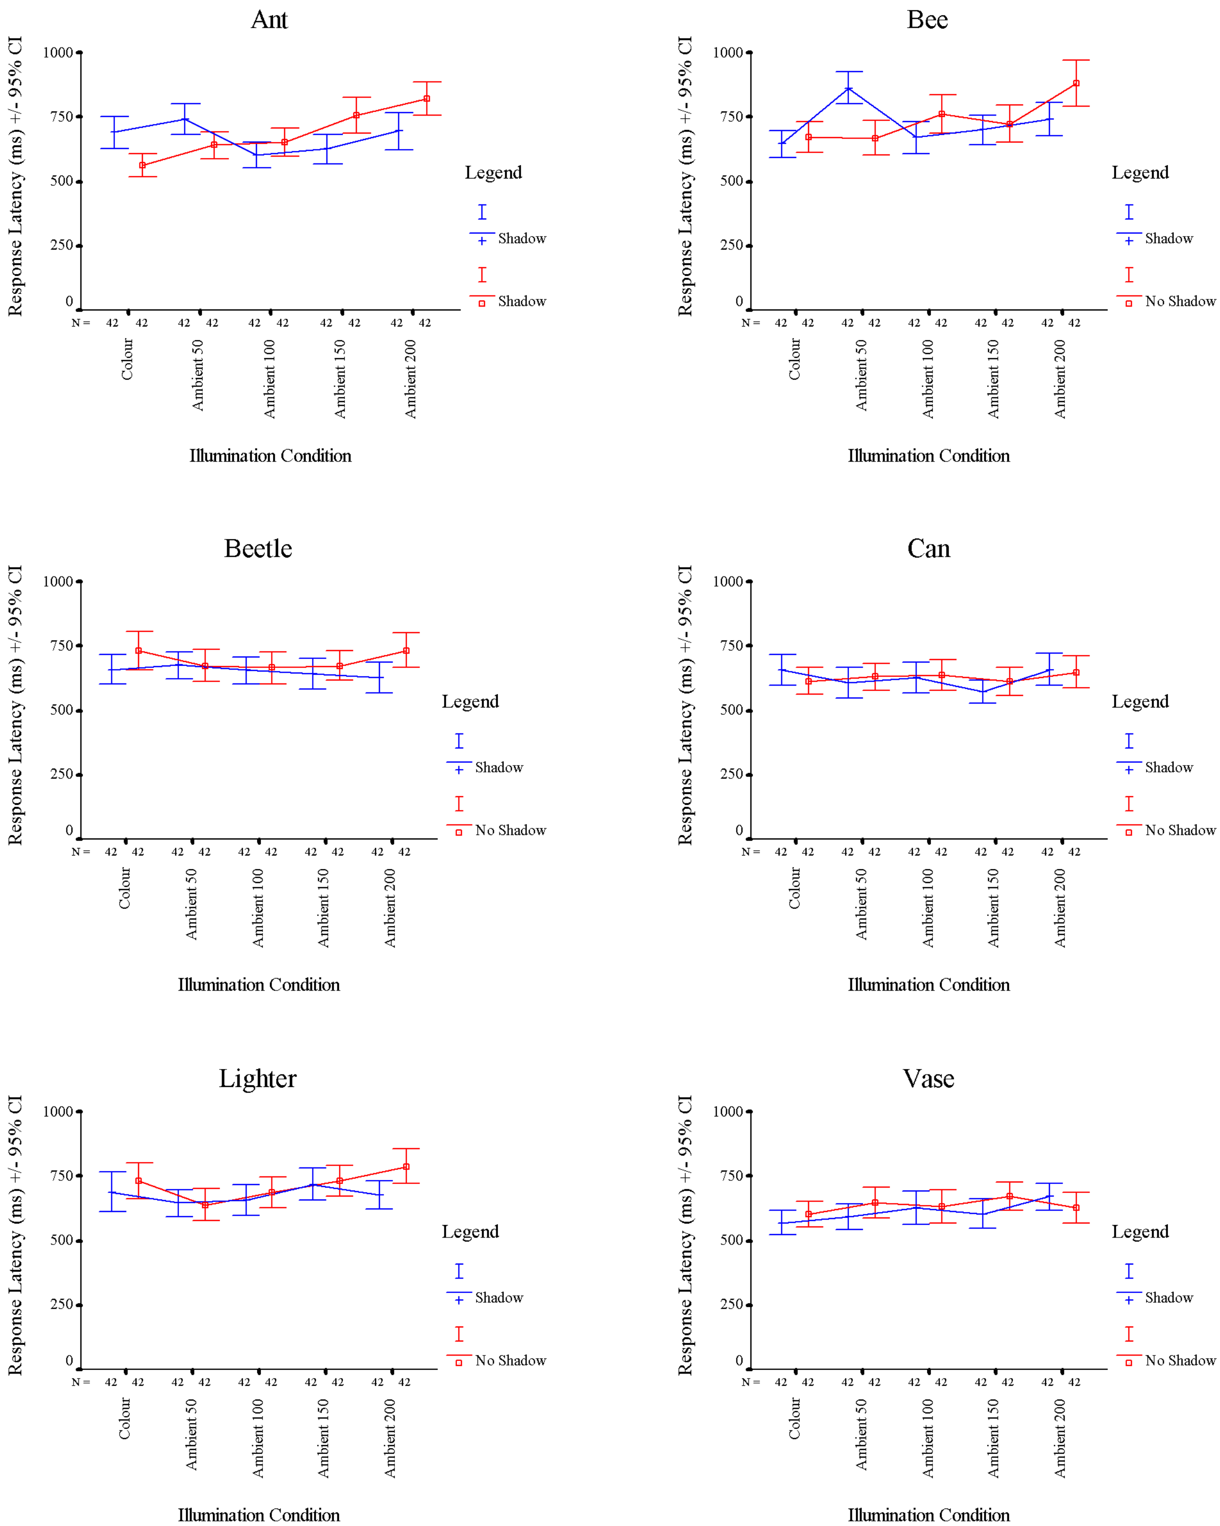

Figure 8.12. Mean response latencies for each Illumination Condition, for each object, and for both Shadow Conditions. There was a significant interaction, mean response latencies for the Ant and Bee image trials increased from the Colour Condition to the Ambient 200 Condition, but only in the No-Shadow Condition. The times remained static in the Shadow Condition. The means are joined as an aid in illustrating the trends across the Illumination Conditions.

*Greenhouse-Geisser adjustment used as the assumption of sphericity was violated) in response latencies being found across the objects (as illustrated in Figure 8.13). Average response latencies for the Bee images were the slowest, at 733 ms (as shown in Table 8.4), while the Can and Vase produced the fastest response latencies, 626 and 624 ms respectively. This pattern matched that found for discrimination (presented earlier in Figure 8.3), where the Bee images were discriminated the worst, and the Can and Vase images were discriminated the best.

The effect of lighting condition upon reaction times was also large (illustrated in Figure 8.14 and Table 8.4) and significant ($F(3.143,128.852)^{*} = 28.430$, $\eta_p^2 = 0.409$, $p < 0.05$, *Greenhouse-Geisser adjustment used as the assumption of sphericity was violated). The Ambient 200 Condition produced response latencies (mean = 713 ms) slower than any of the other conditions (Colour = 652 ms through to Ambient 150 = 669 ms).

The interaction between Object Type and lighting condition accounted for approximately 15% of the variation in response latencies ($F(10.899,446.845)^{*} = 7.212$, $\eta_p^2 = 0.150$, $p < 0.05$, *Greenhouse-Geisser adjustment used as the assumption of sphericity was violated). Figure 8.15 show trends of increasing latencies with increasing illumination levels evident for the Ant, Bee, and Vase, while for the Beetle, Can, and Lighter response latencies were relatively consistent from the Colour to Ambient 200 Conditions. Again, the pattern of increasing response latencies matched decreases in discrimination.

In summary, the results show that the presence of shadows produced discrimination and response latency benefits in comparison to the No-Shadow Condition. However, these benefits were restricted to only two of the six objects, the Ant and the Bee, and to conditions of extreme illumination.

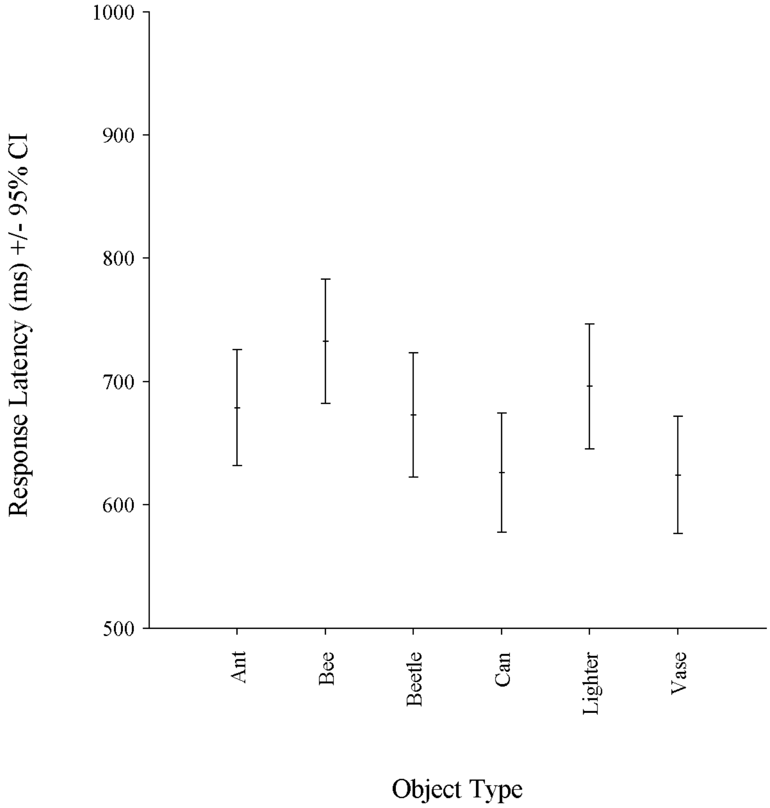

Figure 8.13. Mean response latencies for each Object Type.

Table 8.4.
Mean Response Latencies (ms) For Each Illumination Level, Shadow Condition, and Object Type.

Condition	Level	Significantly Different Levels ($\alpha = 0.05$)	Response Latency (ms)
Illumination	Colour A 50	Ambient 50, Ambient 150, Ambient 200	652
	Ambient 50	Colour A 50, Ambient 100, Ambient 200	668
	Ambient 100	Ambient 50, Ambient 150, Ambient 200	656
	Ambient 150	Colour A 50, Ambient 100, Ambient 200	669
	Ambient 200	Colour A 50, Ambient 50, Ambient 100, Ambient 150	713
Shadow	Shadow	No Shadow	660
	No Shadow	Shadow	684
Object Type	Ant	Bee, Can, Lighter, Vase	679
	Bee	Ant, Beetle, Can, Lighter, Vase	733
	Beetle	Bee, Can, Lighter, Vase	673
	Can	Ant, Bee, Beetle, Lighter	626
	Lighter	Ant, Bee, Beetle, Can, Vase	696
	Vase	Ant, Bee, Beetle, Lighter	624

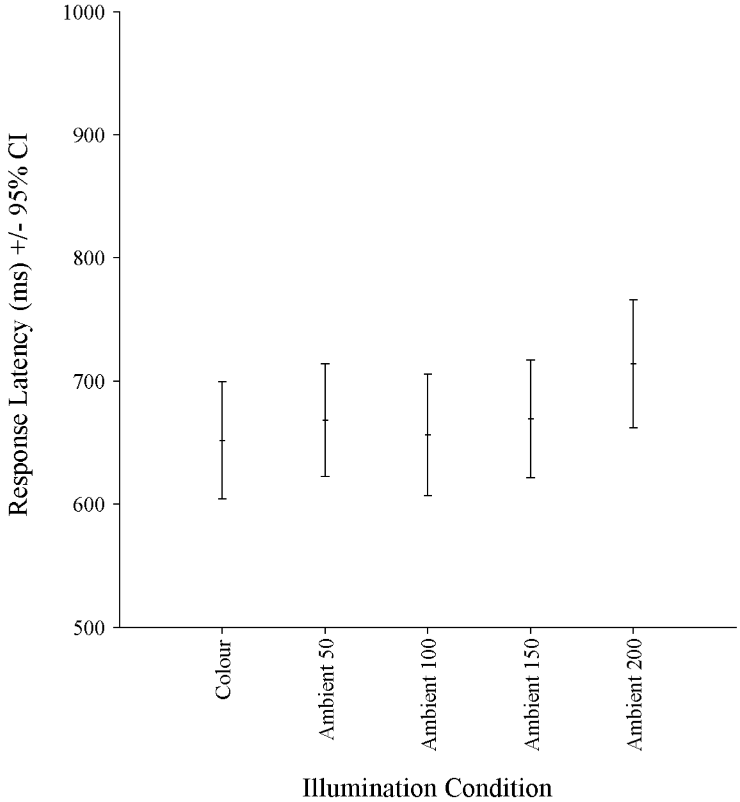

Figure 8.14. Mean response latencies for each Illumination Condition.

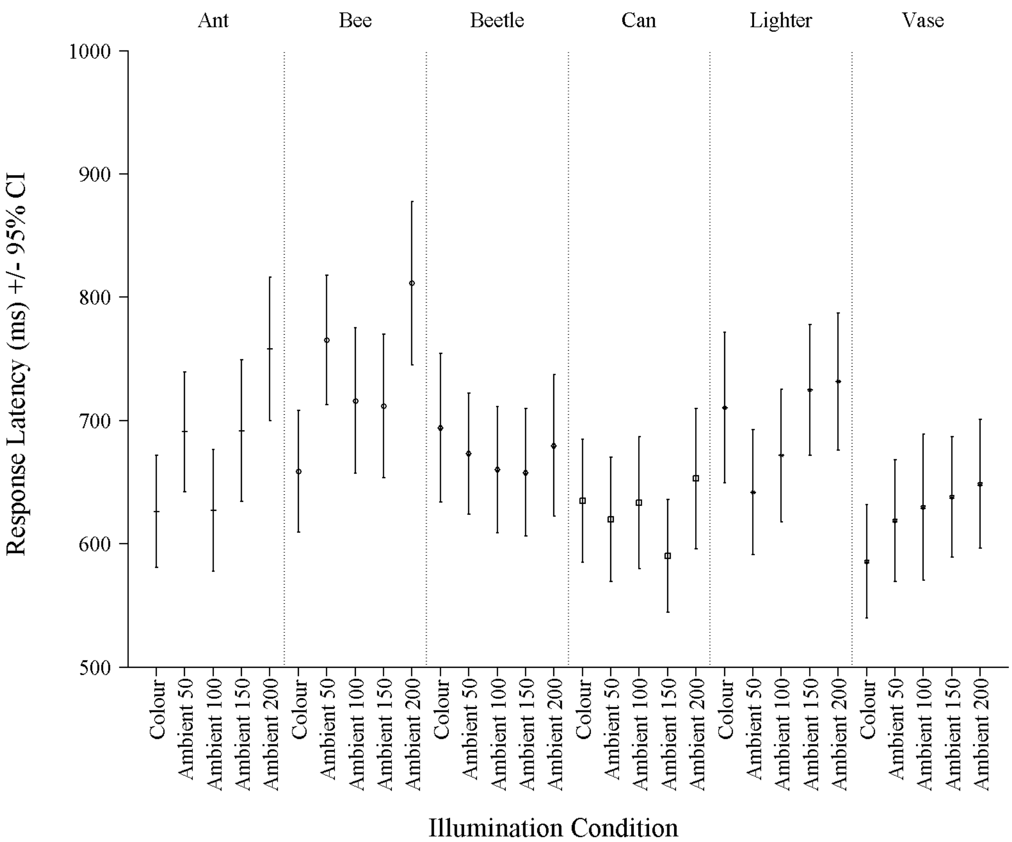

Figure 8.15. The mean response latencies for each Illumination Condition and Object Type. The interaction was significant ($F(10.899,446.845)^* = 7.212$, $\eta_p^2 = 0.150$, $p < 0.05$, *Greenhouse-Geisser adjustment used as the assumption of sphericity was violated). For the Ant images, the Colour and Ambient 100 Conditions produced faster response latencies than the Ambient 200 Condition, and for the Bee images, the Colour Condition produced faster response latencies than the Ambient 50 or Ambient 200 Conditions. A trend of increased response latencies as illumination levels increased is also evident for the Vase images, although there is no significant difference between the means from start to finish. This trend mirrors one of decreased discrimination as illumination increased. For the Beetle, Can, and Lighter, there was also no differentiation in response latency by Illumination Condition.

Discussion

It was hypothesised that shadows would benefit object recognition when shape-from-shading cues were reduced, but that shadows would not be of benefit under normal shading conditions. The results provide partial support for this hypothesis; when assessed across conditions, the presence of shadows was associated with better and faster discrimination. Further analysis showed that this effect was generally restricted to two of the six objects, and, as predicted, only present in the highly illuminated conditions (as shown in Figures 8.5 and 8.12).

In the condition using colour and texture mapping, the presence of shadows produced no effect upon discrimination or response latency (as illustrated in Figures 8.4 and 8.10). A similar result was found under relatively normal shading conditions without colour or texture mapping (the Ambient 50 Condition): shadow presence did not affect response latencies, while there was a discrimination cost associated with the presence of shadows, due to the participants' poor performance on a single object, the Bee. These results support the original hypothesis, in that shadows did not offer any benefit to object recognition under relatively normal shading conditions.

The findings are similar to the results of Freeburg (1966), where accuracy was not improved by shadow presence when the stimuli presented salient cues to recognition (in terms of two prominent features on the textured surface). In Experiment 8, there was no advantage of shadow presence when colour, texture mapping, and normal levels of shading were present. Similarly, Braje et al. (2000), using photos of fruit and vegetables, and Braje (2003) using faces, found that shadow presence did not effect response latency or accuracy.

The result of Experiment 8 differs from the findings of Castiello (2001) who reported a 36 ms response latency advantage for objects presented with congruent cast shadows (a cast shadow that was congruent with the object's shape and direction of illumination, including the lighting direction of any attached shadows), over objects that were presented without any cast shadow. Castiello used voice-activated recording of response latencies, and also concurrently tested conditions in which the object's illumination direction and that of the cast shadow were incongruous, where the shape of the cast shadow was incongruous to the object, and where both cast shadow shape and illumination direction were incongruous to the object's shape and illumination direction. Castiello's results suggested that response times increase when either the direction of lighting is incompatible with the shadow present, or the shadow shape is incompatible with the object shape.

There are several procedural differences between Castiello's (2001) experiment, and Experiment 8 that need to be considered when comparing the results of the two experiments. Castiello's participants verbally named each object. This raises the question of whether a verbal naming task provide more accurate response latency measures than a matching task based upon a manual response. Braje et al. (2000) also used voice-activated recording, and did not find any significant differences between their shadow and no-shadow groups. This implies that the difference in results between Castiello's research and Experiment 8 is not related to the method of response measurement. Furthermore, Davidson and Wright (2002) directly compared verbal-naming latencies with button press latencies in a Stroop task (using a four-choice button-press box). They found button presses to be more accurate than voice activated recording. If the latencies obtained using button pressing are more accurate than those using voice-activated recording, then if an effect is found using voice-activated recording, that effect should be easier to demonstrate using button pressing.

The experiments also differed in terms of how many objects were employed; Castiello (2001) employed 20 objects with one view of each, compared to the six objects employed in this experiment with two views of each. The number of stimuli does not appear to be a critical point in methodology, Braje et al. (2000) also used 20 different sorts of fruit and vegetable, with two views of each, and did not find a benefit of shadow presence.

It is possible that differences between the stimuli may have resulted in the different findings of Castiello (2001) and Experiment 8. However, the experiments share a proportion of similar stimuli: Braje et al. (2000) and Castiello (2001) both used apple, banana, and citrus, although in Braje et al.'s experiment these were in groups, bunches, or sliced; and Experiment 8 and Castiello's experiment both used a can and vase (one third of the objects used in Experiment 8).

The images used by Castiello (2001) were red on a grey background. There is no rationale given for this colour scheme in the paper, nor any assumption that this was critical to the result. Across trials in such an experiment, the contrast of a red object versus a grey background would provide a consistent cue to aid separation of either object from background, or of object from shadow. Therefore, in comparison to a greyscale image, using red objects on a grey background might be expected to reduce the possible influence of any shadow cues.

The last difference between the experiment by Castiello (2001) and Experiment 8, was Castiello's use of trials where either the object and shadow shape were incongruent, the object and shadow illumination directions were incongruent, or both shape and illumination direction were incongruent. Castiello raises the possibility that interference by the incompatible shadows is an explanation for the differentiation he found across the congruent-, incongruent-, and no-shadow conditions. He mentions the Stroop effect (Stroop, 1935) as an example of distracters slowing reaction times, relating that "shadows may act as distracters if they are interpreted as a competing object in the scene" (p. 2308). He suggests that while this would explain a cost in the incongruent-shadow condition versus the no-shadow condition, a cost should also be present in his congruent-shadow condition, as there is still a competing object present, the congruent shadow. However, if an analogy is to be made to the Stroop effect, the congruent shadow would not be classified as a competing object: it would be similar to presenting the word "Green" written in green ink. There are two sources of information, but, they are both sources about the same information. Hayward (1998) has demonstrated that recognition by bounding contour is similar to recognition of shaded objects. Thus, when a cast shadow like that used by Castiello is presented along with a shaded object, the situation is little different to presenting two objects on the screen, where one occludes the other.

In a Stroop experiment using multiple objects, Wühr and Waszak (2003) report latency to responding for incongruous word and colour pairings (e.g., "green" in red ink), neutral word and colour pairings (e.g., "zzzz" in red ink), and congruous word and colour pairings (e.g., "red" in red ink). Congruous word and colour pairings produced the fastest latencies, incongruous pairings the slowest, and neutral pairings produced intermediate values. By considering the similarities of the designs of Castiello's and Wühr and Waszak's experiments, the two sets of results can be compared. Both found that when the two sources of information were congruent, responses were fastest, and when the two sources were incongruent, responses were slowest. When neutral information was presented (as in Wühr and Waszak), or no extra information was presented (as in Castiello), then response latencies were intermediate. Thus, it is possible that the result of Castiello could be a function of using an experimental design that incorporates incongruent trials.

To test this, the procedure of Castiello (2001) could be repeated, but using two objects side by side, one red and one grey, instead of employing an object and shadow. Finding the same results would indicate that the effect seen by Castiello is not because of shadow presence per se, but because of the presence of another information source. Otherwise, a systematic replication of Castiello's study could be performed omitting the incongruent trials, to test if shadow presence still resulted in faster responses than the absence of shadows.

Two systematic replications of Castiello (2001) have been conducted, by Castiello, Lusher, Burton, and Disler (2003), using participants with, and without, left-sided visual neglect. They produced similar results to those obtained by Castiello (2001). Using a procedure involving, congruent and incongruent shadow conditions, plus a no-shadow condition, Castiello et al. (2003) obtained faster reaction times for the congruent shadow condition than the incongruent shadow condition, with the no-shadow condition returning intermediate values across the control subjects in both experiments. Thus, the effect demonstrated by Castiello (2001) has been replicated in each of three experiments with different participants.

Lastly, in Experiment 8, the effect seen for the Ant and Bee did produce a significant main effect (assessed across all objects) of a benefit of shadow presence. At the level of the main effect, the result does match that of Castiello (2001). It would be interesting to know whether individual objects in Castiello's experiment produced different results with respect to the benefit of shadow presence, as Braje et al. (2000) found.

The similarity of the results from the Shadow and No-Shadow Conditions across the Colour and Ambient 50 Conditions suggests that shadows are not of any assistance to object recognition under relatively normal lighting conditions. Aside from the results of Castiello (2001), there is little experimental evidence to suggest otherwise. However, the experimental methodology may not have allowed for the detection of any discrimination benefit due to shadow presence. Accuracy was very high in this experiment, around 90% correct (values presented in Table 8.2). This raises the possibility that there was a ceiling effect present in discrimination. A ceiling in the discrimination measure would prevent the observation of any benefit of shadow presence.

Discrimination was plotted for each object, by Shadow Presence, and for each of the Illumination Conditions, and is illustrated in Figures 8.16 and 8.17. Figure 8.16 shows that there was a marked ceiling in discrimination across the Colour, Ambient 50, and Ambient 100 Conditions. The only object not at a ceiling appears to the Bee in the Ambient 50 Shadow Condition, and in the Ambient 100 No-Shadow Condition. In the Ambient 150 and Ambient 200 No-Shadow Conditions discrimination of both the Ant and Bee images was below a ceiling level (as shown in Figure 8.17). For the Vase images, performance did decrease as illumination increased, but not to the same extent as for the Ant and Bee images: the lowest levels of discrimination for the Vase were about the same as the highest levels for the Ant and Bee images (Figure 8.16).

The presence of the ceiling in discrimination shows that using discrimination measures can be problematic in an object-recognition task: without using reduced presentation times, or reduced image clarity, discrimination is likely to be close to perfect. This thesis has demonstrated that a difference of 200-300 ms in the presentation times used in an experiment could produce results that are either at a ceiling in discriminability, e.g., Experiment 8, or at a floor of chance responding, e.g., Experiment 1. Dependent upon the stimuli and/or procedure, the same presentation times can also produce different levels of discrimination, e.g., compare accuracy in Experiments 3, 5 and 8. It does not appear that using the presentation times of published research will return similar levels of discrimination, even in exactly the same task; compare Experiments 1 to 4 with Tarr et al. (1998). Thus, providing a suitable window of time for participants to view stimuli can be expected to be difficult and time-consuming, in that it may involve multiple sessions of testing or experimentation before suitable presentation durations are established.

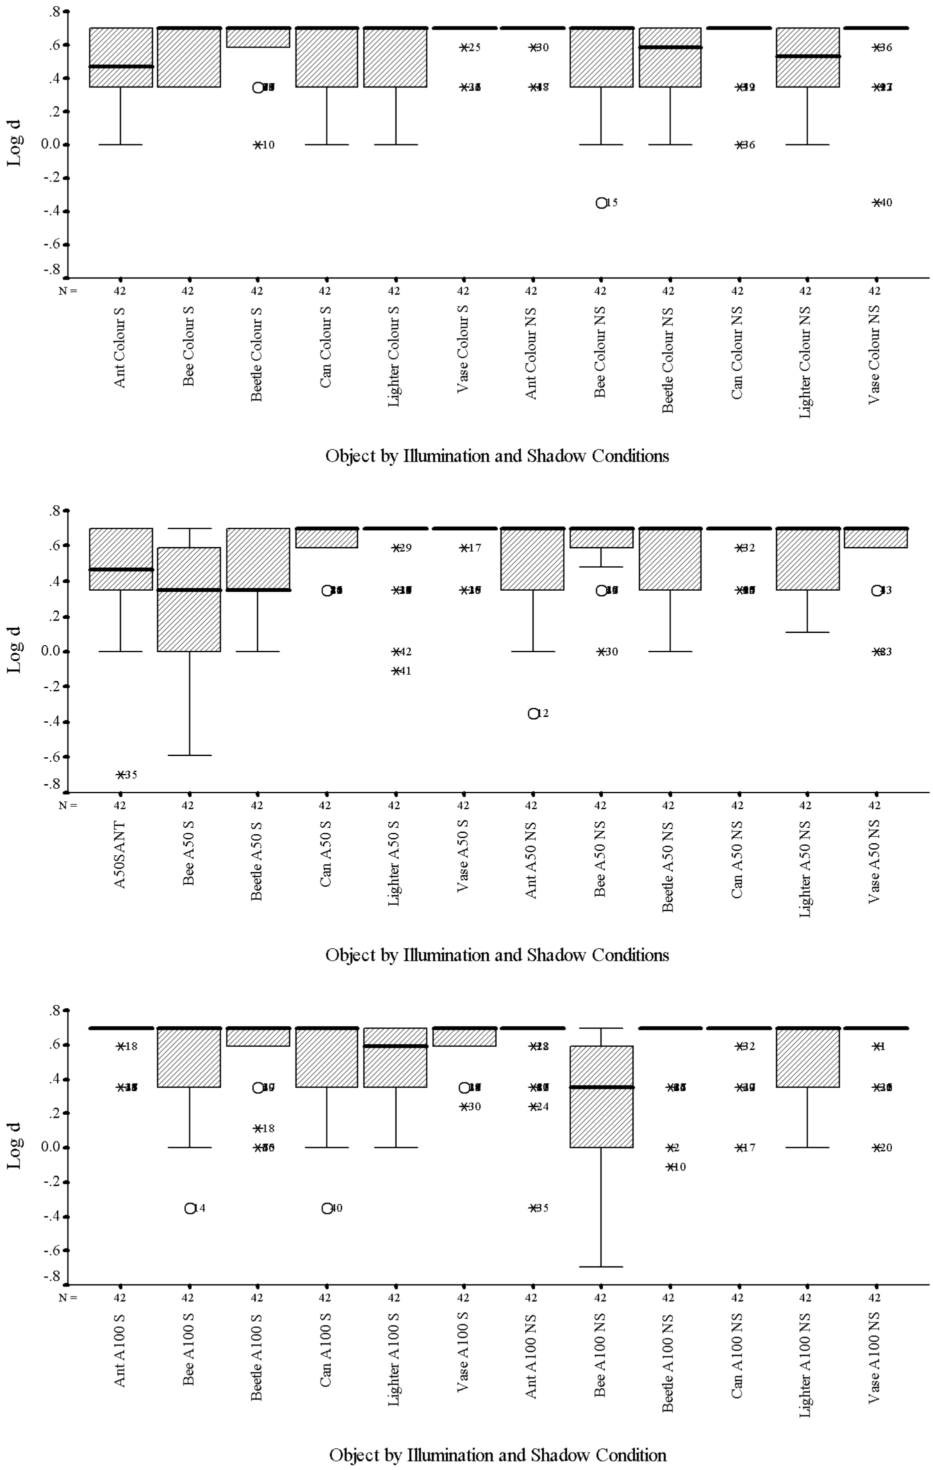

Figure 8.16. Box-plots of discrimination for each object, by shadow condition, for the Colour (top), Ambient 50 (middle), and Ambient 100 (bottom) Conditions. The boxplots show the median (middle bar), 25th and 75th percentiles (box), 25th percentile minus interquartile range and 75th percentile plus the interquartile range or closest data point (whiskers), outliers (circles), and extreme cases (crosses). Outliers are data points between 1.5 and 3 times the interquartile range below the 25th percentile or above the 75th percentile, and extreme values are data points greater than 3 times the interquartile range below the 25th percentile or above the 75th percentile.

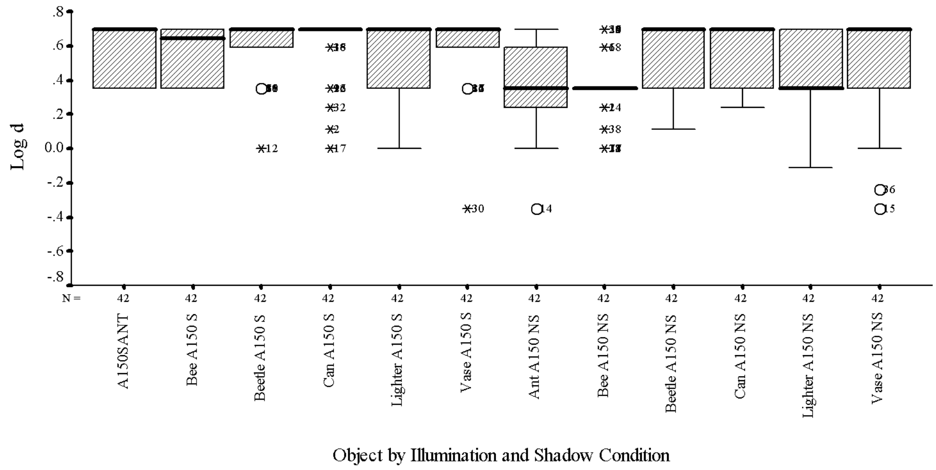

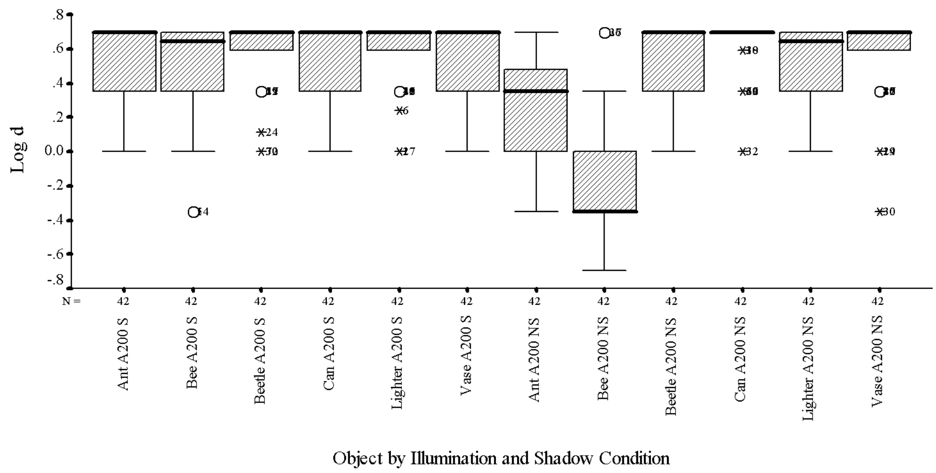

Figure 8.17. Box-plots of discrimination for each object, by shadow condition, for the Ambient 150 (top), and Ambient 200 (bottom) Conditions. The boxplots show the median (middle bar), 25th and 75th percentiles (box), 25th percentile minus interquartile range and 75th percentile plus the interquartile range or closest data point (whiskers), outliers (circles), and extreme cases (crosses). Outliers are data points between 1.5 and 3 times the interquartile range below the 25th percentile or above the 75th percentile, and extreme values are data points greater than 3 times the interquartile range below the 25th percentile or above the 75th percentile.

Even though using severely degraded images would be expected to influence discrimination, in this experiment this was not generally the case: when image clarity was severely reduced in the Ambient 150 and Ambient 200 Conditions, there was not a marked effect upon discrimination for four of the six objects used. The presentation times used here were based upon those employed in previous experiments in this thesis (albeit using a different methodology) and there was no reason to believe that they would produce excessively good discrimination. To reduce discrimination, future experimentation could employ shortened presentation times, yet, for a benefit of shadow presence to become apparent, the utility of either the shadow or non-shadow cues would have to be modulated, with respect to each other, by presentation time. If a benefit of shadow presence was apparent, it would be difficult to establish how the effectiveness of the multiple cues in an image were modulated by stimulus presentation time. Therefore, reducing presentation time in an effort to reduce discrimination may not be a useful change in methodology. To establish how shadow cues interact with other visual cues, the direct manipulation of those cues should provide results that are easier to interpret.

Latency to responding is the other dependent variable employed here, and there was not a pronounced floor evident in response times (as shown in Figures 8.18 and 8.19), as there was a ceiling in discrimination. However, throughout the thesis, response latencies have shown similar patterns of results to discrimination, but these patterns have produced significant results less often. For instance, Figure 8.22 illustrates a trend of increasing response latencies only in the No-Shadow Condition as illumination levels increased, but this increase was not statistically significant. The significant interaction between Object Type, Illumination Condition, and Shadow Presence, did not reveal any consistent differences between the No-Shadow Condition and the Shadow Condition (as shown in Figure 8.12). Thus, based upon reaction times, there is little evidence for a benefit of shadow presence, even under extreme lighting conditions with minimal shading cues present. Thus, even with the problem of a ceiling effect, the discrimination measure provided clearer results than response latency (compare Figures 8.5 and 8.12).

When comparing latency to responding across experiments, there are difficulties in making accurate comparisons. Even systematic replication by

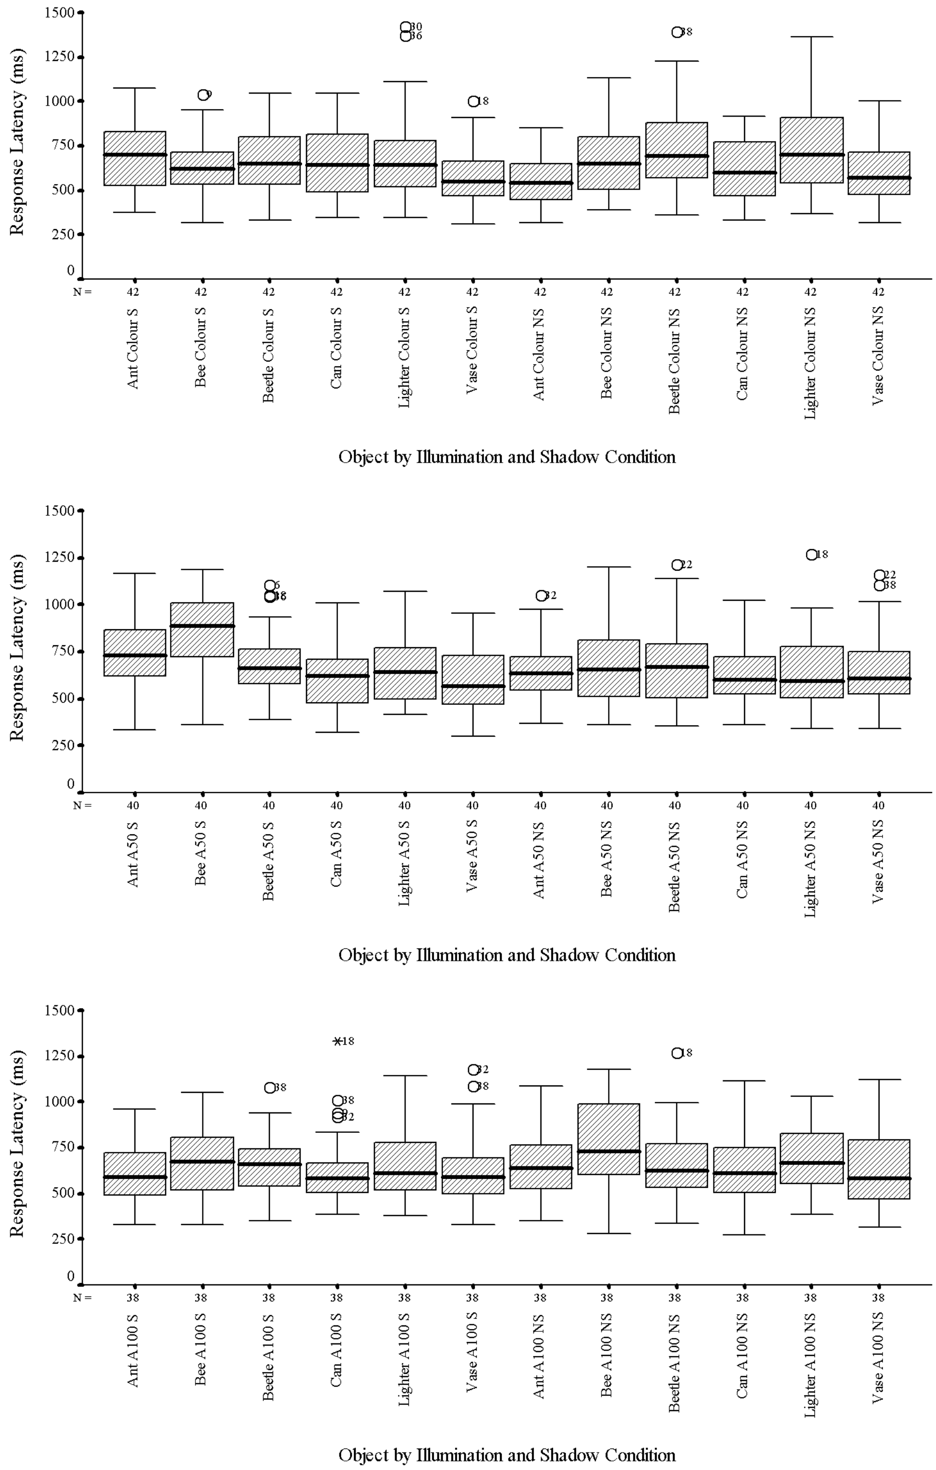

Figure 8.18. Box-plots of response latency (ms) for each object, by shadow condition, for the Colour (top), Ambient 50 (middle), and Ambient 100 (bottom) Conditions. The boxplots show the median (middle bar), 25th and 75th percentiles (box), 25th percentile minus interquartile range and 75th percentile plus the interquartile range or closest data point (whiskers), outliers (circles), and extreme cases (crosses). Outliers are data points between 1.5 and 3 times the interquartile range below the 25th percentile or above the 75th percentile, and extreme values are data points greater than 3 times the interquartile range below the 25th percentile or above the 75th percentile.

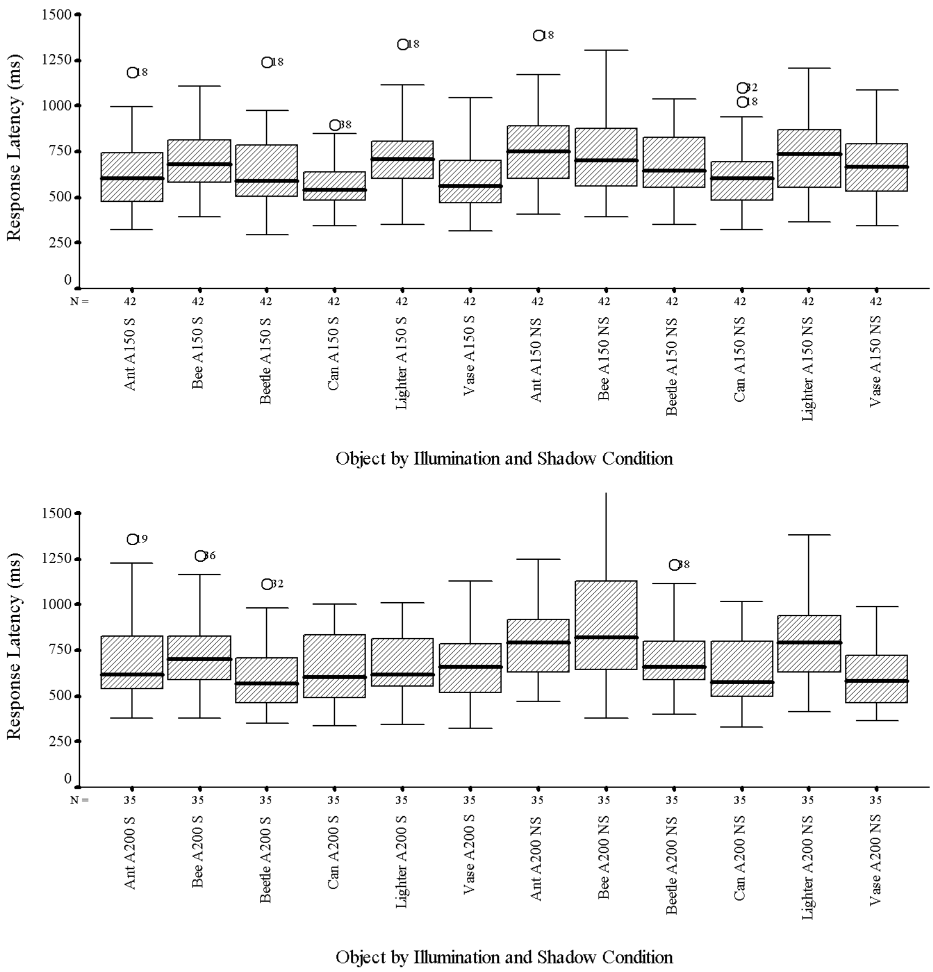

Figure 8.19. Box-plots of response latency (ms) for each object, by shadow condition, for the Ambient 150 (top), and Ambient 200 (bottom) Conditions. The boxplots show the median (middle bar), 25th and 75th percentiles (box), 25th percentile minus interquartile range and 75th percentile plus the interquartile range or closest data point (whiskers), outliers (circles), and extreme cases (crosses). Outliers are data points between 1.5 and 3 times the interquartile range below the 25th percentile or above the 75th percentile, and extreme values are data points greater than 3 times the interquartile range below the 25th percentile or above the 75th percentile.

the same researchers can produce considerably different results. For example, Castiello (2001) reports the average reaction times across conditions as between 444 and 500 ms, while Castiello et al. (2003) report an average reaction time of 778 ms for control subjects. Tarr et al's participants averaged 945-967 ms in the No Shadow Condition across same different object and change/no illumination change conditions, while reaction times in the replication of that work, Experiment 1 of this thesis, averaged 652 ms for the No-Shadow Condition. Furthermore, effects are usually reported in terms of difference in raw scores (e.g., a 50-ms difference), while actual effect sizes are not often cited (e.g., none of: Braje, et al., 2000; Castiello, 2001; Castiello et al., 2003; or Tarr et al., 1998, cite effect sizes). If the reader wishes to calculate an effect size from the results presented in the literature, standard errors are usually presented only graphically, making any calculation only a rough approximation.

Although both discrimination and response latency do have limitations as measures, they were assessed in Experiment 8. The results led to two initial conclusions about the effect of shadow presence on object recognition. Under normal viewing conditions, it seems that one, people cannot do the task any better (they are at a ceiling level of performance anyway), and two, people do not perform the task any faster.

What happens in situations where viewing conditions are not normal? Most of the images used in the experiment were rendered without texture mapping or colour, and there was a paucity of shape-from-shading information in the Ambient 150 and Ambient 200 Conditions. It was in these two conditions only, that discrimination dropped on No-Shadow Trials for the Ant and Bee, while being maintained on Shadow Trials (as depicted in Figure 8.5). For these two objects, shadows had a demonstrated utility: maintaining a high level of discrimination when discrimination dropped in the equivalent No-Shadow Condition.

The results show that people can use the information provided by combined attached and cast shadows to facilitate object recognition. This is a different finding from the results of experimentation using projected shadows or silhouettes, which present the cast contour of a cast shadow (e.g., Hayward, 1998), or experiments using two-tone Mooney images which do not differentiate between shadowing and shading in the two-tone thresholding process (e.g., Moore & Cavanagh, 1998; Moore & Engel, 2001). In this experiment, when shading was insufficient to facilitate recognition, the presence of information from shadows maintained previous high levels of discrimination.

Shadows contain several sources of information, e.g., attached contours, occluded-surface contours, and cast contours. What information is used by the visual system is not known. Figure 8.20 demonstrates the perception of depth that can be achieved by the presence of shadows (also see: Kersten, 1997), although the debate upon whether the perception of depth and 3-D structure is relevant to object recognition is still open (e.g., Hummel, 2000; Tarr, 2003). As well as providing cues to depth, shadow borders may provide cues to shape. Cavanagh (1991) makes a case that cast contours are not likely to be used by the visual system, as their shape will be a function of the casting surface plus the receiving surface. However, Castiello's (2001) research suggests that cast contour information is processed by the visual system: the differences in an image produced by changing between an incongruous and congruous shadow would be most evident in the different cast contours of those shadows, as opposed to changes in extremal borders where the object occludes the shadowed background.

Cavanagh (1991) suggested that attached and occluded boundary contours would be of greater use than cast shadows, and Figure 8.20 illustrates that shadows can enhance differences in brightness across the border between an object and the background, or between separate features of the object (e.g., the wings and thorax of the Bee). Experiment 8 was designed to partially control the brightness gradient across the object-background border. In the No-Shadow Condition, when the objects occlude a white background, the difference in brightness at these borders decreased as the level of illumination increased. In the Shadow Condition, the difference in brightness between object and background remained static where shadows were cast on the background (at least until the object was totally white, i.e., at a brightness ceiling). To further the assessment of any differential contribution of cast versus attached shadows, Experiment 8 could be repeated but using a deformed ground plane to add noise to the cast-shadow contours. If the same effect was found this would support the theory that attached and extremal shadow contours are important sources of information.

Figure 8.20. Illustration of the difference between silhouettes (top) and the images used in Experiment 8 (bottom left = Shadow Condition, bottom right = No-Shadow Condition). Comparison of the bottom left and right images demonstrates the perception of depth due to shadow presence. The shadowing in the bottom left image also highlights many of the same bounding contours as the silhouette.

However, other work by Moore and Cavanagh (1998) suggests that at least the resolution of shape from shadows (if not the recognition of objects) does not require the presence of attached or external borders. Work with Mooney images provides evidence that shadowed areas need to be identified in order for 3-D shape perception to take place. Moore and Cavanagh demonstrated that the presence of shadows in two tone images can reduce discrimination, and Moore and Engel (2001) found that shadow areas are interpreted more slowly, and less accurately, than highlights; the two toning process making the differentiation between object and shadow less likely to occur.

Thus, two-tone images may offer a way to investigate whether attached and external, or cast borders are important for object recognition. Mooney images destroy the penumbra at the cast border of a shadow, removing a vital cue to the identification of a shadowed area (Hering, 1874/1964; MacLeod, 1940), and they mask external borders where the object meets shadow. Similarly, Cavanagh and Kennedy (2000), and Kennedy and Bai (2000), discuss how shape from shadow is destroyed in Mooney images by putting a dark line around the shadow area (as illustrated in Figures 1.8 and 1.9).

Moore and Cavanagh (1998) suggested that a contrast border between object and shadow (e.g., an attached or external border) is not necessary for 3-D shape perception to take place. Their participants were presented with line drawings of the same 3-D shapes used in their two-tone images. In the line drawings, the object-shadow borders were masked by black blobs, Moore and Cavanagh report good 3-D shape perception by their participants. Similarly, modified two-tone figures presented in this thesis (presented in Figures 8.21 and 8.22, objects after Moore & Cavanagh, 1998), illustrate that the presence of a penumbra enables 3-D shape perception of novel objects depicted in two-tone, even though many attached and external shadow borders are obscured by the two-toning process.

Whether any of these cues (depth due to shadowing; cast, attached, or external borders) contribute to discrimination under normal viewing conditions cannot be assessed due to the ceiling effect discussed above: although there is no indication that they enhance the speed of recognition. Thus, a condition upon the finding that combined cast and attached shadows can

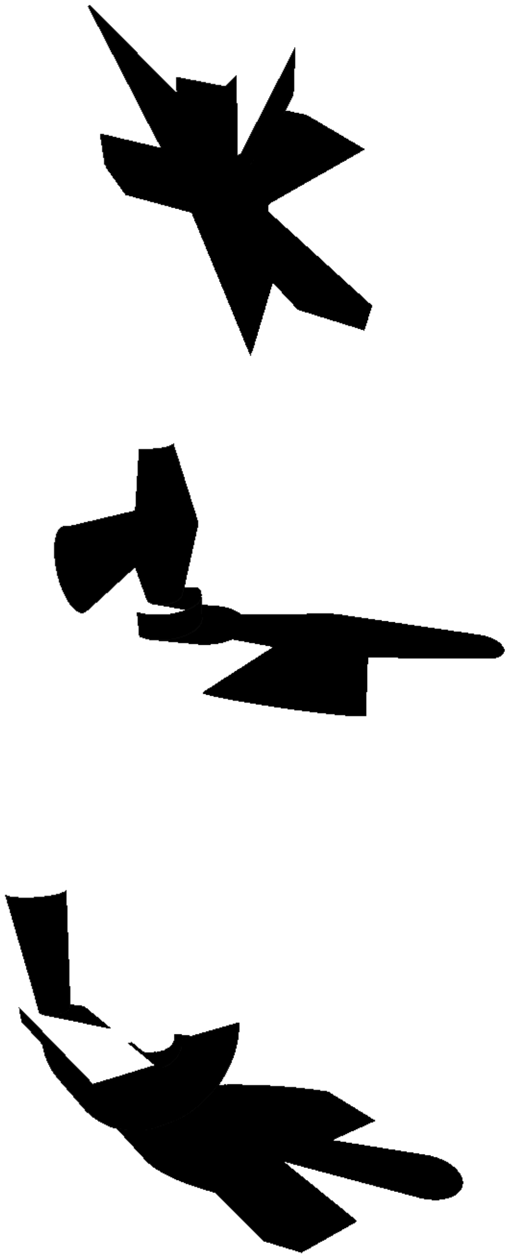

Figure 8.21. Two-tone images similar to those used by Moore and Cavanagh (1998). Moore and Cavanagh found that perception of the 3-D shape of these objects was not well supported by the two-tone images, when the images were novel (also see Moore & Engel, whos' participants performed better than Moore & Cavanagh's in the same shape extraction task).

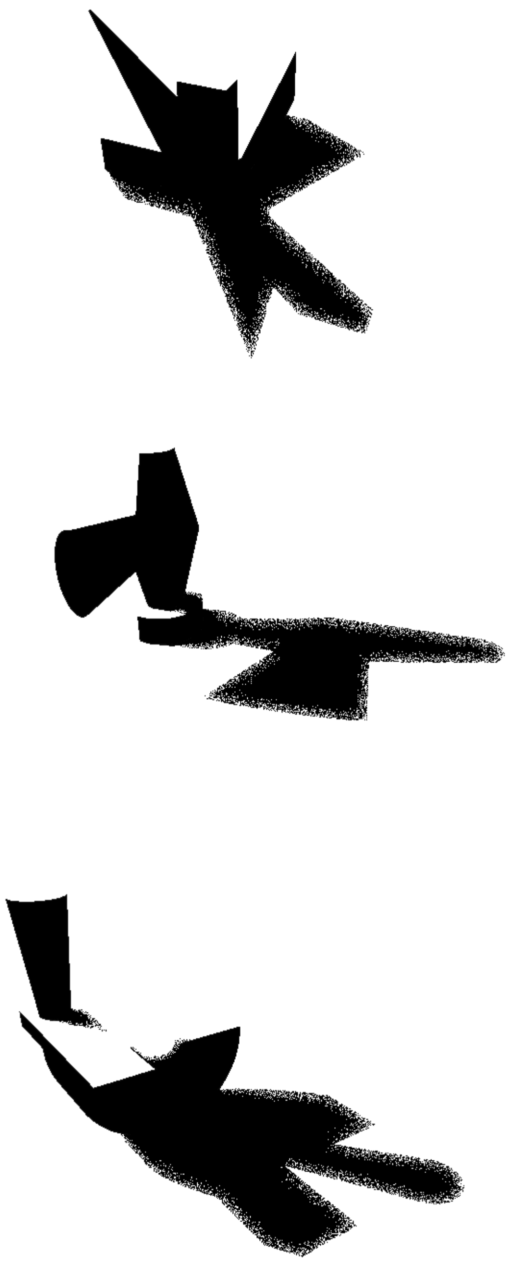

Figure 8.22. Two-tone images similar to those used by Moore and Cavanagh (1998), but where the shadows have penumbra. Contrast between object and shadow at the object's extremal boundary is obscured as in a normal two-tone image, however, perception of 3-D shape appears to be well supported.

facilitate recognition, is that this effect was only demonstrated in conditions where shading was reduced to such an extent that, without shadows, discrimination started to fail.

Why was there an effect for only two objects? As outlined above, it was possible to observe an effect for the Ant and Bee, as they were the only objects for which there was a marked drop in discrimination in the highly illuminated conditions. It is possible that this drop in discrimination was due to these two objects being the most similar and therefore the hardest to discriminate between.

To establish if the Ant and Bee images were more difficult to discriminate between than the other objects (irrespective of shadowing), error rates were calculated for non-matching Word-Picture trials, for each of the different Word-Picture combinations (as shown in Figure 8.23.). For comparison, error rates for matching trials were also calculated. A repeated-measures ANOVA was conducted upon error rate across the different initial word stimuli, and the different picture stimuli. There was a moderate effect of the differing word stimuli ($F(5,200) = 9.663$, $\eta_p^2 = 0.195$, $p < 0.05$), where error rates were higher for the Ant, Bee, and Beetle Trials, than for the Can, Lighter, or Vase Trials. The error rate on the Can Trials was also significantly higher than that for the Vase Trials. Overall, changes in the initial stimuli accounted for approximately 20% of the variation in error rates.

There was a large effect upon error rates of the differing picture stimuli ($F(3.546,141.857)^* = 26.414$, $\eta_p^2 = 0.398$, $p < 0.05$, *Greenhouse-Geisser adjustment used as the assumption of sphericity was violated), with the pictures controlling approximately 40% of the variation in error rate. The Bee picture produced error rates higher than any of the other pictures, at about 20% error, with the Ant and Beetle pictures also producing higher error rates than the Can, Lighter, and Vase.

The interaction between the word stimuli and picture stimuli is used to address whether the participants had difficulty distinguishing between any particular pairings of words and pictures. The interaction was significant ($F(10.044,401.765)^* = 15.526$, $\eta_p^2 = 0.280$, $p < 0.05$, *Greenhouse-Geisser adjustment used as the assumption of sphericity was violated), and is illustrated in Figure 8.32. The error rates of interest are those where the Ant and Bee stimuli appear in combination. When the word “Ant” was the initial

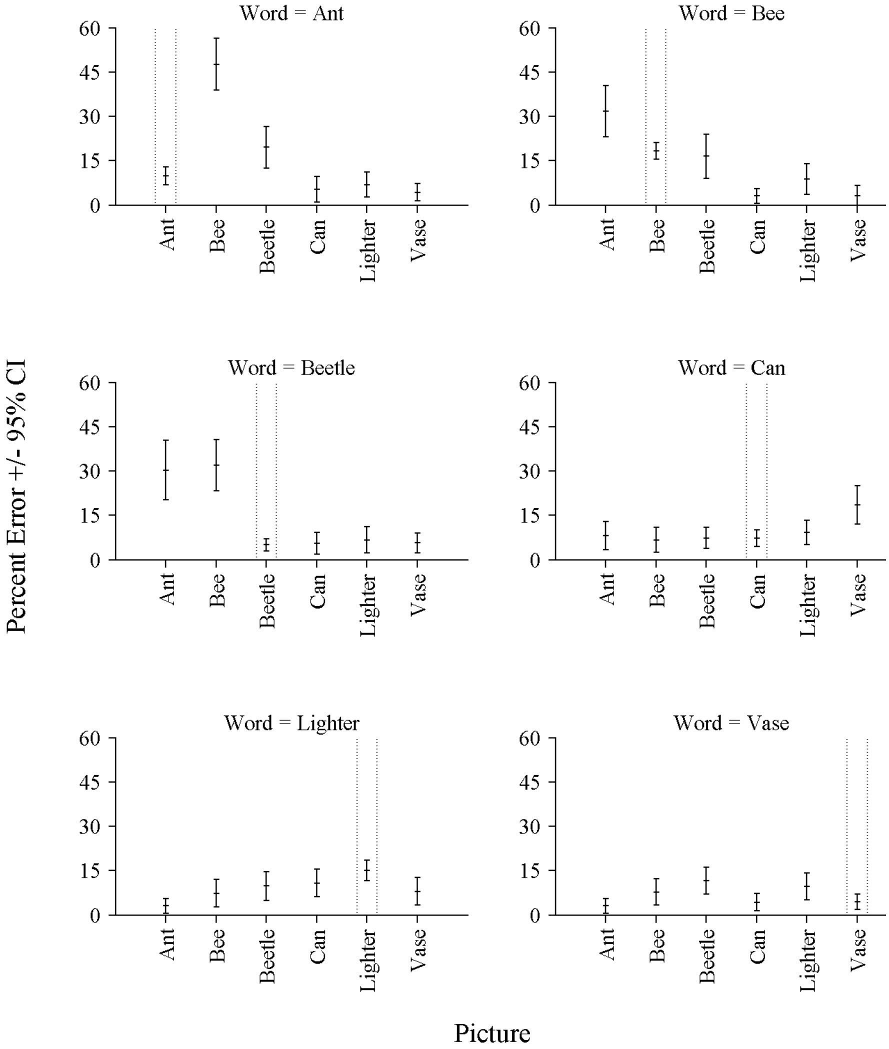

Figure 8.23. The mean error rates for non-matching Word-Picture trials, presented across the different words and for each picture (error rates for matching Word-Picture trials are shown within the dotted lines). Confidence intervals of 95% about the mean error rates are given.

stimulus, the Bee was erroneously said to match 48% of the time, and when "Bee / Fly" was the initial stimulus, the Ant was erroneously said to match 32% of the time. Error rates when either the Ant or Bee were paired with the Can, Lighter, or Vase, ranged from 4 to 9%. However, high error rates also occurred when the word Beetle was paired with the Ant and Bee Pictures, and vice versa: when the word Beetle was the initial stimulus, the error rate on Ant pictures was 30% and on Bee pictures it was 32%. When the Beetle picture was paired with the "Ant" or "Bee / Fly" initial stimulus, error rates were 20% for "Ant" then Beetle image, and 16% for "Bee / Fly" then Beetle image.

Thus, there were three objects that were considerably more difficult to discriminate between than the others were: the Ant, Bee, and Beetle. Yet, only for the Ant and Bee was there an effect of the presence of shadows. For the Beetle there was no effect of the presence of shadows upon discrimination (as illustrated in Figures 8.5 and 8.17). Therefore, the difficulty of the discrimination between the Ant and Bee objects does not explain why there was an effect of shadow presence for only these two objects, as the Beetle initial and comparison stimuli also produced high error rates when combined with the Ant and Bee[1].

Inspection of Figures A6.1 to A6.12 suggests that even where shading was minimal in the Ambient 200 Condition, for a portion of the images there was still a reasonable amount of bounding contour information available from shading, e.g., the Can in view 2, the Lighter in View 2, and the Vase in both views. Without shadows present there is therefore still a contrast border at the object edge. This limited information was generally sufficient to maintain discrimination.

If this shading information was reduced to the most extreme case, where there is no shading information present at all, there is no doubt that an effect of shadow presence could be demonstrated across all objects. Removal of all shading information would make the No-Shadow Condition entirely white and reduce discrimination in this condition to chance levels.

1. Further investigation of the average error rates for each object view did not reveal any notable differences across the Shadow Condition and No-Shadow Condition. This indicated that the difficulty of discriminating the Ant, Bee, and Beetle, was not related to individual images of these objects, either with, or without, shadows present.

Pixel by pixel subtraction of the No-Shadow Condition images from the Shadow Conditions images produces images that only contain shadow information (Figure 8.24 presents examples from Version 1 of each object). The figure demonstrates that object recognition can be facilitated by shadow information in the total absence of shading information[1].

Conclusions

The results suggest that under normal viewing conditions, any cues available from the addition of shadowing are not sufficient to increase the speed or accuracy of object recognition. However, when cues such as shading are too degraded to facilitate normal recognition performance, shadows can provide information that enables similar levels of discrimination and reaction times to those found in normal viewing conditions. A mechanism by which this might occur is where the presence of shadows increases the contrast between the brightness of an object and that of the background, or increases the contrast between different parts of an object. A large gradient in brightness across an attached or extremal shadow border would serve to highlight the casting contour of an object or its parts: this may provide valuable cues to aid object recognition. The size of any gradient in brightness across the discontinuities of an object, or at its bounding contour, would be dependent upon several factors, e.g., the illumination conditions, physical properties of the object and background, and viewpoint. Therefore, producing such a benefit would be very context dependent. The current results indicate that under normal shading conditions such a cue is unlikely to affect the speed or accuracy of object recognition. The results of previous research (Experiments 5 to 7; Castiello, 2001) indicate that how, and when, the visual system employs the information available from a shadow's cast contours needs more investigation.

1. Note: these images differ from traditional Mooney images, in that the thresholding procedure used to create Mooney images incorporates shaded sections of the object into the dark areas, and thus, shadows can be perceived as object surfaces (see Moore & Cavanagh, 1998)

Figure 8.24. Illustration of images that contain shadows but no shading. The images are very similar to those from the Ambient 200 Condition (compare to Figures A6.1 to A6.12), and as such would be expected to produce discrimination close to that obtained for that condition. Images from a comparable No-Shadow condition would be totally white, and therefore could not be discriminated.

General Discussion

Langer (1999) states that there are three sorts of illumination phenomena in the visual world: classical shading, shadowing, and interreflections between surfaces. The experiments conducted as part of this thesis investigated how one of these phenomena, shadowing, can affect human object recognition. Specifically, whether shadow presence contributes to improved accuracy, and/or, faster responding in object-recognition tasks. The results of the experiments reported here show that shadows can be of benefit to object recognition, but that in the majority of situations, shadow presence does not noticeably affect recognition.

When assessing accuracy or discrimination, this finding parallels the results of the research reviewed, where a general benefit of shadow presence, has not been demonstrated in an object-recognition task. Shadows can disambiguate otherwise ambiguous shape, as demonstrated by the experiments of Yonas' (1978), Berbaum et al. (1984), Erens et al. (1993), and Bülthoff et al. (1994), and Freeburg (1966) demonstrated that shadowing can improve matching of textured surfaces under moderate degrees of shadowing and task difficulty. Yet overall, there has been no study into the recognition of objects that has demonstrated a general benefit in accuracy or discrimination due to shadow presence, e.g., Experiments 1 to 8, and Braje et al. (2000) all report no benefit of shadow presence, and Castiello (2001) only reports a 2% error rate across all conditions. The error rate of Castiello's participants, and the results of Experiment 8, highlighted the problem associated with demonstrating a discriminatory benefit due to shadow presence: people are very good at object recognition, usually performing at close to perfect levels of discrimination.

Braje et al. (2000) reported differences in accuracy between shadow and no-shadow conditions for individual objects, in some cases shadow presence was beneficial, in others it was a hindrance, but overall there was no effect of shadow presence. In Experiment 7, the cast borders of cast shadows only benefited recognition when the objects themselves provided few bounding contour cues, and in Experiment 8, shadow presence was beneficial only when cues to shape from shading were negligible.

As well as discrimination, latency to responding, or reaction time, has often been used as the dependent variable in object-recognition experiments.

In object-recognition tasks, there is little association of shadow presence with faster latencies to responding. Throughout the course of this thesis, latencies to responding have provided less differentiation between the shadow and no-shadow conditions than that shown by the discrimination measures. Tarr produced a benefit of approximately 100 ms in the matching of novel objects, which could not be replicated here in four attempts. Castiello (2001) and Castiello et al. (2003) produced response time benefits between 36 ms (2001) and 41 to 74 ms (2003). Neither of these results were directly stated as being statistically significant benefits, although the effect is reliable across experiments.

The experiments by Castiello (2001) and Castiello et al. (2003) provide the only reliable effect of shadow presence upon response times. As discussed in Experiment 8, further systematic replication of these results is required to confirm that the effects are specifically related to shadows, rather than the use of multiple stimuli in the displays. Two possible modifications of the procedure would be to: one, use only a congruent shadow and a no-shadow condition; and two, to perform a similar experiment, but using two objects in each image, rather than an object and a shadow.

However, the research conducted to date does indicate that shadows are cues that can be used to enable recognition in certain situations (e.g., Experiments 5 to 8). The benefits of shadows as cues to object recognition can be compared to the potential benefit of colour to object recognition: colour is not always beneficial. For some objects colour is a highly diagnostic cue, and for others it has low diagnosticity (Nagai & Yokosawa, 2003). As cues, shadows need to be treated in the same manner, with research focusing on the conditions under which they may be useful, as opposed to looking for an all or none benefit to recognition.

The experiments reported in this thesis indicate that these conditions are limited: it seems that shadows may be of benefit when other cues are minimal or absent, i.e., situations where shadows offer the most salient information. For example, in situations where: shading is extremely reduced (Experiment 8); other shape cues are subtle (Experiments 5 to 7 and Freeburg, 1966); or only a silhouette/cast shadow is visible (Warrington & James, 1986). Across Experiments 5 to 8, and in Freeburg's experiment (1966) when more cues were available, there was no effect of shadow presence.

Where shadow presence has had an effect, some clarification is still required as to which shadow borders produced the effect. Experiments 5 to 7 indicate that the cast border of a cast shadow can be used as a source of information regarding the shape of the object. This finding is supported by the results of Castiello (2001) and Castiello et al. (2003), although the degree to which external shadow borders were also present in these experiments is unknown. Experiment 8 and Cavanagh's (1991) work advance the suggestion that attached and external shadow borders can also enable recognition. Future research could focus in this area. Rendering images against irregular background surfaces may enable the separate effects of cast borders, versus attached and external borders, to be evaluated. The creation of shadow-only images has also been discussed in this thesis, and hopefully similar techniques will allow continued research akin to that conducted with Mooney images, but using two-toned, shadow-only images.

The progression of the experiments in this thesis illustrates an important point. To date, conclusions regarding the effects of shadow presence upon object recognition have often been quite generalised, yet research into the effects of shadowing upon object recognition has involved little systematic variation of the amounts of shadowing employed, the types of shadow borders being assessed, or the degree to which other cues to shape are present. This thesis has started to address these areas by assessing the utility of different shadow borders under multiple levels of object foreshortening and shape-from-shading cues. As part of this process, several methodologies have been proposed that should aid future research in evaluating the value of shadow information for object recognition. Methods of quantifying the contribution of shadow to an image have been used in both, the post-hoc testing of the effects of shadow presence (Experiments 1 to 4), and in the selection of stimuli (Experiment 8). In a similar manner, a method of quantifying the change in silhouettes between S1 and S2 was used to establish the contribution to discrimination of bounding-contour change between S1 and S2.

A side note to the thesis, is the reduction in response times across trials noted in the discussion of Experiments 1 to 4. Earlier, the idea that this reduction in response times across trials could have masked the effect demonstrated by Tarr et al. (1998) was ruled out. Each of the next four experiments' results were examined for evidence of the same pattern. For

Experiments 1 to 4, the analysis was conducted upon correct trials, to enable the results to be compared against those of Tarr et al. (1998), but as noted, analyses of the entire data sets revealed the same trends. For Experiments 5 to 8, analyses of response time over trials was limited to the entire data sets. When averaged across participants, the correlations between response time and trial order were: -0.23 in Experiment 5; -0.24 in Experiment 6; -0.19 in Experiment 7; and -0.26 in Experiment 8. In Experiments 1 to 4, the range was -0.19 (Experiment 1 No Shadow) to -0.32 (Experiment 3 No Shadow).

Thus, reductions in response times over trials are evident across all the experiments, and is of an equivalent magnitude (One-way ANOVA: $F = 0.69$, $p > 0.05$). It is present across: one, the different trial orders used for Experiments 1 to 4, Experiment 5, Experiment 6, Experiment 7, and Experiment 8; and two, the different experimental methodologies of Experiments 1 to 4 (sequential picture-to-picture matching), Experiments 5 to 7 (two-alternative forced choice), and Experiment 8 (sequential word to picture matching). Therefore, speculation that the relationship seen in Experiments 1 to 4, of a reduction in response time across trials, was an effect of the particular sequence of trials used can be eliminated. To avoid this speculation in future experiments, it would be advisable to employ different stimulus presentation orders for participants within experimental conditions. The trend of a reduction in responses time across trials, is also a reminder for researchers to ensure that the different trial types they use are evenly distributed across the trial order (and typically unpredictable across the trial order) for each participant.

Theoretical Implications of the Thesis:

The final points of discussion are the implications of the eight experiments' findings upon current object recognition theory. There are presently two significant object recognition theories, structural description theories, and view-based theories. Structural description theories posit that recognition is based upon the representation of objects in the visual system using 3-D shape dimensions (Biederman, Subramaniam, Bar, Kalocsai & Fiser, 1999; Stankiewicz, 2002). Biederman et al. (1999) relate that the qualitative non-accidental properties (NAPs) of 2-D images, such as parallelism and co-linearity, are the basis for our perception of an object's 3-D parts and of the spatial relation between the parts. Stankiewicz (2002) suggests that some 2-D metric properties (primary-axis curvature, cross-section, aspect ratio) are used in the determination of an object's 3-D structural description. View-based theories suggest that recognition is based upon 2-D templates or exemplars of multiple views of a single object, with recognition being achieved by assessing the degree of similarity between the 2-D input and the learnt views (e.g., Tarr and Bülthoff, 1995; Edelman, 1998).

In the last 15 years, the debate over whether recognition is performed on the basis of structural descriptions or multiple views has produced an abundance of research, but no reduction in the number of theories being offered, and little consensus of opinion on how the theories differ, or what the critical test of each would be. Peissig, Wasserman, Young, and Biederman (2002) state that the conceptual differences between the structural description and view-based theories may be more apparent than real, that the theoretical perspectives are often caricatured in order to draw distinctions, and that, at least in their study, the theories make nearly identical behavioural predictions. They also note that the theories continue to evolve in response to the available empirical data.

Stankiewicz (2002) notes that when attempting to demonstrate the validity of one of the theories over the other, the majority of scientific endeavour has been put into researching the effects of viewpoint dependent and viewpoint independent object recognition, even though this is not the critical test of the theories:

> Most object recognition studies addressing the question of whether human object recognition is mediated by a view-based or a structural description have investigated whether human object recognition behaviour is viewpoint sensitive or viewpoint invariant… The research is premised upon the assumption that an object recognition system that is view based will always demonstrate viewpoint-sensitive behaviour and one that is based on a structural description will always demonstrate viewpoint-invariant behaviour. Unfortunately, this is not the case. (Stankiewicz, 2002, p.914)

Stankiewicz (2002) relates that the question of viewpoint sensitivity cannot differentiate between the broad classes of models. He suggests that the critical question is whether human observers can make independent estimates of different three-dimensional shape dimensions (e.g., independently estimate aspect ratio and primary-axis curvature): if they can, then this would support the use of a structural description representation by the human visual system. However, there does appear to be a leap in logic here, even if we can perform independent estimates of three-dimensional shape dimensions, this does not mean that we perform object recognition in this manner.

In considering viewpoint sensitivity as the critical question, Biederman et al. (1999) and Biederman (2001) state that all theories are view-based, and under certain conditions will produce viewpoint sensitivity. Biederman (2001) states that the important question is not whether a representation is view-based (as all representations are view-based), but what that representation is, i.e., is it a qualitative structural description, a feature list, or a set of multiple views.

In interpreting the results of this thesis in light of the different theories of object recognition, three of the points above are particularly salient. The first is that all theories are view-based (Biederman, 2001). The second is that the distinctions between the theories are unclear (Peissig et al., 2002). The third is that the theories continue to evolve in response to the available empirical data (Peissig et al., 2002).

With regard to the first point, that all the theories are view-based, Biederman (2001) (one of the most significant proponents of structural description theories) notes that a template representation may well characterise the earlier stages of cortical processing (V1 to V4) and that metric templates are likely to be the representation mediating face recognition. Stankiewicz (2002) states that all extant structural description models begin with simple image features as their initial input, and they then make use of the vertices and edges present. Given that both sorts of theory propose the same input, and the same initial 2-D view analysis, it is difficult to form hypotheses about how each theory would produce different predictions about the effects of shadow presence. For example: in the case of structural description theories, changes in illumination could either obscure or enhance (dependent upon situation) the discontinuities evident in an image that are necessary for the perception of NAPs and 3-D shape dimensions, e.g., parallel edges, cotermination of edges. In the case of view-based theories, the same object discontinuities are necessary for perception of the 2-D view required to judge similarity to a stored view.

The second point highlighted was that a critical test of the theory types has not been established. In this situation, and given that the experiments were not designed to differentiate between the two theory types, any comment upon how the results may support one theory other the other can only be conjecture. As discussed below, both types of theory can cope with the results of the experiments.

The last point was that the theories continue to evolve in response to the empirical data. The main results (from this thesis) that the theories need to be able to account for, are addressed next with respect to each block of experiments.

In Experiments 1 to 4, accuracy in a sequential matching task was highly correlated with gross differences in silhouette between S1 and S2. Hayward (1998) notes that his finding, that the visual system can match silhouette images virtually as accurately as shaded images, does not exclude either sort of theory, although there is no present algorithm to extract NAPS from silhouettes alone. Hayward's finding parallels the neuro-physiological research by Vogels and Biederman (2002) and Kayaert, Biederman, and Vogels (2003) into the effects of illumination on object coding in Macaque inferior temporal cortex. Vogels and Biederman selected cells that respond to shaded versions of two part objects, and found that these cells also responded (albeit to a lesser degree) to silhouettes of those objects (which are missing the internal contrast borders). This indicated that the external border was sufficient for a selective response. Kayaert et al. tested macaque IT neurons using one and two part shapes as stimuli. In 65% of cases the silhouettes produced a response equally as strong as shaded images. Furthermore, the experiment compared response modulation in the neurons to changes in the NAPS and the metric properties (MPs) of the images. Kayaert et al. found that, even after reduction to silhouette images, changes in NAPs continued to produce significantly greater response modulation than changes in metric properties (MPs). The finding suggests that NAPs (the requirements for Beiderman's structural description theories) are recoverable from silhouettes. However, the lack of distinction between the types of theories is again highlighted when Kayaert et al. (2003) note that their finding is consistent with structural description theories, but could also be incorporated by view-based models:

> The observed NAP advantage is consistent with one of the assumptions of the geon structural description (GSD) model (Biederman, 1987; Hummel and Biederman, 1992). Geons are the shape primitives of the GSD model and are defined by contrasting NAPs, a distinction not incorporated into current viewbased models. View-based models could be modified to incorporate the NAP advantage without including other assumptions of GSDs (or structural descriptions, in general), such as the explicit coding of the relationships among object parts. In general, the present work shows that the use of a computationally inspired parameterization of shapes can provide at least hints of the principles behind shape coding by primates. (Kayaert et al., 2003, p. 3026)

Therefore, the use of silhouettes to perform the matching task in Experiments 1 to 4 can be accounted for by both structural description and view-based theories. The two theories need to also account for the viewpoint sensitivity reported in Experiments 5 to 7.

According to structural description theories, viewpoint sensitivity could be due to more geons becoming visible, those geons already visible becoming less occluded, or the geons evident being poor predictors of object classification (e.g., attempting to use very similar geons to make a sub-ordinate level classification). View-based theories would account for viewpoint sensitivity because as both objects become foreshortened, their 2-D views become increasingly similar. Discrimination (and/or reaction time) may then reduce because the level of similarity of either objects' 2-D input to stored views of the other object is greater. Thus, as discussed above and by other researchers (e.g., Biederman, 2001), both theories can account for a finding of viewpoint sensitivity.

The theories also need to account for the result that the provision of shape cues from two sources (object and cast-shadow) produced an additive benefit to object recognition in Experiment 5 but not in Experiment 6 after familiarisation with an object. In Experiments 5 and 6, cues to shape were restricted to little more than two sets of outer contour, and discriminability was related to the least foreshortened (most informative) outer contour. It was discussed with respect to Experiments 1 to 4, that both structural description and view-based theories can cope with recognition from outer-contour alone. In Experiment 5, when the shadow contour was more informative than the object contour, discriminability was higher: both theories predict that when salient features of an object are obscured (e.g., through foreshortening) recognition will be impaired. In Experiment 6, after familiarisation with the discrimination recognition was high in the foreshortened view prior to the introduction of shadows. Both types of theory can account for this difference between Experiments 5 and 6 through the observer learning to attend to the more subtle differences between the two objects (irrespective of whether the different object shapes are represented in a 2-D or a 3-D fashion).

However, in Experiment 6, for most of the participants, the additional cues added by the shadows reduced discrimination in the more foreshortened shadow conditions (for most participants). Both types of theory could account for this by stating that after the introduction of the shadows the task was essentially a new one. The non-foreshortened shadows would have provided additional information to the memorial representation: if it was 2-D, then new views would have been added, and if it was 3-D then more shape information, that was not available from the foreshortened view, would have been incorporated. In the case of a view-based theory, if participants based each response upon the similarity of the stimulus to the two views of the objects that were the most different (i.e., the two views with the least foreshortening), then the pattern of responding observed would be predicted. In the case of a structural description theory, if the participants based each response on the presence and arrangement of the most salient features in each object's 3-D description, then the pattern of responding observed would also be predicted (because in the foreshortened views the most salient parts of each objects' 3-D description are not visible, and therefore their arrangement is also not visible). In either case, neither theory type has trouble accounting for these findings.

Experiment 8 demonstrated one simple effect that theories of object recognition need to be able to account for. That the combination of cast and attached shows can maintain fast and accurate recognition, when there are not sufficient other cues (e.g., shading) present to do so. Experiment 8 leaves no doubt that shadows alone can maintain recognition (whereas other demonstrations of shape-from-shadow have included shading in the two-toning process, e.g., Cavanagh & Leclerc, 1989). The two major theories do not provide differing predictions about whether shadow contours alone should be able to maintain recognition, but both theories need to be able to account for how this would occur. Given both theories can account for recognition based only upon an objects' silhouette, both theories can account for recognition based upon a mixture of cast and attached shadow borders: the occluded-background and occluded-surface borders present the same 2-D information as the bounding contour of a silhouette, and interior attached-shadow borders offer additional information about internal contours. Thus, for both theories, the 2-D contours highlighted by the shadow borders provide information to the visual system, that, combined with the correct brightness gradients at these borders (for interpretation of dark patches as shadows, Cavanagh & Leclerc, 1989), enables recognition to occur.

While the results across the course of the thesis indicate that shadow presence can be beneficial under certain conditions, e.g., when there are few other cues available, there is some evidence to suggest that shadows could also be detrimental to recognition. The results of Experiment 6 indicate that shadows can act as distracters from other reliable shape cues to object recognition. As discussed above, in Experiment 6 discrimination was sometimes poorer when shadows were present in comparison to when they had been absent from the task. It was suggested that this could be explained by the addition of shadows essentially set up a new task for the participants, one in which the decision criteria were changed given the new information available. In other research, Freeburg (1966) has demonstrated that when shadow borders dominate the borders visible on a surface, recognition of the surface is reduce. Similarly, Moore and Cavanagh (1998) using two-tone images, demonstrated that when shadow and shading are difficult to separate (due to the two-toning process) shape perception can be disrupted. Braje et al. (2000), in their study using photos of fruit and vegetables found that recognition of some images was enhanced by shadow presence, and for other images it was degraded. They speculated that reduced contrast at object borders may have led to some of the decreases in recognition.

The results of Experiment 6 then support other research demonstrating that shadows could be detrimental to object recognition in situations similar to those where they could provide a measurable benefit: those conditions in which other shape cues are limited. When there are few shape cues available to the recognition system, each object border or discontinuity that is visible is important. In these conditions the impact of shadow presence would be greater. Any border highlighted by shadow presence is proportionally more important when the number of borders is reduced, and the loss of any border as it is obscured by shadow is also proportionally more important.

Conclusions

To conclude, shadows are just like any other cue to recognition; their salience will be dependent upon the context of recognition and the other cues present. The research conducted as part of this thesis has demonstrated that even a low level presence of cues such as non-foreshortened bounding contour (Experiments 5 to 7), or shading (Experiment 8), will negate any contribution to recognition speed or accuracy made by shadows. This means that the presence of shadows will rarely provide a measurable effect upon recognition. We are usually functioning at a ceiling level of discrimination due to the abundance of cues present, with any one or two cues likely to be able to produce extremely fast recognition. That said, where few other cues are present, we can still use the information available from shadows to facilitate fast and reliable recognition.

References

Allen, B., P. (1999). Shadows as sources of cues for distance of shadow-casting objects. *Perceptual and Motor Skills, 89*, 571-584.

Allen, B. P. (2000). Angles of shadows as cues for judging the distance of shadow casting objects. *Perceptual and Motor Skills, 90*, 864-866.

Battelli, L., Cavanagh, P., & Thornton, I. M. (2003). Perception of biological motion in parietal patients. *Neuropsychologia, 41*, 1808-1816.

Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behaviour, 22*, 231-242.

Baum, W. M. (1979). Matching, undermatching, and overmatching in studies of choice. *Journal of the Experimental Analysis of Behaviour, 2*, 151-169.

Berbaum, K., Bever, T., & Chung, C. S. (1984). Extending the perception of shape from known to unknown shading. *Perception, 13*, 479-488.

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological review, 94*(2), 115-147.

Biederman, I. (2001) Recognizing depth-rotated objects: A review of recent research and theory. *Spatial Vision, 13*, 241-253

Biederman, I., & Bar, M. (1999). One-shot viewpoint invariance in matching novel objects. *Vision Research, 39*, 2885-2899.

Biederman, I., & Gerhardstein, P. C. (1993). Recognising depth-rotated objects: Evidence and conditions for three-dimensional viewpoint invariance. *Journal of Experimental Psychology: Human Perception and Performance, 19*(6), 1162-1182.

Biederman, I., Subramaniam, S., Bar, M., Kalocsai P., & Fiser J. (1999). Subordinate-level object classification reexamined. *Psychological Research, 62*, 131-153

Braje, W. L., Kersten, D., Tarr, M., & Troje, N. F. (1996). Illumination and shadows influence face recognition. *Investigative Opthalmology & Visual Science, 37*(3), S176.

Braje, W. L., Kersten, D., Tarr, M., & Troje, N. F. (1998). Illumination effects in face recognition. *Psychobiology, 26*(4), 371-380.

Braje, W. L., Legge, G. E., & Kersten, D. (2000). Invariant recognition of natural objects in the presence of shadows. *Perception, 29*, 383-398.

Bülthoff, H. H., Edelman, S., & Tarr, M. (1995). How are three-dimensional objects represented in the brain? *Cerebral Cortex, 3*, 247-260.

Bülthoff, I., Kersten, D., & Bülthoff, H. H. (1994). General lighting can overcome accidental viewing. *Investigative Opthalmology & Visual Science. 35*, 1741 #2257

Cameron, P. A., & Gallup, G. J. (1988). Shadow recognition in human infants. *Infant Behavior and Development, 11*, 465-471.

Castiello, U. (2001). Implicit processing of shadows. *Vision Research, 41*, 2305-2309.

Castiello, U., Lusher, D., Burton, C., & Disler, P. (2003). Shadows in the Brain. *Journal of Cognitive Neuroscience, 15*(6), 862-872.

Cavanagh, P. (1991). What's up in top down processing. In A. Gorea (Ed.), *Representations of vision: trends and tacit assumptions in vision research.* (pp. 295-304). Cambridge: Cambridge University Press.

Cavanagh, P. (1995). A horse of a different colour: shadows have to be darker, but shading does not. *Investigative Opthalmology and Visual Science, 36*(4), S184.

Cavanagh, P., & Kennedy, J. M. (2000). Close encounters: Details veto depth from shadows. *Science* (Letters), 287, 2423-2424

Cavanagh, P., & Leclerc, Y. G. (1989). Shape from shadows. *Journal of Experimental Psychology: Human Perception and Performance, 15*(1), 3-27.

Davidson, E. J., & Wright, P. (2002). Selective processing of weight- and shape-related words in bulimia nervosa: Use of a computerised Stroop test. *Eating Behaviors, 3*, 261-273.

Davison, M. C., & McCarthy, D. C. (1988). *The Matching Law: A research review.* Hillsdale, NJ: Erlbaum.

Davison, M. C., & Tustin, R. D. (1978). The relation between the generalised matching law and signal detection theory. *Journal of the Experimental Analysis of Behaviour, 29*, 331-336.

Day, R. H. (1989). Apparent depth from progressive exposure of moving shadows: The kinetic depth effect in a narrow aperture. *Bulletin of the Psychonomic Society, 27*(4), 320-322.

DeCaro, S., & Reeves, A. (2002). The use of word-picture verification to study entry level object recognition: Further support for view-invariant mechanisms. *Memory and Cognition, 30*(5), 811-821.

Delorme, A., Richard, G., & Fabre-Thorpe, M. (1999). Rapid processing of complex natural scenes: a role of the magnocellular visual pathways? *Neurocomputing, 26-27*, 663-670.

Delorme, A., Richard, G., & Fabre-Thorpe, M. (2000). Ultra-rapid categorisation of natural scenes does not rely on colour cues: A study in monkeys and humans. *Vision Research, 40*(15), 2187-2200.

Edelman, S. (1998). Representation is representation of similarities. *Behavioural and Brain Sciences., 21*, 449-498.

Erens, R. G. F., Kappers, A. M. L., & Koenderink, J. J. (1993) Perception of local shape from shading. *Perception and Psychophysics, 54* (2), 145-156

Etheredge, R. J. M. (1997). *Delayed matching to sample: the effects of delaying reinforcement at a single post-sample stimulus delay.* Unpublished Master of Social Science, The University of Waikato, Hamilton, NZ.

Freeburg, N., E. (1966). Shadow effects in recognition of complex textured surfaces. *Perceptual and Motor Skills, 22*, 251-256.

Gauthier, I., & Tarr, M. J. (1997). Becoming a "Greeble" expert: Exploring mechanisms for face recognition. *Vision Research, 37*(12), 1673-1682.

Gauthier, I., & Tarr, M. J. (2002). Unravelling mechanisms for expert object recognition: Bridging brain activity and behavior. *Journal of Experimental Psychology: Human Perception and Performance, 28*(2), 431-446.

Gauthier, I., Williams, P., Tarr, M. J., & Tanaka, J. (1998). Training 'greeble' experts: a framework for studying expert object recognition processes. *Vision Research, 38*, 2401-2428.

Gegenfurtner, K. R., & Rieger, J. (2000). Sensory and cognitive contributions of color to the recognition of natural scenes. *Current Biology, 10*, 805-808.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics.* New York: Wiley.

Hautus, M. J. (1995). Corrections for extreme proportions and their biasing effects on estimated values of *d'*. *Behavior Research Methods, Instruments, & Computers, 27*(1), 46-51.

Hayward, W. G. (1998). Effects of outline shape in object recognition. *Journal of Experimental Psychology: Human Perception and Performance, 24*(2), 427-440.

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

Hecht, E. (1998). *Optics*. (3rd ed.). Reading: MA: Addison-Wesley.

Hering, E., & (translated by Hurvich, L. M. a. J., Dorothea. (1874/1964). *Outlines of a Theory of the Light Sense*. Cambridge, Massachusetts: Harvard University Press.

Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behaviour, 4*, 267-272.

Herrnstein, R. J. (1990). Levels of stimulus control: A functional approach. *Cognition, 37*, 133-166.

Hummel, J. E. (2000). Where view-based theories break down: The role of structure in shape perception and object recognition. In E. Dietrich & A. Markman (Eds.), *Cognitive Dynamics; Conceptual Change in Humans and Machines* (pp. 157-185). Hillsdale, NJ: Erlbaum.

Johnstone, V. & Alsop, B. (1996). Human signal-detection performance: effects of signal presentation probabilities and reinforcer distributions. *Journal of the Experimental Analysis of Behavior, 66*, 243-263.

Kayaert, G., Biederman, I., & Vogels, R. (2003). Shape tuning in macaque inferior temporal cortex, *The Journal of Neuroscience, 23*(7), 3016-3027

Kanwisher, N., Tong, F., & Nakayama, K. (1998). The effect of face inversion on the human fusiform face area. *Cognition, 68*, B1-B11.

Kennedy, J. M., & Bai, J. (2000). Cavanagh and Leclerc shape-from-shadow pictures: Do line versions fail because of the polarity of the regions or the contour? *Perception, 29*, 399-407.

Kersten, D., Knill, D. C., Mamassian, P., & Bülthoff, I. (1996). Illusory motion from shadows. *Nature, 379*(4), 31.

Kersten, D., Mamassian, P., & Knill, D. C. (1997). Moving cast shadows induce apparent motion in depth. *Perception, 26*, 171-192.

Knill, D., C., Mamassian, P., & Kersten, D. (1997). Geometry of shadows. *Journal of the Optical Society of America, 14*(12), 3216-3232.

Langer, M., S. (1999). When shadows become interreflections. *International Journal of Computer Vision, 34*(2/3), 193-204.

Liu, C. H., Collin, C., A, & Chaudhuri, A. (2000). Does face recognition rely on encoding of 3-D surface? Examining the role of shape-from-shading and shape-from-stereo. *Perception, 29*, 729-743.

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*(4), 489-498.

Logothetis, N. K., & Sheinberg, D. L. (1996). Visual object recognition. *Annual Review Neuroscience, 19*, 577-621.

Logvinenko, A., & Menshikova, G. (1994). Trade-off between achromatic colour and perceived illumination as revealed by the use of pseudoscopic inversion of apparent depth. *Perception, 23*, 1007-1023.

Lucas, B. A., & Taylor, J. S. (1979). Kinetic depth effect examined under conditions of unimodal versus bimodal sensory input. *Perceptual and Motor Skills, 48*, 1315-1349.

MacLeod, R. B. (1940). Brightness-constancy in unrecognized shadows. *Journal of Experimental Psychology., 27*(1), 1-23.

Madison, C., Thompson, W., Kersten, D., Shirley, P., & Smits, B. (2001). Use of interreflection and shadow for surface contact. *Perception and Psychophysics, 63*, 187-194.

McCarthy, D. C., & Davison, M. C. (1979). Signal probability, reinforcement and signal detection. *Journal of the Experimental Analysis of Behaviour, 32*, 373-386.

McCarthy, D. C., & Davison, M. C. (1980a). On the discriminability of the sample stimulus. *Journal of the Experimental Analysis of Behaviour, 33*, 187-211.

McCarthy, D. C., & Davison, M. C. (1980b). Independence of sensitivity to relative reinforcement rate and discriminability in signal detection. *Journal of the Experimental Analysis of Behaviour, 345*, 273-386.

McCarthy, D. C., & Davison, M. C. (1984). Isobias and alloiobias functions in animal psychophysics. *Journal of Experimental Psychology: Animal BEhavior Processes, 10*(3), 390-409.

Mingolla, E., & Todd, J. T. (1986). Perception of Solid Shape from Shading. *Biological Cybernetics, 53*, 137-151.

Moore, C., & Cavanagh, P. (1998). Recovery of 3D volume from 2-tone images of novel objects. *Cognition, 67*, 45-71.

Moore, C., & Engel, S. A. (2001). Mental models change rapidly with implicitly acquired information about the local environment: A two-tone image study. *Journal of Experimental Psychology: Human Perception and Performance, 27*(5), 1211-1228 URLJ: http://www apa org/journals/xhp html.

Moses, Y., Adini, Y., & Ullman, S. (1994). Face Recognition: the Problem of Compensating for Changes in Illumination Direction. *Third European*

Conference on Computer Vision (Stockholm, Sweden, May 2-6, 1994), 1, 286-296.

Nagai, J., & Yokosawa, K. (2003). What regulates the surface colour effect in object recognition: Colour diagnosticity or category? *Technical Report on Attention and Cognition, 28*, 1-4.

Naor-Raz, G., Tarr, M., & Kersten, D. (2003). Is colour an intrinsic property of object representation? *Perception, 32*(667-680).

Norman, J. F., Dawson, T. E., & Raines, S. R. (2000). The perception and recognition of natural object shape from deforming and static shadows. *Perception, 29*, 135-148.

Norman, J. F., Phillips, F., & Ross, H. E. (2001). Information concentration along the boundary contours of naturally shaped solid objects. *Perception, 30*, 1285-1294.

Norman, J. F., & Todd, J. T. (1994). Perception of rigid motion in depth from the optical deformations of shadows and occlusion boundaries. *Journal of Experimental Psychology: Human Perception and Performance., 20*(2), 343-356.

Peissig, J. J., Wasserman, E. A., Young, M. E., & Biederman, I. (2002) Learning an object from multiple views enhances its recognition in an orthogonal rotational axis in pigeons. *Vision Research, 42*, 2051-2062

Perrett, D. I., Oram, M. W., & Ashbridge, E. (1998). Evidence accumulation in cell populations responsive to faces: an account of generalisation without mental transformations. *Cognition, 67*, 111-145.

Poling, A., & Foster, T. M. (1993). The matching law and organisational behaviour management revisited. *Journal of Organisation Behaviour Management, 14*, 83-97.

Price, T. J., O'Toole, A. J., & Dambach, K. C. (1998). A moving cast shadow diminishes the Pulfrich phenomenon. *Perception, 27*, 591-593.

Puerta, A. M. (1989). The power of shadows: shadow stereopsis. *Journal of the Optical Society of America A., 6*(2), 309-311.

Ramachandran, V. S. (1988). Perception of shape from shading. *Nature, 331*(14), 163-166.

Rensink, R. A., & Cavanagh, P. (1993). Processing of shadows at preattentive levels. *Investigative Opthalmology and Visual Science, 34*, 1288.

Rock, I., Halper, F., & Clayton, T. (1972). The perception and recognition of complex figures. *Cognitive Psychology, 3*, 655-673.

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.

Schyns, P. G., & Oliva, A. (1999). Dr. Angry and Mr. Smile: when categorization flexibly modifies the perception of faces in rapid visual presentations. *Cognition, 69*, 243-265.

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*(4), 913-932

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*(3), 732-744.

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology., 29*, 643-662.

Stubbs, D. A. (1976). Response bias and the discrimination of stimulus duration. *Journal of the Experimental Analysis of Behavior, 25*, 243-250.

Tarr, M. J. (2003). Visual Object Recognition: Can a Single Mechanism Suffice? In M. A. Peterson & G. Rhodes (Eds.), *Perception of Faces, Objects, and Scenes: Analytic and Holistic Processes.* (pp. 177-207). London: Oxford University Press.

Tarr, M. J. & Bülthoff, H. H. (1995). Is human object recognition better described by geon structural descriptions or by multiple views? Comment on Biederman and Gerhardstein (1993). *Journal of Experimental Psychology: Human Perception and Performance, 21*(6), 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*(4), 282-289.

Tarr, M. J., Kersten, D., & Bülthoff, H. H. (1998). Why the visual system might encode the effects of illumination. *Vision Research, 38*, 2259-2275.

Ullman, S. (1995). The visual analysis of shape and form. In M. S. Gazzaniga (Ed.), *The Cognitive Neurosciences*. Cambridge, Massachusetts: The MIT Press.

Vecera, S. P. (1998). Visual object representation: An introduction. *Psychobiology, 26*(4), 281-308.

Vogels, R., & Biederman, I. (2002). Effects of illumination intensity and direction on object coding in macaque inferior temporal cortex. *Cerebral Cortex, 12*, 756-766

Warrington, E., K., & James, M. (1986). Visual object recognition in patients with right-hemisphere lesions: axes or features. *Perception, 15*, 355-366.

Weavers, R. (1993). *DMTS: Task acquisition and reinforcer efficacy.* Unpublished Master of Social Science, The University of Waikato, Hamilton.

White, K. G., & Wixted, J. T. (1999). Psychophysics of remembering. *Journal of the Experimental Analysis of Behavior, 71*, 91-113.

Wimmer, F. L. (1994). *The suitability of simplified cast shadows as a visual depth cue for computer graphics applications.* Unpublished Doctor of Philosophy, Carleton University, Ottawa.

Wühr, P., & Waszak, F. (2003). Object-based attentional selection can modulate the Stroop effect. *Memory & Cognition, 31*(6), 983-994.

Wurm, L. H., Legge, G. E., Isenburg, L. M., & Luebker, A. (1993). Color improves object recognition in normal and low vision. *Journal of Experimental Psychology: Human Perception and Performance, 19*(4), 899-911.

Yonas, A., Farr, M., & O'Connor, A. (2001). Seven but not 5-month-old infants extract depth from cast shadows (Abstract). *Journal of Vision, 1(3), 389a, http://journalofvision.org/1/3/389/, doi:10.1167/1.3.389., 1*(3), 389a.

Yonas, A., Goldsmith, L. T., & Hallstrom, J. L. (1978). Development of sensitivity to information provided by cast shadows in pictures. *Perception, 7*, 333-341.

Zayan, R., & Vauclair, J. (1998). Categories as paradigms for comparative cognition. *Behavioural Processes, 42*, 87-99.

Appendix 1 Images used in Experiments 1 to 4

List of Figures:

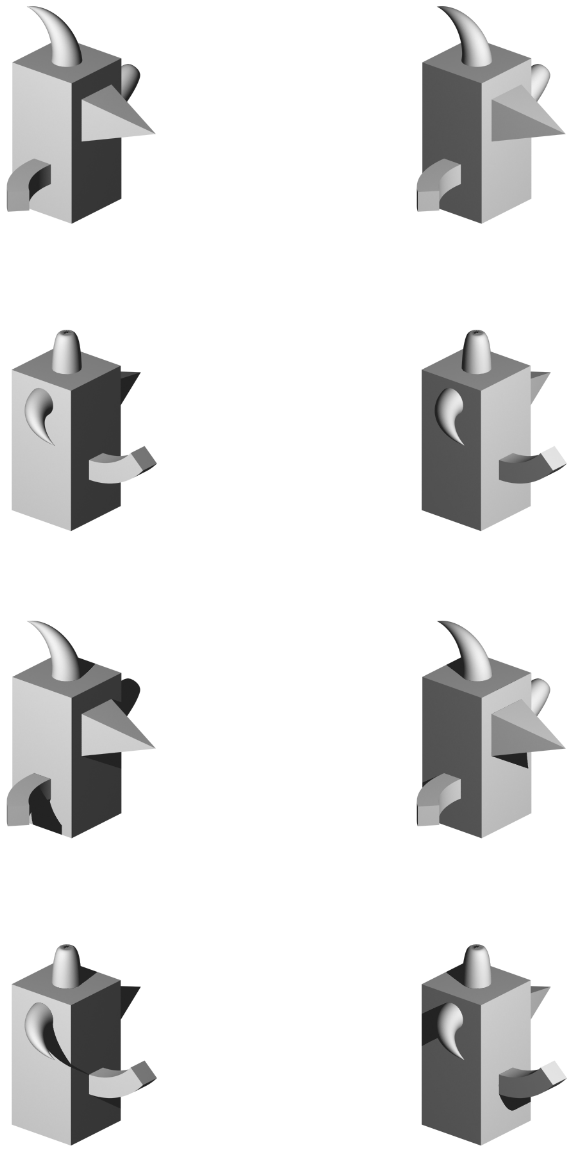

Figure A1.1. Object Base 1: No-Shadow Condition first and second rows, Shadow Condition third and fourth rows. Version 1 first and third rows, Version 2 second and fourth rows. Illumination from left in left column, illumination from right in right column.

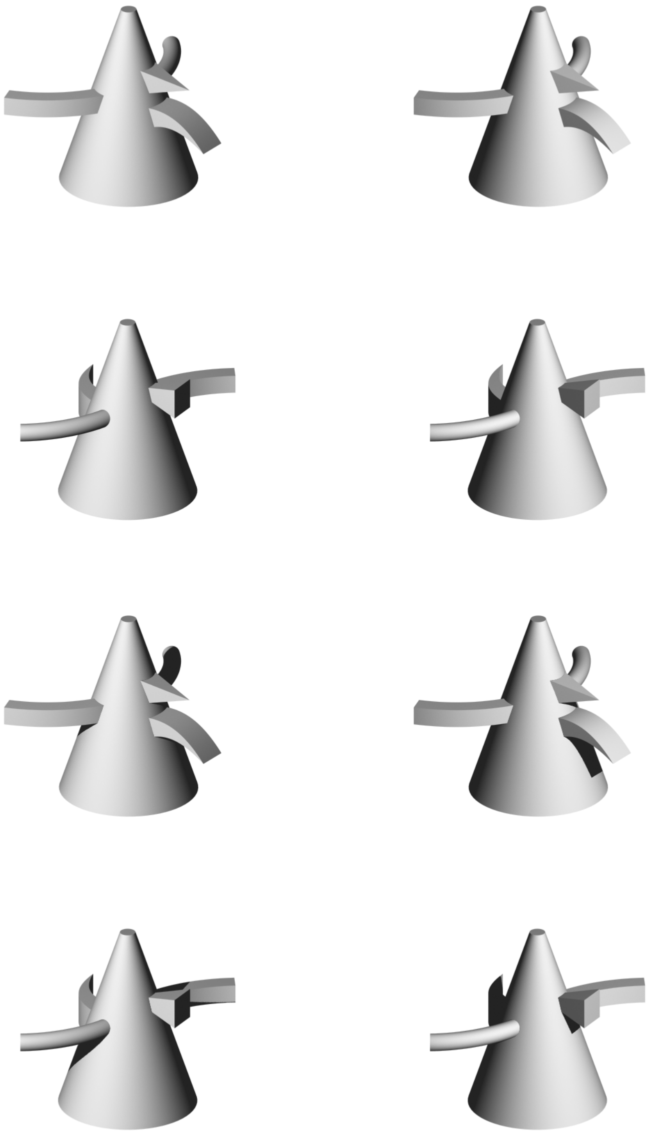

Figure A1.2. Object Base 2: No-Shadow Condition first and second rows, Shadow Condition third and fourth row Version 1 first and third rows, Version 2 second and fourth rows. Illumination from left in left column, illuminati from right in right column.

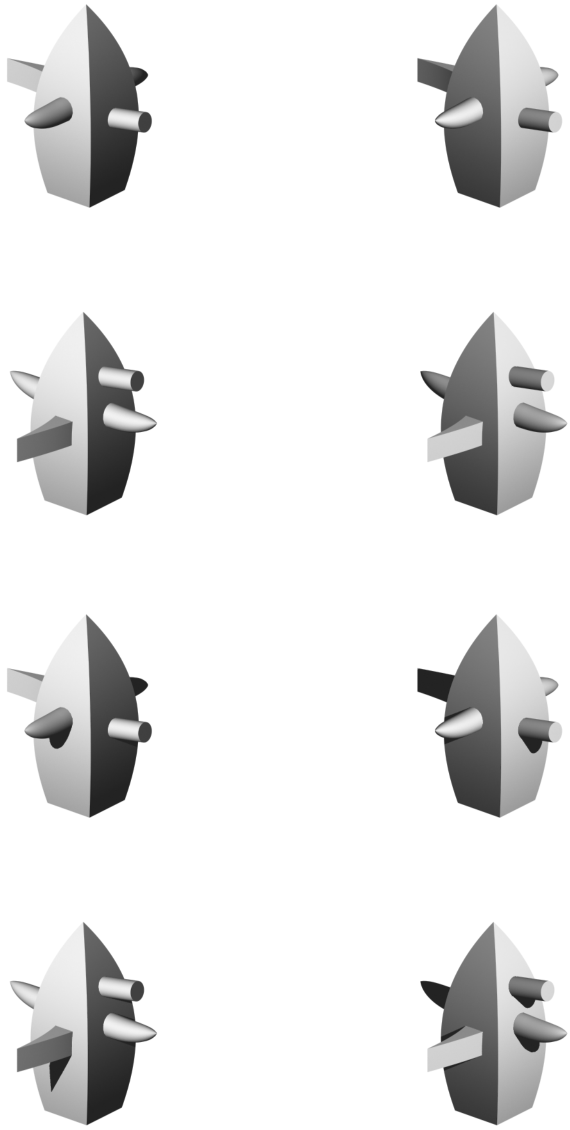

Figure A1.3. Object Base 3: No-Shadow Condition first and second rows, Shadow Condition third and fourth rows. Version 1 first and third rows, Version 2 second and fourth rows. Illumination from left in left column, illumination from right in right column.

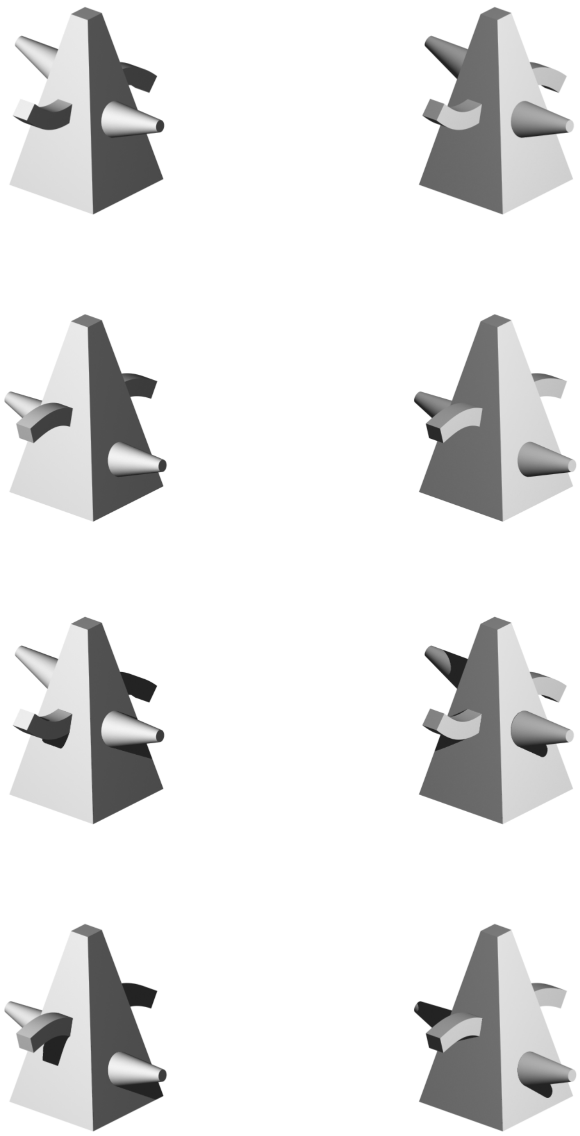

Figure A1.4. Object Base 4: No-Shadow Condition first and second rows, Shadow Condition third and fourth rows. Version 1 first and third rows, Version 2 second and fourth rows. Illumination from left in left column, illumination from right in right column.

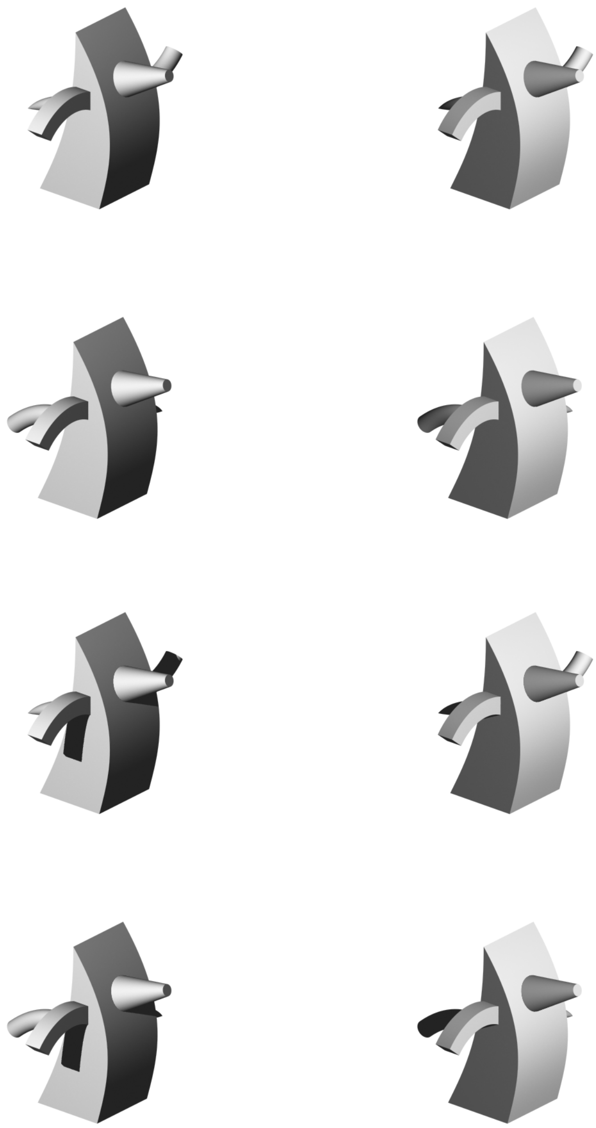

Figure A1.5. Object Base 5: No-Shadow Condition first and second rows, Shadow Condition third and fourth rows. Version 1 first and third rows, Version 2 second and fourth rows. Illumination from left in left column, illumination from right in right column.

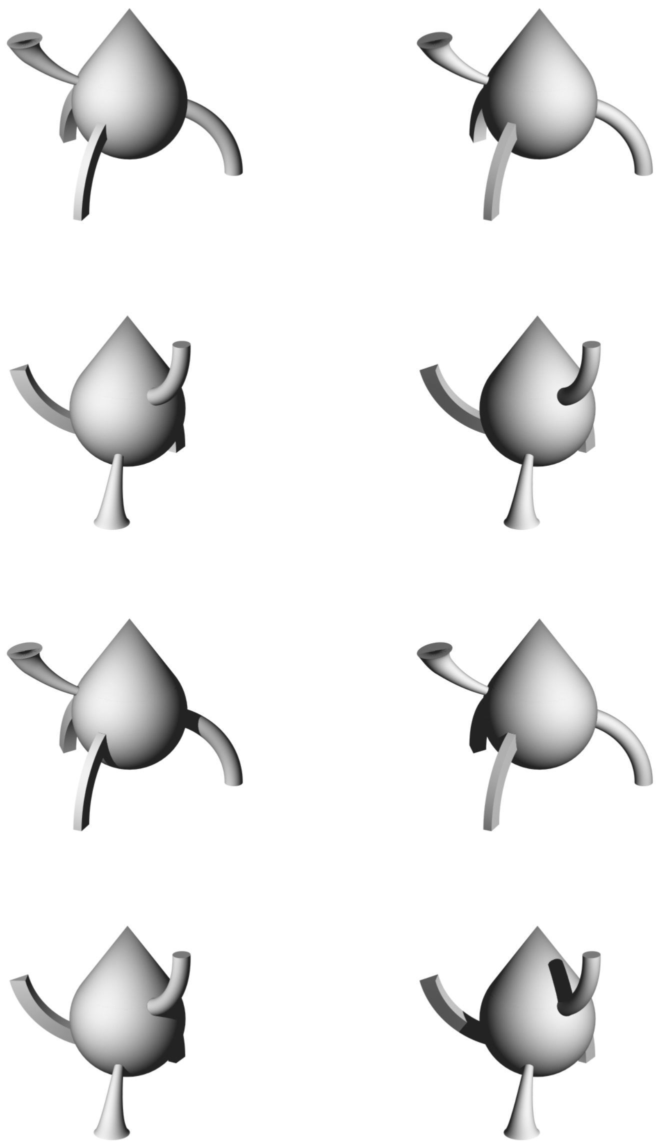

Figure A1.6. Object Base 6: No-Shadow Condition first and second rows, Shadow Condition third and fourth rows. Version 1 first and third rows, Version 2 second and fourth rows. Illumination from left in left column, illumination from right in right column.

Appendix 2 Instructions for Experiments 1 to 4

At the start of the experiments these instructions were presented to the participants over four sequential windows on the computer monitor.

Window 1.

"On each trial in this experiment you will see a crosshair and then four images will be presented sequentially to you."

Window 2.

"The second and fourth images are masking stimuli, you do not need to attend to them. Your task is to determine if the first and third images presented were of the same object."

Window 3.

"Press "Q" for Yes the two images were the same, or press "P" for No they were different. Your reaction time will be given after each trial. Please try to respond as quickly as possibly"

Window 4.

"Pressing "spacebar" once will "o.k." the reaction time and pressing it again will start a new trial. Feel free to have a break between trials if you would like one. Thanks, Richard."

Appendix 3 Measuring Discrimination and Bias

Appendix 3 is a discussion of the use of log *d* and log *c*, in matching-to-sample situations, adapted from Etheredge (1997).

Percent correct is a measure of accuracy that is widely used and understood. An alternative measure is the measure of discrimination log *d*, the major benefit of log *d* over percent correct, is that it is designed to provide a bias free measure of discrimination. It involves viewing a matching-to-sample procedure in a similar manner to a signal detection procedure. In a signal detection experiment the stimulus can be either present or absent and the subject can report either it is present or it is absent. This results in four alternative outcomes as displayed diagrammatically in Figure A3.1.

The four outcomes are termed: a hit, when the stimulus is present and it is reported as such (yes); a false rejection, when the stimulus is present but is reported as absent (no); a false alarm, when the stimulus is absent but is reported as present (yes); and a correct rejection, when the stimulus is absent and is reported as such (no).

Davison and Tustin (1978) suggested viewing signal detection procedures as two separate stimulus situations. In one situation there is a stimulus present. In the other situation there is no stimulus present. Matching to sample (MTS) procedures can be viewed in a similar manner, as shown in Figure A3.2.

Davison and Tustin (1978) pointed out that when viewed this way, it can be seen that under a MTS procedure there are two concurrently available behaviours associated with each stimulus situation. Thus the two stimulus situations could be viewed as two separate concurrent schedules. Under concurrent schedules there are two or more schedules of reinforcement available and the subject can respond on either. In effect under a matching to sample task each stimulus situation has a separate pair of concurrent schedules, with reinforcement for correct and no reinforcement for incorrect responses. In human object-recognition tests, there is often no defined reinforcer given, although sometimes a feedback beep indicating a correct response is given (e.g., Tarr et al. 1998).

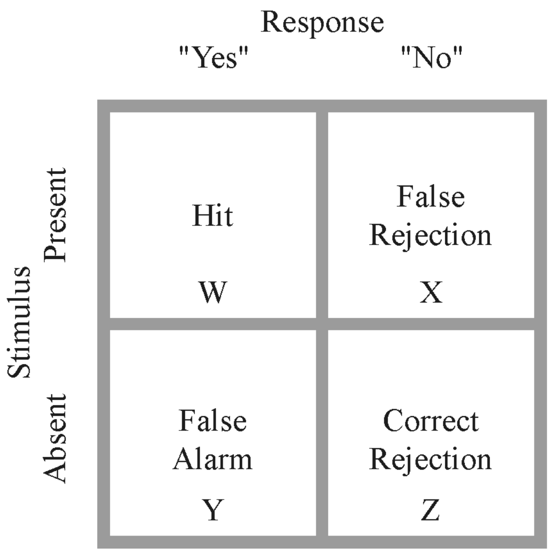

Figure A3.1. The signal detection matrix.

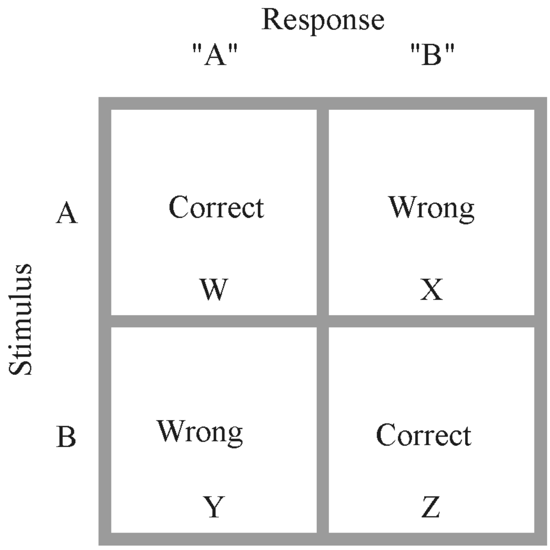

Figure A3.2. The matching-to-sample matrix.

Behaviour under concurrent schedules has been found to be orderly. Baum (1974) reported that concurrent schedule behaviour is generally described by an equation of the form:

(1)

$$\log (B1/B2) = a \log (R1/R2) + \log c$$

where B1 and B2 are the frequencies of responding to the two alternatives, R1 and R2 are the number of reinforcers obtained from alternatives 1 and 2, a is sensitivity to reinforcement and log c is bias in responding to one of the two alternatives. Equation 1 is known as the generalised matching law.

Prior to this Herrnstein (1961) had suggested that behaviour under concurrent schedules conformed to strict matching. He suggested that under variable schedules the proportion of responses to each alternative equals the proportion of reinforcers obtained under that alternative. This strict matching is mathematically described by Herrnstein's (1961) matching law (Equation 2).

(2)

$$B1 / (B1 + B2) = R1 / (R1 + R2)$$

Where B1, B2, R1, and R2 are as given for Equation 1. When log $c = 0$ and $a = 1$ in Equation 1, Equations 1 and 2 are equivalent.

The generalised matching law (Equation 1) extends the matching law (Equation 2) to include occasions where the allocation of behaviour between schedules does not strictly match the allocation of reinforcement between them. Deviations from strict matching are frequent (Baum, 1974). Baum (1974, 1979) discusses three systematic deviations from strict matching that the generalised matching law takes into account. They are undermatching, overmatching and bias.

Sensitivity to reinforcement is, in the logarithmic form of the generalised matching law (Equation 1), represented by the slope of the line, a. Undermatching, where a is less than 1.00, results when relatively more behaviour than would be predicted by the matching law, is allocated to the alternative providing less reinforcement. Thus, the behaviour is said to be less sensitive to reinforcement rate changes than would be predicted by strict matching. Overmatching, where *a* (sensitivity to reinforcement) is greater than 1.00, results when relatively more behaviour than would be predicted by the matching law, is allocated to the alternative providing more reinforcement. In this case behaviour is said to be more sensitive to reinforcement rate changes than would be predicted by strict matching.

Bias is represented by the intercept of the logarithmic equation of the generalised matching law, log *c*. When consistently more or less behaviour is allocated to one alternative than would be predicted by the matching law, irrespective of whether or not that alternative yields greater or lesser reinforcement than the other alternative, the behaviour is said to be biased towards or away from that alternative and log *c* gives a measure of this. Thus, when all other things are equal between the schedules except the rate of reinforcement, log *c* gives a measure of the inherent bias of the subject, and/or uncontrolled aspects of the manipulanda. Inherent bias is differentiated from bias resulting from other sources. Davison and McCarthy (1988) point out there may be a deliberately produced constant bias resulting from differences between reinforcers, which add to log *c* (inherent bias), and bias due to differences in reinforcer frequency between schedules (relative reinforcer-frequency bias) (Davison & McCarthy, 1988). Together, inherent bias and reinforcer bias comprise response bias (Davison & McCarthy, 1988).

Evidence to date on the generalised matching law suggests that it provides a good description of a large body of data collected under concurrent VI VI schedules of reinforcement (Poling & Foster, 1993). Davison and Tustin (1978) showed how the generalised matching law could be applied to MTS data when the procedure is seen as two concurrent schedules. They point out that in a typical signal detection experiment it is usual for every correct response to be reinforced, with a lack of reinforcement for errors, and that this means that W and Z (Figure A3.1) represent the number of reinforcements as well as the number of responses. Thus, as given by Davison and Tustin (1978), if the response ratios were put into the generalised matching law equation, while ignoring bias due to the stimulus' presence or absence, the equation obtained would be:

(3)

$\log (W+Y / X+Z) = a \log (W / Z) + \log c$

If the presence or absence of the stimulus is not discriminated then the same bias will be shown when the stimulus is present and when it is absent. However, it follows that when the subject does discriminate between the presence and absence of the stimulus, the subject will be biased toward saying "yes" when the stimulus is present and "no" when it is absent. Thus bias will consist of any inherent bias, and also any bias resulting from the presence or absence of the stimulus. Davison and Tustin (1978) give two equations that account for this:

(4)

$\log (BW/BX) = ar1 \log (RW/RZ) + \log c + \log d$

(5)

$\log (BY/BZ) = ar2 \log (RW/RZ) + \log c - \log d$

where BW, BX, BY and BZ are the frequencies of responding with respect to the four cells of the signal detection matrix, RW and RZ are the number of reinforcers obtained, taken from cells W and Z of the signal detection matrix, *ar*1 and *ar*2 are the sensitivity to reinforcement values of Equations 4 and 5 respectively, and log *d* is bias caused by the presence or absence of the stimulus[1]. Thus, Equation 4 includes bias toward saying "yes" due to the presence of the stimulus and Equation 5 includes bias toward saying "no" due to the absence of the stimulus. Davison and Tustin (1978) point out that log *d* gives a measure of discriminability of the stimulus. When there is no discrimination between presence and absence of the stimulus log *d* is zero, and Equation 3 applies. When there is discrimination log *d* is non-zero and this can be accounted for by Equations 4 and 5.

Davison and McCarthy (1988) relate that, if *ar*1 = *ar*2, Equation 5 can be subtracted from Equation 4. The assumption that *ar*1 = *ar*2 has been

1. Note that when there is no defined reinforcer provided for correct responding, as is the case in most human research [*ar1* log (RW/RZ)] and [*ar2* log (RW/RZ)] will each equal zero.

shown to hold by Davison and Tustin (1978) in an analysis of data reported by Green and Swets (1966) and Stubbs (1976) and also by McCarthy and Davison (1979, 1980a)[1]. By subtracting Equation 5 from Equation 4 the effects of reinforcer-frequency bias and inherent bias (log *c*) are removed (Davison & McCarthy, 1988). Rearranging gives a point estimate of discriminability that is free of inherent bias and reinforcer effects:

(6)

$$\log d = 0.5 \log (BW.BZ/BX.BY)$$

The assumption that discriminability is independent of response bias (reinforcer bias and inherent bias) was made by Davison and Tustin (1978) and subsequently tested by McCarthy and Davison (1980b). They concluded that there was no interaction between discriminability and response bias, and that they are additive logarithmic quantities. McCarthy and Davison (1980b) also showed that discriminability can be varied widely, in a signal detection procedure, without affecting sensitivity to reinforcement.

It is also possible to provide point estimates of inherent bias (log *c*). By adding Equations 4 and 5, log *d* is removed leaving a bias function (Davison & McCarthy, 1988), Equation 7.

(7)

$$\log (BW/BX) + \log (BY/BZ) = 2\ ar \log (RW/RZ) + 2 \log c$$

Rearranging gives:

(8)

$$0.5 \log (BW.BY/BX.BZ) = ar \log (RW/RZ) + \log c$$

The measure on the left side of the equation is response bias, termed log *b*, and it is comprised of inherent bias and relative reinforcer-frequency bias (Davison & McCarthy, 1988). When relative reinforcer-frequency variation is one (or when reinforcers are not used), its log is zero, and inherent

1. If no reinforcers are given, then sensitivity to reinforcement will equal zero, and *ar1* and *ar2* will both equal zero.

bias equals response bias. From this, a point estimate of inherent bias can be obtained.

(9)

$$\log c = 0.5 \log (BW.BY/BX.BZ)$$

Using this equation, inherent bias can be calculated for an individual stimulus. To do this the cells, from the matrix in Figure A3.2, must be arranged in the form, responses to the stimulus for which bias is being measured, over responses to the other stimulus/position (Weavers, 1993).

Thus, it is possible to derive, using the generalised matching law and signal detection methodology, both a bias free measure of a subject's performance in discriminating between two stimuli (or presence/absence of a stimulus), log *d*, and an estimate of inherent bias, log *c*. The values obtained for both log *d* and log *c* can range from negative infinity to positive infinity. A large value indicates either high accuracy, or extreme bias, respectively. A value of positive infinity occurs for log *d* when there are no errors in one, or both, stimulus situations, i.e., when either BX or BY is zero, or both are zero. A log *d* value of zero indicates a lack of discrimination between the sample stimuli. In a matching-to-sample procedure this means that the participant is responding at a level equivalent to chance. For log *c* a value of infinity occurs when all responses are allocated to one response alternative, while a value of zero indicates that there is no bias towards one or other of the stimuli.

To avoid problems calculating log *d* and log *c* due to infinite values, a correction is often used (Hautus, 1995), where a value of 0.5 is added to every cell in the matching to sample matrix. This correction value produces an artificial ceiling in the measures dependent upon the number of trials used in their calculation. That is, with a number of trials close to infinity, the maximum value of either measure will be close to infinity, but with a smaller number of trials, the maximum value obtainable will be restricted to a smaller value.

Appendix 4 Response latencies from Experiments 1 to 4.

List of tables:

Table A4.1.

Average Reaction Times For Experiments 1 to 2, Analysed With Within-Subjects ANOVAs on Illumination Change and Version Change Within a Trial, With Experimental Group, No-Shadows Versus Shadows, as a Between-Subjects Factor.

	Expt. 1, All Data	Expt. 1, Corrects	Expt. 1, Incorrects	Expt. 2, All Data	Expt. 2, Corrects	Expt. 2, Incorrects
Group means						
No Shadows	637 ms	653 ms	668 ms	704 ms	690 ms	761 ms
Shadows	636 ms	645 ms	670 ms	752 ms	763 ms	789 ms
F Expt. 1 (1,21), Expt. 2 (1,22)	0.001	0.034	0.002	0.582	1.295	0.123
η_p^2	0.000	0.002	0.000	0.026	0.056	0.006
Illumination change means						
Same illumination direction	632 ms	649 ms	661 ms	717 ms	725 ms	766 ms
Different illumination direction	641 ms	649	677 ms	739 ms	729 ms	785 ms
F Expt. 1 (1,21), Expt. 2 (1,22)	1.682	0.000	1.155	7.652*	0.185	0.764
η_p^2	0.074	0.000	0.052	0.258	0.008	0.037
Group*Illumination change means						
No Shadows, same illumination direction	636 ms	655 ms	661 ms	700 ms	687 ms	759 ms
Shadows, same illumination direction	628 ms	643 ms	661 ms	734 ms	762 ms	7649 ms
No Shadows, different illumination direction	638 ms	650 ms	674 ms	707 ms	694 ms	7729 ms
Shadows, different illumination direction	644 ms	647 ms	679 ms	767 ms	765 ms	8069 ms
F Expt. 1 (1,21), Expt. 2 (1,22)	1.038	0.447	0.027	3.474	0.035	0.398
η_p^2	0.047	0.021	0.001	0.136	0.002	0.019
Object change means						
Same objects in a trial	620 ms	610 ms	696 ms	705 ms	692 ms	803 ms
Different objects in a trial	654 ms	687 ms	641 ms	751 ms	762 ms	748 ms
F Expt. 1 (1,21), Expt. 2 (1,22)	8.043*	14.351*	6.263*	10.541*	18.307*	9.305*
η_p^2	0.277	0.406	0.230	0.324	0.454	0.318
Group*Object means						
No Shadows, same objects in a trial	619 ms	606 ms	696 ms	687 ms	671 ms	781 ms
Shadows, same objects in a trial	621 ms	615 ms	696 ms	723 ms	714 ms	824 ms
No Shadows, different objects in a trial	656 ms	699 ms	639 ms	721 ms	710 ms	741 ms
Shadows, different objects in a trial	651 ms	676 ms	644 ms	782 ms	813 ms	755 ms
F Expt. 1 (1,21), Expt. 2 (1,22)	0.113	0.611	0.016	0.762	3.283	0.645
η_p^2	0.005	0.028	0.001	0.033	0.130	0.031
Illumination*Object means						
Same illumination, same objects in a trial	614 ms	606 ms	688 ms	694 ms	692 ms	798 ms
Different illumination, same objects in a trial	625 ms	615 ms	704 ms	716 ms	693 ms	807 ms
Same illumination, different objects in a trial	650 ms	692 ms	633 ms	741 ms	758 ms	733 ms
Different illumination, different objects in a trial	657 ms	682 ms	650 ms	761 ms	765 ms	763 ms
F Expt. 1 (1,21), Expt. 2 (1,22)	0.130	1.824	0.003	0.027	0.020	0.145
η_p^2	0.006	0.080	0.000	0.001	0.001	0.007
Group*Illumination*Object means						
No Shadows, same illumination, same objects	617 ms	604 ms	691 ms	682 ms	668 ms	773 ms
Shadows, same illumination, same objects	611 ms	607 ms	686 ms	706 ms	715 ms	824 ms
No Shadows, same illumination, different objects	655 ms	706 ms	631 ms	719 ms	706 ms	745 ms
Shadows, same illumination, different objects	646 ms	678 ms	635 ms	763 ms	809 ms	721 ms
No Shadows, different illumination, same objects	619 ms	608 ms	702 ms	692 ms	673 ms	790 ms
Shadows, different illumination, same objects	631 ms	622 ms	706 ms	740 ms	714 ms	824 ms
No Shadows, different illumination, different objects	657 ms	692 ms	647 ms	722 ms	715 ms	738 ms
Shadows, different illumination, different objects	657 ms	673 ms	653 ms	800 ms	816 ms	789 ms
F Expt. 1 (1,21), Expt. 2 (1,22)	0.107	0.001	0.020	0.143	0.003	0.627
η_p^2	0.005	0.000	0.001	0.006	0.000	0.030

Note: Response times are given for all the data, and for correct responses and incorrect responses separately. Any differences that are significant at an alpha level of 0.05 are indicated by an asterix (*).

Table A4.2.

Average Reaction Times For Experiments 3 to 4, Analysed With Within-Subjects ANOVAs on Illumination Change and Version Change Within a Trial, With Experimental Group, No-Shadows Versus Shadows, as a Between-Subjects Factor.

	Expt. 3, All Data	Expt. 3, Corrects	Expt. 3, Incorrects	Expt. 4, All Data	Expt. 4, Corrects	Expt. 4, Incorrects
Group means						
No Shadows	759 ms	756 ms	758 ms	719 ms	714 ms	777 ms
Shadows	723 ms	725 ms	759 ms	758 ms	757 ms	790 ms
F Expt. 3 (1,26), Expt. 4 (1,24)	0.401	0.319	0.000	0.174	0.238	0.015
η_p^2	0.015	0.012	0.000	0.007	0.010	0.001
Illumination change means						
Same illumination direction	738 ms	740 ms	761 ms	720 ms	709 ms	793 ms
Different illumination direction	744 ms	742 ms	756 ms	757 ms	762 ms	775 ms
F Expt. 3 (1,26), Expt. 4 (1,24)	0.880	0.092	0.040	19.516*	22.997*	0.405
η_p^2	0.033	0.004	0.002	0.448	0.489	0.017
Group*Illumination change means						
No Shadows, same illumination direction	749 ms	750 ms	756 ms	699 ms	686 ms	810 ms
Shadows, same illumination direction	726 ms	730 ms	767 ms	740 ms	733 ms	776 ms
No Shadows, different illumination direction	768 ms	762 ms	760 ms	740 ms	742 ms	744 ms
Shadows, different illumination direction	720 ms	721 ms	752 ms	775 ms	782 ms	805 ms
F Expt. 3 (1,26), Expt. 4 (1,24)	2.956	2.804	0.137	0.114	0.112	2.855
η_p^2	0.102	0.097	0.006	0.005	0.005	0.106
Object change means						
Same objects in a trial	709 ms	706 ms	767 ms	732 ms	733 ms	766 ms
Different objects in a trial	773 ms	776 ms	750 ms	745 ms	738 ms	801 ms
F Expt. 3 (1,26), Expt. 4 (1,24)	23.607*	12.854*	0.275	3.831	0.687	5.201*
η_p^2	0.476	0.331	0.012	0.138	0.028	0.178
Group*Object means						
No Shadows, same objects in a trial	709 ms	704 ms	742 ms	716 ms	713 ms	782 ms
Shadows, same objects in a trial	709 ms	708 ms	792 ms	748 ms	754 ms	750 ms
No Shadows, different objects in a trial	808 ms	808 ms	774 ms	722 ms	715 ms	771 ms
Shadows, different objects in a trial	737 ms	743 ms	727 ms	767 ms	760 ms	831 ms
F Expt. 3 (1,26), Expt. 4 (1,24)	7.355*	3.076	2.229	1.074	0.221	9.066*
η_p^2	0.221	0.106	0.092	0.043	0.009	0.274
Illumination*Object means						
Same illumination, same objects in a trial	703 ms	699 ms	787 ms	707 ms	701ms	787 ms
Different illumination, same objects in a trial	715 ms	713 ms	748 ms	758 ms	766 ms	745 ms
Same illumination, different objects in a trial	772 ms	781 ms	736 ms	732 ms	718 ms	798ms
Different illumination, different objects in a trial	773 ms	770 ms	764 ms	757 ms	758 ms	804 ms
F Expt. 3 (1,26), Expt. 4 (1,24)	1.205	4.157	2.996	2.435	3.091	0.965
η_p^2	0.044	0.138	0.120	0.092	0.114	0.039
Group*Illumination*Object means						
No Shadows, same illumination, same objects	701 ms	695 ms	757 ms	693 ms	678 ms	849 ms
Shadows, same illumination, same objects	705 ms	702 ms	816 ms	721 ms	723 ms	726 ms
No Shadows, same illumination, different objects	798 ms	805 ms	755 ms	705 ms	694 ms	771 ms
Shadows, same illumination, different objects	747 ms	758 ms	717 ms	759 ms	742 ms	825 ms
No Shadows, different illumination, same objects	717 ms	713 ms	728 ms	740 ms	748 ms	716 ms
Shadows, different illumination, same objects	713 ms	713 ms	768 ms	775 ms	784 ms	774 ms
No Shadows, different illumination, different objects	818 ms	812 ms	793 ms	739 ms	736 ms	772 ms
Shadows, different illumination, different objects	728 ms	729 ms	736 ms	775 ms	779 ms	837 ms
F Expt. 3 (1,26), Expt. 4 (1,24)	2.597	1.354	0.000	0.647	0.010	2.923
η_p^2	0.091	0.050	0.000	0.026	0.000	0.109

Note: Response times are given for all the data, and for correct responses and incorrect responses separately. Any differences that are significant at an alpha level of 0.05 are indicated by an asterix (*)

Appendix 5 Instructions for Experiment 5

The instructions for Experiment 5, as given on the computer, were:

Window 1.

"On each trial in this experiment you will see a crosshair and then two images will be presented sequentially to you."

Window 2.

"The second image is a masking stimulus, you do not need to attend to it. Your task is to determine if the first image matches the picture to the left or right of your screen."

Window 3.

"Press "Q" if you think it matches the left picture, or press "P" if you think it matches the right picture. Your reaction time will be given after each trial. Please try to respond as quickly as possibly."

Window 4.

"Pressing the "spacebar" once will remove the reaction time window and pressing it again will start a new trial. Feel free to have a break between trials if you would like one. Thanks, Richard."

Appendix 6 Images used in Experiment 8

Colour images are provided in Appendix 8 (CD)

Figure A6.1. Images of the Ant view 1.

Figure A6.2. Images of the Ant view 2.

Figure A6.3. Images of the Bee view 1.

Figure A6.4. Images of the Bee view 2.

Figure A6.5. Images of the Beetle view 1.

Figure A6.6. Images of the Beetle view 2.

Figure A6.7. Images of the Can view 1.

Figure A6.8. Images of the Can view 2.

Figure A6.9. Images of the Can view 1.

Figure A6.10. Images of the Can view 2.

Figure A6.11. Images of the Vase view 1.

Figure A6.12. Images of the Vase view 2.

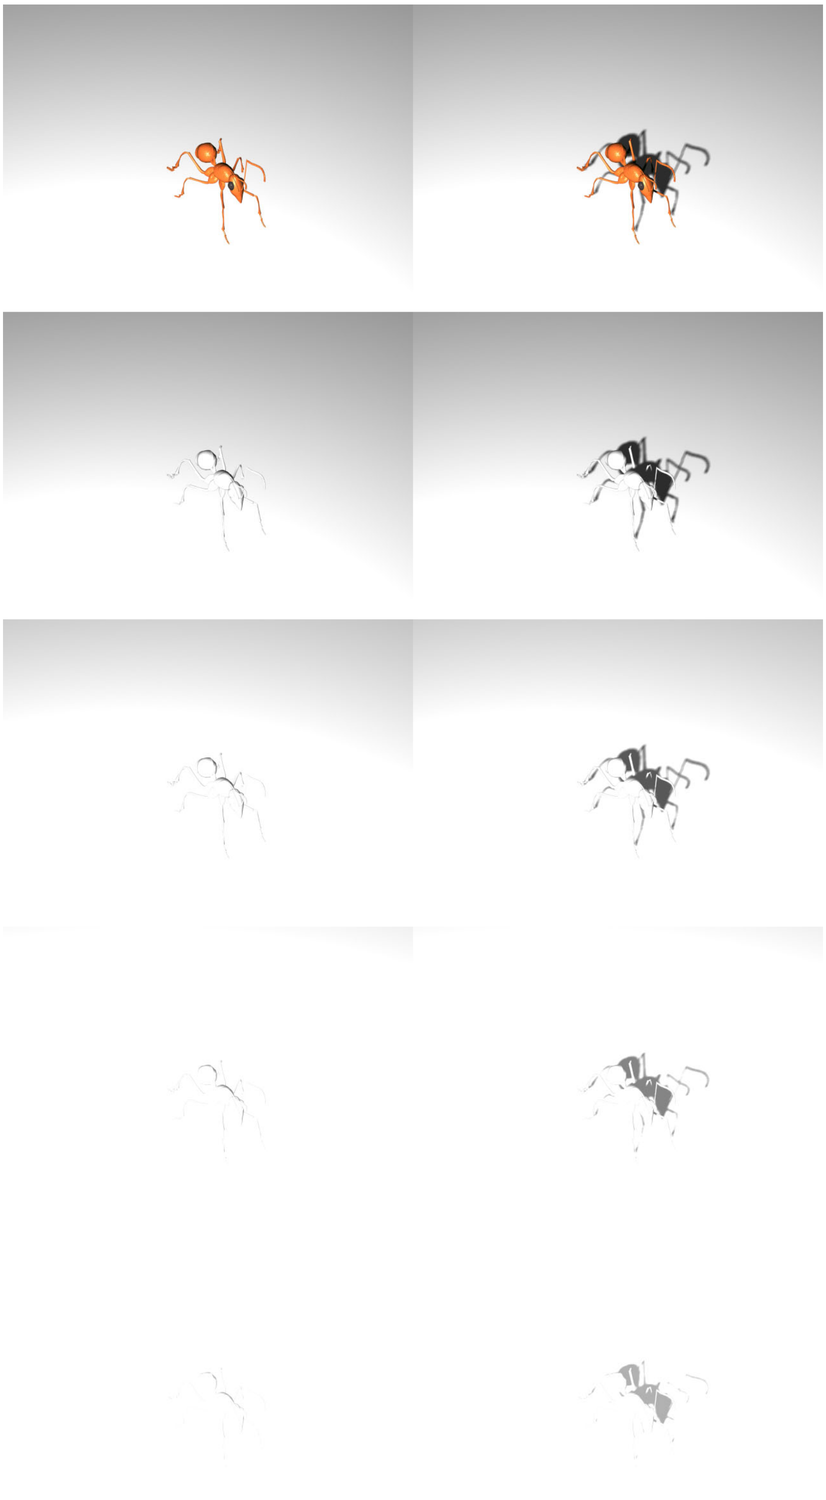

Figure A6.1. Images of the Ant view 1. The No-Shadow Condition images are on the left and the corresponding Shadow Condition images are on the right. The Illumination Conditions, from top to bottom, are: Colour Ambient 50; Ambient 50; Ambient 100; Ambient 150; and Ambient 200.

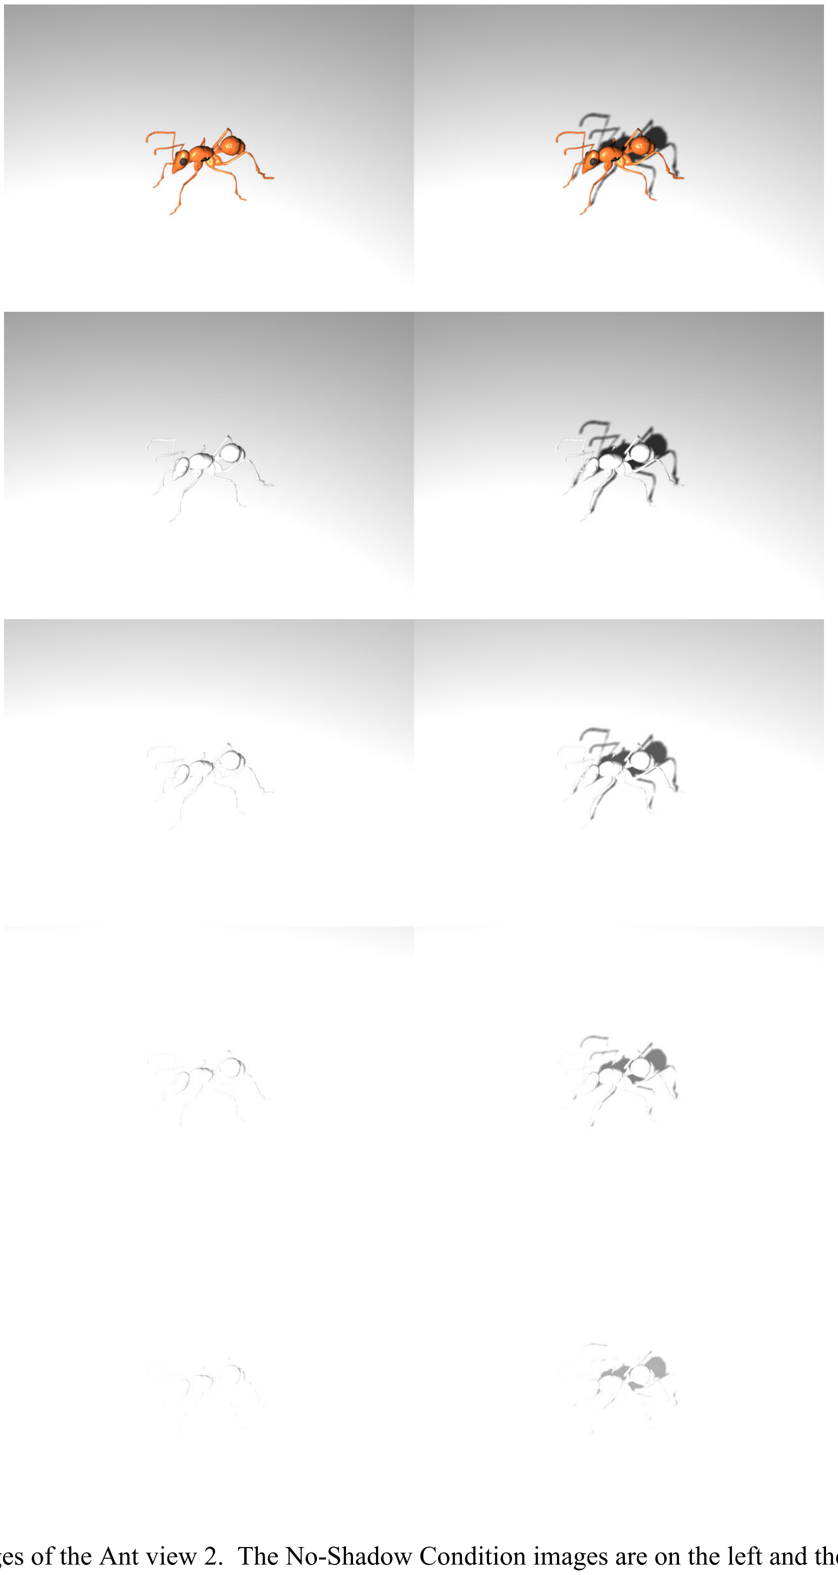

Figure A6.2. Images of the Ant view 2. The No-Shadow Condition images are on the left and the corresponding Shadow Condition images are on the right. The Illumination Conditions, from top to bottom, are: Colour Ambient 50; Ambient 50; Ambient 100; Ambient 150; and Ambient 200.

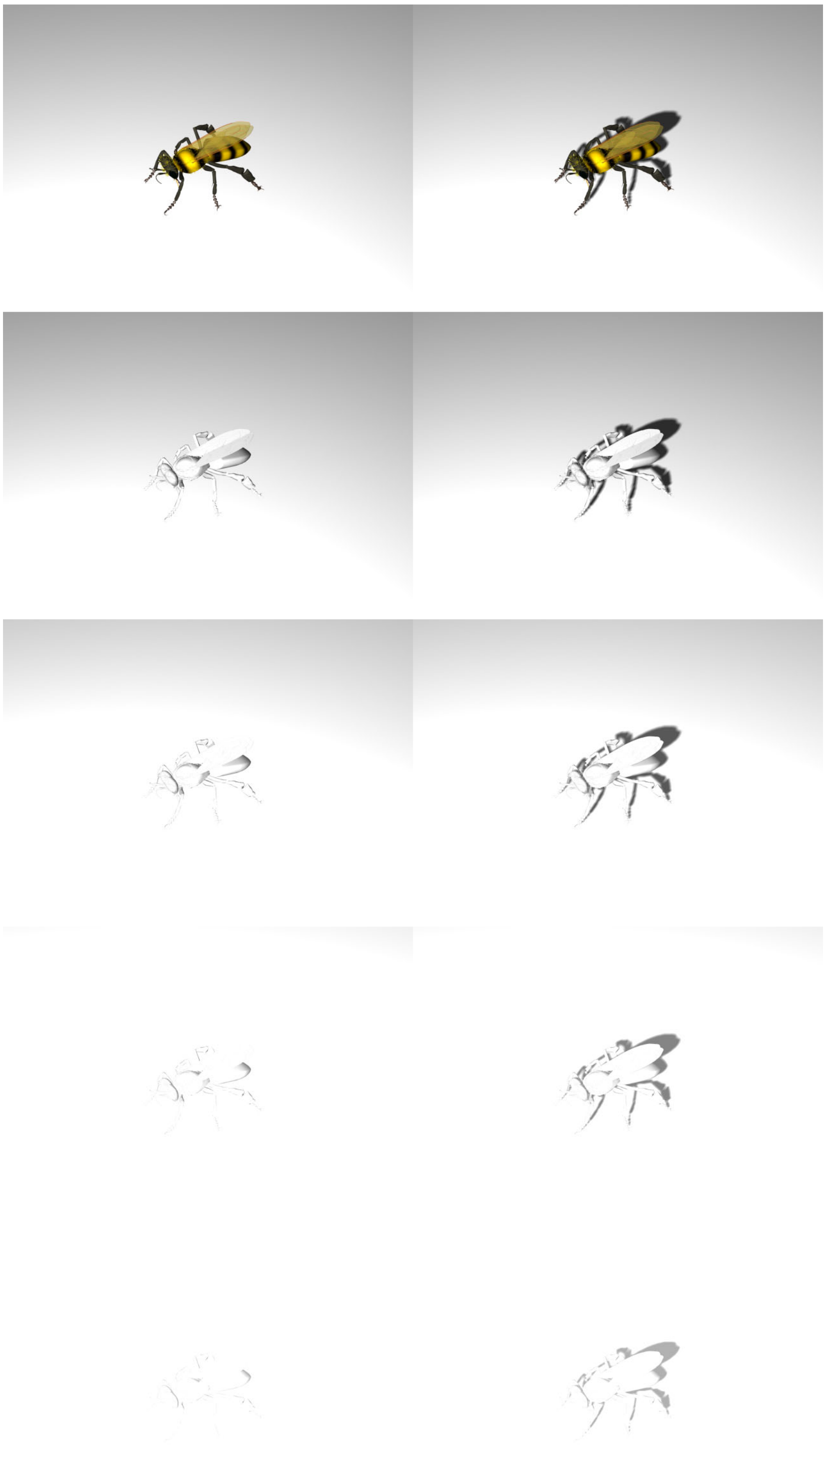

Figure A6.3. Images of the Bee view 1. The No-Shadow Condition images are on the left and the corresponding Shadow Condition images are on the right. The Illumination Conditions, from top to bottom, are: Colour Ambient 50; Ambient 50; Ambient 100; Ambient 150; and Ambient 200.

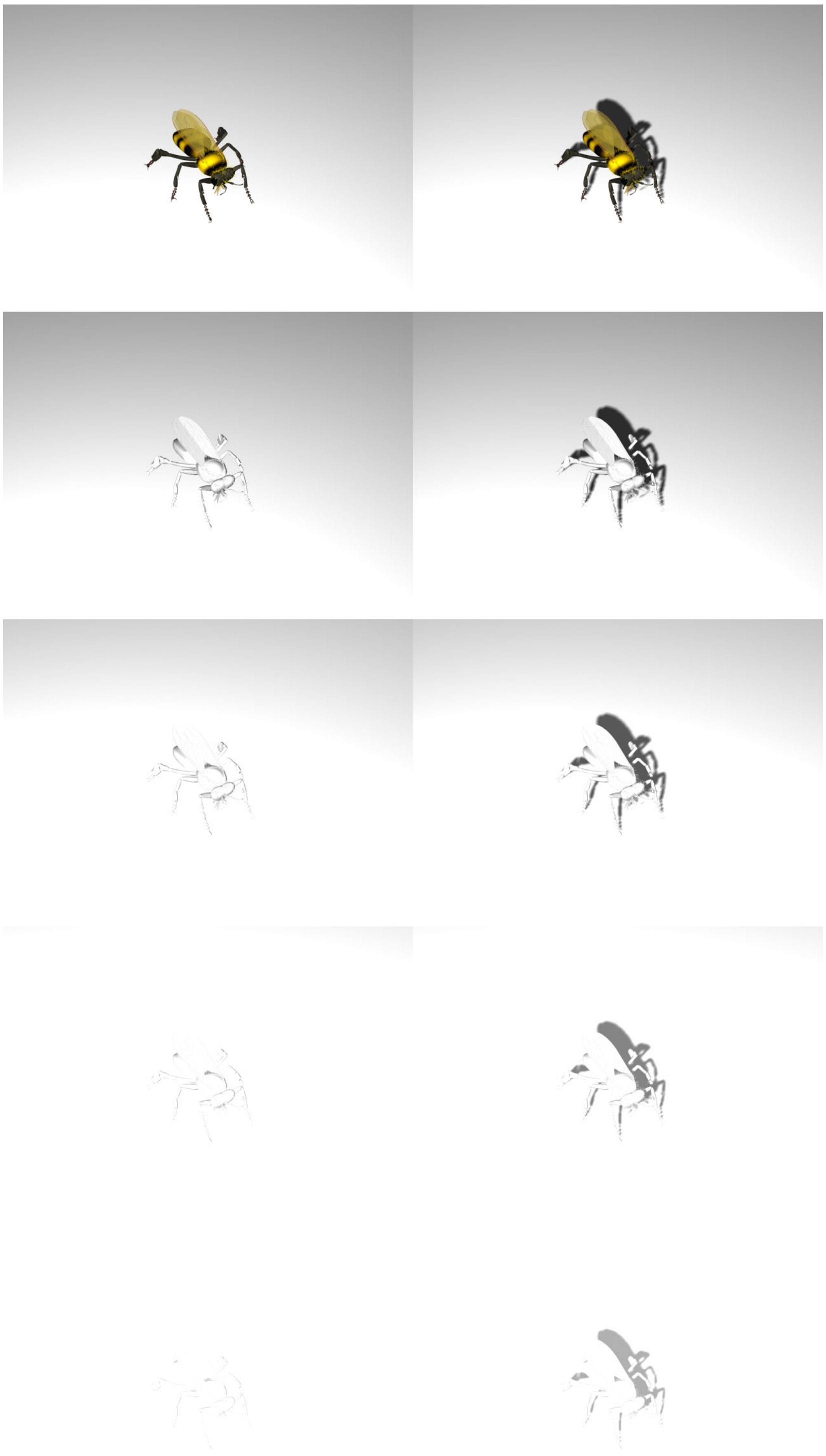

Figure A6.4. Images of the Bee view 2. The No-Shadow Condition images are on the left and the corresponding Shadow Condition images are on the right. The Illumination Conditions, from top to bottom, are: Colour Ambient 50; Ambient 50; Ambient 100; Ambient 150; and Ambient 200.

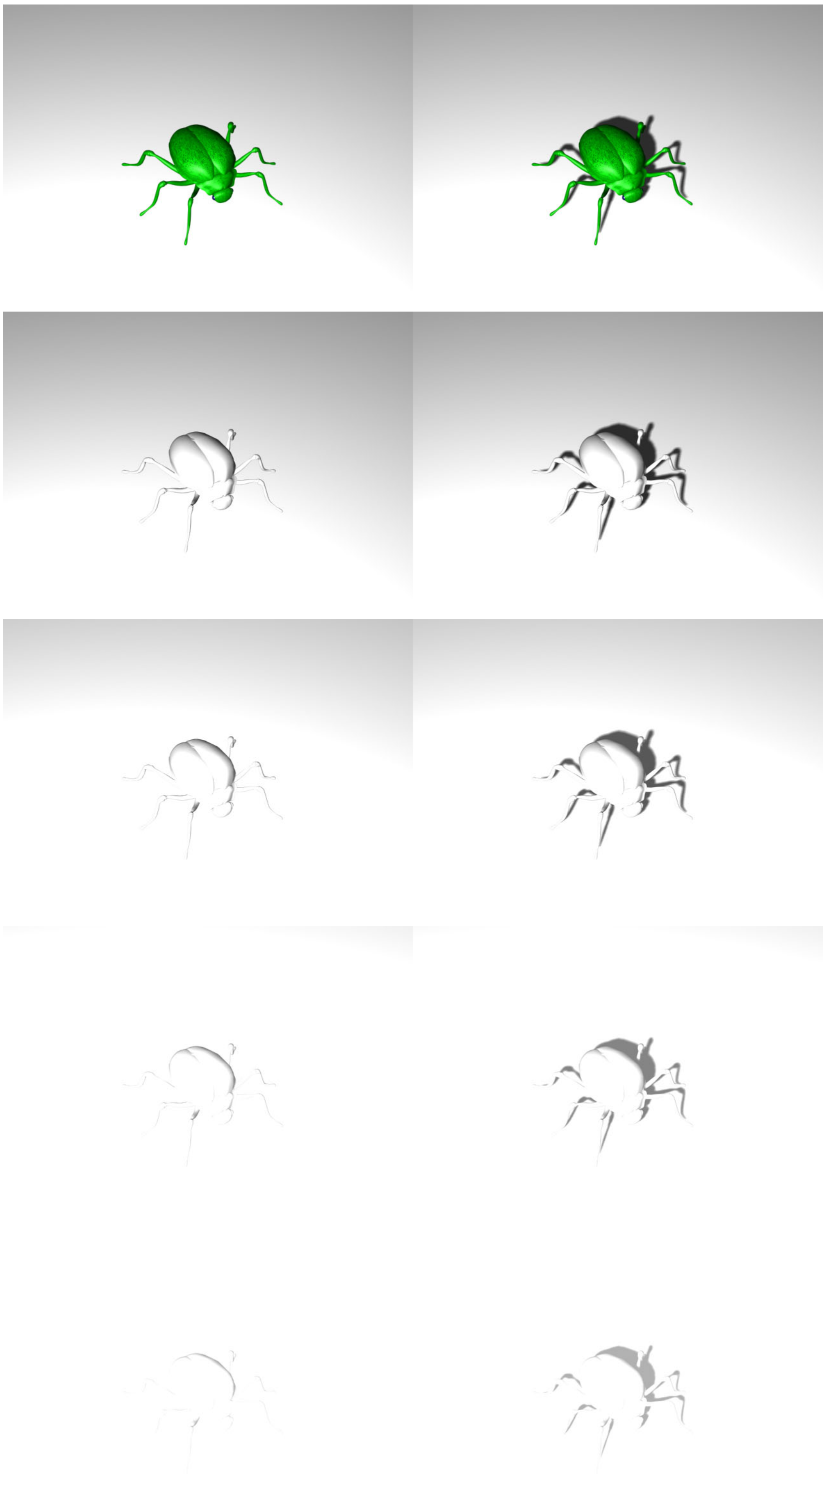

Figure A6.5. Images of the Beetle view 1. The No-Shadow Condition images are on the left and the corresponding Shadow Condition images are on the right. The Illumination Conditions, from top to bottom, are: Colour Ambient 50; Ambient 50; Ambient 100; Ambient 150; and Ambient 200.

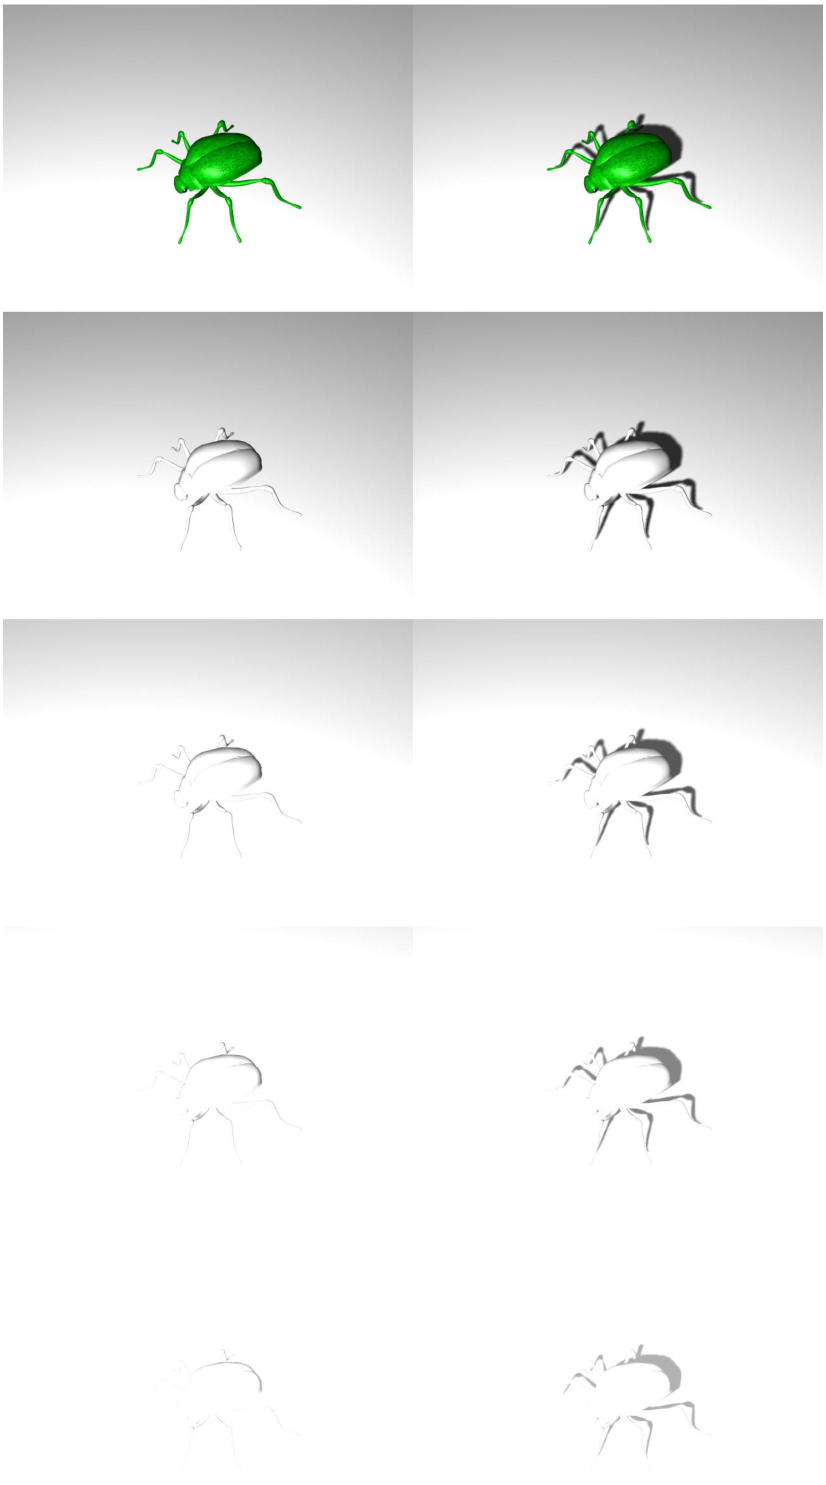

Figure A6.6. Images of the Beetle view 2. The No-Shadow Condition images are on the left and the corresponding Shadow Condition images are on the right. The Illumination Conditions, from top to bottom, are: Colour Ambient 50; Ambient 50; Ambient 100; Ambient 150; and Ambient 200.

Figure A6.7. Images of the Can view 1. The No-Shadow Condition images are on the left and the corresponding Shadow Condition images are on the right. The Illumination Conditions, from top to bottom, are: Colour Ambient 50; Ambient 50; Ambient 100; Ambient 150; and Ambient 200.

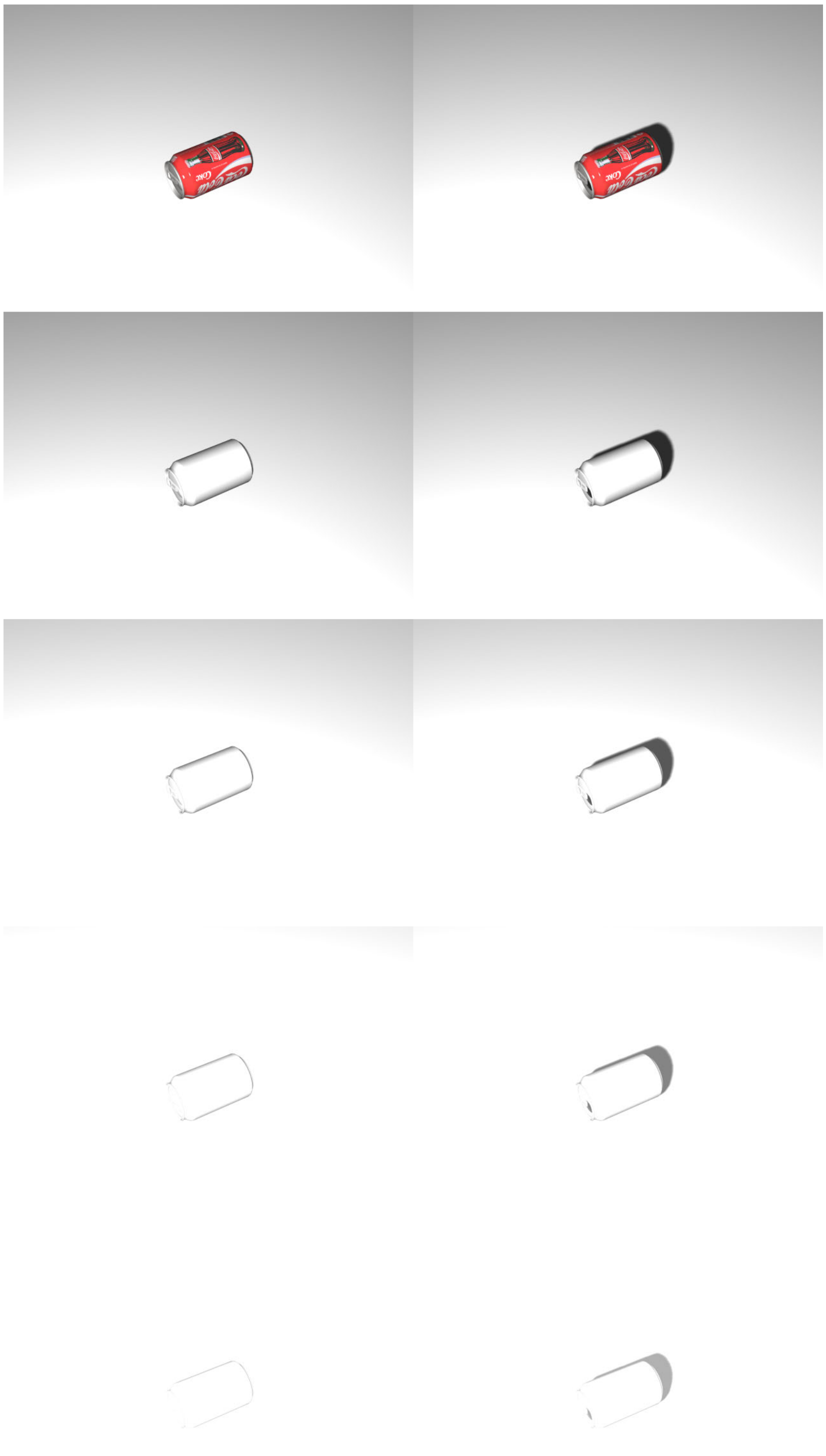

Figure A6.8. Images of the Can view 2. The No-Shadow Condition images are on the left and the corresponding Shadow Condition images are on the right. The Illumination Conditions, from top to bottom, are: Colour Ambient 50; Ambient 50; Ambient 100; Ambient 150; and Ambient 200.

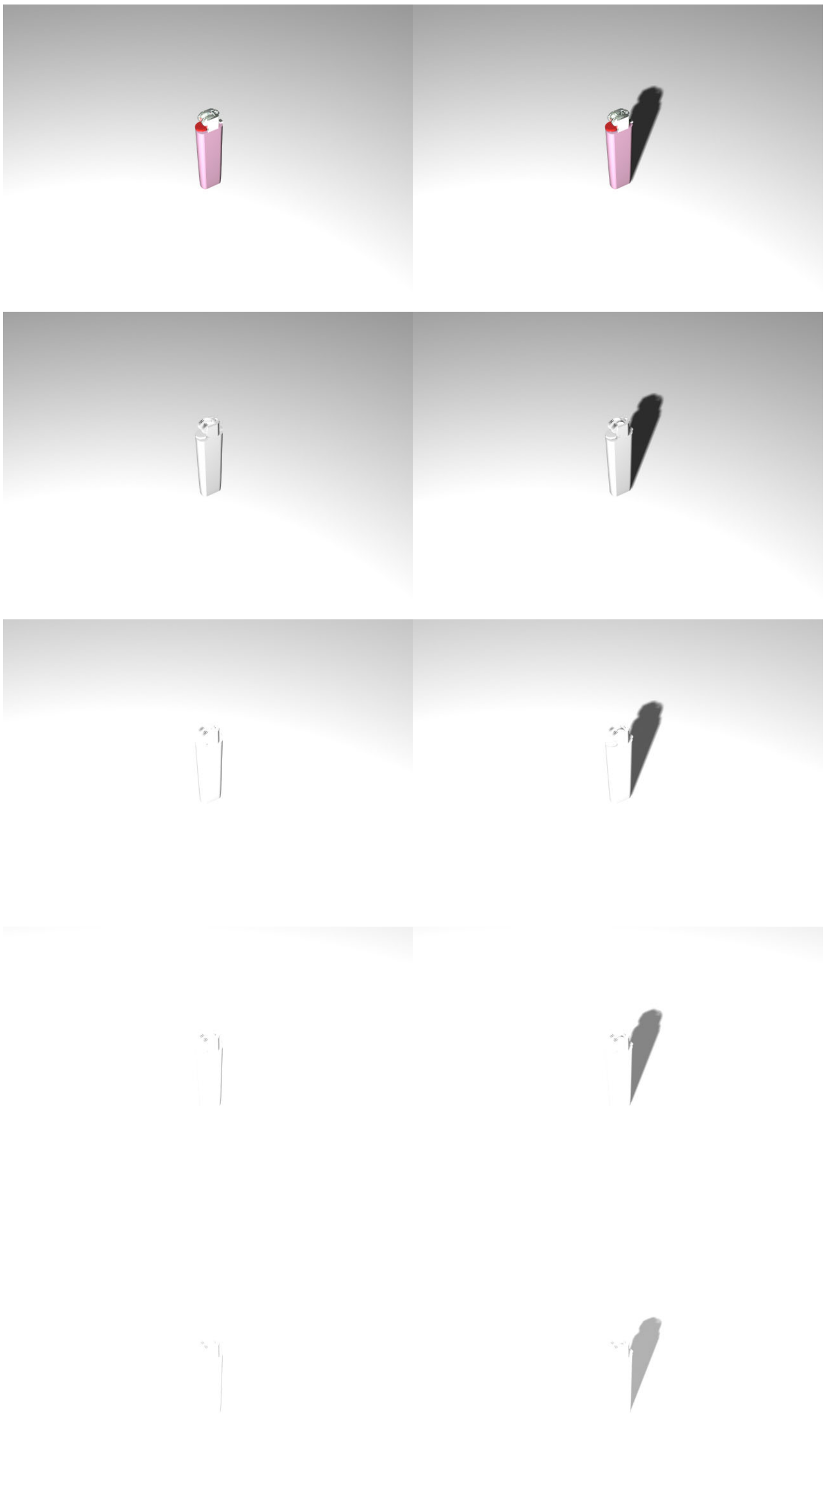

Figure A6.9. Images of the Lighter view 1. The No-Shadow Condition images are on the left and the corresponding Shadow Condition images are on the right. The Illumination Conditions, from top to bottom, are: Colour Ambient 50; Ambient 50; Ambient 100; Ambient 150; and Ambient 200.

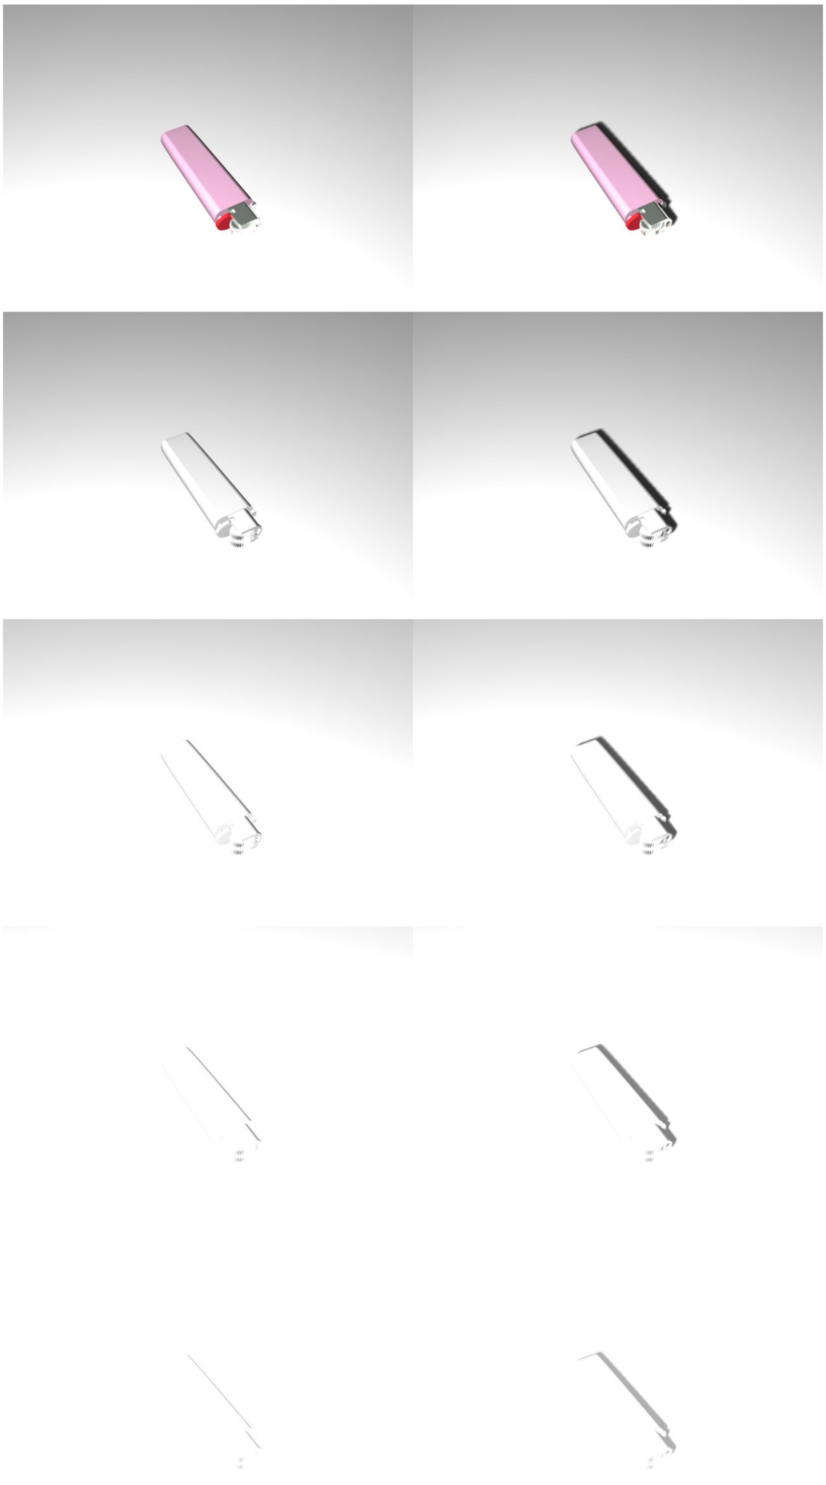

Figure A6.10. Images of the Lighter view 2. The No-Shadow Condition images are on the left and the corresponding Shadow Condition images are on the right. The Illumination Conditions, from top to bottom, are: Colour Ambient 50; Ambient 50; Ambient 100; Ambient 150; and Ambient 200.

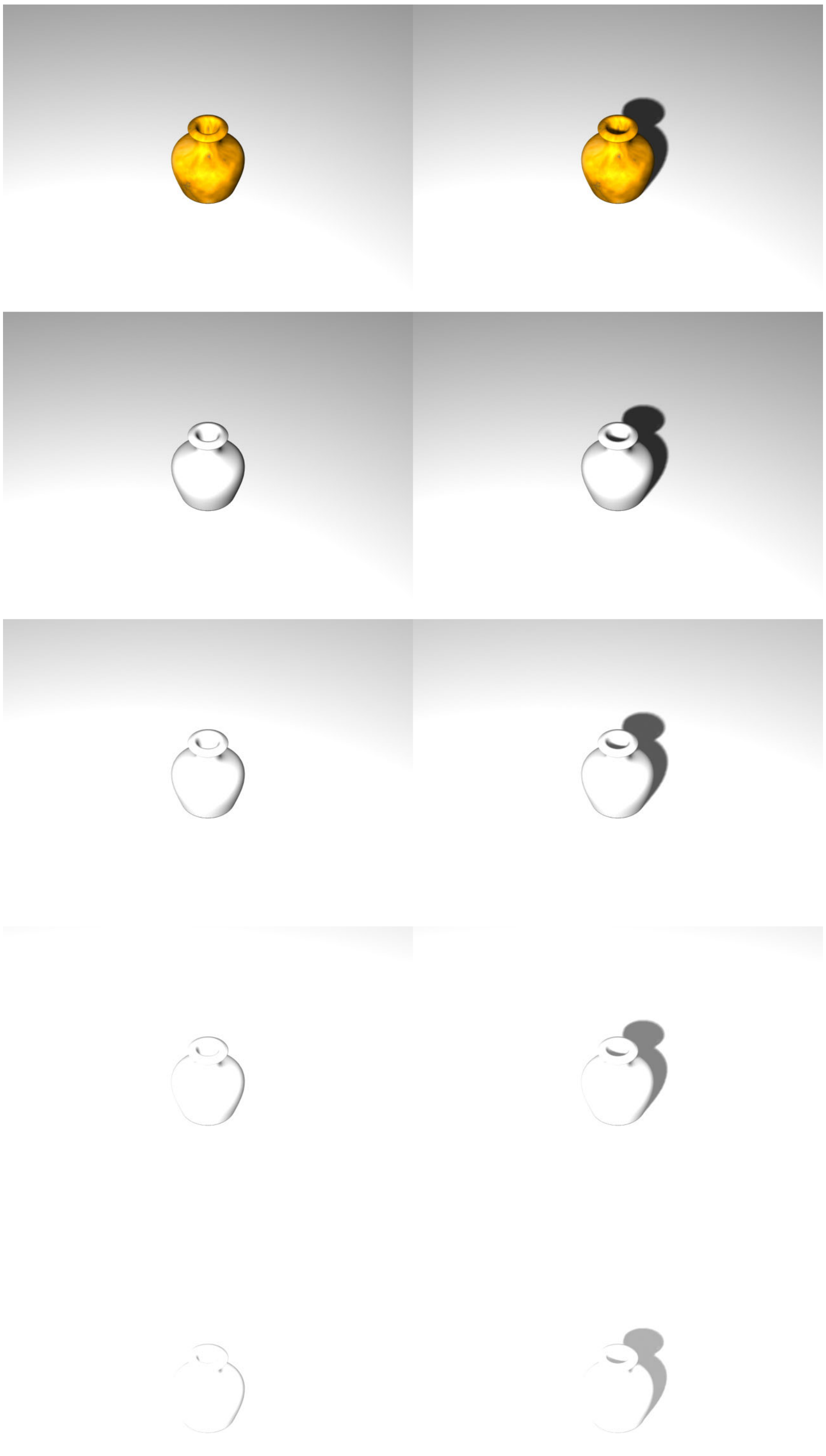

Figure A6.11. Images of the Vase view 1. The No-Shadow Condition images are on the left and the corresponding Shadow Condition images are on the right. The Illumination Conditions, from top to bottom, are: Colour Ambient 50; Ambient 50; Ambient 100; Ambient 150; and Ambient 200.

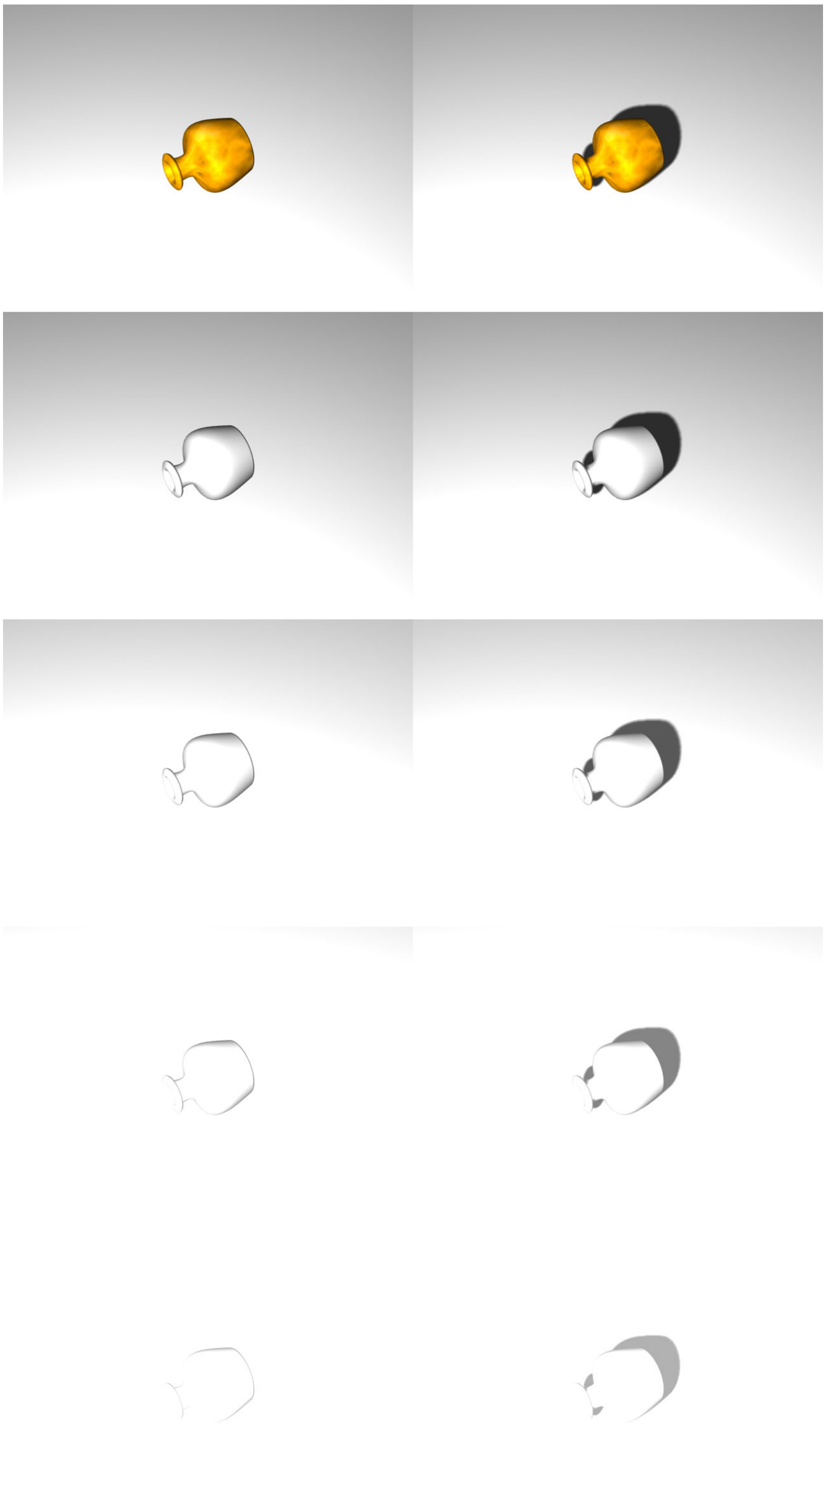

Figure A6.12. Images of the Vase view 2. The No-Shadow Condition images are on the left and the corresponding Shadow Condition images are on the right. The Illumination Conditions, from top to bottom, are: Colour Ambient 50; Ambient 50; Ambient 100; Ambient 150; and Ambient 200.

Appendix 7 Instructions to the participants for Experiment 8

The experiment contains 240 trials. On each trial you will see the sequential presentation of: 1, a cross hair; 2 the name of an object; 3 an image of an object; and 4 a masking stimulus (see: Figure 1). The images you see will only be visible for a very short period of time. Your task is to choose whether you think the object you saw matched the word that preceded it. If you think the word and picture matched, then you should press the "Q" key. If you think the word and picture did not match, then you should press the "P" key. Please respond as quickly as possible.

You can start each trial in your own time by pressing the "spacebar" key. Pressing the space bar also removes the reaction time window that appears after each trial. You have about 1 ½ seconds to respond on each trial. If you don't respond in this time, then you can just move onto the next trial.

There is no deception involved in this experiment, that is, the experiment is exactly as outlined above. If at any stage you feel that you do not wish to continue with the experiment then feel free to discontinue your participation. Course credit will still be given for any honest attempt at participation in the experiment.

Thanks,

Richard Etheredge.

Appendix 8 Data

Appendix 8 is a CD containing:

1. The raw data from all the experiments.

2. The thesis in Adobe PDF format, containing colour images where appropriate (e.g., stimuli used in Experiment 8).