EFFECTS OF LIGHTING ON THE PERCEPTION OF FACIAL SURFACES.

HAROLD HILL (B.Sc.)

Thesis submitted to the University of Stirling
for the degree of Doctor of Philosophy
(September 1993)

TABLE OF CONTENTS.

ABSTRACT.

The problem of variable illumination for object constancy has been largely neglected by "edge-based" theories of object recognition. However, there is evidence that edge-based schemes may not be sufficient for face processing and that shading information may be necessary (Bruce, 1988). Changes in lighting affect the pattern of shading on any three-dimensional object and the aim of this thesis was to investigate the effects of lighting on tasks involving face perception.

Effects of lighting are first reported on the perception of the hollow face illusion (Gregory, 1973). The impression of a convex face was found to be stronger when light appeared to be from above, consistent with the importance of shape-from-shading which is thought to incorporate a light-from-above assumption. There was an independent main effect of orientation with the illusion stronger when the face was upright. This confirmed that object knowledge was important in generating the illusion, a conclusion which was confirmed by comparison with a "hollow potato" illusion. There was an effect of light on the inverted face suggesting that the direction of light may generally affect the interpretation of surfaces as convex or concave. It was also argued that there appears to be a general preference for convex interpretations of patterns of shading. The illusion was also found to be stronger when viewed monocularly and this effect was also independent of orientation. This was consistent with the processing of shape information by independent modules with object knowledge acting as a further constraint on the final interpretation.

Effects of lighting were next reported on the recognition of shaded representations of facial surfaces, with top lighting facilitating processing. The adverse effects of bottom lighting on the interpretation of facial shape appear to affect within category as well as between category discriminations. Photographic negation was also found to affect recognition performance and it was suggested that its effects

may be complimentary to those of bottom lighting in some respects. These effects were reported to be dependent on view.

The last set of experiments investigated the effects of lighting and view on a simultaneous face matching task using the same surface representations which required subjects to decide if two images were of the same or different people. Subjects were found to be as much affected by a change in lighting as a change in view, which seems inconsistent with edge-based accounts. Top lighting was also found to facilitate matches across changes in view. When the stimuli were inverted matches across changes in both view and light were poorer, although image differences were the same. In other experiments subjects were found to match better across changes between two directions of top lighting than between directions of bottom lighting, although the extent of the changes were the same, suggesting the importance of top lighting for lighting as well as view invariance. Inverting the stimuli, which also inverts the lighting relative to the observer, disrupted matching across directions of top lighting but facilitated matching between levels of bottom lighting, consistent with the use of shading information. Changes in size were not found to affect matching showing that the effect of lighting was not only because it changes image properties. The effect of lighting was also found to transfer to digitised photographs showing that it was not an artifact of the materials. Lastly effects of lighting were reported when images were presented sequentially showing that the effect was not an artifact of simultaneous presentation.

In the final section the effects reported were considered within the framework of theories of object recognition and argued to be inconsistent with invariant features, edge-based or alignment approaches. An alternative scheme employing surface-based primitives derived from shape-from-shading was developed to account for the pattern of effects and contrasted with an image-based account.

ACKNOWLEDGEMENTS.

I would like to thank everybody who has helped with the writing of this thesis, especially Vicki Bruce and Mike Burton, my supervisors. Alan Johnston got me started and has kept in touch.

Other members of department at both Nottingham and Stirling have provided invaluable assistance. In particular, thanks to Claire O'Malley who wrote the program for the matching experiments. Thanks to all those contemporaries at both Universities who provided stimulating discussion and concrete support. At the risk of omitting some, but wanting to name some names: Neal Dench, Elias Hanna, Oli Mason, Richard Wright, Derek Carson, Pat Waterson, Jim Bright, Andy Waters, Steve Dakin, Trish Carlin, Ian Paterson and all my other contemporaries at both Stirling and Nottingham.

Thanks too to all the technical staff at both universities especially Steve Tristram and Bob Lavery for their photographs and John Russell for his hollow objects.

Lastly thank to the subjects and stimuli without whom the experiments would not have been possible.

The work was supported by a SERC earmarked studentship to Vicki Bruce.

1. CHAPTER 1: VISUAL ASPECTS OF FACE PROCESSING

1.1 INTRODUCTION

Faces are the bumpy front surfaces of heads. They contain the input channels for air, water, food and several modalities of sense data. They also serve as both input and output channels for many types of communication; verbal and non-verbal. From their appearance alone we derive information about their owner's identity, age, sex , "race", emotional state and attractiveness. Some of these abilities including the recovery of identity, recognition, may only have evolved because of the high visual salience of faces due to their primary role in communication.

This thesis will be primarily concerned with the visual processes underlying perception of a face as a face and which allow recovery of some of these different types of information from faces. In particular the recovery of identity specific visual information from faces despite changes in viewing conditions, particularly in lighting, will be considered. Such changes introduce variation into the retinal image projected by a face from which information that reliably identifies the individual must be derived. The problem of the recovery of visual information despite variations in viewing conditions is common to all face processing tasks and tasks involving other classes of objects. The input is the retinal image(s) but the output depends on the particular task, different tasks requiring different information. For example, the recovery of identity from the face requires information that is unique to that face while the appreciation of expression requires that equivalent information be recovered from different faces. Previous research about visual aspects of face and object processing will be reviewed in this chapter, while other chapters will report experimental evidence about the effects of variable viewing conditions including illumination.

1.2 PREVIOUS RESEARCH

Research into face processing has been conducted in diverse areas including psychology, neuropsychology, anatomy, neurophysiology, computing and engineering. In psychology there has been a lot of work on memory for faces as well as some into expression perception, eye gaze perception and lip-reading. The evidence from these different sources has been integrated into functional models of the processes involved which provide a theoretical framework for the interpretation of experimental results (eg. Hay & Young, 1982; Bruce & Young, 1986). Such models are based on evidence from "normal" subjects and brain damaged patients, and draw on reports of clinical conditions and everyday errors as well as the results of experimental investigations. A complete review of the large and expanding psychological literature concerning all aspects of face processing will not be attempted here. Instead the review will be limited to aspects of the models and evidence relevant to understanding visual aspects of the processing of faces.

Many early psychological studies investigated recognition memory for previously unfamiliar faces, probably because of the relevance of this task to eyewitness testimony. Recognition memory is a specific task also used in the study of word recognition. Subjects are first shown a set of unfamiliar faces during the learning phase which are subsequently presented along with unfamiliar faces which have not been learnt in the test phase. The subject's task is to distinguish between the faces they have seen once before and the faces seen for the first time at test. Measures of both hits, correct recognition of faces seen previously in the learning phase, and false positives, mistaken "recognition" of faces never seen before, are normally used and subjects typically perform at a high level, often over 95% correct. However, many such experiments are confounded by using identical pictures at study and test. This means that the task could be performed by recognising details of the particular pictures used rather than of the faces themselves. Some aspects of picture recognition may be shared with face recognition, for example early stages of visual processing, but in general knowledge about picture recognition will not inform about face recognition as faces are rarely, if ever, seen under and identical conditions and in the same context twice.

Real faces are three-dimensional objects, not fixed flat patterns of intensities like photographs, and project different retinal images depending on the angle of view, the viewing distance and the lighting conditions. Further, real faces undergo a number of natural transformations in structure such as those that characterise different expressions, for example, or the changes associated with aging. Everyday face processing must recover the information necessary for particular tasks despite such changes. Empirically this means that in order to demonstrate that face, as opposed to picture, processing is involved it is necessary to transform pictures between study and test in recognition memory or other experimental paradigms (Bruce, 1982). Realisation of possible differences between face and picture processing did lead to more detailed consideration of the visual processing of faces, including consideration of what information is used for what tasks and how this information might be derived.

1.3 STRUCTURAL ENCODING

Any particular picture of a face will be dependent on the viewing conditions, including lighting, as well as on the structure of the face itself. For most tasks involving faces it is not the details of the particular viewing conditions that are relevant, but aspects of the structure of the face itself. The processes by which information about the structure of the face, as opposed to aspects of the viewing conditions, is derived from the retinal image for particular tasks have been referred to as "structural encoding" (Bruce & Young, 1986). The products of structural encoding would be structural codes, descriptions of those aspects of the structure of the face necessary for particular tasks. For example, recognising a familiar face would require information sufficient to distinguish that face from other faces. Structural codes can be contrasted with pictorial codes, descriptions of a particular visual stimulus which specify details of the viewing conditions as well as of the object viewed. It is argued that the task involved in the experiments described in the last section could be done on the basis of pictorial codes, but that these would not be sufficient for everyday face recognition where faces have to be processed under a wide variety of conditions (Bruce, 1982).

The nature of the structural codes for faces is as yet unknown, as is the nature of the structural encoding processes, and this was one of the main reasons for doing the work reported in this thesis. Different tasks require different information and therefore would probably be supported by different structural codes. For example "configural" information, that is information about the spatial relationships between "features", is known to be important for many tasks including recognition (Sergent, 1984; Young, Hellawell & Hay, 1987) and so would be expected to be contained in the structural code(s) used for those tasks. An understanding of structural encoding would require that the processes that allowed the recovery of the configuration of a face be specified, as well as the form that the description takes.

As well as being different for different tasks, structural codes for the recognition of familiar and unfamiliar faces might be different. As new faces become familiar those aspects that are less variable, and thus more reliable signals of identity, appear to become more important. For example there is evidence, that the internal features of faces are more salient for familiar faces and external features for unfamiliar faces, and that this is a property of structural, not pictorial, codes (Ellis, Shepherd & Davies, 1979; Endo, Takahashi & Maruyama, 1984; Young, 1984; Young, Hay, McWeeny, Flude & Ellis, 1985).

In summary, the precise nature of structural codes is not known but they are thought to represent a more abstract, derived level of description of the stimulus than that of pictorial codes, describing details of the structure of the face, for example, rather than the particular pattern of image intensities. Different levels of description may be needed for different tasks including perhaps object-centred codes for recognition, a description of the face within its own coordinate system which is completely independent of the viewer and viewing conditions (Marr & Nishihara, 1978). The work reported in this thesis sought to investigate how structural encoding processes are affected by lighting conditions and also whether the structural codes themselves contain lighting dependent information. However, the nature of structural codes and structural encoding processes is likely to depend on what information is important for particular tasks and evidence about this will be reviewed next.

1.4 CONFIGURATIONS vs FEATURES

Faces can be thought of as being composed of the set of features or parts for which we have linguistic handles; eyes, eyebrows, ears, nose, mouth, cheeks, chin and forehead for example. It is possible that facial processing is based on the independent processing of these features, and this was the assumption that underlay the development of the photofit facial reconstruction system (Penry, 1971). People could, for example, be recognised on the basis of one or two distinctive features, without reference to other aspects of their face that are shared with many other people. However, there is evidence that features are not processed independently but rather that faces are processed holistically, as gestalts, with the relationships between features being as important as the individual features themselves (Sergent, 1984; Young, Hellawell & Hay, 1987). Of particular importance appears to be 'second order relational' information, that is information about the spatial relations between features within a shared configuration (Diamond & Carey, 1986).

Good evidence for the importance of configural information is the demonstration that halves of familiar faces are more difficult to recognise when they are presented as mismatched pairs aligned so as to form new, composite faces than when the same halves are presented misaligned (Young et al., 1987). The features depicted in each case are the same, but when they are seen within the context of a plausible face the new configural information generated appears to interfere with the processing of individual features independently. This finding seems to be independent of the way that the face is divided in half as it has been demonstrated with faces divided into top and bottom halves or into internal and external features. Either of these divisions will preserve some of the local spatial relationships but the interference demonstrated when the face is shown as a whole suggests that there are aspects of the face seen as a whole that are important. It is interesting to note that turning the stimuli upside-down, a transformation known to disrupt face processing (Yin, 1969), actually facilitates the recognition of the individual halves of composite faces. This manipulation, often referred to as figural inversion and discussed in detail below, is thought to disrupt the encoding of configural information (Diamond & Carey, 1986) but seems in this case to facilitate the processing of features independently. Other evidence for the importance of processing faces as wholes rather than parts is the 'face superiority effect' (Homa, Haver & Schwartz, 1976) which shows that the processing of individual features may actually be enhanced if they are part of a facial configuration with a nose, for example, being better remembered when it is learnt as part of a face.

The precise nature of configural information, like that of the hypothesized structural codes, is not established. Even what constitutes a facial feature is not known, as the features for which we have words are not necessarily the ones that are important for visual processing. The first experiments reported investigated effects of viewing conditions on the perception of the hollow face illusion (Gregory, 1973; see also figures 2.2 & 2.2) as a face, that is as a globally convex whole with local convex and concave features. The possible role of lighting and viewpoint in highlighting particular features in a way that affects identification is considered in Chapters 3 and 4. Figural inversion is used to investigate the relationship between viewing conditions and configural information.

1.5 OBJECT CONSTANCY

The general problem for the structural encoding of both faces and other classes of objects is that of object constancy, the recovery of the constant structure of the object under variable viewing conditions. Objects with different structures will give rise to different retinal images but so too will any one object seen under different viewing conditions. Recognising and encoding aspects of structure that distinguish an object despite differences that arise because of variable viewing conditions is the problem of object constancy. With faces, as with some other classes of object with moveable parts, this is further complicated by variations in the structure of the objects themselves.

It is normally assumed that the recognition of an object involves the matching of a description derived from the input image to a stored description of the object. This raises questions about the derivation of descriptions from the image, questions general to vision research, and also about the form of the stored representations, questions that overlap with memory research. A common problem for object constancy is the recovery of three-dimensional object shape despite variations in the two-dimensional image resulting from changes in viewpoint. Most three-dimensional objects, including faces, project different images depending on the angle of view, the image variation itself being a function of the structure of the object.

A straightforward but computationally expensive solution to this problem would be to represent all objects from all viewpoints from which they could be seen. In this vein, instance-based models of memory that could be applied to object recognition have been proposed (eg. Hintzman, 1986; Nosofsky, 1986), where a trace would be stored for each "instance" of an object. This approach has been further developed through the use of neural networks which can be trained to learn different instances of a particular object, recognition of which then generalizes to novel instances. Different methods of storing and accessing traces lead to different patterns of performance, but seemingly abstractive properties like prototype learning have been demonstrated. Particular types of connectionist network solve learning and recognition problems in different ways, though their performance is often describable statistically. For example the auto associator network which stores instances in that it associates a particular pattern with itself, can be set up so that it effectively performs principal components analysis on the input - a popular solution to the problem of variability of input images (eg. Turk & Pentland, 1991).

Other "invariant features" schemes of object recognition attempt to find view invariant features in the image that provide sufficient information for recognition. Taking the example of faces certain characteristics of the hair, like colour and texture, might serve as invariant features. These properties are view invariant in that they are recoverable from any orientation where the hair can be seen. Consistent with this, hair is known to be important to memory for faces, particularly previously unfamiliar ones (Shepherd, Davies & Ellis, 1981). However hair is highly variable in other respects - it can be cut, brushed, dyed and it also grows, falls out and changes colour - and so would not always be a reliable cue for distinguishing between faces.

Other schemes suggest, using arguments similar to those applied to structural codes, that the stored representation for a familiar object will characterise that object without reference to particular instances. For example a single three-dimensional object centred model of the object could be stored (Marr & Nishihara, 1978). Such representations are characterised by object centred coordinate systems, derivable from a two-dimensional image, and by being composed of parts at different scales arranged in an hierarchical way. If such stored three-dimensional models exist then the problem of matching the input image to a particular representation can be shown to be over constrained provided that at least three non-colinear features can be found in common between the input image and the model (Lowe, 1987). The question remains as to how the three-dimensional stored model is derived in the first place (Poggio & Edelman, 1990). However, an object centred coordinate system at least appears plausible for faces, with perhaps a vertical axis given by symmetry and a horizontal axis defined by the eyes.

Most such schemes represent the object in terms of component parts and are therefore referred to as object decomposition methods (Ullman, 1989). Objects are decomposed into a limited set of generic parts, for example generalized cones, volumes produced by sweeping a cross-section of constant shape but varying size along an axis (Binford, 1971), or geons, another set of volumetric primitives (Biederman, 1987). Such parts can be defined in terms of image properties such as occluding contours or edges. Some image properties, including colinearity, curvilinearity, cotermination, parallelism and symmetry, bear non-accidental relationships to properties of the three-dimensional world and thus could be used to define the three-dimensional parts present from the information in the two-dimensional image (Lowe, 1984). Geons and generalized cones were both designed to describe objects composed of clearly separated and often articulated parts and are not necessarily appropriate for faces. An alternative set of component parts that have been proposed specifically for faces are eight primitive surface types (Bruce, Coombes & Richards, 1993). These are defined according to their mean and Gaussian curvatures (Koenderink,1990), though how these would be recovered from the two-dimensional image has not yet been specified. It is claimed that such parts capture psychologically relevant dimensions of facial variation.

Object decomposition and invariant features methods are not mutually exclusive. The parts recovered from object decomposition, or their relationships, may themselves constitute invariant features. Both the three-dimensional model representation (Marr & Nishihara, 1978) and the recognition by components (Biederman, 1987) scheme make use of relationships between parts that are likely to remain invariant over a large range of views in which those objects are likely to be encountered. Most object decomposition schemes are models of between class distinctions and it is not clear how within class distinctions would be made. For example, (almost) all faces would decompose into the same parts having the same approximate relationships making it unclear on what basis different examples within the class of faces would be distinguished.

Another solution to the problem of object constancy for recognition is the alignment approach (eg. Ullman, 1989). Again single models of each object are stored. The difference from object decomposition models is that rather than decomposing the input image into parts for matching, the model is transformed as a whole so that it is aligned to the image. For example the catalogue of models might be transformed so that their appearance as seen from the same viewpoint as the input image would be made explicit. Axes of symmetry and labelling of the top and bottom of objects might help such a task. Such prior alignment simplifies the second stage of recognition, the matching of input to the correct stored model, as all the objects have to be checked but only as they appear from the appropriate view. Matching can then be accomplished by directly matching view dependent object features without the necessity of deriving view independent parts or other features. Transformations could also allow for other sources of variation in input, such as size or even constrained object variability such as that allowed by the articulation of parts. This is the part of the power of the approach but begs the question as to how the models and transformations are acquired.

The experiments reported in this thesis, investigated the plausibility of such alternative systems as explanations of the workings of the human visual system by exploring the effects of variations in input image. These effects can reveal clues about the nature of stored models and also the processes applied to the input which allow access of such models.

1.6 FACES vs OBJECTS

Faces are three-dimensional objects and as such pose many of the same problems for any recognition system as other classes of three-dimensional object. However, there are also problems more or less specific to face recognition which may be reflected in their visual processing and in the recognition strategies used and these possibilities are considered next. The neuropsychological evidence about possible impairments to the human brain suggests that object and face recognition are clinically dissociable from each other and from word recognition.

Many of the differences between objects and faces are a result of the different functions they support. For example, the need to make within category distinctions within the highly homogeneous class of faces may impose particular demands on any system for face recognition and lead to particular solutions (Hay & Young, 1982). Faces all share the same parts or features arranged in much the same way - eyes above nose above mouth etc. - and so these properties will not be sufficient to distinguish between individuals. This problem is further complicated by the variability in the structure of faces required for their communicative functions. The mouth is a good example of a highly deformable facial feature important for communication. For the purposes of recognition what distinguished any particular mouth would need to be recovered while for lip-reading or expression properties of changes common to all mouths would be the important information. These seemingly contradictory requirements will have shaped both the structure of the face and the systems evolved for processing it. If the communicative functions are primary this will have led to selective pressure for equivalence or homogeneity within the class of faces, making the problem of recognition more difficult and error prone.

As well as differences in function the structure of faces may also differ in important ways from that of other objects. As was noted in the section on object constancy, faces are not readily decomposable into parts separated by clearly defined contours. Such edges play an important part in many theories of *object* recognition (eg Biederman, 1987; Marr & Nishihara, 1978) but may not be applicable to faces.

Consistent with this there is evidence that faces are not as well recognised from line drawings as are objects and that for faces but not objects there is an advantage for photographs over line drawings (Davies, Ellis & Shepherd, 1978; Biederman, 1987).

An alternative to edge detection specifically developed for encoding faces and hands is valledge detection (Pearson & Robinson, 1985). Following work which had shown that the application of edge detectors to facial images produced cluttered representations of the internal facial features (Pearson & Six, 1982; Pearson, 1983) an alternative operator was developed sensitive to luminance valleys in the image. These luminance valleys occur in regions of the image corresponding to areas of the surface of the object or face where the surface slopes sharply away from the line of sight. This operator is also sensitive to large edges, hence reference to it as a valledge detector. When combined with a thresholding operation which adds "mass" to the images produced the "cartoons" produced have been found to be nearly as well recognised as full grey level images (Bruce, Hanna, Dench, Healey & Burton, 1992). It has even been speculated that the system is effective because it mimics the operation of the early visual system; the cartoon of a cartoon is itself and it may be that the cartoon and full grey level image produce the same neuronal response. The features extracted by the composite operator could be used for economical representation and storage, fulfilling many of the functions required by the object recognition schemes discussed above. The thresholding component of the composite detector is necessary for good recognition of faces probably because it provides the sort of information normally provided by shading or shadow. Faces are composed of the sort of relatively smooth featureless surfaces for which other sources of shape information, especially shading (Watt, 1991), may be as important as edges. Shading information will be considered in more detail in the next section.

Thus the functions and structure of faces may pose problems for their processing that are not shared with other classes of object and empirical findings may not generalize. However, it is also possible that once the ability to make the fine distinctions in structure necessary for distinguishing faces had evolved, the same system could be adapted for other purposes including general object recognition.

1.7 SHAPE-FROM-SHADING

Shading is the spatial variation of intensity as a function of illumination, surface orientation and viewing angle and shape-from-shading is the process by which three-dimensional shape is recovered from this information. The recognition of faces from line drawings compared to photographs, along with the importance of lighting and make-up and the existence of the hollow face illusion, have all been cited as evidence that shading information is important for face perception (Bruce, 1988). This section will outline what is known about shading information and its use and will consider why it might be important for face perception.

The optics of image irradiance are well understood including the relationship between changes in surface orientation and changes in image irradiance (Horn, 1977). However the human visual system has to recover scene structure from image irradiance, the reverse problem. Image intensities are available but they are determined not only by the shapes of surfaces present in the scene but also by properties of the source of illumination, the position of the viewer and the reflectance of the surface. As all these are initially unknown the problem is effectively under constrained and its solution is likely to require the use of simplifying assumptions. Specification of possible assumptions and investigation of the assumptions that the human visual system actually employs has been one goal of shape-from-shading research.

Shading may not be as effective a cue to shape as motion or stereopsis (Cavanagh & Leclerc, 1989). However the fact that we can interpret static two-dimensional images containing little else than shading information (eg. photographs of sculptures) together with the widespread exploitation of shading in representational art suggests that the human visual system has evolved to make use of this source of information. Indeed the apparent advantage in providing camouflage of counter shading, the use of pigment changes to cancel out the revealing effects of shading, in many species suggests that shading may be a widely used and primitive source of shape information.

The use of shading by the human visual system has been studied psychophysically (Pentland, 1982; Todd & Mingola, 1983; Mingola & Todd, 1986; Todd & Reichel, 1989; Johnston, Passmore & Morgan, 1991; Erens, Kappers & Koenderink, 1993) however the conclusions drawn seem contradictory or very task dependent. These studies do suggest that the visual system does not invert the image formation process at each point in the image. For example lighting direction is an important determinate of point intensities but does not seem to be recovered prior to the interpretation of shape (Mingola & Todd, 1986). Where the studies differ is in how accurately they conclude that surface shape can be specified by shading. Studies which measure subjects' impressions of local surface orientation show high levels of errors (eg. Mingola & Todd, 1986). However when the task is to judge relative curvature subjects can perform quite accurately (Johnston, Passmore & Morgan, 1991). While curvature is the derivative of orientation, it may be that it is recovered directly from shading and used as a metric for quantifying surface shape. So it seems possible that humans can recover surface shape from shading, and that shape information may be coded in terms of curvature.

As well as testing the ability of humans to make use of shading information there has been other work on the relationship between shading and shape (Horn, 1977; Koenderink & van Doorn,1980). These approaches seek to specify how shape might be recovered from shading and what assumptions would be required. One of these approaches has attempted to solve the inverse optics problem by making assumptions about the illumination and surface reflectance and shape (Horn, 1977). Psychological validity is not claimed for this approach and there is no evidence that the human visual system requires the same prior knowledge or makes the same assumptions. Another approach is to relate global properties of the image intensities to surface shape. For example, it has been shown that the positions of luminance singularities and the topography of isoluminance contours are lawfully related to object structure (Koenderink & van Doorn, 1980). In this case the constraints required are global rather than local. The effects of stimulus manipulations on human performance can be used to test the psychological validity of the assumptions invoked by such approaches.

For example another set of experimental studies have used variants on the crater illusion to investigate the assumptions and processes the visual system uses to interpret shape-from-shading (Ramachandran, 1988a & 1988b). The basic stimulus used, the crater illusion, is a circular pattern of shading that can be interpreted as a convex bump or a concave dent (See figure 1.1). The different interpretations of the stimuli entail different lighting directions and it is sometimes possible to change one's perceived interpretation by mentally shifting the light source. The tendency to perceive the interpretation consistent with light being from above is cited as evidence that a light-from-above assumption is employed by the human visual system in the interpretation of shape-from-shading. Indeed, when the image of a top lit bump is rotated by 180° it is seen as a dent lit from above, rather than as a bump lit from below, again consistent with a light-from-above assumption.

Another demonstration shows that the interpretation of displays consisting of two configurations of such stimuli that are 180° rotations of each other, are mutually constrained - one type are seen as bumps and the other as dents. The two configurations would imply two opposite lighting directions if they were perceived as the same shape, but interpreting one type as bumps and one as dents allows an interpretation consistent with a single lighting direction for the entire image. This is interpreted as evidence that shape-from-shading is a global operation incorporating a single light source assumption. Other demonstrations show that such stimuli, when interpreted as three-dimensional surfaces, can also support grouping, segregation and apparent motion phenomena implying that their three-dimensional interpretations can act as elementary features or primitives for subsequent processes.

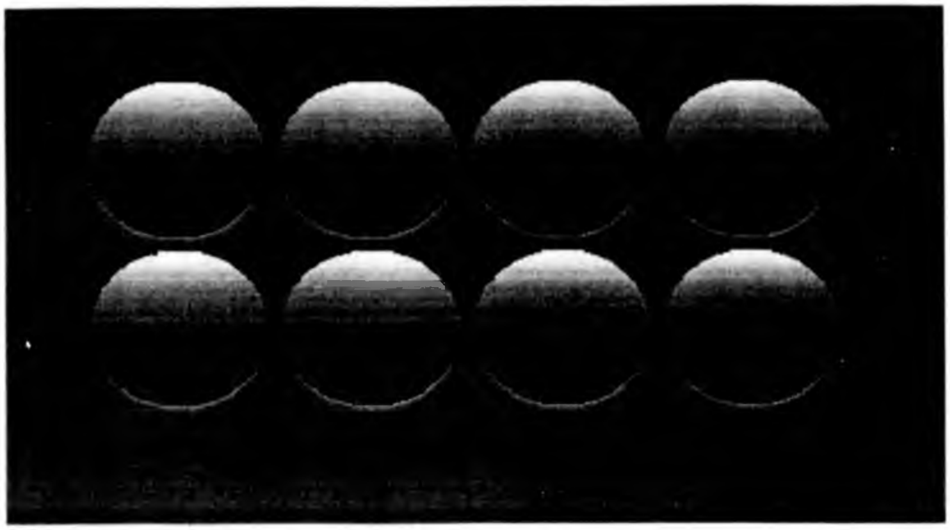

Figure 1.1: Examples of the type of stimuli used in the experiments reported by Ramachandran (1988a & 1988b). (See text for discussion.)

Experimental evidence that shape-from-shading is particularly important for face processing has come indirectly from the effect of photographic negation (Phillips, 1972). This manipulation, like figural inversion, affects visual information in a clearly defined way - it reverses intensity values about the mean. Negation impairs performance on many tasks involving face perception (Galper, 1970; Phillips, 1972; Luria & Strauss, 1978; Kemp, McManus & Pigott, 1990). This may be largely due to the effects of negation on the appearance of pigmented areas - eg. light hair becomes dark and visa versa. However negation may also disrupt the interpretation of shading and shadows (Cavanagh & Leclerc, 1989). Taking the negative of a photographic image of a face in effect results in a set of image intensities that are incompatible with any illuminated three-dimensional object and this may be partly why negatives are difficult to process (Horn, 1990). Interestingly there is other information that negation does not affect, like the position of isoluminance contours. The adverse effects of negation do not seem consistent with accounts of face processing based solely on such features. Negation is also interesting in that it effects low spatial frequencies more than high spatial frequencies, consistent with the importance of the former for face perception (Hayes, Morrone & Burr, 1986).

Another interesting and important feature about research on shape-from-shading concerns how information from this source is combined with other sources of shape information like contours, specularities and texture (Bulthoff & Mallot, 1987). For example information from occluding contours has been used in computational approaches to provide boundary conditions for the differential equations to be solved (Ikeuchi & Horn, 1981). The same information has also been shown to affect human interpretation of patterns of shading (Ramachandran, 1988a & 1988b).

The work reported in this thesis both investigated the ability of humans to make use of shading information under different conditions and also the relationship of shading information to other sources of information. One stimulus condition that was varied in many of the experiments reported was lighting and this will be considered next.

1.8 LIGHTING

The retinal image projected by any object is affected by the number, properties and arrangement of the light sources illuminating that object. The discussion of lighting in this section and in the thesis in general will be limited to the effects of changing the **direction** of a single light source. The physics of illumination is well enough understood to allow the effects of multiple light sources to be modelled including the effects of mutual illumination, the illumination of one surface by light reflected off another. However, it appears that the human visual system normally assumes a single light source, at least in the interpretation of shape-from-shading

(Ramachandran, 1988a & 1988b). Studying the effects of variation introduced in the relatively simple case of one light source varying in direction may generalize to predict the effects of other sources of variation. For example if lighting invariant features are used these would remain invariant whatever the number and arrangement of light sources.

The basic problem for processing faces under conditions of variable lighting is that of object constancy, the same problem as has already been discussed in the context of variable viewpoint. Changing the direction of a single light source affects the retinal image projected by an object just as changing the viewpoint does. Changing lighting direction determines which surfaces receive direct illumination and which are in shadow in the same way as changing viewing direction determines which surfaces are visible and which hidden. Indeed the information provided by a particular lighting direction is in some ways analogous to that provided by a different viewing direction - the shadow areas correspond to what would be hidden surfaces and the shadow boundaries to what would be occluding contours if the object were viewed from the position of the light source. In the particular case where light source and viewer are in the same position no shadows will be visible as the overlap with hidden surfaces will be precise.

Like changes in viewing direction, changes in lighting direction do not change the underlying structure of the face or object viewed but both do change the retinal projections of all objects including faces. For most face processing tasks it is aspects of the structure of the face that needs to be recovered from the retinal image, not the particulars of illumination or viewing direction. The variability introduced by variable viewing conditions can often be considered as a source of noise in the recovery of the structure of the object, the signal. However, any retinal projection has an associated viewpoint and lighting direction and these determine the surfaces visible and the particular pattern of shadows, shading and contours that provide cues to shape. So, while precise details of a particular image may be accidents of the viewing condition those details also reflect properties of the underlying structure. How well human performance is predicted by details of particular images or whether it seems dependent on derived properties is an important question in this thesis as well as how information about structure might actually be derived from image properties.

Many of the problems and solutions that have been discussed in the context of variable viewing direction also apply to variable lighting direction. In particular the relationship of the lighting dependent retinal image to the stored representations thought to be required for recognition must be considered.

All objects, including faces, could be stored under all lighting conditions within an instance based account, although this additional source of variability emphasises the possible computational expense in terms of memory required by such a solution. An instance based store could contain traces for instances with different directions of lighting and some of these might be sufficiently similar to a novel retinal image, as lit from a particular direction, to allow matching for recognition (ie. generalization across lighting conditions might emerge from an instance based model). Alternatively lighting invariant features, such as high contrast pigment differences, might be used for recognition. However, such information may be relatively sparse and itself variable under other changes in viewing condition (eg. occluding contours are lighting independent but change with changes in viewpoint). The parts used by object decomposition schemes (eg. Marr & Nishihara, 1978; Biederman, 1987) may be lighting independent once recovered. Indeed the problem of lighting has not been widely considered in the object recognition literature because most object recognition schemes are based on edges, edge detection being insensitive to most changes in lighting (Poggio & Edelman, 1990). Many edges, including occluding contours and high contrast borders between areas of different reflectance, are lighting invariant features and may be sufficient to support matching. However faces are not composed of parts separated by clear contours and edges may not be sufficient for face perception. Instead other aspects of the image, for example shading information, may be necessary to support face processing. This information, its processing and even the derived descriptions may be lighting dependent in a way that edges are not.

Previous research into the effects of lighting on face perception has included reported effects of lighting on the recognisability of photographs of faces (Johnston, Hill & Carman, 1992) and on the recognition of facial surfaces portrayed by shading with no information about pigmentation (Bruce, Healey, Burton, Doyle, Coombes & Linney, 1991). The recognition of photographs has been shown to be disrupted by lighting faces from below the horizontal meridian, bottom lighting (Rock, 1973; Johnston et al, 1992). This effect may be associated with difficulties in recovering shape-from-shading which may incorporate a light-from-above assumption (Ramachandran, 1988a & 1988b). When images of faces lit from below are turned upside-down the light source appears to be from above relative to the observer and this appears to offset some of the difficulties in their recognition resulting from the face being upside-down (Johnston et al, 1992). Interestingly the valledge detector mentioned earlier (Pearson & Robson, 1985) also breaks down under conditions of bottom lighting as luminance valleys can become luminance ridges. Despite the apparent difficulties normally associated with bottom lighting, this direction was found to facilitate the recognition of shaded representations of facial surfaces, "laser heads" of the type used in many of the experiments reported here (Bruce et al, 1991). This apparently anomalous result may have been an artifact of the task used, matching, perhaps the result of bottom lighting highlighting more informative surfaces or surface details in these particular representations. This previous work did highlight the importance of the precise form of illumination used and it was a major aim of the work reported in this thesis to systematically investigate such effects on both matching and recognition tasks as outlined in the overview below. Laser heads themselves will be more fully described in chapter 3, the first experimental chapter in which they were used.

1.9 OVERVIEW OF EXPERIMENTAL WORK

Experiments are reported on the effects of lighting direction on the perception of the hollow face illusion (Gregory, 1973; see also figure 2.1), and on the matching and recognition of laser heads (cf. Bruce et al, 1991).

The work on the hollow-face illusion is reported in Chapter 2. Perception of this illusion may involve some of the same processes that are involved in perception of a face as a face, that is as a globally convex object with appropriate local concave and convex features. The illusion reflects the ambiguity of shading which, as has already been discussed, is always consistent with two possible interpretations of surface geometry, convex or concave, with their opposite associated lighting directions. Indeed the illusion has been cited as evidence of the importance of shading for face perception (Bruce, 1988) as it appears to be the shading that defines the face as a face. Differences between lighting the illusion from above or below are tested, in order to establish the role of a light-from-above assumption in the interpretation of shape-from-shading. The integration of shape-from-shading with other sources of depth information, particularly stereopsis (experiment 4), is also considered. The possible importance of object specific knowledge in generating the illusion is also investigated, especially whether such knowledge is dependent on lighting, through comparison of the illusion upright and inverted and comparison with the hollow potato illusion (experiments 1, 2 & 3).

The other experiments reported in the thesis investigated effects of lighting on the recognition and matching of laser heads (cf. Bruce et al. 1991; see also figure 3.1). Laser heads are shaded representations of faceted facial surfaces and so especially appropriate for investigating the effects of lighting. Differences between lighting the face from above or below on the recognition of laser heads are tested and the relationship of the effects of lighting and viewpoint (experiments 5 & 6). Experiment 6 also included a manipulation of photographic polarity, the effect of which is also thought to be related to shadow and shading information (Phillips, 1972; see also section 1.6). The use of laser heads allowed investigation of the effects of polarity on the perception of shape without its confounding effects on the perception of pigmentation.

The experiments reported in Chapter 4 (experiments 7 - 12) examined subjects' abilities to match heads shown under different viewing conditions. How common structure is recovered from differing images was of particular concern. Differences between top and bottom lighting were again tested with upright (experiment 7) and inverted faces (experiment 8), together with the effects of less extreme changes in lighting that were either both from above (experiments 9 & 10) or both from below the face(experiments 11 &12). The matching task allowed investigation of perceptual effects independently from the effects of memory. In Chapter 5 control experiments are reported that were intended to test the generalizability of the effects of lighting found (experiments 13 - 15).

The final chapter attempts to integrate the results of the different sets of experiments and to summarise and explain the effects of lighting found. Image and edge based explanations are considered but an explanation within a surface based framework is offered as the most parsimonious explanation of the pattern of results reported.

2. CHAPTER 2: EFFECTS OF LIGHTING ON THE HOLLOW FACE ILLUSION

2.1 GENERAL INTRODUCTION

The work reported in this chapter is concerned with possible effects of lighting on the perception of a face as a face, a basic level category decision (Rosch, 1978). This is the type of decision modelled by most of the object recognition schemes discussed in Chapter 1. If categorisation can be achieved on the basis of simple outline or edge information, as many such schemes suggest, there would be no reason to expect effects of lighting. However, if shading information is more important for faces than for other classes of objects, as suggested in Chapter 1, this might be reflected by effects of lighting at the level of perceiving a face as a face. This was investigated using the hollow face illusion, where a hollow mask is wrongly seen as a convex face when viewed from beyond a certain distance (Gregory, 1973; see also figures 2.1 & 2.2). The illusion will be described more fully next followed by a description of the method used to assess its strength under different viewing conditions.

The illusion is an example of depth reversal occurring with a real object under dynamic, binocular viewing conditions. A mask or mold of a face appears as a normal convex hemispherical face with appropriate local convex and concave surface features when viewed from beyond a certain distance. It has been argued that it is the overall pattern of *shading*, a so-called pictorial cue, that helps reveal the face in such a mask (Bruce, 1988). The mask presents the observer with "a visual conflict between different sources of depth information" (Georgeson, 1979) and the experiments reported sought to provide cues as to how such conflicts are resolved. All the normal sources of depth information, including shading, motion, stereoscopic disparities and texture, are present but they are giving contradictory or ambiguous information. It is

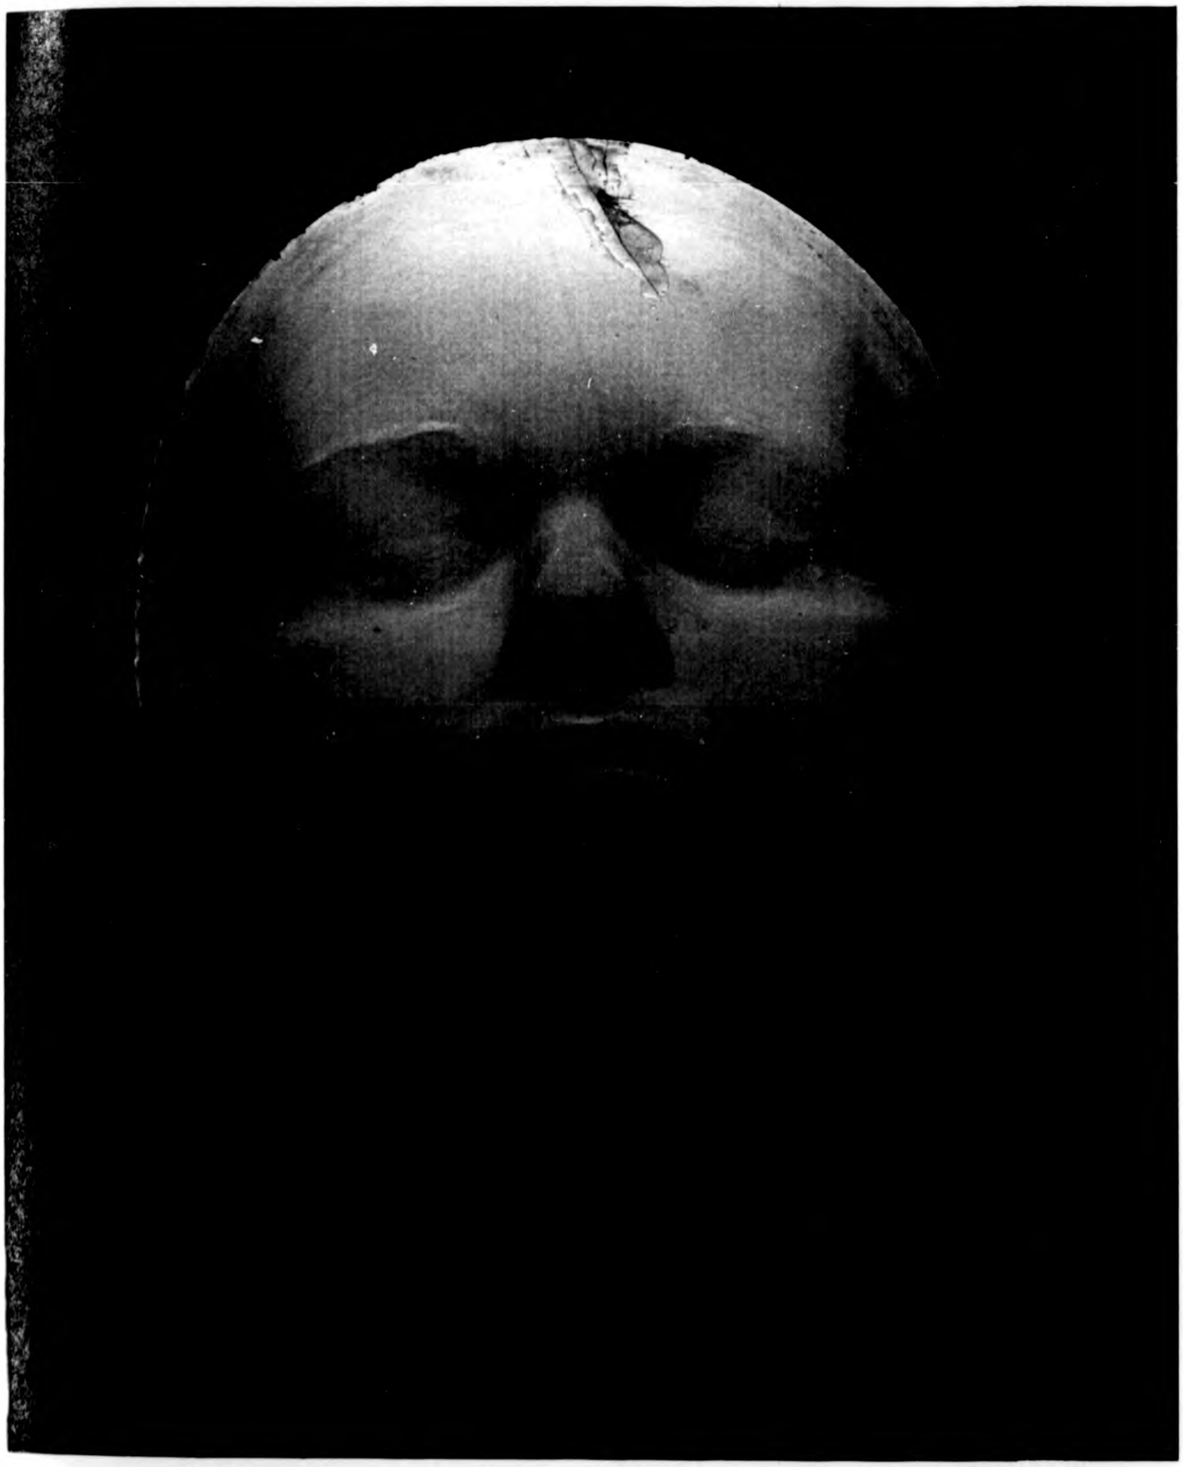

Figure 2.1: The hollow face used at Nottingham, lit from above and behind.

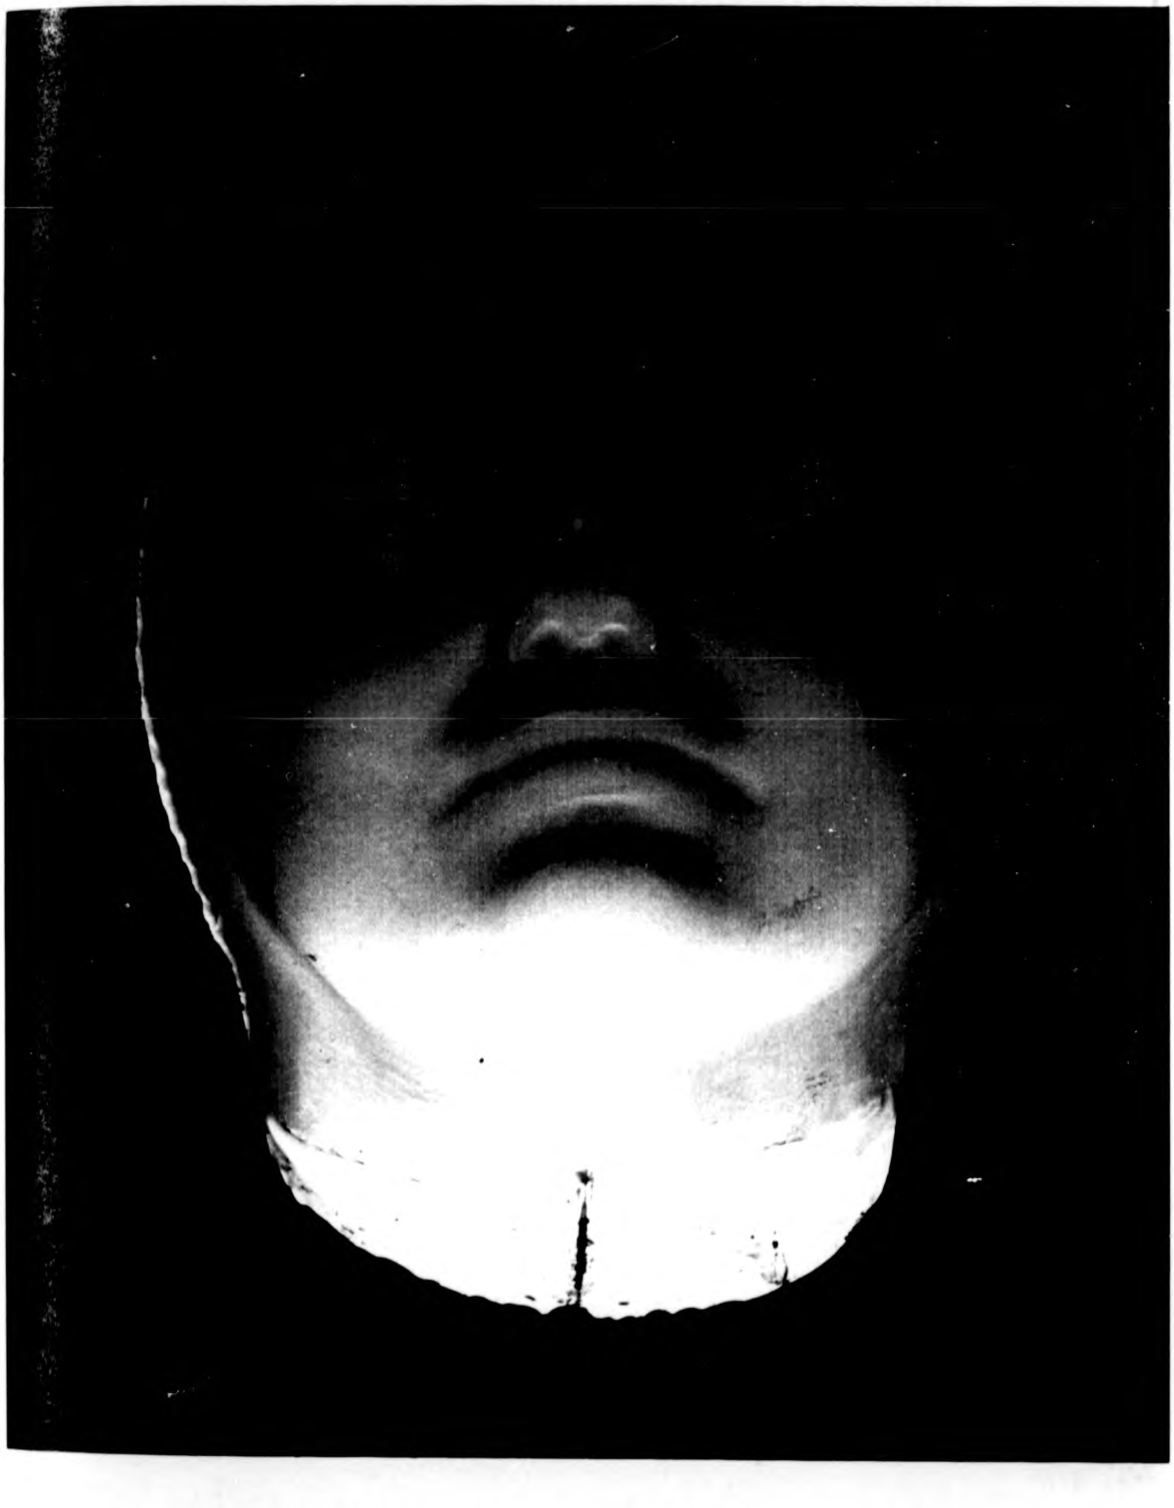

Figure 2.2: The hollow face used at Nottingham, lit from below and behind.

now commonly thought that such alternative sources of information are analysed by separate modules (eg. Marr, 1982) and investigation of the illusion may give clues as to the operation of separate modules as well as how information from the different sources is integrated to give a unitary percept.

The hollow face has been categorised as an illusion generated by ambiguity (Gregory, 1973). Shading itself is ambiguous because there are always two possible interpretations for a given pattern of shading, a convex surface apparently lit from one direction or a concave surface apparently lit from the opposite direction (Horn, 1977). In the case of the hollow face illusion, the pattern of shading is produced by a mask lit from the actual direction of the light source but could also be produced by a convex face apparently lit from the opposite direction.

The novel method that was developed to assess the strength of the illusion was based on the observation that at close distances the mask can be seen to be concave, but at greater distances the illusory percept is of a convex face. At some distance in between perception must change between the two solutions and the novel measure was the distance at which this change took place. On each trial subjects either approached or retreated from the mask, judging when their percept changed from the convex to the concave solution or vice versa respectively. This distance was assumed to be inversely related to the strength of the illusion; the larger the distance the weaker the illusion and vice verse. At close distances many cues to the actual concavity of the mask would become more pronounced, for example retinal disparities, so maintenance of the illusory convex percept at these short distances is consistent with a stronger illusion.

One aim of the first two experiments reported was to establish why it is that the concave mask is perceived as a convex face when viewed from beyond a certain distance, despite the availability of depth cues to the contrary. One explanation is that alternative perceptual hypotheses have associated probabilities, and that it is the familiarity or high probability of faces as convex objects that favours this as the interpretation of the ambiguous data (Gregory, 1973). An alternative explanation, not requiring reference to such object specific knowledge (cf. Marr, 1982), would be that the illusion results from a *general* preference for convexity in the interpretation of shading information. The existence of a "hollow potato" illusion where an irregular, *unfamiliar*, smooth, potato-shaped mould is seen as convex when viewed from beyond a certain distance, is consistent with this explanation (Johnston, Hill & Carman, 1992). Experiment 3 involved a direct comparison between the strengths of the hollow face and hollow potato illusions to determine what, if any, is the contribution of the familiarity of faces as convex objects to the strength of the illusion.

The two possible solutions of shape-from-shading, convex or concave, entail opposite lighting directions (Horn, 1977). When the concave mask is lit from in front and below the illusory percept is of a convex face lit from in front and above. However, if a row of bumps, convexities, lit from above are rotated 180° in the image plane which would normally also involve a rotation of apparent lighting direction, they are perceived as a row of dents, concavities, still lit from above rather than as bumps lit from below (Ramachandran, 1988a & 1988b; See also fig. 1.1). This is interpreted as evidence for a "light-from-above assumption" but seems at odds with a preference for convexity. It does suggests that the hollow face illusion would be stronger when the convex face appears to be lit from above rather than below. This was tested in experiments 1 & 2. That the illusion is still seen as a convex face even when this entails it appearing to be lit from below suggests that in this case, unlike the bumps, a convex interpretation can "win" over a light-from-above assumption.

If the illusion were solely the result of a general preference for convexity then no effect of the orientation of the mask would be expected. Changing orientation only affects the perception of the mask as a face, not convexity or concavity. However, if the illusion is dependent upon the nature of "the face" as familiar object (Gregory, 1973), then inversion might be expected to reduce the strength of the illusion, or eliminate it altogether, as upside-down faces are much less familiar objects. Inverting faces is known to have a disproportionate effect on many aspects of face processing (Yin, 1969; Valentine, 1988), probably because it disrupts configural encoding (Young et al, 1987; see Chapter 1). One would not expect any effect of orientation on the "hollow potato" illusion, a prediction that is directly tested in experiment 3.

If the fact that the object seen is an upright face is critical, an interaction between the effects of lighting and orientation might also be expected; there might be different effects of lighting when the face was upright from when it was upside-down. The object specific hypotheses that have been invoked to explain the hollow face illusion should hold "with any lighting" (Gregory, 1973), and should mean little or no effect of lighting direction. When the mask is upside-down a "light-from-above assumption" might be predicted to give an effect of lighting direction, with the illusion stronger when lighting appears to be from above. Previous research investigating effects of orientation and lighting on the hollow face illusion by Sakurai, Endo and Maruyama (1985) reports that highest "depth reversibility", their evidence for a weak illusion, was obtained for upside-down faces apparently illuminated from below the chin. This is consistent with an advantage for upright faces but not what would be expected from a light-from-above assumption as, in this condition, the lighting would be from above relative to the observer. Their experiments confounded the effects of lighting and orientation by including them as a single factor, stimulus category, making the individual effects of lighting and orientation and any interaction between them unrecoverable. The design used here allowed the effects of lighting and orientation, including any interactions, to be assessed independently using a clearer method than that of Sakurai et al (1985). The presentation of the mask as a real three-dimensional object also allows investigation of the illusion in a situation where all the non-pictorial cues to depth are also available.

As already mentioned, the novel measure we used to assess the strength of the illusion was the distance from the mask at which the subjects' perceptions changed between convex and concave. Subjects started by seeing the mask as a face and then approached it until they could see it as a mask, or started by seeing the mask and retreated until they could see the illusory face. All movement was perpendicular to the mask. It was predicted that subjects would tend to maintain whatever initial perceptual interpretation they made (ie. concavity when retreating from the mask and convexity when approaching it) with the result that the distance at which the change in perception occurred should be greater when retreating than when approaching the mask whatever the other conditions. In general the strength of the illusion in any condition was taken to be inversely related to the distance from the mask at which the change between convex and concave occurred; short distances implying a strong illusion and long distances a weak illusion. The smaller the distance, the stronger other cues like stereopsis and the effects of motion are (retinal disparities are larger and effects of movement greater), but perception of the illusion implies that the convex interpretation of shading is still dominating these other, unambiguous, cues to shape. This method was used to assess the strength of the illusion under different conditions which are described in the introductions to the individual experiments.

2.2 EXPERIMENTS 1 & 2

INTRODUCTION.

In the first two experiments the effects of orientation and lighting direction on the strength of the illusion, as measured by the novel method, were investigated. The illusion was presented either upright or inverted and lit from above or below. As explained more fully in the introduction, it was predicted that the illusion would be strongest when presented upright, due the familiarity of upright faces, and also that it would be stronger when the illusory convex interpretation was consistent with light appearing to be from above relative to the observer, consistent with a light-from-above assumption. However, if the illusion is the result of an object specific hypothesis "which holds with any lighting" (Gregory, 1973), then no effect of lighting would be expected at least when the face was upright. It was expected on the basis of the importance of object specific knowledge and on the basis of previous work, which reported the weakest illusion with the face inverted and lit from above (Sakurai et al, 1985), that he effect of inverting the illusion would be greater than the effect of changing lighting direction. Therefore the upright illusion lit from below was expected to be stronger than the inverted illusion even when this was lit from above.

Experiments 1 and 2 differed in that for experiment 1 light was from behind and through the semi-translucent mask (see materials section for a full description of the mask), while for experiment 2 it was from in front of the mask. When lit from behind, the illusion involves perceptual reversal of the apparent direction of lighting, from behind to in front - but actual top lighting still appears to come from above the head. When lit from in front the illusion involves a reversal from actual top to apparent bottom lighting and versa. It is when light *appears* to be from above relative to the observer, whatever its actual direction, that it was expected that the illusion would be stronger on the basis of a light-from-above assumption.

METHOD

Subjects.

Two sets of twenty-four subjects recruited by advertisement took part in experiments 1 and 2. All subjects had normal or corrected to normal vision and were tested to ensure that they had functioning stereoscopic ability (see procedure section). Subjects were paid one pound for their participation in an experiment which lasted approximately ten minutes.

Materials.

The hollow mask used for all experiments was made of semi-translucent monochrome plastic and was housed in a wooden cabinet with a glass front and back (Figures 2.1 and 2.2). The light source was an angle-poise lamp with a sixty watt bulb directed from above or below the mask. Distances were measured with a tape measure. Both experiments took place in a long windowless room with no lighting other than the angle poise illuminating the illusion.

A random dot stereogram of a spiral was used for testing stereopsis, taken from Frisby "Seeing" (1979, p. 81). Red/green spectacles were used to view it.

Design.

Experiments 1 & 2 were both 2(Orientation) x 2(Lighting) x 2(Movement) within subjects designs. The factors were the Orientation of the mask (Upright/Upside-down), the Lighting direction (Top/Bottom) and the subject's direction of Movement with respect to the mask (Approaching/Retreating). There was one trial in each cell giving a total of eight trials for each subject. For experiment 1 lighting was from behind (and therefore through) the mask while for experiment 2 it was from in front of the mask (through the glass). The dependent variable used was the distance away from the mask at which the subjects judged their perception to have changed between convex and concave.

Procedure.

Subjects were first tested to ensure that they had stereopsis, the ability to perceive depth via the fusion of disparate retinal images. Subjects were told that viewing the 17 x 18 cm image (Frisby, 1979) at a distance of approximately 0.5 with the red-green spectacles should enable them to see a three-dimensional spiral. They were asked to report whether the centre of the spiral appeared to be in front of or behind the page. Potential subjects who failed to report seeing the expected three-dimensional percept took no further part in the experiments. Remaining subjects were then shown the hollow face display in a room lit only by the angle-poise lamp used to illuminate the illusion and their task explained to them. They were told that the mask would be presented to them under various conditions and that for each condition two measures would be made; one where they would start by seeing the illusion, that is the hollow-mask as a convex face with the nose pointing towards them, and then approach it until they saw it as concave, and another where they would start by seeing the concave mask and retreat until they saw the convex face. They were asked to stop when this change had taken place so that their distance from the mask could be measured. Subjects were asked to be consistent across conditions in the subjective criterion they used for judging when the change between the convex and concave interpretations had taken place. It was further explained and demonstrated that the illusory percept of the face, but not the verdical percept of the mask, would appear to follow them when they moved from side to side and they were asked to use this as a check for when the change had occurred. It was requested that all other movement should be directly towards or away from the mask, that is perpendicular to it, though movement was not constrained.

The measure of the strength of the illusion was then taken under the different conditions which were presented in a different random order for each subject.

RESULTS.

Experiment 1: Back Lighting.

The mean distances, in centimetres, at which subjects judged their percept to have changed are shown for each condition in table 2.1 and, averaged across direction of approach, in figure 2.3. As explained in the introduction to this chapter, the larger the distance the weaker the illusion.

Analysis of variance was conducted on the results. With lighting behind the mask there were main effects of Orientation ($F_{1,23}=57.3$, $p<<0.05$), Lighting ($F_{1,23}=11.0, p<<0.05$) and direction of Movement ($F_{1,23}=56.9, p<<0.05$) with no significant interactions (All p's >0.1). Upright, top lit and approaching trials all gave smaller distances than their counterparts, upside-down, bottom lit and retreating trials.

Fig. 2.3: Mean Distances For Experiment 1.

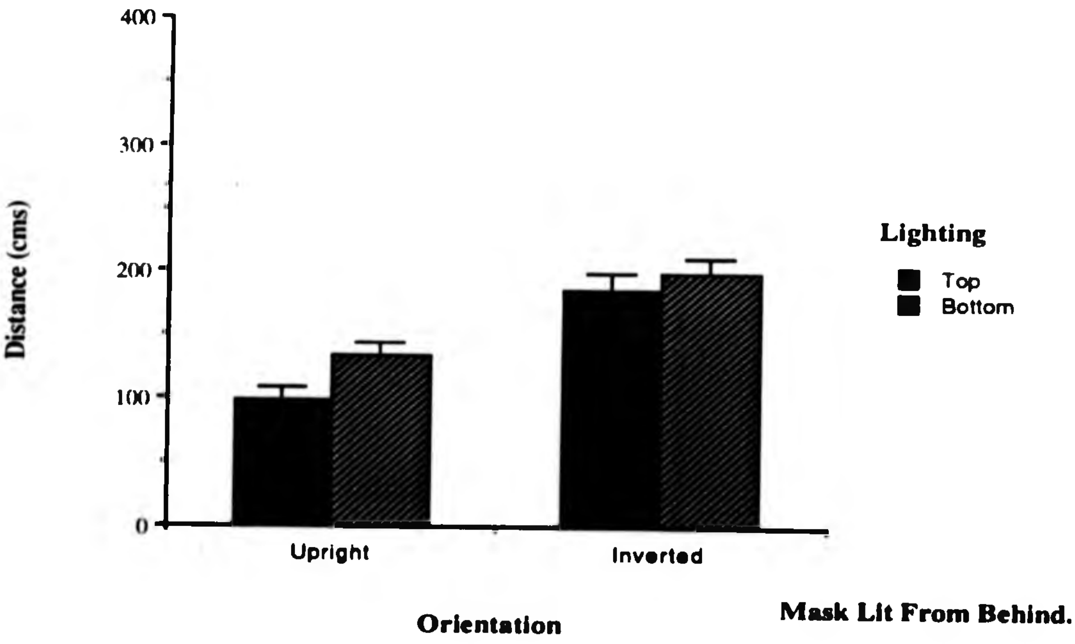

Table 2.1: Mean Distances, Experiment 1.

Mean Distance (cms)		ORIENTATION			
		Upright		Inverted	
MOVEMENT		App.	Ret.	App.	Ret.
LIGHTING	Top	71	127	140	230
	Bottom	96	172	163	236

App. - Approaching; Ret. - Retreating.

Experiment 2: Front lighting.

The corresponding means for experiment 2 are shown in table 2.2 and figure 2.4. The overall mean for experiment 2, with the light source in front of the mask was 223 cms, higher than the equivalent mean for experiment 1, 154 cms. This suggests that the illusion was more effective when lit from behind.

An ANOVA gave a main effect for Orientation ($F_{1,23}=92.1, p<<0.05$) and an interaction between Lighting and direction of Movement ($F_{1,23}=8.3, p<0.05$). The interaction between Orientation and direction of Movement just missed significance ($F_{1,23}=3.5$, $p=0.075$). Upright and approaching trials both again gave smaller distances than upside-down and retreating trials. Apparently top lit trials, though actually bottom lit, resulted in smaller distances than apparently bottom lit trials, though analysis of simple main effects showed this difference was significant only when subjects were retreating ($F_{1,46}=7.3$, $p<0.05$). The interaction between lighting and movement reflects the absence of any effect of lighting for approaching trials. It is not clear why this should be.

The significant interaction between lighting and movement does not have a obvious explanation, nor did it have a parallel in experiment 1. The other interaction between orientation and movement, which was not quite significant, may have been an artifact of scaling as in more familiar configurations of the mask distances are smaller and there is less room for approaching and retreating trials to differ. Neither of the interactions involving both lighting and orientation approached significance (F's < 1).

Fig. 2.4: Mean Distances For Experiment 2.

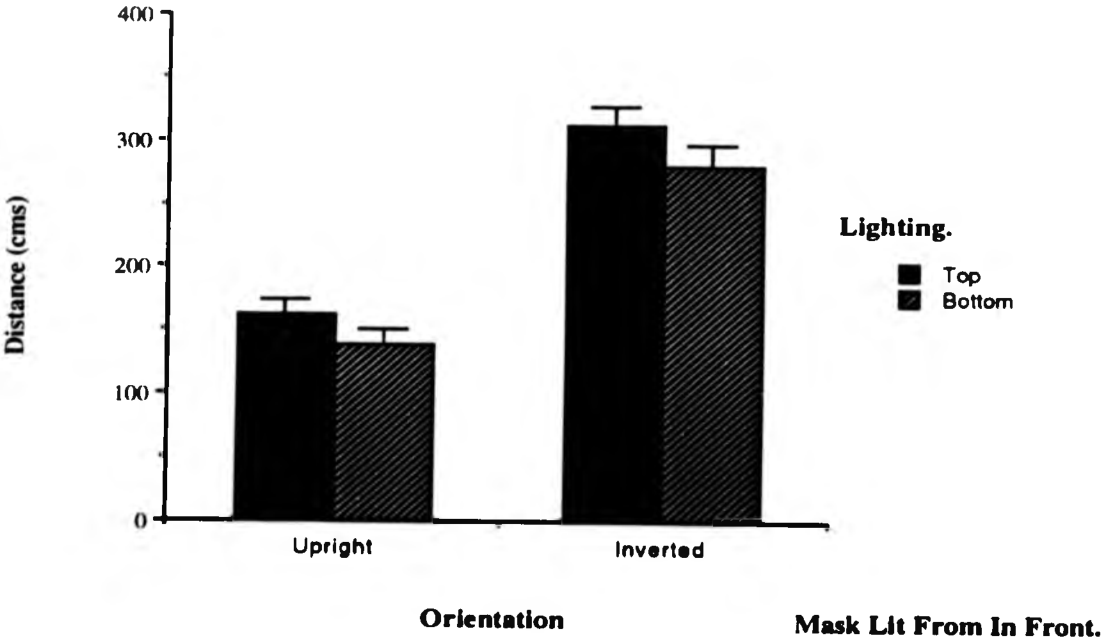

Table 2.2.Mean Distances. Experiment 2.

Mean Distance (cms)		ORIENTATION Upright		Inverted	
MOVEMENT		App.	Ret.	App.	Ret.
LIGHTING	Top	116	209	245	378
	Bottom	110	167	242	318

DISCUSSION.

In both experiments 1 and 2 there were independent effects of orientation and of lighting. The illusion was stronger when the face appeared upright and lit from the top, consistent with a light-from-above assumption and the importance of knowledge of upright faces.

If the familiarity of faces was the only factor, the illusion might not have worked when the mask was upside-down. That it did, although less effectively, suggests that the availability of an object specific hypothesis is not essential for the illusion (cf. Gregory, 1973). It is possible that the hypothesis of an upside-down face would be sufficient, a possibility further investigated in the next experiment. There an unfamiliar object, the "hollow potato" (Johnston et al, 1992), is used for which there would be no available object specific hypothesis. However, the consistent effect of orientation in both experiments, with the illusion less strong when the mask was inverted, suggests that the familiarity of upright faces is an important factor in determining the strength of the illusion and may be important in explaining its existence.

The existence of the illusion when the mask is upside-down is consistent with a general preference for convexity as a solution to shape-from-shading. However, a convexity preference alone would not be consistent with the effect of orientation reported as it would apply equally to the upright and inverted orientation. When images of bumps lit from the top are inverted they are seen as dents also lit from the top (Ramachandran, 1988a & 1988b), reflecting a light-from-above assumption but apparently violating any convexity preference. However, the inverted hollow face looks convex even when it also appears lit from below. Thus whether a convexity preference or light-from-above assumption 'wins' when they are in conflict appears to depend on the situation and perhaps the stimulus type and neither are absolute.

Object specific hypotheses were also thought to be lighting independent but there were effects of lighting both when the mask was upright and when it was upside-down. If a lighting independent object specific hypothesis was the sole cause of the illusion (Gregory, 1973) no effect of lighting on the upright illusion would have been expected. The finding suggests that the knowledge of faces, or access to that knowledge, is lighting dependent, being disrupted by bottom lighting. The result is consistent with a "light-from-above" assumption being used in the interpretation of shape-from-shading (Ramachandran, 1988a & 1988b) and shading information being critical to the generation of the illusion. The top lit advantage for the inverted condition is inconsistent with the result reported by Sakurai et al (1985) who found highest reversibility, implying a weak illusion, with upside-down stimuli lit from the chin, that is lit from above relative to the observer. It should also be noted that in experiment 2, when lighting was from the front, apparent lighting direction is opposite to actual lighting direction and it was when lighting *appeared* to be from above that the illusion was stronger, even though it was actually from below.

With both experiments the effect of direction of movement of the observer relative to the mask was as expected. The distance at which perception changed was significantly higher when retreating from the mask than when approaching it. This is consistent with a subject's initial perceptual interpretation having some stability. The results imply a range of distances from the mask for each condition over which *both* of the two interpretations, convex and concave, are made, depending on the direction of movement. This is an hysteresis like effect. Within this range the image formed at the retina must be the same whether the subject is 'seeing' the concave mask or the convex face. A similar hysteresis like effect is also found with stereopsis (Fender and Julesz,1967) and here as there it could be the result of a cooperative algorithm or of the operation of a short term memory buffer (Marr, 1982). At a different level of explanation the effect would be consistent with a hypothesis testing model of vision with whatever hypothesis was initially accepted having a certain stability.

In experiment 2 the direction of movement factor was not completely independent of the other two factors. There was a significant interaction of direction of movement with orientation and a marginal interaction of direction of movement with lighting. Reasons for these interactions are not clear. What is more important is that there were no interactions involving both lighting and orientation.

Overall the patterns of results between experiments 1 and 2 are consistent, with the exception of the marginal interactions just mentioned. The overall mean for experiment 1 (154 cms) was lower than that for experiment 2 (223 cms). With the mask and lighting used, the illusion was more effective when lit from the back than when lit from the front. The mask available was housed in a display cabinet with a glass front and the specular effects this introduced may have interfered with the illusion. Specular effects are known to be important to the visual system's interpretation of surface orientation (Todd & Mingola, 1983) and specular highlights can force an ambiguous convex/concave figure to change state (Blake and Bulthoff, 1990). When a concave surface is lit from the front, there are also greater possibilities for cast shadows and the effects of mutual illumination to disambiguate depth as convex. These may be the reasons why the illusion was less effective when the mask was lit from the front compared to when it was lit from the back and may also account for the interactions with direction of movement found only in experiment 2, though how remains unclear.

2.3 EXPERIMENT 3

INTRODUCTION.

This experiment involved a direct comparison between the hollow face illusion (see fig. 2.5) and a "hollow potato" illusion (see figs. 2.6 & 2.7) in order to further investigate the role of face specific knowledge in generating the illusion. The effect of orientation reported for the previous experiments suggests that the familiarity of faces as upright convex objects can make the illusion stronger but that even the inverted mask is seen as a face from beyond a certain distance. However, even inverted faces are clearly faces and this knowledge may be contributing to the generation of the illusion under these circumstances.

Alternatively, the possibility of a general convexity preference in the interpretation of shape-from-shading may account for the inverted illusion. Evidence for this includes the existence of a "hollow potato" illusion, where a unfamiliar, smooth, concave surface with local bumps and dents (which is not actually a potato!) can appear convex when photographed or viewed from beyond a certain distance (Johnston et al, 1992). For such an object there would be no object specific knowledge about its shape available and the simplest explanation of the illusion therefore seems to be a general bias to perceiving objects defined by shading as convex.

This experiment sought to assess the role of object specific knowledge by comparing a situation where such knowledge was readily available, the hollow face, with one where there was no such knowledge, the hollow potato. The orientation effect found on the previous experiments suggests that such knowledge does strengthen the illusion, and therefore that the hollow potato illusion should be less strong than the hollow face illusion. If inverting the face disrupts important aspects of face processing (Carey & Diamond, 1977; Young, Hellawell & Hay, 1987) then the hollow potato illusion might be expected to be as strong as the inverted hollow face, assuming that object specific knowledge is not available in either case.

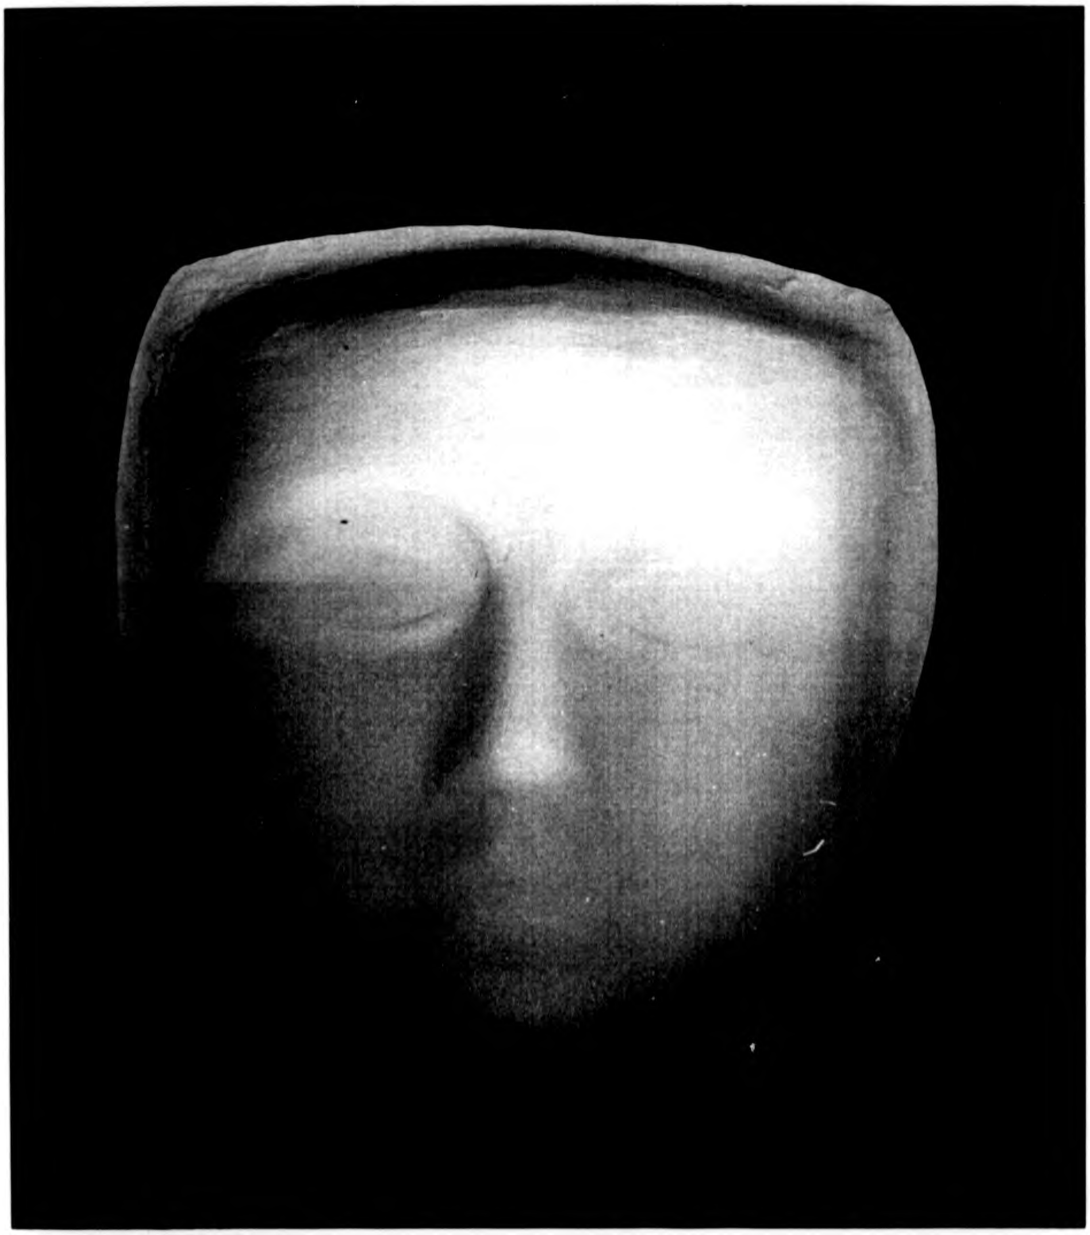

Figure 2.5: The hollow face used at Stirling as it was lit for experiment 3.

Figure 2.6: The "upright" hollow potato.

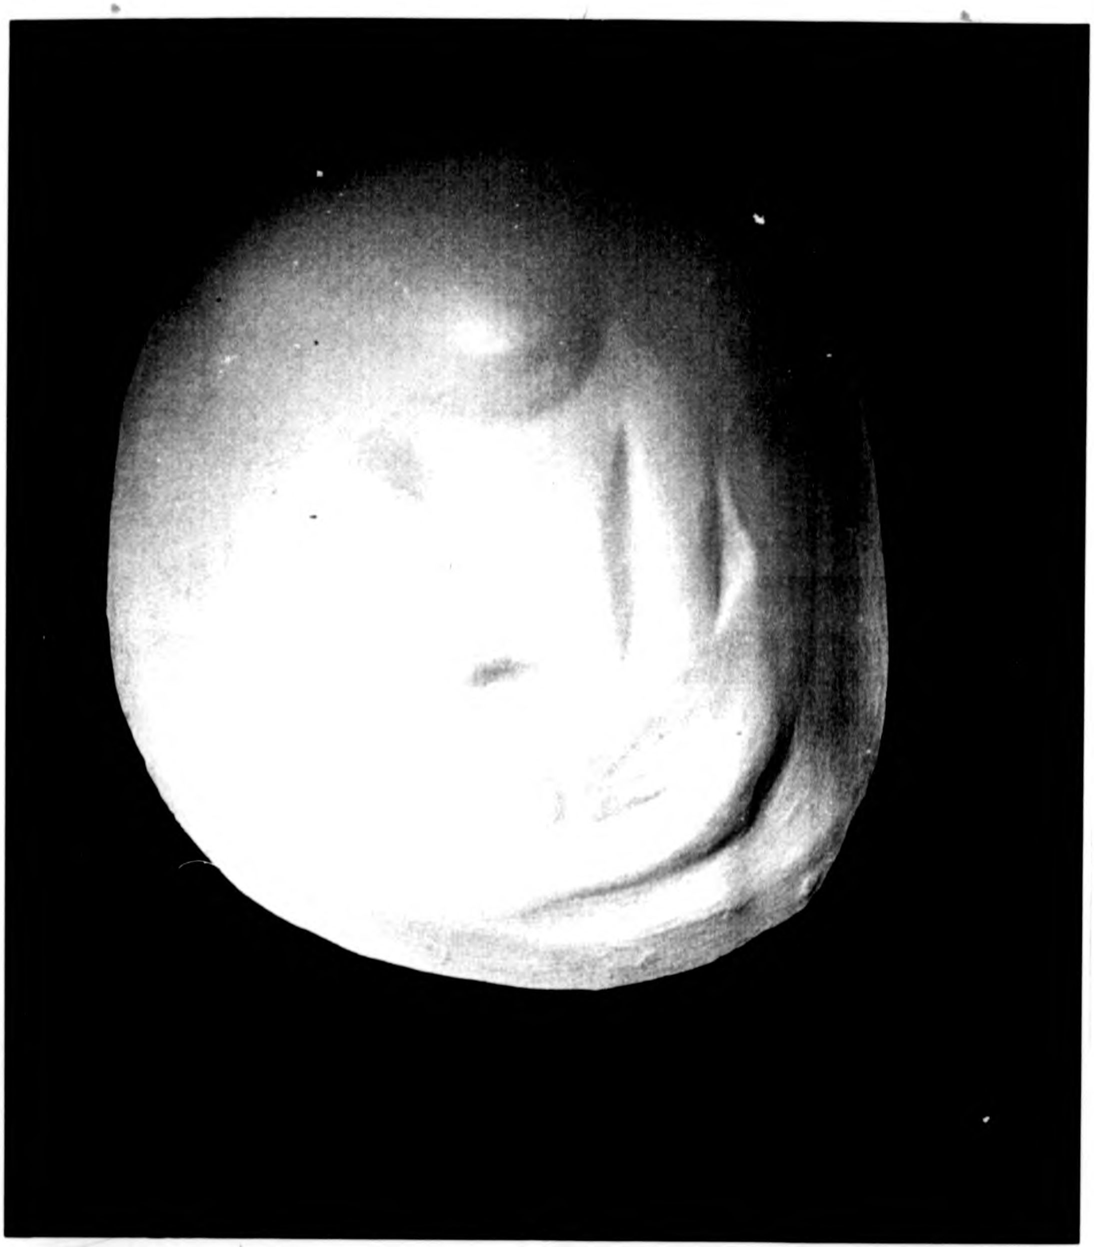

Figure 2.7: The "inverted" hollow potato.

The inverted hollow face and hollow potato should also produce equivalent effects of movement. It was further predicted that, as the "upright" and "inverted" hollow potato are equally unfamiliar, the hollow potato should show no effects of orientation.

Finally, in this experiment two trials were included per cell. This served as test for order effects if, for example, the increasing familiarity of the potato affected the illusion and as a general test of the reliability of the measure. The experiment was also conducted at a different university and with a different mask and lighting thus testing the generalizability of the method. In particular the mask was made of different, non translucent material, and had to be lit from in front though it had no glass front. These differences are described more fully in the methods section.

METHOD.

Subjects.

Twelve subjects were recruited from around the University of Stirling and were paid one pound for their participation. All subjects had normal or corrected to normal vision and were tested for stereopsis in the manner described for experiments 1 and 2.

Materials.

The hollow face and the hollow potato used for the experiment were made at the University of Stirling for the purposes of the experiment (see figures 2.5, 2.6 and 2.7). The hollow face was produced by taking a mold of a mask using molding plaster. The same plaster and casting method was used for the hollow potato, the "potato" itself having been made from plaster and newspaper. The potato was no specific shape but a convex hemisphere of the same size as the face. It had arbitrarily distributed concave and convex features of approximately the same depths and heights as features found on the face. This was to ensure that the degree of depth reversal required to see the illusory percept was as great for the hollow potato as for the hollow face. One orientation of the potato was labelled upright as a few people reported perceiving a "potato face" in that orientation. This accidental face like configuration was not intentional and reflects the general tendency to perceive face like structures in many smooth objects including clouds and rocks. Any significant effect of this apparent faceness would be reflected by an effect of orientation on the hollow potato illusion.

All other details of the illusion were made as similar as possible. Both were painted matte white when completed and housed in boxes with black surrounds. Both were presented at approximately head height and both were lit by the light from a slide projector directed from slightly below and to the right of the illusion. Due to the reversal of apparent lighting direction the illusory percepts appeared lit from above and to the left. The light source was approximately 2.5 m from the mask.

Design and Procedure.

Testing took place in a long room with windows blocked out by dark blinds. Other details of the procedure were the same as for previous experiments.

The experiment was a 2(Illusion) x 2(Orientation) x 2(Movement) repeated measures design. The levels of Orientation and direction of Movement were the same as for the other experiments. The factor Illusion refers to which illusion was used in the condition, the hollow potato or the hollow face. Other details of the design were the same as for previous experiments, except that each cell contained two trials. The order of trials, including whether it was the first or second measure in any particular cell, was randomised for each subject. Whether each measure was the first or second in any particular condition was recorded so that it could be included as a factor, Order, with two levels, first or second, in an initial analysis. This allowed the reliability of the measuring method to be assessed but was not included as part of the main design as it did not bear on the relationship of the hollow potato to the hollow face illusion.

RESULTS.

As can be seen from figure 2.8, which shows the mean distances collapsed across the Movement and order factors, the upright hollow face gave the strongest illusion with inverted hollow face and either orientation of hollow potato equivalent in strength. Table 2.3 presents means distances for all conditions. There was clearly the expected effect of Movement, with retreating measures greater than approaching measures.

Initially results were analysed with the order factor included but, as there were no significant effects of this factor or any interactions involving it (all p's > 0.1), this analysis is not reported.

The 2(Illusion) x 2(Orientation) x 2(Movement) analysis gave a significant three-way interaction ($F_{1,11}$= 12.0, $p<<0.05$). This seems largely due to the effect of Movement being suppressed for the upright face condition, perhaps because performance was approaching ceiling - distances were small so there was less room for approaching and retreating measures to differ. Analysis of simple interactions showed a Movement x Orientation interaction for the face ($F_{1,22}$= 17.8, $p<<0.05$) but not for the potato ($p>0.1$). Consistent with the explanation offered for the three-way interaction there were simple simple main effects of Movement for the inverted face ($F_{1,44}$= 39.7, $p<<0.05$), the upright potato ($F_{1,44}$= 59.7, $p<<0.05$) and the inverted potato ($F_{1,44}$= 50.5, $p<<0.05$) but not for the upright face ($F_{1,44}$=3.5, $p=0.068$).

There were simple simple main effects of Orientation for both approaching ($F_{1,44}$= 28.3, $p<<0.05$) and retreating ($F_{1,11}$= 101.7, $p<<0.05$) measures for the hollow face, consistent with the effect of orientation found on the previous experiments using the hollow face. There was no simple simple main effect of orientation for the hollow potato in either approaching or retreating conditions ($p>0.1$).

In order to test whether the hollow potato in either orientation was equivalent to the inverted face a further analysis was carried out with just these three conditions. The three conditions were included as a single three level factor, Type, together with the Movement factor. The analysis showed a main effect of Movement ($F_{1,11}$= 61.8, $p<<0.05$) but no effect of Type or any interaction involving this factor (p's >0.1). This confirmed that the hollow potato produced equivalent effects to the inverted hollow face.

Fig. 2.8: Mean Distances For Experiment 3.

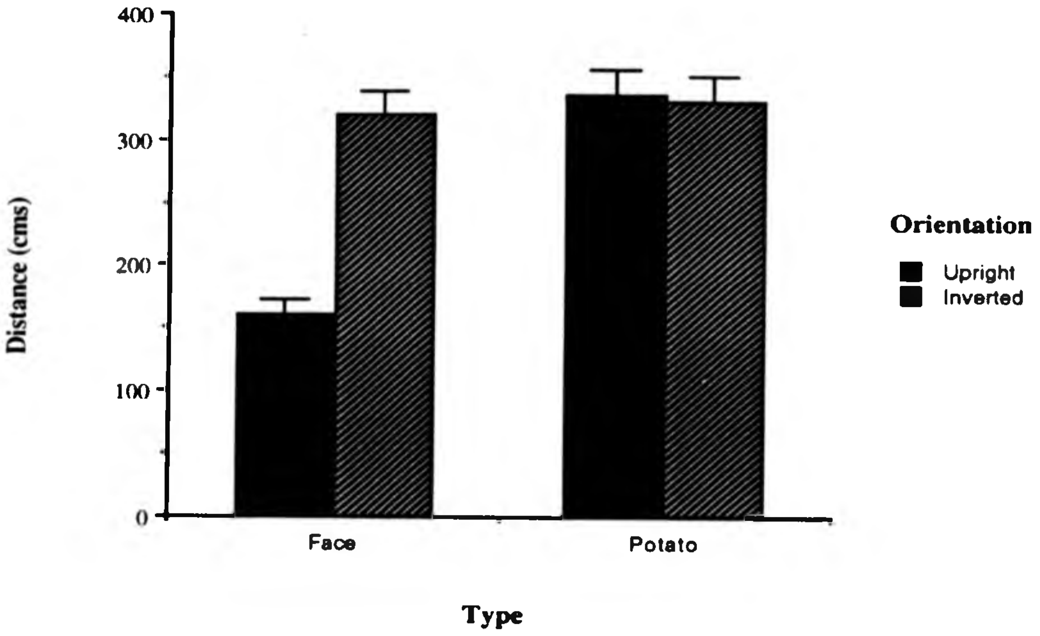

Table 2.3: Mean Distances For Experiment 3.

Mean Distance (cms)		ORIENTATION			
		Upright		Inverted	
MOVEMENT		App.	Ret.	App.	Ret.
TYPE	Face	145	176	254	383
	Potato	256	414	258	414

DISCUSSION.

The results showed that the upright hollow face gave a stronger illusion than either the inverted hollow face or either orientation of the hollow potato. This is again consistent with face specific knowledge contributing to the strength of the illusion. The hollow potato illusion was found to behave like the inverted hollow face, confirming the use of the latter as the control condition for the effects of object specific knowledge in the previous experiments. An effect of orientation was only found for the hollow face illusion, consistent with object specific knowledge strengthening the illusion but only being available when it appears as an upright face. The parallel between the hollow potato and inverted hollow face suggests that both would show the same effects of lighting direction as reported for the inverted face, although this remains to be tested.

The existence of the hollow potato illusion is cited as evidence for a convexity preference in the interpretation of shape-from-shading (Johnston et al, 1992). There seems no other reason the hollow potato should appear convex, particularly as other sorts of information including stereopsis are unambiguously consistent with the concave interpretation. Shading information makes available the convex interpretation, due to its inherent ambiguity (see Chapter 1, section 1.7), and the perception of the potato as convex suggests there may be a bias towards the convex interpretation. The illusion again suggests the general importance of shading as a source of shape information, and how it can override other sources of such information.

This experiment also showed that the method developed in the other experiments transferred to new materials and a new population of subjects. Two measures were also made in each condition for this experiment and an analysis showed that there were no differences between the first measure and the second. A single measure, as used in the previous experiments, seems sufficient to give reliable differences between conditions. A possible problem with the method, one that may also have been reflected in the results of experiment 2, was that in conditions where the illusion is strong differences may be suppressed as the lower limit on distances begins to be important.

2.4 EXPERIMENT 4

INTRODUCTION.

The fourth experiment sought to assess the role of stereoscopic depth information in the perception of the mask. The "visual conflict between different sources of depth information" (Georgeson 1979) is between so-called pictorial cues, which are often depth reversible, and other cues including binocular stereoscopic depth information. The nature of stereopsis means that the disparities between the observer's two retinal images are not ambiguous and, in this case, are only consistent with the concave solution; yet the mask is seen as convex except at very close distances. The role of stereoscopic depth information in perception of the mask was investigated by comparing monocular viewing, where no stereoscopic information would be available, to binocular viewing. This experiment was carried out at Nottingham University with the same hollow face as used for experiments 1 and 2.

The precise role of stereoscopic depth information in perception of the mask is unclear. It has been demonstrated that the perception of the hollow face illusion cannot be explained by the complete loss of stereoscopic depth information nor by the pseudoscopic combination of retinal disparities with their signs reversed (Yellot & Kaiwa, 1979). It is known that information from both eyes is used because when an observer is facing the illusion the nose on the illusory face points straight ahead and not towards one eye, as it would if the information from only one eye were being used. A remaining possible interpretation is that of "fusion without stereopsis" , that is that the images from the two eyes can be combined into a unitary percept but without the depth implied by the magnitude of their disparities (Yellott and Kaiwa, 1979).

The point at which stereoscopic depth information 'wins' over the pictorial cues might still be dependent upon the magnitudes of the disparities between the retinal images reaching a threshold beyond which they can no longer be fused without their depth implications being realised. If this is so then one would expect subjects to see the mask as concave over a shorter range when viewing monocularly, as stereoscopic depth information is not available to disambiguate the illusion. Other cues, perhaps motion parallax, self occlusion or mutual illumination, must fulfil this role.

As well as comparing monocular to binocular viewing, upright and upside-down orientations were again included. These two orientations differ trivially in terms of stereoscopic depth information and so any effect of orientation would be expected to be independent of viewing condition. A replication of the effect of direction of movement was also predicted.

METHOD.

Twelve subjects were run in this experiment using the same method as for experiment 1 and 2 except for the following details. All lighting was from above and behind the cabinet as this had previously been found to give the strongest illusion. The lighting variable was replaced by a new factor, called Eyes with levels monocular (mono) or binocular (bino). The design was therefore a 2(Orientation) x 2(Eyes) x 2(Movement) within subjects design with the levels of the other factors as for experiments 1 and 2. Subjects wore an eye-patch covering their non-dominant eye in the monocular condition.

RESULTS.

The mean distances for experiment 4 appear in Table 2.4, and figure 2.9 shows means collapsed over direction of movement.

Analysis of variance on these distances gave independent main effects of Orientation ($F_{1,11}=19.1$, $p<<0.05$), Eyes ($F_{1,11}=14.7$, $p<<0.05$) and direction of Movement ($F_{1,11}=7.8$, $p<0.05$). The illusion was stronger for monocular and upright conditions and distances were again smaller when subjects were approaching than when they were retreating.

There were no significant interactions though Orientation x Movement was close ($F_{1,11}=4.1$, $p=0.067$). As with experiment 2 this non significant interaction was probably an artifact of scaling. All other F's were less than 1.

Fig. 2.9: Mean Distances For Experiment 4.

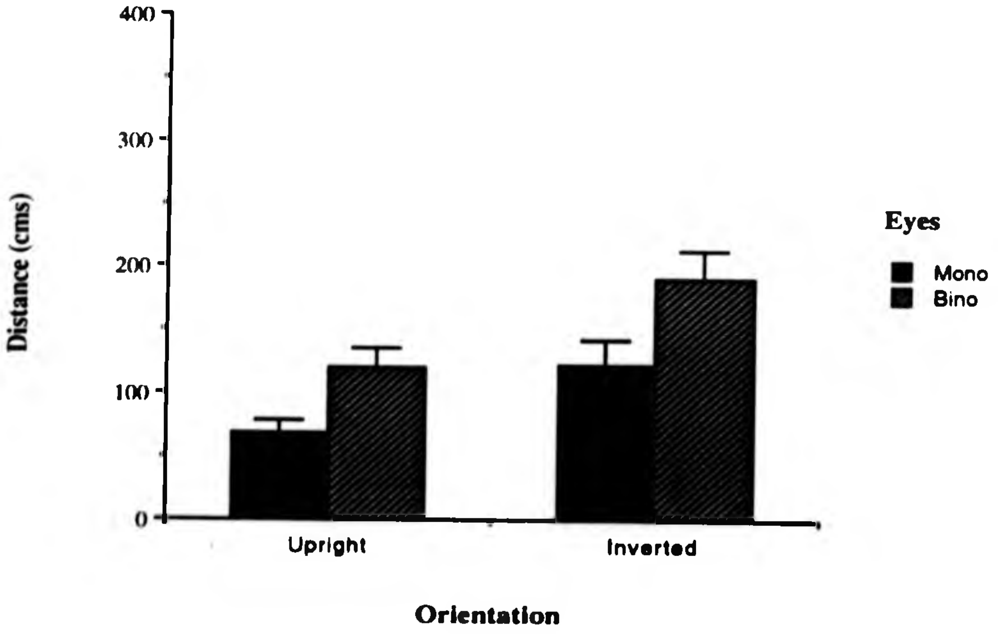

Table 2.4: Mean Distances, Experiment 4.

Mean Distance (cms)		ORIENTATION			
		Upright		Inverted	
MOVEMENT		App.	Ret.	App.	Ret.
EYES	Mono	49	87	84	161
	Bino	101	140	159	221

DISCUSSION.

The effects of orientation and direction of movement shown in experiment 4 were consistent with those obtained in experiments 1 and 2 and the explanations offered there apply. In addition, experiment 4 gave an independent effect of viewing condition with the illusion stronger when viewed monocularly than when viewed binocularly.

This difference between monocular and binocular viewing suggests that stereoscopic depth information is important in resolving the illusion at close distances. In the monocular condition, when no disparity information was available, monocular motion cues, the effects of mutual illumination or pictorial cues must have served this function. The effect of viewing condition suggests that stereoscopic depth information, when available, is the most effective cue to concavity as a solution at short distances.

When the illusion is viewed binocularly information from both eyes is incorporated but the magnitudes of disparities are not used to determine perceived depth, or else it would appear concave. The results reported are consistent with the suggestion that the mask may be perceived as convex until the retinal disparities, as determined by the distance from the illusion, reach a threshold of magnitude beyond which they can no longer be fused without their depth implications being realised.

The effect of monocular compared to binocular viewing reported here further justifies the use of real rather than screen based presentation of the illusion (cf Sakurai et al 1985). Screen based presentation restricts one to investigation of the effects of pictorial depth cues. The demonstration, in experiment 4, of the importance of stereoscopic depth information highlights this limitation.

2.5 GENERAL DISCUSSION

The experiments reported have shown effects of orientation, lighting direction, number of eyes used and direction of observer movement on the strength of the hollow face illusion as measured by the distance from the mask at which its appearance changed between the convex and concave interpretations. The effects of lighting and orientation were found to be independent of each other in experiments 1 and 2 as were the effects of orientation and viewing condition in experiment 4. Experiment 3 showed that the hollow potato behaved in the same manner as the inverted hollow face, suggesting that in neither case was object specific knowledge contributing to the illusion.

These results provide clues about the visual system's interpretation of shape-from-shading and the relationship of shape-from-shading to other depth cues in determining the perceived depth and orientation of surfaces and objects. These issues are discussed within the framework of the two-and-a-half dimensional (2 1/2-D) sketch, a description hypothesized as representing the orientations and distances of visible surfaces relative to the viewer (Marr, 1982). The proposed inputs to the 2 1/2-D sketch are the outputs of independent visual processing modules including shape-from-shading, stereopsis, motion, texture and contour analysis. Such a representation would be the last stage of pure data-driven perception before a final stage of perceptual interpretation which might involve accessing knowledge of objects.

In general the independence of the effects of lighting and viewing condition is consistent with a view that posits independent parallel modular processing of the information encoded at the retina without reference to object specific knowledge. The effect of lighting would be on a shape-from-shading module, while viewing condition (binocular/monocular) would affect the recovery of shape-from-stereopsis. The effect of lighting is consistent with a light-from-above assumption in the interpretation of shape-from-shading. However, the hollow face illusion, in either orientation, can still be perceived as a convex surface even when this entails it appearing to be lit from below suggesting such an assumption is not absolute. The existence of the hollow potato and inverted face illusions also suggests that there is a convexity preference or bias in the interpretation of shape-from-shading.

Independence of processing by different modules like shape-from-shading requires consideration of the problem of how such modules are combined to give a percept of the world that incorporates all the information available. The complicated nature of this combination of cues is evidenced by the way in which the output of shape-from-shading appears to be able to provide the input for apparent motion (Ramachandran, 1988a & 1988b). Alternating the positions of a sphere and a hollow defined by shading of the type shown in figure 1.1 can give rise to apparent movement of the sphere. With the hollow face changes in shading due to sideways movement relative to the mask results in apparent movement of the illusory face; the face appears to follow the observer as they move to the left or right. The illusion arises out of a *conflict* between different sources of information (Georgeson, 1979), perhaps as processed by different modules. In order to give rise to a consistent unitary percept including all the different available information the different modules must **constrain** each other's solutions (Barrow & Tenenbaum, 1981). The effects produced by lighting and the number of eyes used further suggest that the interaction between modules is not fixed but can vary according to the availability and associated uncertainty of the outputs of the parallel modules.

The effect of direction of observer movement, whereby the initial perceptual interpretation made by subjects seems to be maintained, must also be affected by the way in which the different sources of orientation and depth information are combined. Information consistent with the initial percept may be used to maintain the initial solution. When a change between the convex and concave interpretations occurs, with its implications for perceived distance, it may be reflecting a new solution to the combination of different modules becoming the most probable.

The independent effects of orientation reported and the comparison between the hollow face and the hollow potato in experiment 3 is consistent with face specific or top-down information acting as another constraint on the solution. This would explain why the convex solution is maintained at closer distances when the appearance is of an upright face. The familiarity of upright faces may also act to produce a bias towards this as the initial interpretation. Once any particular initial interpretation of the illusion has been accessed, it may feed back to constrain the solutions of the processing modules, consistent with the effect of direction of observer movement.

The results of experiment 3 that the hollow potato produces an illusion of the same strength as the inverted hollow face is consistent with no object specific knowledge being available in either case. Both also showed the same effect of observer movement. This apparent equivalence suggests that the hollow potato would show the same effects of lighting and monocular or binocular viewing condition as have been shown for the inverted hollow face.

Finally the results suggest a convexity preference which can dominate the light-from-above assumption, as the face mask appears convex even when this requires it to appear to be illuminated from below, and even when the mask is inverted. However, evidence from flat patterns of shading that are interpreted as dents when this is consistent with light appearing to be from above (Ramachandran, 1988a & 1988b) shows that any convexity preference is also not absolute and can 'lose' to a light-from-above assumption in some circumstances. The result here does suggest that Ramachandran's (1988a & 1988b) stimuli might be most effective when perceived as bumps illuminated from above above compared to when they appear as dents or as illuminated from below. This prediction could be tested using computer based displays of the type used by Ramachandran if methods were developed to measure their degree of depth reversibility (cf Sakurai et al 1985).

Experiments 1 and 2 showed that whether the illusory face appeared lit from above or below was important in determining the strength of the illusion. This effect probably reflects the importance of shape-from-shading for face perception and the use of a light-from-above assumption in the interpretation of this information. If seeing the illusory face entails lighting appearing to be from below, the face interpretation appears to become less likely. The work in the next chapter investigated whether this effect of lighting on seeing the face as a face generalizes to the ability to distinguish between faces. If bottom lighting disrupts the interpretation of shape-from-shading and this information is important for distinguishing between faces, it would be expected to affect the ability to distinguish faces lit from below.

3. CHAPTER 3: EFFECTS OF LIGHTING ON THE RECOGNITION OF LASER HEADS

3.1 GENERAL INTRODUCTION

In the last chapter a clear difference between top and bottom lighting on the perception of the hollow face illusion was shown. This suggested that lighting may be important to the processes involved in seeing a face as a face. A similar difference between these lighting directions has been reported on the recognition of photographs (Johnston, Hill & Carman, 1992) suggesting that lighting can also affect the ability to recognise particular, familiar, faces. In this chapter experiments are reported on the recognition of "laser heads" (eg. see figure 3.1) under conditions of top and bottom lighting to investigate the basis of this difference.

Laser heads are shaded representations based on a database of three-dimensional coordinate points derived by laser scanning the surface of a head as described more fully in the general methods section below. For laser heads, as with the hollow face, much of the information about their shape is only apparent from the overall pattern of shading. The perception of both types of stimuli may inform about aspects of the visual processing of faces, including the possible importance and use of lighting dependent shading information. Laser heads derived from different individuals, unlike the hollow face, contain individuating information about the person from whom they were derived (see Appendices B & C). This allowed the effects of lighting on the visual processes involved in within category distinctions, recognising individual faces, to be investigated as reported here.

Recognition requires that laser heads be distinguished from each other, rather than just all seen as convex, despite changes in lighting and viewpoint. If lighting, in particular, affects shading information this may be used to distinguish between individual faces as well as for conveying the overall shape of the face. A previous study using laser heads showed that laser heads can be recognised at levels that are above chance but far from perfect (Bruce, Healey, Burton, Doyle, Coombes &

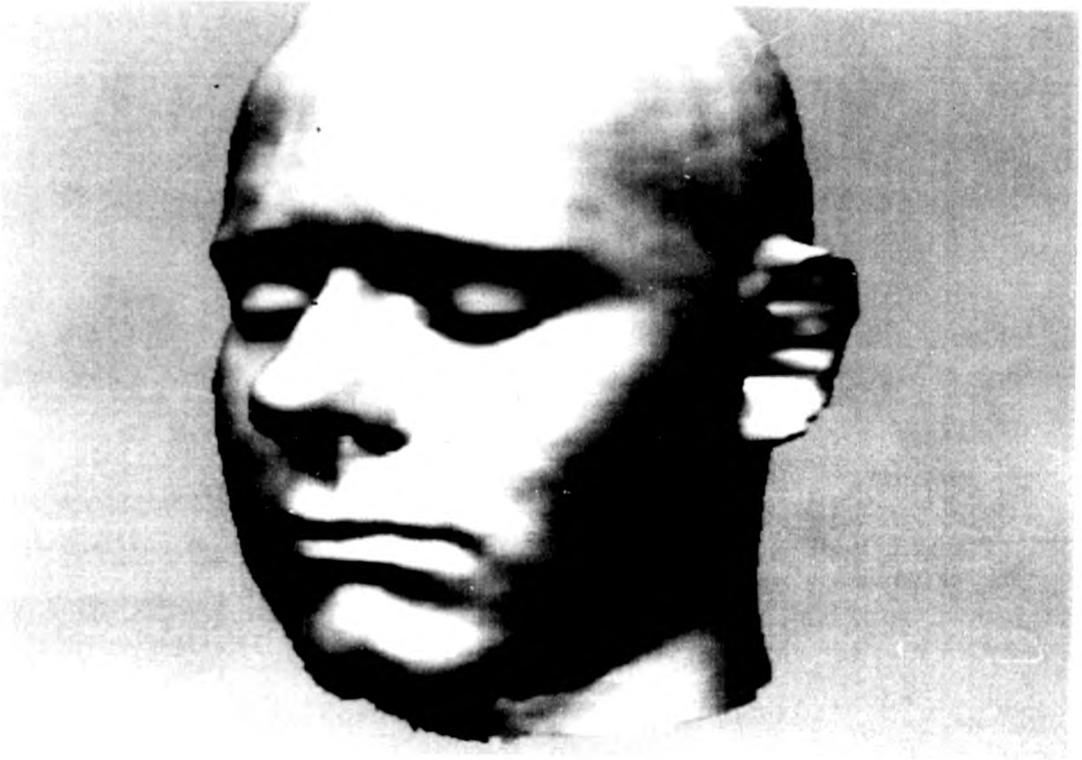

Figure 3.1: An example of a laser head shown lit from above in three-quarter view.
(Lighting direction top1 - see general methods)

Linney, 1991). Laser heads do contain information other than shading, for example they preserve two- and one-dimensional metric information like lengths and aspect ratios, and it may be this information that supports their recognition. Laser heads also provide contour cues to shape when presented as shaded representations. However, an effect of the resolution with which laser heads were portrayed has been interpreted as evidence that information about surface structure, that might be conveyed by shading, is important (Bruce et al, 1991). Effects of other image manipulations, like lighting and polarity reversal, known to affect shading but not many of the alternative sources of information are used to investigate the use of shading information. Other useful sources of information are absent from laser heads, for example pigment differences and texture (see Appendix D), and this allowed the role of shading to be examined in relative isolation. A comparison between laser heads and black and white (grey!) photographs of people with eyes closed and hair covered has shown that photographs are better recognised, suggesting that any problems associated with the recognition of laser heads are not solely the result of the absence of those particular facial features. The comparison suggests that the texture and pigment based information absent from laser heads is useful and there appear to be particular problems with the recognition of female targets from their laser heads for reasons which were not clear. The recognition of females from laser heads was investigated in the first experiment reported here (see also Appendix C).

The major aim of the work reported was systematic investigation of the effects of lighting and viewpoint, two variables which laser heads allow to be controlled more accurately and easily than is possible with natural images. Previous research has shown the importance of these variables for performance of tasks involving laser heads (Bruce et al, 1991) and they are also known to be important for tasks involving photographic stimuli (eg. Bruce, Valentine & Baddeley, 1987; Johnston et al, 1992). For example full-face and three-quarter views are generally found to produce more accurate recognition than profile views, and top lighting to be better than bottom lighting, at least when photographs are used. However, with laser heads some anomalous results have been reported such as advantages for bottom lighting and profile views (Bruce et al, 1991). It may be that these were a function of the information absent from laser heads; for example the absence of eyes and lip colouring may have reduced the importance of the full-face view and heightened the importance of contour information derivable from the profile view. The apparent advantage for bottom lighting was explained in terms of the highlighting of more useful surface features and this effect of lighting has to be considered alongside its effect on shading. The experiments reported here sought to test if these anomalous effects were reliable and therefore evidence for important differences between laser heads and photographic stimuli. This raises the issue of the generalizablity of findings using laser heads as stimuli, but knowledge of the effects of view and lighting on laser heads may anyway be useful for their applications although it was hoped that the results would be generalizable.

In the second experiment reported, experiment 6, the well known adverse effects of photographic negation were investigated using laser heads as stimuli. These effects have been attributed to the effect of negation on the interpretation of shape-from-shading and shadows (Phillips, 1972; Pearson, Hanna & Martinez, 1990; Cavanagh & Leclerc, 1991; see also chapter 1) although negation also importantly affects information about pigment differences. The use of laser heads allowed investigation of the former in the absence of the latter. The effects of negation on shading may not be independent of the effects of changing lighting direction; negatives look similar to bottom lit faces in some ways because both reverse the brightness of light and dark areas such as, for example, the forehead, the eye sockets and the bridge and nostril area of the nose at least for full-face views (Johnston et al, 1992; see also figs. 3.2d & 3.5a). Thus some of the effects of negation and bottom lighting may be complimentary and taking the negative of a bottom lit face, like taking the negative of a negative, results in a stimulus that is similar in some respects, including apparent lighting direction, to a positive of a face lit from above (see figures 3.2a & 3.5d). However, a previous study failed to show anomalous results have been reported such as advantages for bottom lighting and profile views (Bruce et al, 1991). It may be that these were a function of the information absent from laser heads; for example the absence of eyes and lip colouring may have reduced the importance of the full-face view and heightened the importance of contour information derivable from the profile view. The apparent advantage for bottom lighting was explained in terms of the highlighting of more useful surface features and this effect of lighting has to be considered alongside its effect on shading. The experiments reported here sought to test if these anomalous effects were reliable and therefore evidence for important differences between laser heads and photographic stimuli. This raises the issue of the generalizablity of findings using laser heads as stimuli, but knowledge of the effects of view and lighting on laser heads may anyway be useful for their applications although it was hoped that the results would be generalizable.

In the second experiment reported, experiment 6, the well known adverse effects of photographic negation were investigated using laser heads as stimuli. These effects have been attributed to the effect of negation on the interpretation of shape-from-shading and shadows (Phillips, 1972; Pearson, Hanna & Martinez, 1990; Cavanagh & Leclerc, 1991; see also chapter 1) although negation also importantly affects information about pigment differences. The use of laser heads allowed investigation of the former in the absence of the latter. The effects of negation on shading may not be independent of the effects of changing lighting direction; negatives look similar to bottom lit faces in some ways because both reverse the brightness of light and dark areas such as, for example, the forehead, the eye sockets and the bridge and nostril area of the nose at least for full-face views (Johnston et al, 1992; see also figs. 3.2d & 3.5a). Thus some of the effects of negation and bottom lighting may be complimentary and taking the negative of a bottom lit face, like taking the negative of a negative, results in a stimulus that is similar in some respects, including apparent lighting direction, to a positive of a face lit from above (see figures 3.2a & 3.5d). However, a previous study failed to show that top lit positives and bottom lit negatives are equivalent stimuli for recognition (Johnston et al, 1992), although this is likely to have been largely due to differences in the effects of negation and lighting on factors other than shading, particularly the brightnesses of pigmented areas of the kind absent from laser heads. Experiment 6 used laser heads to test whether the effects of lighting and negation on shading were equivalent in the absence of these other factors.

The methods common to the two experiments reported in this chapter, some of which are also relevant to the experiments on matching reported in the next, will be outlined next followed by the experiments themselves.

3.2 GENERAL METHODS

Materials.

All but one of the remaining experiments reported in this thesis made use of surface representations, "laser heads" (see figures 3.1, 3.2 & 3.5 and Appendices B & C), and the technique by which these were produced will be described here. Laser heads allowed the effects of direction of lighting on a simple shading model to be investigated in isolation from many other sources of information. Laser heads also allow lighting and viewpoint to be controlled and changed independently as is described below.

Laser heads were developed as accurate representations of the geometry of facial surfaces for use in planning facial surgery (Linney, 1992). The subject of the scan, wearing a stocking to conceal their hair and with their eyes closed, is seated in a rotating chair in front of the scanning device. This was designed and built at University College London Department of Medical Physics and this is where scanning took place (Moss, Linney, Grindrod, Arridge & Clifton, 1987). The laser source projects a line onto the surface of the face which follows the contours of the profile of the face. This is apparent when the line is viewed from an oblique angle by a video camera. The deviations of the points is a function of the shape of the face and, after calibration using known angles, trigonometric principals can be used to define the relative positions of points on the surface. This is done for successive profiles as the head is rotated in front of the devise. The distribution of these profiles is not even, with a greater density at the front of the head giving greater resolution in the area of the face. The number of points on each profile varies slightly between heads but the result is a database of approximately 20,000 surface points.

This data was then transferred to a Sun 3/110 grey level work station at Nottingham University where it was sampled and transformed to convert it into a database of approximately 16,000 x, y, z coordinate points. These were then joined to produce a wire frame model with approximately 16,000 quadrilateral facets using an autofaceting program written specifically for the purpose. The facets were then shaded using standard computer aided design techniques for a particular viewpoint and direction of lighting. The shading algorithm used for all the experiments reported in this thesis was the Phong diffuse shading algorithm (Bui-Tuong, 1975). This shading model was chosen as it most accurately models the physics of the illumination of matte surfaces. Gouraud shading is used for many applications and produces more accurate performance with laser heads under some conditions (Bruce et al, 1991) but this model makes a number of approximations which might make it less generalizable to real conditions of illumination. Phong shading models the effect of a single distant light source on a matte or Lambertian surface, that is one reflecting light equally in all directions. For Phong shading an intensity is calculated for each facet which is determined according to Lambert's law by the cosine of the angle of incidence, the angle between the incident lighting and the local surface normal (calculated from the vertices of that facet). Surface intensity therefore decreases as the angle between the normal of a patch and the incident light increases and patches not receiving direct illumination are rendered uniformly dark. There is no ambient component in this model. As the surface modelled is Lambertian, that is reflecting light equally in all directions, surface luminance is independent of viewing direction, although this determines which surfaces are visible. Lambertian shading would be only one factor determining the intensity of a point under ordinary conditions of illumination, but it does provide a strong cue to the three-dimensional structure of the surface of the face and is probably always an important cue when available (see chapter 1). Phong shading does not, for example, implement ray tracing and so the effects of mutual illumination, which may be important naturally, are absent. The position of the light source and viewing direction are specified relative to the centre of the head prior to shading as described next.

The available Sun software was adapted to displaying laser heads to allow both viewing and lighting directions to be defined within a head centred world coordinate system. The origin for this coordinate system is in the centre of the head. The x-axis runs from left to right, from ear to ear, the y-axis top to bottom, crown to neck, and the z-axis from front to back. Three viewpoints were used for the recognition experiments reported in Chapter 3, full-face, left three-quarter and left profile. These had viewing directions as defined by coordinates (0, 0, -z), (-x, 0, -z) and (-x, 0, 0) respectively. Values of x, y and z were chosen so that images were of uniform size. To produce a particular view of a head under different lighting conditions exactly the same viewing positions were used to ensure only the lighting changed. For matching experiments only three-quarter and profile views were used, with viewing positions defined in the same way. Figures 3.1 and 3.2 show one head used in the first recognition experiment, experiment 5, in all three views under top and bottom lighting respectively.

The alternative light source positions were specified within a normalised device coordinate system. The z-axis corresponded to the viewing direction, the y-axis was vertical and perpendicular to this and the x-axis was orthogonal to both these. The basic "Top" and "Bottom" lighting directions used in both recognition and matching experiments were from 45° above or below the viewing direction, that is as defined by lighting coordinates (0, 1, -1) and (0, -1, -1) respectively. Figure 3.2 shows examples of these two directions of lighting. With these lighting directions a change in viewing direction involved rotation of the head within a constant light field. Both light source and viewing position coordinate systems were developed in accordance with ACM core (*Status Report of the Graphics Standards Planning Committee*. Computer Graphics. Volume 13, Number 3, August 1979).

The images produced on the Sun were saved from the screen as raster files and transferred to a Macintosh IIcx computer. Here they were converted to PICT files and "tidied up" using an image processing package. This involved trimming excess areas of neck and also removing noise that was produced by reflection of the laser at the top of the head.

3.3 EXPERIMENT 5

INTRODUCTION

Experiment 5 investigated the recognition of the laser heads of people familiar to the subjects as seen from different views and with different directions of lighting. Three views were used, full-face, three-quarter or profile and two directions of lighting, from above or from below (see figure 3.2). This allowed differences between top and bottom lighting to be tested for on the recognition of laser heads.

An effect of viewpoint was expected in this task as faces, at least when photographed, are normally found to be better recognised in full-face and three-quarter than in profile. There is also some evidence for a general three-quarter view advantage (Bruce et al, 1987). Previous work with laser heads has also shown effects of viewpoint (Bruce et al, 1991). However, as mentioned in the general introduction, there was an unexpected advantage associated with profile views, at least for matching to a line-up, in that work perhaps reflecting the more limited sources of information available from laser heads. This experiment was intended as a systematic investigation of the effects of viewpoint on the recognition of laser heads.

a) Top Lit, Full-Face b)Top Lit, Three-Quarter c)Top Lit, Profile

d) Bottom Lit, Full-Face e)Bottom Lit, Three-Quarter f)Bottom Lit, Profile

Figure 3.2: Examples of stimuli used in experiment 5.

As well as varying viewpoint, lighting was also varied with faces in each view lit from above or below. In the last chapter subjects appeared more likely to see a face as a face when it appeared lit from above and previous work with photographs has shown that faces are better recognise when lit from above (Johnston et al, 1992). However, again an anomalous result has been reported with laser heads with an advantage for bottom lighting when matching to a photographic line-up (Bruce et al, 1991). This experiment was intended to systematically investigate the effects of lighting on the recognition of laser heads. It was not known what the relationship between the effects of lighting and viewpoint would be, though as both combine in highlighting certain surfaces and not others, their effects might be expected to interact.

Male and female laser heads were used in this experiment although there appears to be particular problems with the recognition of female faces from laser heads (Bruce et al, 1991). The main measure of performance used was recognition accuracy but a second measure was included, a rating of "likeness" obtained in the second phase of the experiment after subjects had been informed whom each laser head was actually of.

METHOD

Subjects

Twelve subjects took part in the experiment, all members of the same department as the people whose heads were used as stimuli. All potential subjects were shown a list of names of people who might appear in the experiment to ensure that they were familiar with the target heads and knew their names. All subjects had normal or corrected to normal vision.

Materials.

The laser scans of eight members of staff from the University of Nottingham Psychology Department were used as stimuli. Four were male and four female. For full details of the laser scanning technique and the computer graphics used to display them, please see the general methods section. For this experiment three views of each head were used; full face, left three-quarter and left profile. Lighting was from either 45° above or below the viewer's line of sight, that is directions "top" and "bottom" (see general methods section). Different views of the head were in effect rotations within a uniform light field which itself was symmetrical about the line of sight. This arrangement maximises shading information by maximising the amount of visible surface area receiving direct illumination.

Stimuli were presented on a Macintosh IIcx computer using a program written in Supercard. All images included a "button" which provided a stimulus number. Stimulus numbers were assigned randomly and provided no clues to identity. Responses were made on response sheets which included spaces for the stimulus number as displayed on the button, the name of the person portrayed and the likeness rating.

Design.

A within subjects 3(View) x 2(Light) x 2(Sex) design was used. The levels of view were 'full-face', 'three-quarters' and 'profile', of lighting 'top' and 'bottom' and of sex 'male' and 'female'. Eight faces were used under all conditions of view and light, thus each face was presented a total of six times. Four of the face used were male and four female. This gave a total of 8 x 3 x 2 = 48 trials for each subject. The order of trials was randomised for each subject. Subjects were provided with a list of fourteen names, seven male and seven female, which was used to ensure that they knew all the target stimuli. They were also encouraged to guess from the list of names for any stimulus they were not sure of, giving a chance rate of 1/14, 7%.

Two measures of performance were used, percentage correct on the naming task and the likeness ratings given in the second phase of the experiment.

Procedure.

The experiment took place in a small windowless room lit by a single fluorescent bulb. An Apple Macintosh 11cx computer was used to present the stimuli. Subjects were first shown a list of 14 members of the department, including the 8 used as stimuli, to ensure that they knew the people used as target faces. Subjects were then shown the stimulus faces in a random order and asked to name each face on the response sheet next to the appropriate stimulus number. Subjects were asked to guess if they were not sure from the 14 names given on the list. The provision of a list of possible names together with the use of accuracy rather than response latency as a measure of performance was thought to circumvent most of the problems normally associated with the use of naming as a task, namely that the face may be recognised without the name being accessed and that naming takes longer than actual recognition.

After subjects had attempted to put a name to all the images the stimuli were all presented a second time. In this phase subjects were told who the image was derived from and asked to rate its likeness to that individual. Independence of this judgement from their previous attempt to recognise the stimulus was stressed. Subjects were asked to rate likeness on a scale of 1 to 5 with 1 indicating 'not at all like' and 5 indicating 'very like'.

Results

The percentage correct recognition scores for male and female items are shown in figures 3.3 a) and 3.3 b) respectively and figures 3.4 a) and 3.4 b) show the likeness ratings. As can be seen both naming accuracy and likeness ratings were much lower for female than male stimuli. The averages percent correct were 75% for males and 41% for females and the average likeness ratings 3.6 and 2.3 respectively. For the male stimuli, figure 3.3 a), there was a clear interaction between the effects of view and light; there was most effect of light for full-face views, then three-quarters and none for profiles. For female stimuli there are few clear differences and performance may be have been close to floor.

Due to the clear differences in performance analysis of variance was performed separately on male and female items. Analysis of variance on percentage correct recognitions for male items gave a view x light interaction ($F_{2,22}$= 7.2, $p<<0.05$). Analysis of the simple main effects of light showed effects of light for full-face ($F_{1,33}$= 18.9, $p<<0.05$) and three-quarter views ($F_{1,33}$= 6.4, $p<0.05$) but not for profiles ($p>0.1$). There was a simple main effect of view for top lit stimuli ($F_{2,4}$= 9.1, $p<<0.05$) but not for bottom lit stimuli ($p>0.1$). Planned t-tests showed that for top lit stimuli full-face and three-quarter views produced more accurate naming than profiles. With the equivalent analysis for female items there were no significant differences (All p's>0.1).

The pattern of ratings data shown in figures 3.4 a) and 3.4 b) was similar to that for naming accuracy. Male items were rated better likenesses, average 3.6, than female items, average 2.3. For male items there was an interaction view x light ($F_{2,22}$= 19.3, $p<<0.05$). Again there were simple main effects of light for full-face ($F_{1,33}$= 47.7, $p<<0.05$) and three-quarter views ($F_{1,33}$= 43.9, $p<<0.05$), but not for profiles ($F<1$, ns). For female items there was a main effects of light ($F_{1,11}$= 10.3, $p<<0.05$), with top lit faces rated better likenesses than bottom lit faces. There was also a main effect of view ($F_{2,22}$= 5.7, $p<0.05$), with planned comparisons showing three-quarters to be rated better likenesses than either full-face or profile views.

FIG. 3.3 a): Percentage Correct Recognition Performance, Male items, experiment 5.

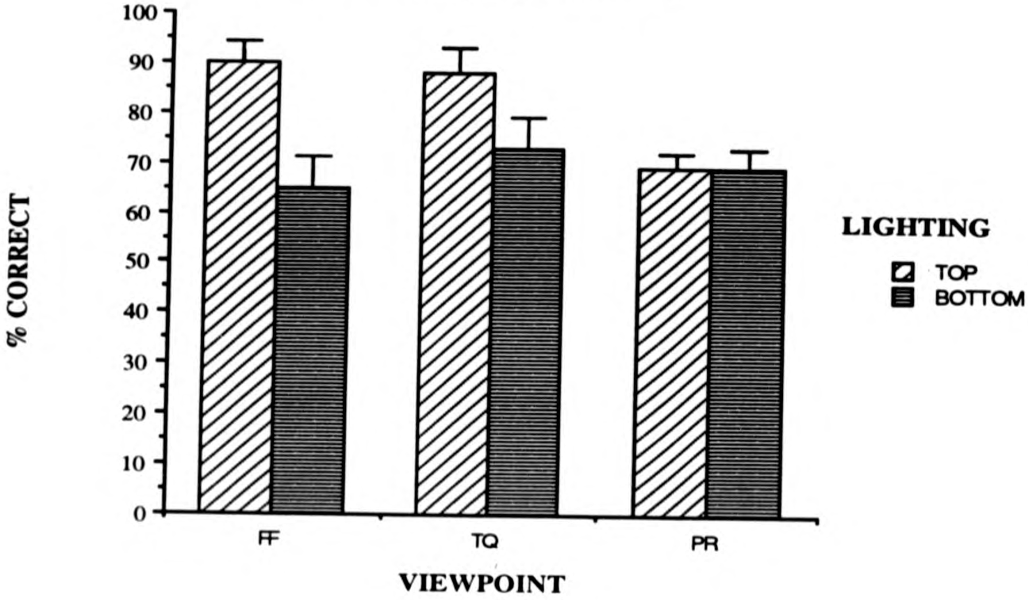

Fig 3.3 b): Percentage Correct Recognition Performance, Female Items Experiment 5.

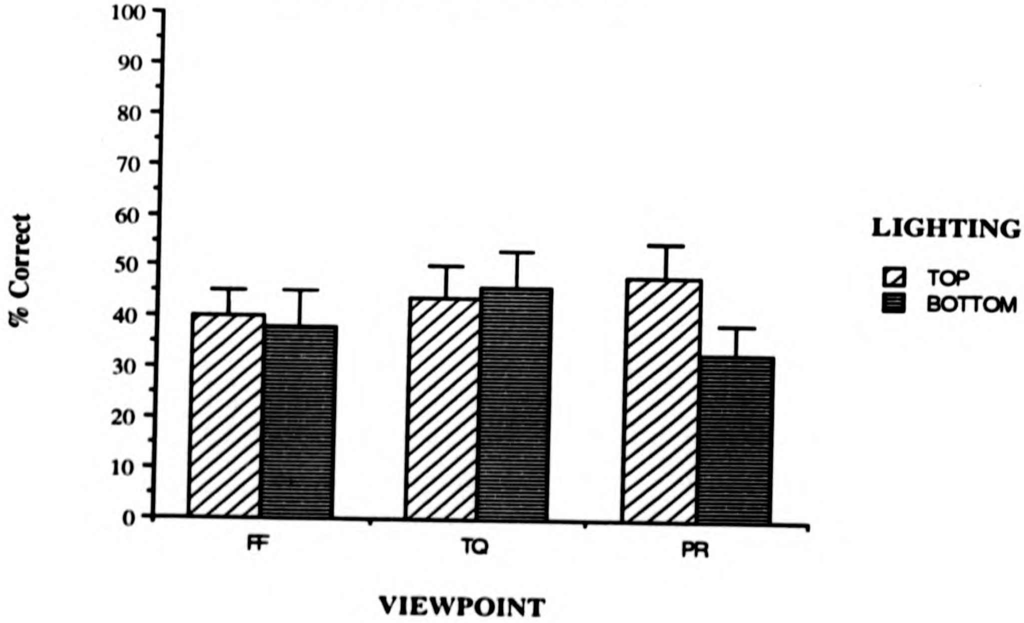

Fig. 3.4 a): Likeness Ratings, Male Items Experiment 5.

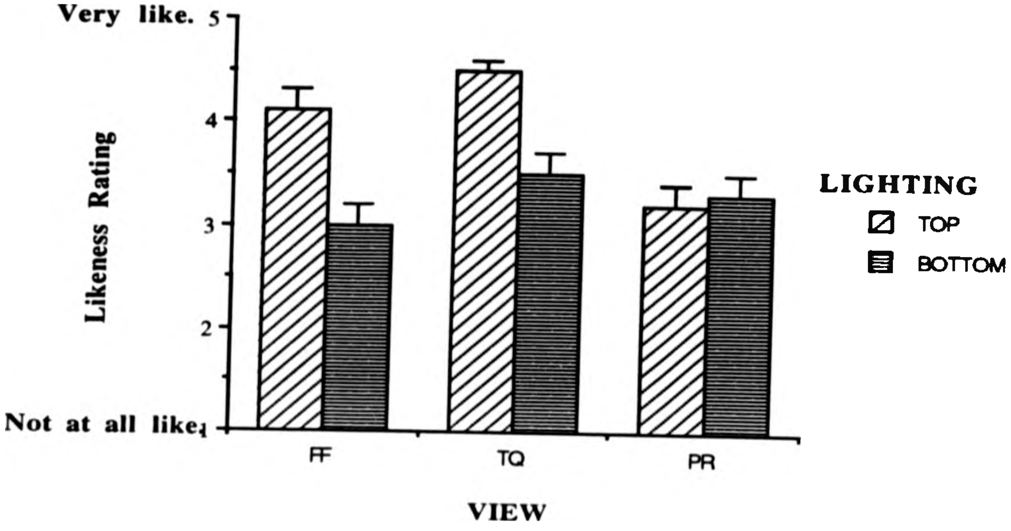

Fig. 3.4 b): Likeness Ratings, Female Items Experiment 5.

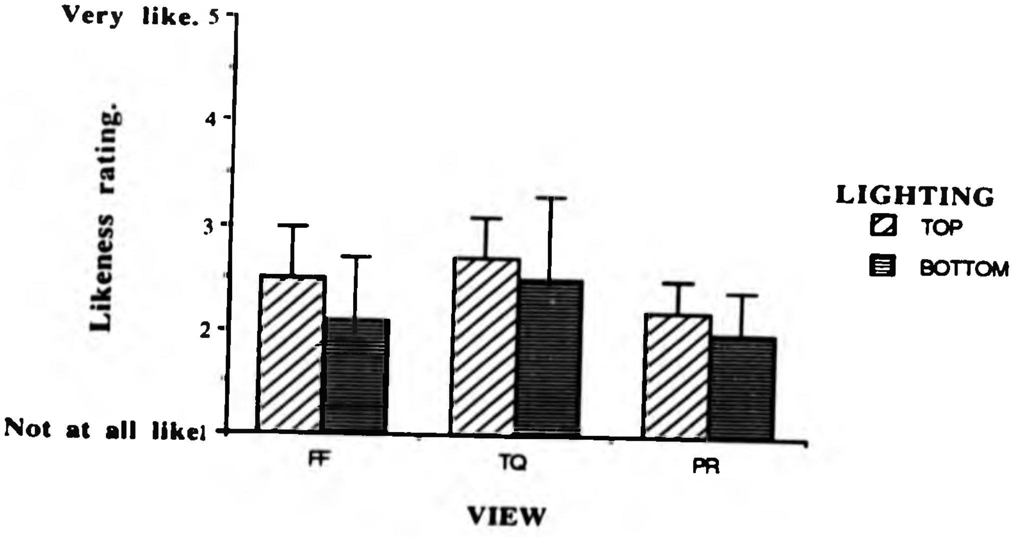

DISCUSSION.

For both measures subjects performed better with the male than female items and there were different patterns of effects of light and view for the sexes. For male items there was a clear interaction between the effects of light and view, with only full-face and three-quarter views being benefited by top lighting. For female items there were no clear effects of light and view on recognition, perhaps because recognition performance was close to floor. There was an advantage for top lighting and three-quarter views on ratings data. The difference between male and female items will be discussed first, followed by a fuller discussion of the pattern of results for male items.

The advantage for male laser heads over female laser heads is a replication of a previous result with this type of stimuli (Bruce et al, 1991). In that experiment a comparison with black and white photographs where eyes were closed and hair concealed showed that the recognition of female faces was particularly impaired in laser scans. Performance with photographs of faces with eyes closed and hair concealed was at a comparable level for male and female items showing that the absence of these features could not explain the difference in performance. It suggests instead that other features, for example skin texture or eyebrow shape, which were present in the photographs but not laser scans must explain the decrement perhaps with these features becoming particularly important for female faces when hair and eyes are concealed (Bruce et al, 1991). Performance with female items on this experiment was higher than on previously reported experiments, though chance levels were also different, but it was still well below that for male items. Performance with female items did not appear to be affected by lighting or viewpoint changes, though these did affect ratings, although this may have been because performance was at floor and what accurate recognition there was may have been was based on local invariant features unaffected by the manipulations of viewing conditions. The real reason for the apparently reliable difference between performance on male and female laser heads remains unknown and may be worthy of further investigation given its theoretical implications for the question of how sex judgements are made (Bruce, Burton, Hanna, Healey, Mason, Coombes, Fright & Linney,1993; Burton, Bruce & Dench, 1993) and because of the applied uses of laser heads. One possibility is that shape information as derived from shading is just less useful for distinguishing female heads than male heads, perhaps because other sources of information are more variable and thus more informative for female heads. Rather than try to determine the reason for the poor recognition of laser heads of females only male heads were used in the remaining experiments reported in this thesis.

For male items in this experiment the effects of view and light were found to interact. This appears to be mainly due to advantages for top lit full-face and three-quarter views. Recognition accuracies for these views were around 90% suggesting these are the best conditions for recognition of laser heads. The previous advantage of bottom lighting for three-quarter views (Bruce et al, 1991) was not replicated in this experiment and may have been an artifact of the task previously used. The advantage for top lighting found here can, as with the previous work using photographs (Johnston et al, 1992) and the hollow face illusion (see Chapter 2), be explained in terms of the importance of shape-from-shading for face perception. Shape-from-shading may be less important for more angled views like profiles because the shape of the internal features in this view are specified by the shape of the occluding contour rather than shading. That the occluding contour projected by a face is lighting invariant may explain the lack of an effect of lighting in this view.

An alternative explanation for the advantage of top lighting would be that it highlights more salient features for recognition. This is the same argument as was used to explain the advantage found for bottom lighting in a previous study! In general top and bottom lighting highlight the same proportion of the total facial surface, just different areas which may be more or less useful (the proportion of lighted surface is only reduced when the lighting direction is at a very different angle from the viewing direction). In general, arguments in terms of feature saliency are likely to have some validity in that different features are differentially useful for different tasks (eg. Shepherd, Davies & Ellis, 1981; Haig, 1986; Fraser & Parker, 1986). However what features are useful appears to vary from face to face (Haig, 1986) and even task to task (Roberts & Bruce, 1988) and is not clearly established. This makes such explanations liable to circularity in that feature salience is used to explain data at the same time as the data is being used as evidence of feature salience. It is not clear how any feature salience argument would explain the effect of lighting direction on the perception of the hollow face illusion. The shape-from-shading explanation of the advantage for top lighting would be supported if an advantage could be shown for the bottom lit condition when the stimuli were presented upside-down. Such presentation entails a rotation of the apparent direction of lighting so that the "bottom" lit condition now satisfies the light-from-above assumption. This has been shown to reduce or eliminate the adverse effect of figural inversion on recognition of photographs (Johnston et al, 1992). However, recognition performance with laser heads is generally lower than with photographs and inverting laser heads might reduce performance to floor causing any such reversal of the advantages associated with top and bottom lighting to be masked. Experiments using inverted laser heads are reported in Chapter 4 which provide some evidence of the effect of reversal of lighting direction with inversion that would be predicted by an account in terms of shape-from-shading.

Previously performance has been shown to be better when laser heads are shown in profile than in three-quarter view on a matching task, but a non-significant trend for a three-quarter advantage when the task was recognition (Bruce et al, 1991). As was the case with the bottom lit advantage already discussed, this apparently anomalous result was not replicated in the current experiment. In this experiment there was an advantage for three-quarter and full-face views over profile views in line with previous findings using photographs (eg Bruce et al, 1987). There appears to be information useful for recognition available from full-face or three-quarter views of laser heads absent from profiles, at least when top lit. Profiles only show half the face but the advantage for top lighting again suggests that shape-from-shading may be important.

The next experiment reported further investigated the effects of lighting and viewpoint on the recognition of laser heads but using a slightly different task. The effects of photographic negation was also investigated as this is also thought to be related to shading and shadow information.

3.4 EXPERIMENT 6

INTRODUCTION.

This experiment has two main aims; to replicate the effects of light and view found in experiment 5 and also to investigate any effects of photographic negation on the recognition of laser heads. A familiarity decision task was used instead of naming, with subjects being presented with familiar and unfamiliar faces. Heads were again presented in full-face, three-quarter or profile view and top or bottom lit. However, faces in this experiment were also presented as photographic positives and negatives for the reason discussed in the general introduction (see figures 3.2 and 3.5). Only male heads were used as stimuli because of the apparent difficulties with the laser heads of females (Bruce, Healey, Burton, Doyle, Coombes & Linney, 1991; see also experiment 5).

The positive stimuli provided a direct replication of experiment 5 although the task was different and the addition of negatives could affect overall performance (cf. Johnston et al, 1992). However, top lighting should still favour the processing of full-face and three-quarter views, with less of a difference expected on profiles.

As explained in the introduction to this chapter, negation is thought to affect the interpretation of shading and shadow information (Phillips, 1972), although its primary effect may be on information about pigment differences available in photographs but not laser heads. There should be effects of negation on the recognition of laser heads, even despite the absence of information about pigmentation, if it affects shading information important for recognition. Effects of negation on the perception of laser heads have been previously shown using a sex judgement task consistent with shading being important (Bruce, Burton, Hanna, Healey, Mason, Coombes, Fright & Linney, 1993).

The effects of polarity on shading may not be independent of the effects of lighting (Johnston et al, 1993; see also the introduction to this chapter). While both negation and bottom lighting reverse the polarity of some areas, the precise nature of

a) Top Lit, Full-Face. Negative

b)Top Lit, Three-Quarter Negative

c)Top Lit, Profile Negative

d) Bottom Lit, Full-Face Negative

e)Bottom Lit, Three-Quarter Negative

f)Bottom Lit, Profile Negative

Figure 3.5: Examples of negative stimuli used in experiment 6.

the relationship is not clear as the two manipulations have different effects on other areas. Areas slanted with zero tilt with respect to the light source, for example the centre of the forehead, change in luminance when lighting changes from top to bottom but surfaces like the sides of the head do not (see figure 3.2 a & d). For the latter areas their luminance is affected by negation (see figure 3.5 a) but not by a change in lighting direction. The relationship between light and polarity is further complicated because the orientation of surfaces relative to the light source, which determines the effect of a change in lighting, will also be dependent on viewpoint. However, it was expected that bottom lit negatives would produce better performance than top lit negatives, more similar to top lit positives in some respects for the reason outlined in the introduction (see figure 3.5 d, e & f).

Experiment 5 showed that the effects of lighting and viewpoint are not independent, perhaps because different sources of information are important for the perception of different views; the shape of the occluding contour being more useful for profile views but shading information being useful for less angled views. If this is the case the effects of all three factors used in this experiment, lighting, viewpoint and polarity, might be expected to interact. For example, negation, like lighting, would not be expected to affect the perception of profile views as much as that of less angled views as it does not affect the shape of the occluding contour.

Therefore an effect of negation was expected along with a replication of the effects of lighting and view reported in experiment 5. The three effects were expected to interact, with the effects of lighting and polarity being complimentary in some cases but perhaps different for different viewpoints.

METHOD.

Subjects.

Sixteen subjects recruited from the same department as the people whose heads were used as stimuli took part in the experiment. All had normal or corrected to normal eyesight.

Materials.

The laser heads were produced in the same manner as described previously. Laser heads of people from University College Hospital were used as distractors, all of whom were assumed to be unfamiliar to the subjects. Negative images were obtained using an image processing package on the Macintosh computer which reverses grey level intensities around the middle value; 0<->255, 1<->254, 2<->253 etc. There was an additional stage in the processing of the laser heads for this experiment in that black and white slides of the stimuli were taken from the screen to allow slide presentation using a Rockwell Aim computer. This controlled recording of responses, latencies and also controlled presentation durations and intervals. Responses were made using a box with two buttons labelled "familiar" and "unfamiliar". Slides were projected at a viewing distance of approximately 350 cms and a size of seventy x forty centimetres, subtending a visual angle of 12° x 7°.

This extra stage of processing necessarily degraded the stimuli introducing additional non-linearities to grey level reproduction and noise which may have affected the generalizability of the results. However, at the time this was the only way to allow presentation of stimuli and recording of responses to be controlled automatically.

Design.

The design of the experiment was a 3(View) x 2(Light) x 2(Polarity) within subjects design. The levels of view were full-face, three-quarter and profile, of light top and bottom, and of polarity positive and negative. Ten heads were used as stimuli, six targets and four distractors. There were therefore 10 x 3 x 2 x 2 =120 trials for each subject. Two random orders were produced with the restriction that no two consecutive trials showed the same person and there were no more than four "familiar" or "unfamiliar" trials in a row. Each pseudo random order was used for eight subjects and the slides were further split into two blocks of sixty, the order of the blocks being alternated between subjects. The order of slides within each block was also reversed after every four subjects, meaning that only two subjects had any one order.

Hits and false positives were recorded, together with reaction times though these were not subsequently analysed due to high error rates.

Procedure.

Testing took place in a dimly lit large windowless room. Stimuli were presented on a blank wall approximately three and a half metres from the subject. Subjects were told that they would be shown 120 head models some of which would be of familiar members of the department and some of which would be of unfamiliar strangers. They were told that their task was to decide for each one which of these categories it belonged to and respond familiar or unfamiliar with the buttons provided. Subjects were asked to respond as quickly and as accurately as possible. The actual experiment was preceded by ten practice trials involving stimuli of people not used in the actual experiment to familiarise the subjects with the stimuli and task.

Slides were presented for 4 seconds each, separated by a gap of 1 second. Each slide was preceded by a short tone half a second before presentation.

RESULTS

Treatment of Results

In order to accurately reflect sensitivity to different viewing conditions it is necessary to combine measures of the numbers of both hits and false positives. A', a non parametric equivalent to d', was used for this purpose in this experiment (McNicol, 1972; Rae, 1976; Valentine & Bruce, 1986). A' allows sensitivity to be calculated from a single pair of hit and false alarm rates using a graphical method to approximate the area under the ROC curve. An A' of 0.5 represents chance performance and an A' of 1 perfect performance. A' tends to be right skewed and is therefore transformed using $2\arcsin\sqrt{(A')}$ prior to analysis of variance (McNicol, 1972). While response latencies were recorded, they are not reported as the large number of errors meant that there were not enough correct responses to give meaningful averages.

Results

The mean A' s for conditions of lighting and polarity for different levels of viewpoint are shown in figure 3.6. As can be seen from the graphs the effects of lighting and polarity were found to be different at different levels of view. For full-face views there was a clear interaction between polarity and lighting, with performance best with top lit positives or bottom lit negatives. For three-quarter views there does not appear to be much effect of either variable, although recognition rates were fairly high overall. For profiles there appears to be an effect of lighting, with performance better with bottom lit stimuli, but not of polarity.

Analysis of variance was carried out on transformed A' s which gave a three-way interaction of View x Lighting x Polarity ($F_{2,30}= 8.9$, $p<<0.05$). Analysis of simple interactions gave a polarity x lighting simple interaction for full-face views ($F_{1,45}= 30.7$, $p<<0.05$) but not for three-quarters or profiles (p's > 0.1). Further analysis of the effects for full-face views showed simple simple main effects of lighting for positive full-face views ($F_{1,90}$= 23.2, $p<<0.05$), with top lighting better than bottom, and for negative full-face views ($F_{1,90}$= 4.5, $p<0.05$), bottom better than top. The simple simple main effect of negation did not quite reach significance for top lit full-face views ($F_{1,90}$=3.6, $p=0.061$) but did for bottom lit full-faces ($F_{1,90}$=24.0, $p<<0.05$). The only other significant difference was a simple simple main effect of lighting for negative profile views ($F_{1,90}$= 8.0, $p<<0.05$), with bottom lighting better than top.

Fig. 3.6 a): A' For Full-Face View Stimuli, Experiment 6.

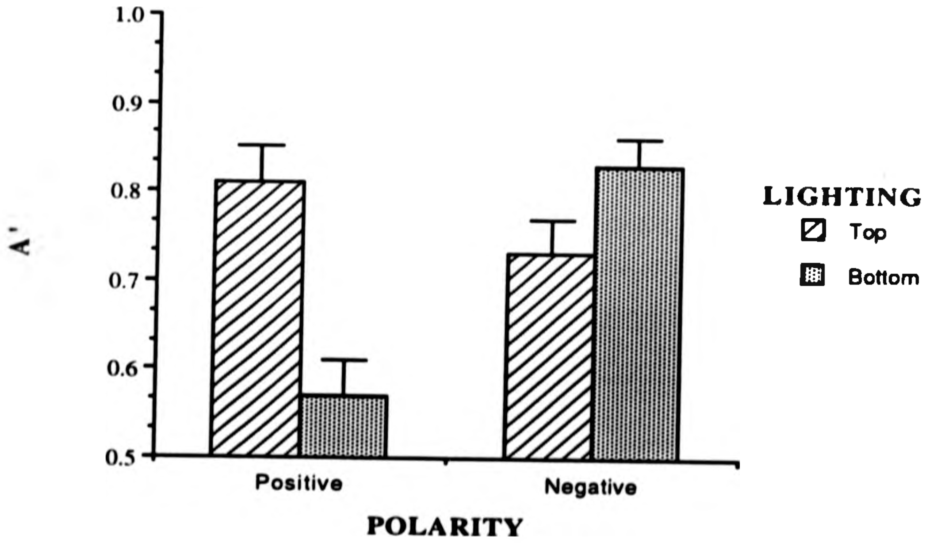

Fig. 3.6 b): A' For Three-Quarter View Stimuli, Experiment 6.

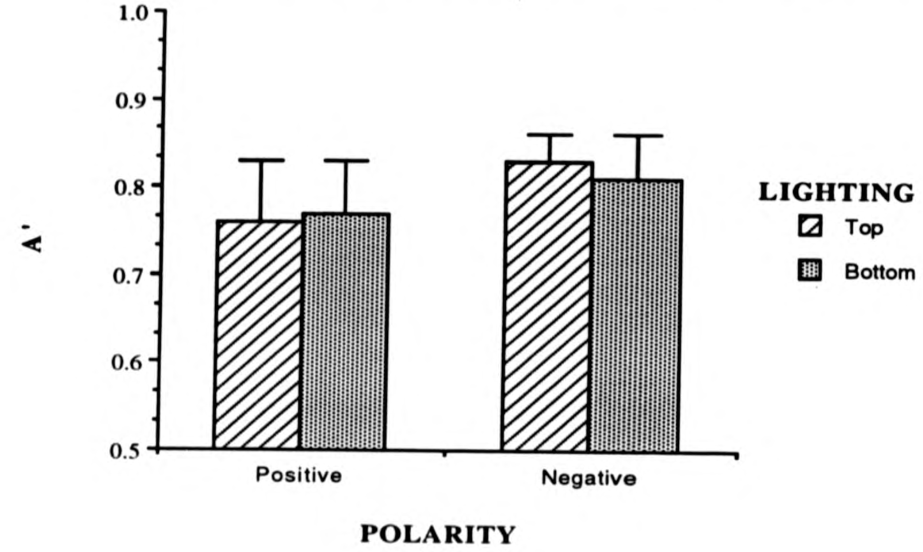

Fig. 3.6 c): A' For Profile View Stimuli, Experiment 2.

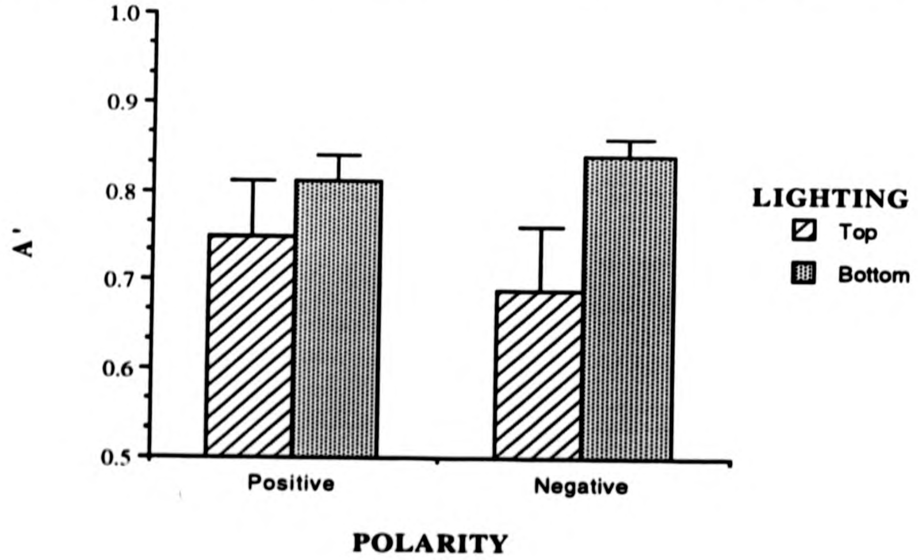

DISCUSSION.

The effects of lighting, polarity and negation were all found to interact. For full-face views the effects of lighting and polarity were as expected, with both bottom lighting or negation reducing performance alone but cancelling each other out when applied together. The lack of effects of either manipulation on three-quarter views was surprising as was the effect of lighting on profiles. Possible explanations of this pattern of data will be discussed first followed by some general conclusions the results suggest.

For full-face views the effects of view and lighting interacted. Bottom lit positives were particularly difficult to recognise (cf Johnston et al, 1992; experiment 5). The effect of negating top lit stimuli did not quite reach significance. It appears bottom lighting disrupts recognition more than does negation for laser heads. This may in part reflect the lack of information about pigmentation which would have been affected by negation but not by lighting. There was a clear effect of negation for bottom lit stimuli, with negatives better recognised than positives. This effect was consistent with negation and bottom lighting being complimentary functions. The effect also shows that negation does affect information other than pigment differences. Negation appears to offset many of the problems associated with bottom lighting, perhaps because bottom lit stimuli do appear to be lit from above (see figure 3.5). It appears that information is not "lost" when faces are bottom lit (as subsequent negation improves performance) but perhaps is not recovered because a light-from-above assumption is not satisfied. This would also be true if figurally inverting bottom lit stimuli improved performance (cf. Johnston et al, 1992).

The lack of effects of lighting and polarity on the three-quarter view was surprising, and in the case of lighting seems inconsistent with the results of experiment 5. The explanation for this is not clear but may have in part reflected changes introduced by the extra stage of processing necessary to produce slides which may have resulted in a loss of grey level information. It is also possible that three-quarter views contain local information useful for recognition not affected by lighting or inversion that the other views do not contain. Information from the shape of internal contours is one possibility.

For profiles there was an effect of lighting when stimuli were presented in the negative, with bottom lighting better than top lighting. As with the full-face view negatives of bottom lit faces were well recognised. Polarity was not expected to affect profiles as it does not change the occluding contours specifying the shape of the internal features. However, when stimuli are bottom lit negation appears to allow the recovery of useful information. There was no effect of lighting when profiles were presented as positives, consistent with the finding reported in experiment 5. The effect of lighting on negative profiles may again be because bottom lit negatives appear top lit. In addition the internal features in the profile are highlighted in the bottom lit negative but not the positive and this area may provide useful information for recognition.

In general the results for stimuli presented as positives in experiment 6 were consistent with the results of experiment 5, an exception being the results for three-quarter views. The inclusion of negative stimuli is likely to have made the task harder and may have suppressed some effects. For full-face and profile views the patterns were much as expected and can be explained in terms of shape-from-shading being more important for less angled views and perhaps being facilitated by lighting from above. Lighting independent contour information is more important for profiles, and the result of experiment 6 show that this information is not affected by negation. Negation does affect apparent lighting direction and appears often to facilitate performance with bottom lit stimuli. However possible loss of shading information through slide based presentation and the absence of pigment information from laser heads have served to reduce the detrimental effects of negation (cf. Phillips, 1972).

In general laser heads are not that well recognised although performance is above chance. Subjects also appeared to adopt fairly strict criteria for judging laser heads to be familiar and were more likely to respond "unfamiliar". In experiment 6 they made fairly few false positives but hits were correspondingly depressed. This is perhaps not surprising given that they lack many sources of information including hair, pigment differences, texture and eyes all of which may be important for everyday face recognition and be part of our memories of familiar faces (Shepherd et al, 1981). That they can be recognised at all suggests that information about the shape of the surface from contours and shading is important, and sometimes sufficient, for recognition even when rendered using a fairly simple shading algorithm. The effects of lighting suggest that it is shading and shadow information is important and not just the two-dimensional metric properties that laser heads also provide. However lighting direction may also serve to highlight different features in the same way as different views reveal different surfaces and part of the explanation of its effect may be in terms of feature saliency.

The main problem found with these studies on the recognition of laser heads is that recognition rates are low, especially when the heads are presented under non ideal viewing conditions. Another problem with recognition studies is that they require that subjects are familiar with the people used as stimuli and there was only a limited pool of such people available. To circumvent these problems a matching task was used for the experiments reported in the next chapter. Matching requires the visual processing of the images presented and measures their discriminability under different viewing conditions, both essential components of recognition and other processes. Thus the effects of lighting and viewpoint and their relationship could be further investigated but without some of the problems associated with the use of laser heads.

4. CHAPTER 4: EFFECTS OF LIGHTING ON THE MATCHING OF LASER HEADS

4.1 GENERAL INTRODUCTION

In earlier chapters effects of lighting have been reported on the general task of seeing a face as a face and on the more specific task of recognising individual faces. In this chapter an essential component of the recognition process, the ability to distinguish between faces despite changes in viewing conditions, was tested using a face identity matching task. Laser heads were again used, as it was the perception of shape that was of interest, and lighting and viewpoint were again varied. This introduction will outline the matching task used and predictions of matching performance derived from the alternative theories of object constancy described in Chapter 1 and from previous experimental evidence.

In all the experiments reported subjects were presented with pairs of faces and asked to decide whether they were of the same or different people. As well as varying identity, the members of each pair could be shown under the same or different viewing conditions (see figure 4.1 and Table 4.1; see also figure 3.2 b, c, e & f). Three-quarter and profile views were used for all experiments and lighting was from above or below. Specific details for each experiments will be given in the appropriate introductions and the general methods section gives full details of the materials and design.

Matching tasks may seem remote from the normal demands of face processing and there is evidence that the matching of unfamiliar faces may be dissociable from the recognition of familiar faces (reviewed Young & Bruce, 1991). However, matching tasks have been successfully used to show effects of different manipulations on face processing (eg. Young, Hay, McWeeny, Flude & Ellis, 1985; Bruce, Valentine & Baddeley, 1987; Kemp, McManus & Pigott, 1990) and have also been used in a clinical context to assess patients' abilities to integrate and discriminate information about faces (Benton & Van Allen, 1968). Of particular relevance to the

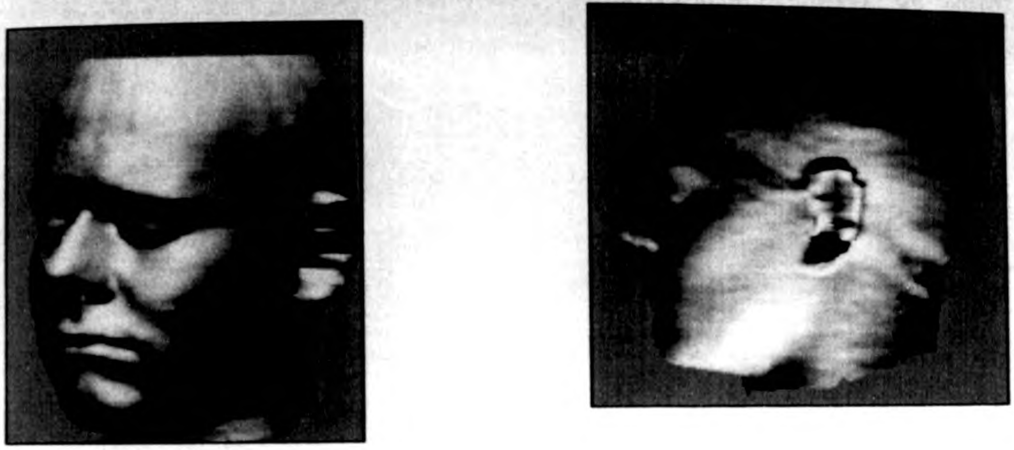

a) "Same" pair, different view and different light (Top Lit Three-Quarter and Bottom Lit Profile)

b) A "Different" pair, same view and same light (Both bottom Lit Three-Quarters)

Figure 4.1: Examples of stimuli from Experiment 7.

experiments reported here is the fact that matching has previously been shown to be affected by changes in viewpoint (Bruce et al, 1987) and by figural inversion of the stimuli (Kemp et al, 1990).

Simultaneous matching, as used in all but one experiment reported here, allows visual processing to be assessed independently of any load on memory (Kemp et al, 1990). Matching was intended to overcome two problems associated with the studies of the recognition of laser heads reported in the last chapter; low recognition rates under non ideal viewing conditions and a limited pool of subjects familiar with the people used as stimuli. As regards the first problem, matching requires faces to be discriminated on the basis of shape, a necessary component of recognition, but does not require actual recognition. The low recognition rates reported in the last chapter may have resulted because the unusual looking laser heads poorly accessed memories derived from normal stimuli. Matching performance does not require access to memory and so performance may not be so low. A group of subjects who knew the people used as stimuli was run as a comparison in experiment 7 to check for any effect of familiarity on performance. However, another reason for using a matching task was that it allowed subjects unfamiliar with the people used as stimuli to be run as subjects.

It was hoped that matching would inform in particular about visual aspects of face processing including possible levels of description derived from the image. For example, in Chapter 1 a distinction was made between pictorial and structural codes (Bruce, 1982). Two faces presented under identical viewing conditions could be matched on the basis of pictorial codes alone but these would not be sufficient when the faces were shown under different viewing conditions. For this structural codes capturing information sufficient to distinguish between faces despite changes in viewing condition (Bruce & Young, 1986) would be needed. The alternative object recognition schemes discussed under the section on object constancy in Chapter 1 make different proposals about possible candidates for structural codes and these lead to different predictions for matching performance which will be reviewed next.

It should first be noted that most accounts of object recognition are accounts of between category decisions and it is not always clear how they would support the within category decisions characteristic of faces and required for the matching task used here. For example instance based models are models of how new stimuli access long term memory and might not be appropriate for explaining the matching of unfamiliar faces. One scheme based directly on raw image properties is the eigen-face approach (Turk & Pentland, 1991). Once faces have been learnt via a process akin to principal components analysis of instances of each face they can be recognised despite changes in, for example, lighting. However, it is not clear what predictions to make from this about the matching of unfamiliar faces though the inputs for this particular scheme have to be standardised for view suggesting that the effects of variable viewpoint may present particular problems.

The implications of invariant features recognition schemes for matching are clearer in that if such features are sufficient for recognition they should also be sufficient to support matching. For example the occluding contour projected by a face seen in a particular view will be invariant across changes in lighting and may be sufficient to support matching when both faces are shown in the same view. It was suggested in the last chapter that such occluding contours might be more salient for profiles than less angled views, explaining the lack of an effect of lighting on their recognition (see experiment 5, figure 3.2), and any such difference might transfer to matching.

Many of the object decomposition schemes described in Chapter 1 are based on information about edges including occluding contours (eg. Biederman, 1987) and this has implications for expected matching performance. For example the problem of variable illumination is normally considered a problem for edge detection, edges which would then support recognition or matching (Poggio & Edelman, 1990). However, changes in view have more profound effects on edge information, including changing the occluding contours, and thus would be expected to affect matching performance. Viewpoint invariant properties of edges (Lowe, 1985; see also Chapter 1, section 1.5) can be used to overcome this problem or different combinations of edges can be used to access common parts which themselves are viewpoint invariant (Biederman, 1987; see also chapter 1). In general changes in viewpoint would be expected to have greater effects on matching than changes in illumination according to edge based schemes.

Alignment approaches to object recognition take explicit account of viewing conditions in order to align a stored description to the input image (Ullman, 1989) but if no stored description was available it is not clear how matching would be accomplished. 2-D alignment schemes might allow the images to be matched to be aligned with each other using linear interpolation between views (eg. Poggio & Edelman, 1990) and a general object model might support 3-D alignment if it were the input images which were transformed when no stored description was available. Many alignment approaches make use of contour information and so, like object decomposition methods, might be expected to be more sensitive to changes in viewpoint than changes in illumination.

Experimentally, effects of viewpoint on matching performance have been shown before (eg. Bruce et al, 1987) and this taken together with the relative neglect of lighting in the object recognition literature led to the expectation of greater effects of viewpoint than lighting on matching performance. Other evidence consistent with this comes from the physiological work which suggests a pattern of responses of single cells selective for faces that is independent of lighting, size and retinal position but modulated by view (Perrett, Rolls & Caan, 1982; Perrett, Hietanen, Oram & Benson, 1992; Hietanen, Perrett, Oram, Benson & Dittrich, 1992). However, effects of light on face perception were reported in the previous two chapters and these may transfer to matching. For example, if shape-from-shading is important and does incorporate a light-from-above assumption at least a difference between matching top and bottom lit stimuli might be expected.

Experiments 7 to 10 investigated the effects of lighting and viewpoint on matching and predictions specific to particular experiments will be discussed in the relevant introductions. The general methods common to all the matching experiments will be described next.

4.2 GENERAL METHODS

Subjects.

Subjects were recruited by advertisement from the University of Nottingham (experiments 7, 9-12, 14 & 15) and the University of Stirling (experiments 8, 13) and were paid at the rate of £3 an hour for their participation. All subjects used had normal or corrected to normal eyesight. All subjects were unfamiliar with the laser head format.

Materials.

The laser scans used in this experiment were in the same format as for the experiments involving recognition and presentation was screen based as it was for experiment 5. The laser heads of eight people were used in each experiment, two for the practice session and six for each actual experiment (see Appendix B). The views used were three-quarter and profile views, that is from viewing directions -C, 0, -C and -C, 0, 0 respectively. C, a constant, was again chosen so that the images used were of approximately uniform size.

For each experiment two lighting directions were used, levels varying for different experiments. Specification was in the same manner as for experiments 5 and 6. The directions top1 and bottom1 used in those experiments were again used, which had light source coordinates (0,1, -1) and (0, -1, -1) respectively.

Presentation of stimuli and recording of responses was controlled by the Macintosh computer using a program written for the purpose in Hypercard. The program presented separately framed faces side-by-side in pairs. The frames were randomly sized to avoid the use of framing cues by the subjects. Examples of the stimuli which were presented in pairs can be seen in figure 4.1. Each image measured 9 cm x 9 cm and was viewed from a distance of approximately 1m thus subtending a visual angle of approximately 5° x 5°.

Stimuli remained on the screen until subjects made a response which was followed by a gap of approximately one second before the next pair. Responses were made using the "shift" key on the right of the keyboard and the "apple"/"command" key on the left.

Design.

The core design for all experiments except one was a 2(View) x 2(Light) within subjects design, with the levels of each factor being "same" or "different". The levels of view and light refer to properties of each pair, as defined by the correspondence of their individual viewing conditions. Half the trials were "Same" trials, in the sense of being composed of two images of the same person, and half were "Different" trials, being composed of images of different people. Capital letters will be used for Same and Different when referring to identity and response and lower case letters when referring to viewing conditions. For each subject there were 144 trials in an order randomised with the restrictions that no two consecutive "Same" trials were of the same person and that there were no more than three "Same" or "Different" trials in a row. Both hits and false alarms were recorded; that is "Same" responses to "Same" pairs and "Same" responses to "Different" pairs. The "Different" trials used the same set of faces as were used for the "Same" trials, paired in a pseudo-balanced way to avoid any biases due to effects of different degrees of similarity between different faces.

Combinations of three-quarter and profile views were used for all experiments but the lighting directions used varied between experiments. Which were used will be specified in individual method sections. Both same view and same light conditions could be further subdivided. For example for same view pairs both faces could be presented as profiles or both as three-quarter views. In order to investigate if there were important differences between these alternatives a 3(View) x 3(Light) analysis was planned with three-levels of both light and view variables. The levels of view for this analysis were profile/profile, three-quarter/three-quarter and profile/three-quarter. The levels of light were Light1/Light1, Light2/Light2 and Light1/Light2 or Light2/Light1 with the directions used for Light1 and Light2 differing between experiments. The full design is show in table 4.1.

There were twice as many same view as different view trials but equal numbers of same light and different light trials. This ensured that all combinations of view and light were balanced while keeping the number of trials per subject as low as possible. More different light than different view trials were used because the effects of changing view have already been more widely investigated than the effects of changing light.

Table 4.1: The Full Within Subjects Design For Matching Experiments.

LIGHTING VIEW	Light1/Light1 (same light)	Light1/Light2 (different light)	Light2/Light1 (different light)	Light2/Light2 (same light)
PRF/ PRF (same view)	6 "Same" & 6 "Different"	6 "Same" & 6 "Different"	6 "Same" & 6 "Different"	6 "Same" & 6 "Different"
PRF/ T-Q (different View)	6 "Same" & 6 "Different"	6 "Same" & 6 "Different"	6 "Same" & 6 "Different"	6 "Same" & 6 "Different"
T-Q/ T-Q (same view)	6 "Same" & 6 "Different"	6 "Same" & 6 "Different"	6 "Same" & 6 "Different"	6 "Same" & 6 "Different"

Table 4.1: The table shows the within subjects design for the matching experiments with the different levels of lighting and view. The same viewpoints were used for all experiments, PRF-Profile and T-Q - Three-quarter. These were combined to give "same view" and "different view" trials as shown. The particular directions of lighting used varied between experiments but for each there were two, here labelled Light1 and Light2. These were combined to give "same light" and "different light" trials also as shown. Within each cell there were twelve trials, six showing the same individual (The "Same" trials) and six showing different individuals (The "Different" trials.).

Procedure.

The experiment took place in a small windowless room lit by a single fluorescent light. Subjects were told that they would be presented with pairs of head models and that their task was to decide if both images were of the same or different people. It was stressed that the "Same" representations would not necessarily be identical but might vary in lighting and/or viewpoint but that "Different" trials would show different people. The keys with which the responses were made were the "Apple" key on the left of the keyboard and the "Shift" key on the right of the keyboard. Subjects chose which key to use for which responses.

Ten practice trials preceded each experiment, with the same combinations of stimuli as the subsequent experiment but using different individual faces. These were to familiarize subjects with the nature of the stimuli and the task. Images during both the practice session and the experiment remained on the screen until a response was made, although subjects were encouraged to respond as quickly and as accurately as possible and were told that their response latencies would be recorded. The inter stimulus interval averaged one second, although it varied slightly preventing subjects from anticipating their responses. In both practice and experimental sessions a single buffer trial, to which subjects responded as normal, marked the beginning of testing. Subjects were informed that they could have a break during testing by holding down the key that they had just pressed.

TREATMENT OF RESULTS

For all experiments hits and false alarms were combined into a single measure of performance, A' (Norman, 1964; McNicol, 1972), the same measure of sensitivity as was used in experiment 6. All the effects reported were present when hits were analysed separately but the use of A', a criterion free measure of sensitivity (Rae, 1976), took account of the possibility of a response bias.

A' was again transformed using 2arcsin√(A') prior to analysis of variance. Both 2(View) x 2(Light) and 3(View) x 3(Light) ANOVA's, with levels as described in the design section, were carried out. Planned pairwise comparisons were used to test for differences between the additional levels included in the 3(View) x 3(Light) analysis. Pairwise comparisons were carried out on the means from these analyses to see if there were important differences between top and bottom lighting and between three-quarter and profile views. Planned t-tests were used for this purpose with the significance levels divided by the number of comparisons in each case.

Graphs will be presented of raw A' s to show the pattern of performance, together with tables of transformed data for more formal comparisons.

Response latencies were recorded but are not reported. "Same" response latencies were analysed for some experiments and showed that subjects were slow in the same conditions as they made most errors, suggesting that the results shown in the accuracy data are not the result of a speed accuracy trade-off.

4.3 EXPERIMENT 7

INTRODUCTION

In this experiment subjects' abilities to match across changes in light and view were investigated. Profile and three-quarter views were used along with top and bottom lighting (see figures 4.1 and 3.2 b, c, e & f). Different effects of lighting and viewpoint were expected as explained in the general introduction. Two groups of subjects were run, the members of one of which were familiar with the people used as stimuli. The lighting and viewpoint manipulations will be discussed first followed by the between groups factor.

Differences between top and bottom lighting have already been reported on tasks including seeing a face as a face (Chapter 2) and recognising individual faces (Chapter 3). If the explanation offered in terms of shape-from-shading being important and incorporating a light-from-above assumption is correct then this difference may transfer to the matching task. In particular if bottom lighting interferes with the encoding of the three-dimensional shape of the face and this is necessary for matching between views a difference between top and bottom lighting would be expected in the different view condition. However, if bottom lighting merely highlights less salient features then there might be no difference on a matching task where salience would presumably only be determined by the features visible in the other image.

As already discussed, effects of changing viewpoint on matching were expected. As changing viewpoint within a uniform light field entails a change in lighting relative to the object and effects of viewpoint and lighting would not be expected to be independent, the effects of lighting and viewpoint might also be expected to interact as they did in the recognition experiments reported in the last chapter. In particular lighting might have less of an effect on profile views than on three-quarter views as was found in experiment 5, reflecting the relative salience of occluding contour information.

Two groups of subjects were run to investigate if prior knowledge affected performance. Subjects familiar with the people used as stimuli might be able to match the faces by recognising both, a dual recognition strategy not available to the members of the unfamiliar group. Such alternative strategies might affect relative sensitivity to changes in viewpoint and lighting. In general familiar faces are known to be processed differently in some ways from unfamiliar faces with, for example, internal features being more salient for familiar faces (Ellis, Shepherd & Davies, 1979). Any differences in the effects of lighting and viewpoint on the perception of familiar and unfamiliar faces should be shown by a difference between the groups.

METHOD

Subjects & Materials

Two groups of subjects took part in the experiment, a "familiar" group and an "unfamiliar" group. All members of the "familiar" group were from the same department as the people used as stimuli. The two levels of lighting used were top1 and bottom1 which are described, along with all other details of subjects and materials, in the general methods section.

Design and Procedure

The basic design reported in the general methods section was used with an additional between subjects factor of Familiarity with the people used as stimuli. Thus the experiment was a 2(Familiarity) x 2(View) x 2(Light) mixed factorial design. The levels of Familiarity were familiar/unfamiliar and the levels of lighting top/bottom. All other details of the design and details of the procedure were as given in the general methods section.

RESULTS.

Mean A' s for the 2 (View) x 2(Light) analysis are shown in figure 4.2, collapsed across familiarity factor. It can be seen by inspection that matching was best when pictures were identical and that changing either light or view reduced performance. Table 4.2 gives the means for the transformed data for the full 3 x 3 analysis. Table 4.2a gives means for the familiar group, Table 4.2b means for the unfamiliar group and Table 4.2c the means collapsed across groups. Subjects familiar with the faces used as stimuli performed better overall, but familiarity did not affect relative sensitivity to changes in view or light.

The 2(Familiarity) x 2(View) x 2(Light) analysis of variance on the transformed data gave independent main effects of Familiarity ($F_{1,22}=8.7$, $p<<0.05$), View ($F_{1,22}=15.7$, $p<<0.05$) and Light ($F_{1,22}=27.3$, $p<<0.05$). The interaction between the within groups factors, View x Light just missed significance ($F_{1,22}=4.2$, $p=0.052$). Analysis of simple main effects was carried out and showed effects of light for same view pairs ($F_{1,44}=21.7$, $p<<0.05$) but not for matches between views ($p > 0.1$). There were simple main effects of view when light was the same ($F_{1,44}=17.6$, $p<<0.05$) but not when light was different ($p > 0.1$). This suggests that changing either light or view reduced performance but that changing both did not produce a significant additional effect.

The 2(Familiarity) x 3(View) x 3(Light) ANOVA gave an independent main effect of familiarity ($F_{1,22}=8.2$, $p<<0.05$) and a view x light interaction ($F_{4,88}=4.1$, $p<0.05$). There were simple main effects of light for prf/prf pairs ($F_{2,132}=5.6$, $p<<0.05$), t-q/t-q pairs ($F_{2,132}=15.9$, $p<<0.05$) and prf/t-q pairs ($F_{2,132}=5.5$, $p<<0.05$). There were simple main effects of view for top/top pairs ($F_{2,132}=7.3$, $p<<0.05$) and for bottom/bottom pairs ($F_{2,132}=17.1$, $p<<0.05$) but not for top/bottom pairs ($p>0.1$).

Planned pairwise comparisons between the means shown in Table 4.2c did not show any direct differences between profile and three-quarter views. However, when both heads were lit from above only performance with pairs of three-quarters differed significantly from performance matching between views. There was a difference between top and bottom lighting in that different view pairs were matched significantly more accurately when both faces were lit from above than when both were lit from below.

Fig. 4.2: Mean A' s For Experiment 7

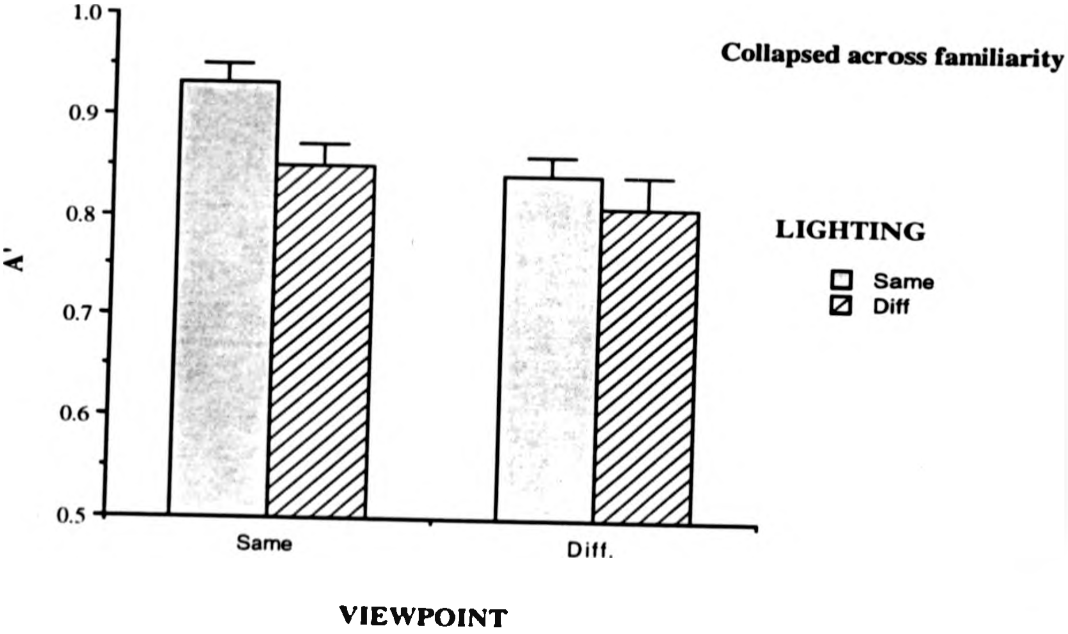

Table 4.2: Means of Transformed A' s for the 3(View) x 3(Light) Analysis. Experiment 7.

a) Familiar Subjects.

VIEW	LIGHT Top/Top	 Bott/Bott	 Top/Bott
Prf/Prf	2.81	2.70	2.49
T-Q/T-Q	2.87	2.87	2.50
T-Q/Prf	2.63	2.46	2.42

b) Unfamiliar Subjects.

VIEW	LIGHT Top/Top	 Bott/Bott	 Top/Bott
Prf/Prf	2.38	2.62	2.28
T-Q/T-Q	2.83	2.58	2.25
T-Q/Prf	2.37	1.99	2.16

c) Means Collapsed Across Familiarity.

VIEW	LIGHT Top/Top	 Bott/Bott	 Top/Bott
Prf/Prf	2.60	2.66	2.38
T-Q/T-Q	2.85	2.73	2.38
T-Q/Prf	2.50	2.22	2.29

DISCUSSION

The results of the experiment show that there were effects of lighting, viewpoint and familiarity on matching performance. Subjects familiar with the people used as stimuli performed better overall, but their knowledge did not affect their relative sensitivity to changes in viewpoint and lighting. Identical images were matched best, as would be expected, with changes in either light or view reducing performance. Changing both light and view did not produce an additional effect compared to changing either variable alone.

No effects of changing lighting were expected on the basis of evidence from physiology and edge based object recognition schemes cited in the introduction. However it seems that, for a task that requires human observers to distinguish people on the basis of information derived from the three-dimensional shape of their faces, lighting is an important variable. Even when viewpoint remained the same, making available lighting invariant occluding contour information for matching, there was an effect of changing lighting on matching performance. The effect of lighting on pairs of profile views was particularly surprising (given the lack of an effect of lighting on profiles in experiment 5) and suggests that lighting is important for the perception of profiles and occluding contour information not always sufficient. When lighting was changed there was no additional effect of changing viewpoint, again highlighting the importance of lighting dependent information over occluding contours affected by view but not lighting. These results suggest that the level of description being used for matching is affected by lighting and that lighting invariant information, like the shape of occluding contours, is not sufficient for this task.

As suggested in the introduction bottom lighting may present particular problems for deriving information about three-dimensional shape. When view and light were the same there was no difference between top and bottom lighting, but under these conditions pictorial codes would be sufficient for matching and derived descriptions would not be required. However, when subjects had to match between views performance was better if both faces were lit from above than if both faces were lit from below although the change in view was no different. Matching between views requires the derivation of structural codes, and the results suggest that this may be affected by whether lighting is from above or below. The advantage for top lighting is consistent with the use of shape-from-shading information as this is thought to incorporate a light-from-above assumption. It was suggested that bottom lit faces may be more poorly recognised because this direction of lighting highlights less salient features. However, this seems unlikely to account for difficulties matching between views as the same features would be highlighted in both images, those facing downwards, and thus be salient. Differences between top and bottom lighting are further investigated in experiments 9, 10, 11 and 12.

Effects of changes in viewpoint on matching performance were expected. Most three-dimensional objects project a very different retinal image depending on the angle from which it is viewed and different faces viewed from the same viewpoint will look more similar in many respects than the same face viewed from different viewpoints. What the results show is that the effects of viewpoint do not appear to be independent of the effects of lighting. This is perhaps not surprising as the viewpoint changes used were rotations within a uniform light field, and so the direction of lighting was in effect varying with respect to the face. However, that there was no additional effect of changing viewpoint when lighting was different was unexpected as changing viewpoint undoubtedly affects information, such as occluding contours, which changing lighting alone leaves unchanged. This, together with the effect of lighting even when viewpoint was unchanged, suggests that occluding contour or silhouette information is not sufficient for discriminating between individuals. This must be reconciled with the absence of an effect of lighting on the recognition of profiles (experiment 5).

Lastly with regards view, previous research has suggested that some tasks involving face processing might show an advantage for three-quarter views as these might serve as canonical representations (Bruce et al, 1987). This experiment suggested that pairs of three-quarter views may be marginally better matched than pairs of profiles when lighting is from above. However, any three-quarter view advantage did not generalize to when lighting was from below or different suggesting that the difference between views is not that great. Profiles would have been expected to contain more salient information about the shape of the face in their occluding contours but were found to be as much affected by changes in lighting direction as pairs of three-quarters.

Individuals familiar with the people used as stimuli performed better overall showing that prior knowledge facilitated the task. There are many possible explanations for this, including dual recognition as discussed in the introduction. However, both familiar and unfamiliar faces appeared to be affected by changes in viewpoint and lighting in much the same way. This could be because the processes of visual encoding are largely the same but with familiarity making the matching process more reliable because of the availability of additional knowledge.

The matching of laser heads was found to be sensitive to lighting and viewpoint whether or not the stimuli used are familiar to the subjects. The effects of changes in lighting direction seem inconsistent with edge based accounts of object recognition which might be applied to matching. It may be that matching is image based, with any changes in the image disrupting performance. However, absolute grey level differences do not easily account for differences between, for example, top and bottom lighting when matching across the same change in view. These effects seem more consistent with an explanation in terms of more abstract visual codes derived from the image perhaps involving shape-from-shading.

In the next experiment reported the same stimuli were presented inverted, a manipulation which would not affect grey levels differences between images but which would affect the derivation of shape-from-shading as top lighting becomes bottom lighting and vice versa (Johnston, Hill & Carman, 1992). Inversion is also thought to disrupt the derivation of configural information (Young, Hellawell & Hay, 1987; Valentine, 1988; see also Chapter 1) and the next experiment served as a test of whether the task was reflecting properties of face processing or general pattern matching. The effect of familiarity was not investigated further as it had not been found to affect relative sensitivity to lighting and view changes, but also because there was only a limited pool of subjects familiar with the people used as stimuli available.

4.4 EXPERIMENT 8

INTRODUCTION

In the last experiment effects of both lighting and viewpoint on matching performance were reported. This experiment was intended as a test of whether these results reflected general properties of pattern matching or aspects of face processing in particular. To this end this experiment was a repeat of experiment 7 but with all the stimuli inverted, a common control condition for experiments concerning face processing (eg. Yin, 1969; Diamond & Carey, 1986; Young et al, 1987; Johnston et al, 1992; see also experiments 1-4, chapter 2).

The stimuli used in this experiment were exactly the same ones as used in experiment 7, the only difference between the two experiments being the orientation in which they were presented. Thus any differences found in the effects of lighting and viewpoint cannot be attributed to differences in grey level descriptions, as these would have been unaffected by the change in orientation. General knowledge about three-dimensional surfaces and objects would also not explain any such differences as inverted faces are also three-dimensional. If, for example, "allowable" transformations specify possible image variations that can result from changes in viewing conditions rather than differences in structure these would be equally applicable to inverted faces. Alternatively, if invariant features are used to match between viewing conditions these too should be recoverable independently of orientation and support matching even when the images are inverted. A study using Identikit faces concluded that both upright and inverted faces could be matched using a serial self-terminating comparison of facial features (Bradshaw & Wallace, 1971) but the effects of inversion on matching between views or lighting conditions have not been assessed. Only if performance on this matching task requires knowledge of upright faces and/or configural information (cf. Diamond & Carey, 1986) would a different pattern of effects would be expected in this experiment compared to the results of experiment 7. There are many ways in which knowledge of upright faces may support matching performance. For example, it is possible that matching faces across viewing conditions is dependent on accessing a stored object model which is aligned to the input image (Ullman, 1989) and that this process might break down for inverted faces. A clear difference in the pattern of effects between this experiment and the last would imply that the matching task was reflecting aspects of face processing, and not just general pattern matching, and may provide clues as to the nature or access of the information required.

The views used in these experiments differed by 45° in both upright and inverted orientations and an effect on performance of presenting the faces in different views would again be expected. If the ability to match between views is dependent on knowledge of upright faces, in the ways suggested for any changes in viewing conditions above, then subjects may be less able to match between views when the stimuli are presented inverted. This would suggest that performance on matching between views is not just determined by the magnitude and type of resulting image changes but is also determined by knowledge of upright faces.

In experiment 7 lighting was also found to have an effect on the matching of upright faces. If this was the result of the grey level changes that lighting changes introduce to the image an equivalent effect would be expected with inverted stimuli, as such differences will be as great with inverted images. Also, if the effect of lighting is because different lighting directions highlight different features then the effect would also be expected to transfer to inverted images. Only if matching across lighting conditions in some way requires knowledge of upright faces would different and more extreme effects of changing lighting condition be expected. In chapter 2 the effects of lighting and inversion on perception of the hollow face illusion were found to be additive, but this was on a very different task involving perception of the face as a face and not requiring discrimination between faces. However, inversion may affect the ability to integrate information about the face and see it as a whole, a gestalt, and this may be necessary for matching between lighting conditions. Evidence from patients with right hemisphere damage suggests that this may be necessary (Benton & Van Allen, 1968).

As well as testing whether the effects of lighting on inverted faces were different, the experiment was designed as a test of the relationship between top and bottom lighting. The results of experiment 7 suggested a difference between top and bottom lighting especially for matching between views and this may reflect a role for shape-from-shading and/or a light-from-above assumption. Figural inversion entails a reversal of apparent lighting direction relative to the viewer, top becomes bottom and vice versa, and so "bottom" lighting would satisfy the light-from-above assumption in the inverted orientation (cf. Johnston et al, 1992). This might be reflected in an advantage of bottom lighting over top lighting in this experiment.

In experiment 7 three-quarter views were found to be better matched than profile views when stimuli were top lit. Differences in the information that is important for the perception of different views were also suggested by experiments 5 & 6, with shading information likely to be more important for three-quarter views and contour information for profiles. Some of these differences may hold when stimuli are inverted, the contours of profiles remaining more salient for example, and so the effects of light may be different for pairs of profiles or pairs of three-quarters.

METHOD.

All details of the methods used for this experiment were the same as for experiment 7 except that all stimuli were presented upside-down and only twelve subjects were run, recruited from the University of Stirling, all of whom were unfamiliar with the people used as stimuli.

RESULTS.

The treatment of results was the same as for experiment 7. Mean A' s for the 2 x 2 analysis are presented graphically in Fig. 4.3 and Table 4.3 shows the A' s for the full 3 x 3 analysis. These values were transformed before analysis of variance. As can be seen from Fig. 4.3 inverting the stimuli appears to have caused an overall drop in performance, overall mean A' 0.71, compared to the unfamiliar group in experiment 7, mean 0.82. The effects of changing light and particularly view appear more pronounced, with changes in view reducing performance to near chance.

Analysis of variance on transformed data confirmed this pattern of results. In the 2(View) x 2(Light) analysis the interaction just missed significance ($F_{1,11}= 4.3$, $p=0.06$) but there were main effects of light ($F_{1,11}= 5.9$, $p<0.05$) and view ($F_{1,11}= 43.2$, $p<<0.05$). In the 3(view) x 3(light) ANOVA on the transformed data presented in Table 4.3 below the view x light interaction was significant ($F_{4,44}= 3.4$, $p<0.05$). Analysis of simple main effects showed effects of light when both views were profiles ($F_{2,66}= 3.4$, $p<0.05$) and when both views were three-quarters ($F_{2,66}=6.9$, $p<0.05$) but no additional effect of light when view was different ($F_{2,66}=1.21$, ns). There were simple main effects of view at all levels of light; for "top" lit (from below relative to the subjects) pairs ($F_{2,66}=22.3$, $p<<0.05$), for "bottom" lit (from above relative to the subjects) pairs ($F_{2,66}= 8.2$, $p<<0.5$) and for differently lit pairs ($F_{2,66}= 5.3$, $p<<0.05$).

Planned comparisons on the means in table 4.3 did not show any direct differences between the pairs of profiles and pairs of three-quarters or between top lit pairs and bottom lit pairs. There was a slight advantage for "bottom" lighting for matching between views as expected but this did not reach significance. The only evidence for an advantage for "bottom" lighting was for pairs of three-quarter views, when only this same light condition was significantly better than the different light condition. In the different light condition only pairs of profile were matched better than different views, perhaps reflecting the use of occluding contour information.

Fig. 4.3: Mean A' s Experiment 8

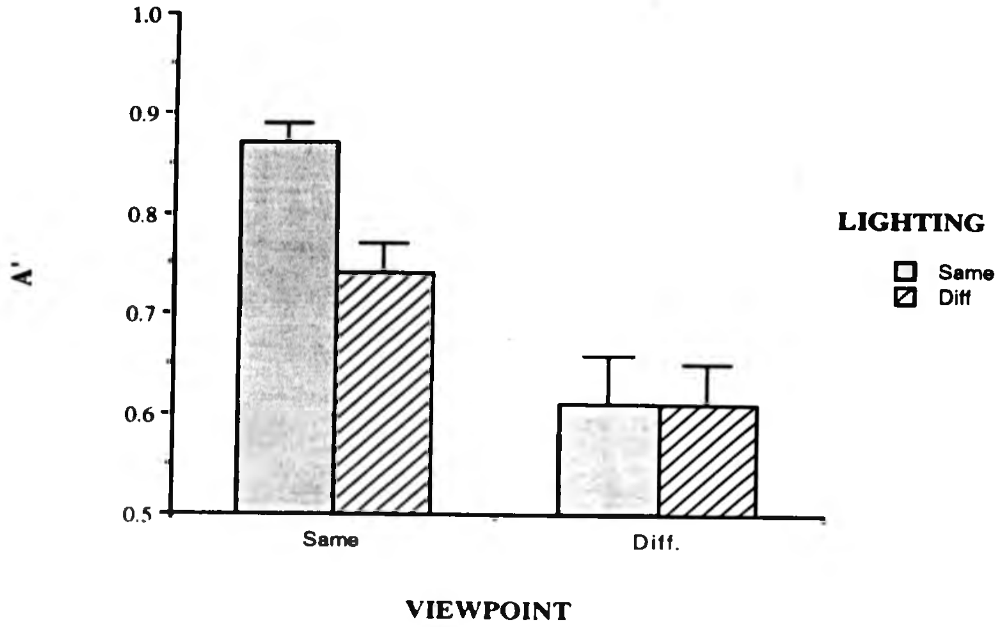

Table 4.3: Means of Transformed A' s for the 3(view) x 3(light) Analysis, Experiment 8.

VIEW		LIGHT	
	Top/Top	Bott/Bott	Top/Bott
Prf/Prf	2.57	2.25	2.23
T-Q/T-Q	2.39	2.48	1.98
Prf/T-Q	1.71	1.93	1.79

DISCUSSION.

The results of this experiment show that inverting the stimuli did affect subjects' abilities to match across changes in viewing condition. Changes in view reduced performance almost to chance but there was also a clear effect of changing light alone. Comparison with experiment 7 shows that performance was lower overall and more mistakes were made even when the pictures were presented under identical viewing conditions (the raw data showed that this was due to a number of false positives in this condition). The main conclusion was that matching performance is affected by knowledge of upright faces or configural information and cannot be explained simply in terms of changes in raw grey level information.

There was a clear identical picture advantage showing that pictorial codes are sufficient to distinguish faces presented inverted so long as viewing conditions are the same. The fact that some errors were made even in this condition suggests that laser heads may be confused when presented inverted, perhaps more so than more natural images as they can vary in fewer ways. Inversion must have disrupted the derivation of information that allows the laser heads of different people to be distinguished, suggesting that configural information may be important for this purpose.

Inversion also affected subjects' abilities to match between viewing conditions. Matching between views appeared particularly affected. The difference between the two views, 45°, was no greater when stimuli were inverted than when they were upright. Differences between grey level descriptions would also be no greater when stimuli were inverted suggesting that the ability to match between views does not have a simple explanation at this level. Instead it seems likely that the problem is with deriving view insensitive descriptions, and the effect of inversion again suggests that configural information may have a role in this. Configural information itself may be largely independent of view.

As with upright stimuli there was a clear effect of changing lighting when view was unchanged. Again lighting independent occluding contours do not appear to be sufficient to distinguish between faces. An informal comparison with experiment 7 shows that performance in this condition was worse when stimuli were inverted (0.74 vs 0.81) and also that the drop in performance caused by changing lighting alone was greater (-0.13 vs -0.09). This suggests that matching across changes in lighting may be facilitated by knowledge about upright faces. However, inversion affected the ability to match across changes in view more suggesting that there is useful lighting but not view invariant information available when stimuli are presented upside-down. A possible candidate is occluding contour information perhaps sufficient to distinguish between at least some pairs of faces. Consistent with this there appears to be a trend for pairs of profiles to be least affected by lighting condition. In this experiment there was no effect of light when view was different but this probably reflected a floor effect. This may also have been the reason for the lack of the expected advantage of "bottom" lighting for matching between views when the stimuli were inverted although the trend was in the right direction.

In summary, matching performance was found to be much poorer when stimuli were inverted, suggesting that performance with upright stimuli does reflect aspects of face processing. Matching between viewpoints seemed particularly impaired by inversion suggesting that face specific knowledge may be particularly important for this task. Performance was above chance for matching between lighting conditions when view was not changed, probably achieved because of lighting independent features including occluding contours are useful even when inverted. Matching performance was best for matching inverted profiles across lighting conditions, consistent with this view providing most useful contour information. Evidence for the importance of shading information, particularly for three-quarter views, was more equivocal probably because the difficulty of the task suppressed sensitivity to such differences.

4.5 EXPERIMENT 9

INTRODUCTION.

Experiment 7 showed large effects of changing lighting direction when it was changed from above to below the face, top to bottom. The same change has been shown to have an effect on the perception of the hollow face illusion (see Chapter 2) and on the recognition of laser heads (see Chapter 3) and photographs (Rock, 1973; Johnston et al, 1992). The aim of the next two experiments reported was to test whether these effects were the result of violating a light-from-above assumption used in the interpretation of shape-from-shading (Ramachandran, 1988a & 1988b) or whether they are a function of the image changes resulting from any change in lighting direction.

This experiment investigated the effects of changing between two different directions of top lighting (see figure 4.4). Any effect of this change cannot be the result of a light-from-above assumption as this was satisfied by both levels of lighting. A change between directions of top lighting does still result in changes to the image which might still produce an effect on matching performance. In order to equal the magnitude of the change in lighting direction the two directions of top lighting should have been 90° apart from each other as the levels of top and bottom lighting used for previous experiments had been. For profiles a 90° rotation in lighting in the same direction as rotation of the head resulted in the face being lit from the direction in which it was looking. This meant that the image produced was different from when lighting was from the original direction, as intended, but that the face still received direct illumination. However, for three-quarter views a light source in the same position would result in surfaces not visible to the viewer being highlighted, increasing the chance of an effect of lighting. As a result it was decided to also light three-quarter views from the direction in which they were facing, a difference from the original direction of lighting of only 45° but one which ensured that visible surfaces were highlighted. The image projected was still different from when a three-quarter view was lit from the original direction of top lighting and so if image changes are the cause of the effect of lighting they will again apply here. The new direction of lighting for both three-quarter and profile views is fully described in the methods section.

The new lighting direction also provided a test of the relationship between lighting and viewpoint. Under the original directions of lighting a change in view entailed a change in lighting *with respect to the face*. The pattern of shading and shadows on the surface of a three-quarter and a profile view lit from the original direction of lighting would be different. However, when both three-quarter and profile views were lit from the direction in which they were facing the pattern of shading and shadows should remain constant between views. If it is the effect on shading and shadows that is the cause of the effect of view then there should be less of an effect of changing view when both heads are lit from the new direction of light.

This experiment also provided a test of the effect of lighting condition on matching between views reported in experiment 7. There matching across views was better when light was from above than when it was from below. All light was from above in this experiment and so matching between views would be expected to be accurate, although still significantly lower than when viewing conditions are identical.

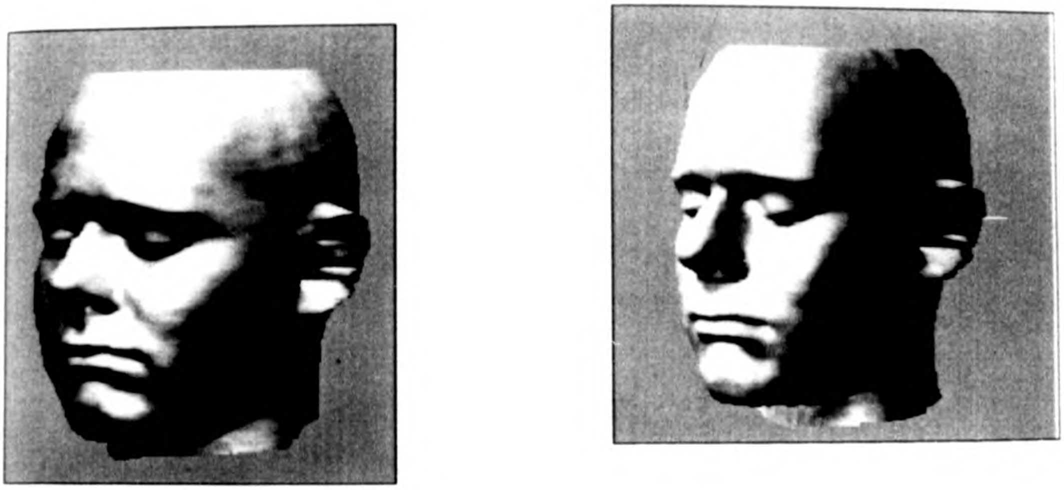

a) Top1 Lit, Three-Quarter.

b)Top2 Lit, Three-quarter.

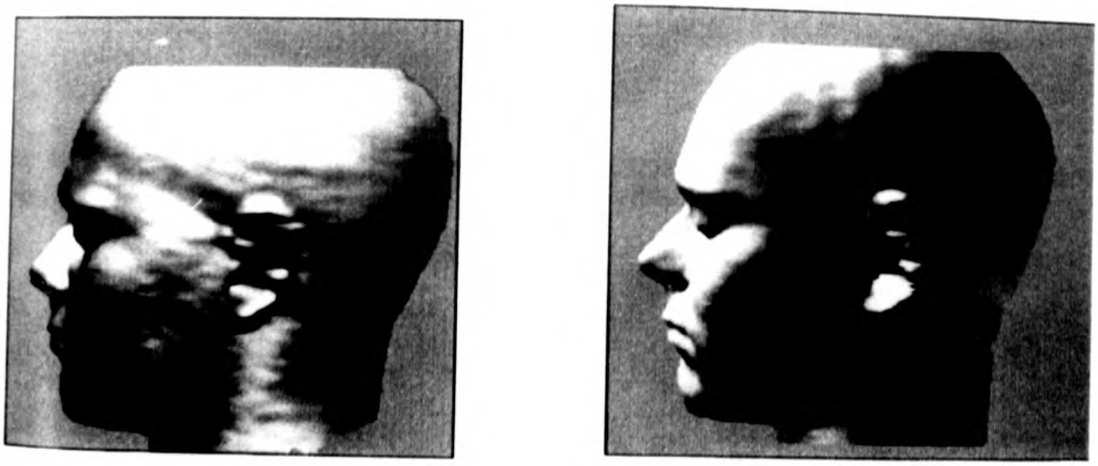

c) Top1 Lit, Profile

d)Top2 Lit, Profile.

Figure 4.4: Examples of the two directions of top lighting used in experiments 9 & 10. Same view, different light pairs.

METHOD.

The method for experiment 9 was the same as for experiment 7 except for the following points. Only twelve subjects were tested, all unfamiliar with the people used as stimuli and the laser head format. The design of the experiment was the same except that there was no between subjects familiarity factor and the levels of the lighting factor were different reflecting the new direction of lighting used. All conditions involving bottom1 lighting in experiment 7 were replaced by top2 lighting in experiment 9. The new direction of lighting was from 45° above or below the direction in which the face was looking, that is 45° above a vector normal to the plane of the face (see figure 4.4). For profile views the light source coordinates were (-1, 1, 0) for top2 lighting and for three-quarter views (-1, 1, -1). With the model of lighting used there is no difference in intensity for lighting sources that are at a distance greater than 1, and therefore no difference between the two used here.

RESULTS.

Treatment of data was the same as for experiment 7. Raw A' s for the 2(View) x 2(Light) analysis are presented graphically in Figure 4.5 and table 4.4 gives the transformed means for the full 3(View) x 3(Light) analysis. Performance was better in this experiment than for the equivalent, unfamiliar, group in experiment 7 (Average A' 0.89 for experiment 9 and 0.82 for unfamiliar group in experiment 7), probably because all stimuli were top lit.

The 2(View) x 2(Light) analysis of variance gave a main effect of view ($F_{1,11}$= 13.3, $p<<0.05$) but no other effects (all p's > 0.1). In the 3(View) x 3(Light) analysis of variance the view x light interaction just reached significance ($F_{4,44}$= 2.8, $p<0.05$). Analysis of simple main effects showed effects of light for pr/pr pairs ($F_{2,66}$=8.7, $p < 0.05$) but not, quite, for t-q/t-q pairs ($F_{2,66}$= 2.4, $p=0.099$) or for different view pairs ($p > 0.1$). There were effects of view when lighting was from the same direction for both faces, both top1 lit ($F_{2,66}$=3.3, $p<0.05$) or both top2 lit ($F_{2,66}$=12.0, $p<<0.05$) but not when lighting was different ($F<1$).

Planned t-tests on the means shown in Table 4.4 showed differences between the paired profiles and different view conditions when light was from the new direction, top2, and pairs of profiles lit from this direction were also significantly better matched than pairs of profiles lit from different directions. Interestingly there was no difference between profile pairs both lit from direction top1 and differently lit pairs.

Fig. 4.5: Mean A' s Experiment 9

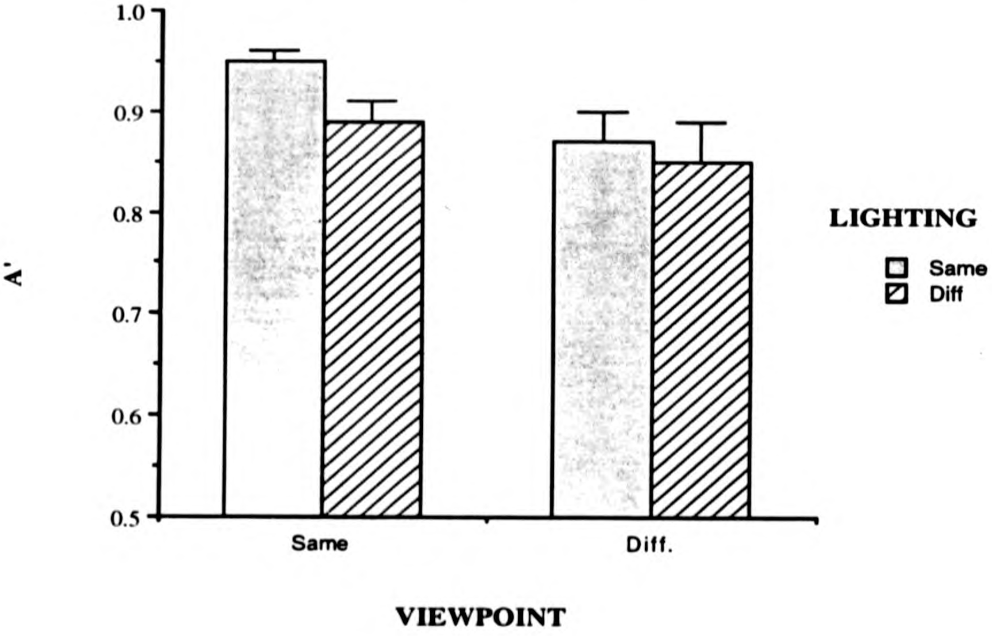

Table 4.4: Means of Transformed A' s for the 3(view) x 3(light) Analysis, Experiment 9.

VIEW	LIGHT Top1/Top1	Top2/Top2	Top1/Top2
Prf/Prf	2.66	2.92	2.48
T-Q/T-Q	2.74	2.68	2.52
T-Q/Prf	2.49	2.43	2.39

DISCUSSION.

There was again a clear effect of changing view in this experiment but the effect of changing lighting direction was reduced or absent when the change was between two levels of top lighting. In particular there was no difference between levels of lighting when view was different, unlike experiment 7. Overall levels of performance were higher in this experiment compared to the equivalent, unfamiliar, group in experiment 7 - mean total A' s 0.89 and 0.82 respectively (2 x 2 analysis). The lack of effects of light and the higher overall performance are both consistent with an advantage for top lighting.

The pattern of data revealed by analysis of variance was also different from experiment 7. There was no main effect of light in the 2(View) x 2(Light) analysis or any interaction suggesting that the effect of light found on experiment 7 was largely a result of violation of a light-from-above assumption. In this experiment there was only an effect of light for pr/pr pairs and this seems to have resulted from an advantage for profiles when lit from the new direction of light. The new level of lighting highlighted the 'face' in the profile view, rather than the sides of the face, and the relatively high level of performance in this condition suggests that this provides useful information for matching.

This advantage did not transfer to three-quarter views. Indeed there was no significant effect of light for this view although this might be explained by the smaller change in light, 45° not 90°. Even so, the lack of an effect showed that not all image changes introduced by changes in lighting affect matching performance. There were also two results inconsistent with the change in angle determining performance; firstly there was no difference matching across changes in lighting for pairs of three-quarters or pairs of profiles although the difference in directions was greater for profiles and secondly there was no difference between top1/top1 lit pairs and differently lit pairs of profiles although the change between the lighting directions was 90°. Instead, the effect of lighting in the 3(View) x 3(Light) analysis reflects a particular advantage for profiles lit from the new direction which highlights their internal features rather than the sides of the face.

There was no difference between levels of light for matches between views. In experiment 7 there was a clear advantage for top lighting but when all lighting was from above it did not matter whether lighting was the same or different. Again this suggests that the results cannot be explained in terms of image differences - changing lighting would have increased these and there might have been differences between top1 and top2 lighting. The results are consistent with top lighting facilitating matching between views. The effect of changing view was at least as great when light was from the new direction although this did not involve a change in the direction of light relative to the face. This may be because we are more used to changes in view within a uniform light field or because the gross effects of view in determining which surface and contours were visible masked any effects of equivalences in surface brightness.

It appears from this experiment that performance was better and that changes in lighting are reduced or absent so long as all lighting is from above. Matching between views also appears better when light is from above, but an effect of viewpoint is not dependent on a rotation with respect to lighting direction To further test differences between top and bottom lighting and to control for the effects of different magnitudes of lighting differences this experiment was repeated inverted and with two equivalent levels of bottom lighting.

4.6 EXPERIMENT 10

INTRODUCTION.

This experiment used the same stimuli as experiment 9 but, like experiment 8, all were presented inverted. It was expected that this would again reduce matching performance across changes in viewing condition, especially viewpoint. It was not known what the effects of changing between the two directions of lighting would be when the stimuli were inverted. In experiment 9 this change was found to produce less effect than changing between top and bottom lighting and this would transfer to inverted stimuli if the change produces less effect on the image. However, when inverted all the stimuli used were lit from below relative to the observer, and violation of a light-from-above assumption may effect matching between lighting directions. Bottom lighting does appear to disrupt matching between viewpoints (see experiment 7) and this is another reason for expecting different view matches to be poor in this experiment.

METHOD.

The method for this experiment was the same as for experiment 9 except that all stimuli were presented inverted.

RESULTS.

Mean A' s for the 2(Light) x 2(View) analysis are presented in figure 4.6 and means of transformed A' s for the 3(View) x 3(Light) analysis are shown in Table 4.5. As can be seen from figure 4.6 changing either light or view reduced matching performance to near chance when top lit stimuli were inverted.

A 2(View) x 2(Light) analysis of variance gave a View x Light interaction ($F_{1,11} = 15.9$, $p << 0.05$). Changing either view or light reduced performance but changing both produced no additional effect. Analysis of simple main effects confirmed this pattern of data with simple main effects of view when light was the same ($F_{1,22} = 32.7$, $p << 0.05$) but not when light was different ($p > 0.1$) and simple main effects of light when view was the same ($F_{1,22} = 35.8$, $p << 0.05$) but not when view was different ($p > 0.1$).

The 3(View) x 3(Light) analysis of variance also gave an interaction between the effects of view and light ($F_{4,44} = 4.6$, $p < 0.05$). Analysis of simple main effects showed effects of light for pairs of profiles ($F_{2,66} = 8.0$, $p << 0.05$) and for pairs of three-quarters ($F_{2,66} = 15.9$, $p << 0.05$) but not for different view pairs ($p > 0.1$). There were also simple main effects of view when both faces were lit from direction top1 ($F_{2,66} = 18.2$, $p << 0.05$) or both from direction top2 ($F_{2,66} = 11.0$, $p << 0.05$) but not, quite, when the faces were differently lit ($F_{2,66} = 2.6$, $p = 0.08$).

Planned comparisons did not show any significant differences between the same view or the same light conditions.

Fig. 4.6: Mean A' s For Experiment 10.

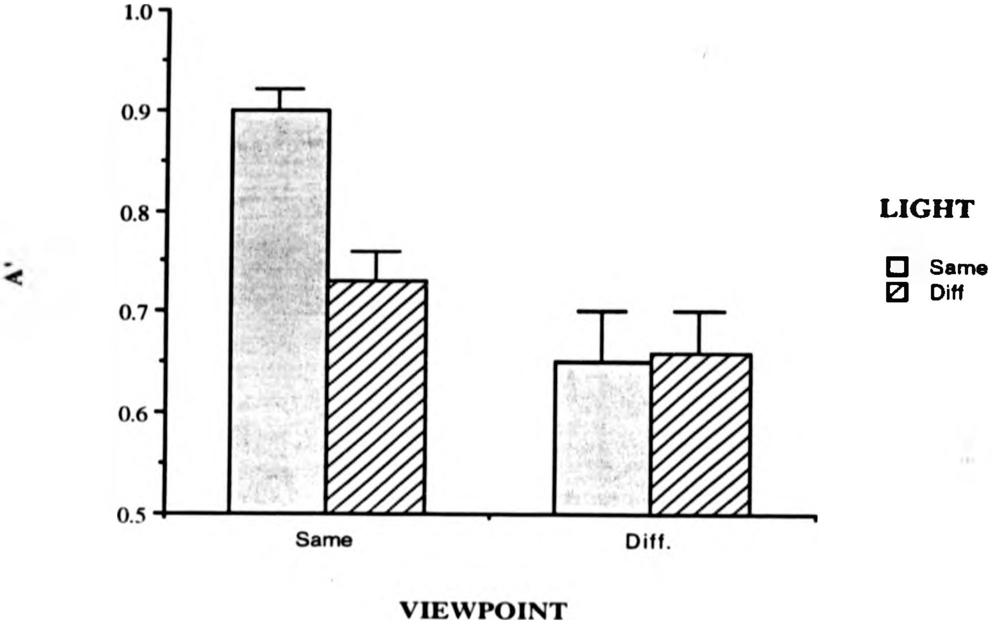

Table 4.5: Means of Transformed A' s for the 3(view) x 3(light) Analysis, Experiment 10.

VIEW		LIGHT	
	Top1/Top1	Top2/Top2	Top1/Top2
Prf/Prf	2.56	2.68	2.18
T-Q/T-Q	2.53	2.51	1.87
T-Q/Prf	1.78	2.07	1.92

DISCUSSION.

Inverting the stimuli reduced subjects' ability to match between faces under different conditions of view or light. As with experiments 7 and 8, exactly the same stimuli are matched more poorly when presented inverted, although image differences are no greater. Again it seems likely that inversion is disrupting the derivation of some sources of information from the image perhaps especially configural information. These effects may have been exaggerated in this experiment because all lighting was from below.

It was expected that matches across changes in viewpoint would be disrupted by inversion (see experiment 8) but subjects also appeared to be worse at matching across the change in lighting when stimuli were presented inverted (see figures 4.5 & 4.6). This suggests that the effects of lighting changes are not determined by the magnitude of the image changes they introduce. In this experiment all stimuli were lit from below relative to the observer and this may be the cause of the low level of matching across changes in lighting, rather than the fact that the stimuli were inverted. This was tested in the next two experiments reported where faces lit from below were presented upright and inverted.

4.7 EXPERIMENT 11

INTRODUCTION.

The design was the same as for experiment 9 except that all lighting was from below the head rather than above (see figure 4.7). The differences in angle between the directions of lighting used were the same and any difference between the pattern of results for the two experiments would have to be explained in terms of a difference between top and bottom lighting.

The results of experiment 7 suggested that matching between views is poorer when stimuli are lit from below. If this is so all different view matches should be especially poor in this experiment. Matches between lighting conditions were also examined. The differences between directions was no greater than for experiment 9 but, as was suggested in the discussion to the last experiment, matching between lighting conditions may also be dependent on a light-from-above assumption. If this is so a strong effect of lighting direction would also be expected in this experiment.

METHODS.

The methods for experiment 10 were the same as for experiment 9 except that two levels of bottom lighting were used instead of two levels of top lighting. The bottom lit direction from experiment 7 was used along with a new direction of bottom lighting, from 45° below the direction in which the face was looking. The new direction of lighting will be referred to as bottom2 lighting. For profile views the coordinates of the light source were (-1, -1, 0) and for three-quarter views, (-1, -1, -1).

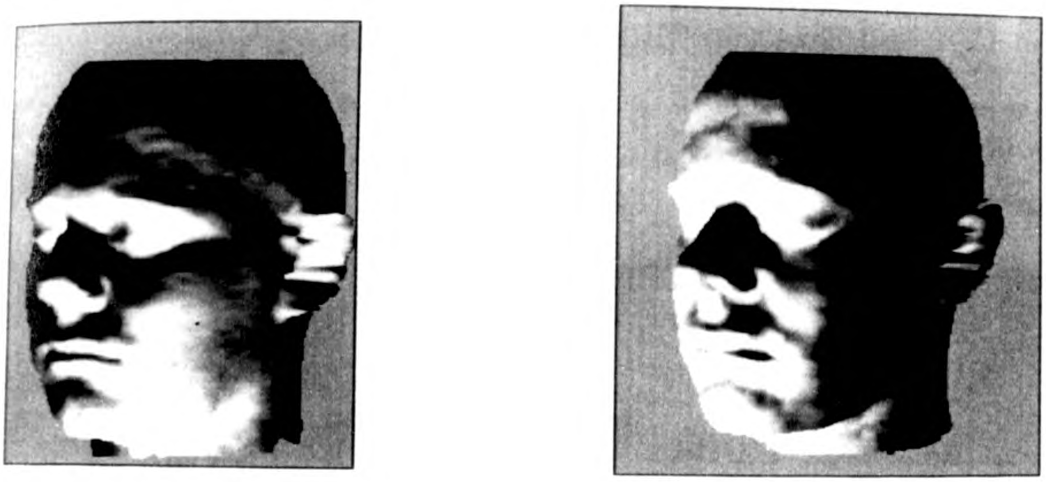

a) Bottom1 Lit, Three-Quarter

b)Bottom2 Lit, Three-Quarter

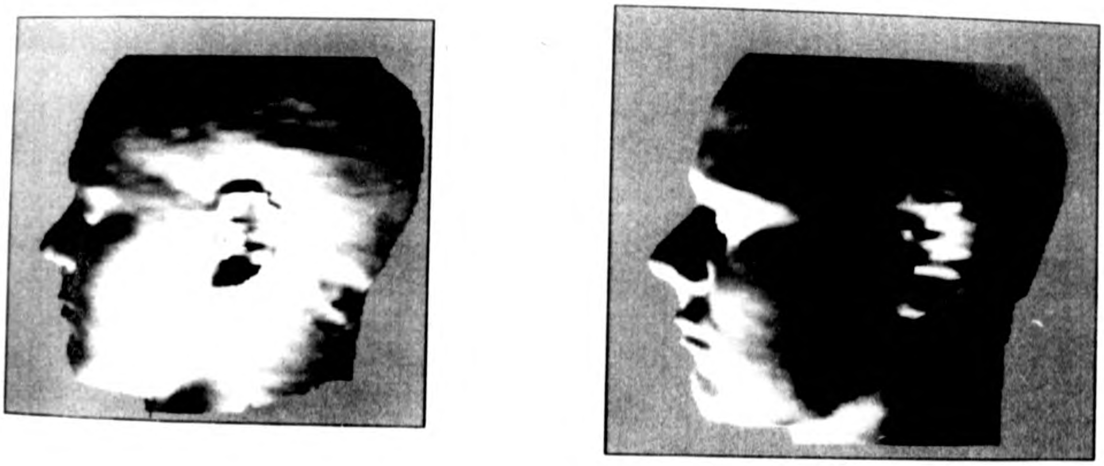

c) Bottom1 Lit, Profile

d)Bottom2 Lit, Profile

Figure 4.7: **Examples of the two directions of bottom lighting used in experiments 11 & 12. Same view, different light pairs.**

RESULTS.

Treatment of data was the same as for experiment 9. Mean A' s for the 2(View) x 2(Light) analysis are shown in figure 4.8 and the means of transformed data for the full 3(View) x 3(Light) analysis appear in Table 4.6. As can be seen from figure 4.8 any change in viewing conditions disrupted performance when stimuli were lit from below. Overall performance was lower when all stimuli were presented lit from below, mean A' 0.81 compared to 0.89 for experiment 9.

A 2(View) x 2(Light) analysis of variance showed a view x light interaction ($F_{1,11} = 27.4$, $p<<0.05$). Analysis of simple main effects showed effects of light when view was the same ($F_{1,22} = 26.1$, $p<<0.05$) and when view was different ($F_{1,22} = 11.3$, $p<<0.05$). The effect of light when view was different was the opposite to what would have been expected, with performance better when light was different than when it remained the same. A possible explanation that can be offered for this is a switch in subject strategy when stimuli are different in conditions of both view and light, perhaps to looking for evidence of similarity rather than difference. There was a simple main effect of view when light was the same ($F_{1,22} = 60.4$, $p<<0.05$) but not when light was different ($F<1$).

The full 3(View) x 3(Light) analysis on the data shown in table 4.6 below also showed a view x light interaction ($F_{4,44} = 11.0$, $p<<0.05$). Analysis of simple main effects showed effects of light when both faces were in profile ($F_{2,66}=15.0$, $p<<0.05$), when both faces were presented in three-quarter view ($F_{2,66}=12.2$, $p<<0.05$) and when the faces presented were seen from different views ($F_{2,66}=4.4$, $p<0.05$). There were simple main effects of view when light remained constant, both lit from direction bottom1 ($F_{2,66}=27.2$, $p<<0.05$) or from direction bottom2 ($F_{2,66}=25.1$, $p<<0.05$) but not when the faces in a pair were lit from different directions ($F<1$, ns).

Planned comparisons showed no differences between same light or same view conditions; bottom1 lit pairs did not differ from bottom2 lit pairs and three-quarter pairs did not differ from profile pairs. Different view pairs were worse matched than same view pairs so long as light remained the same and different light pairs were worse matched than same light pairs so long as view remained the same. When view was different, different light pairs were, surprisingly, significantly better matched than pairs which were both lit from direction bottom1. When light was different there was no difference between conditions of view.

Fig. 4.8: Mean A' sExperiment 11.

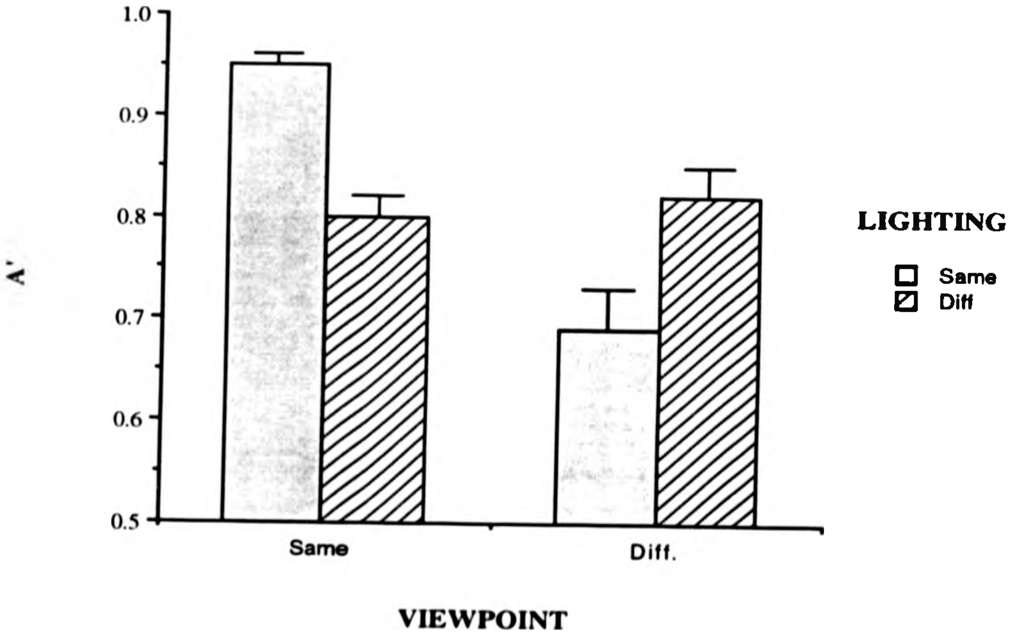

Table 4.6: Means of Transformed A's for the 3(View) x 3(Light) Analysis, Experiment 11.

VIEW	LIGHT Bott1/Bott1	Bott2/Bott2	Bott1/Bott2
Prf/Prf	2.84	2.80	2.29
T-Q/T-Q	2.78	2.63	2.19
T-Q/Prf	2.04	1.93	2.29

DISCUSSION

In this experiment there were clear effects of both view and light which interacted. Identical picture matching performance was high, as expected, despite all stimuli being lit from below. However, changes in either light or view reduced performance significantly. These effects were much larger than those experiment 9, although the only difference was that both lighting directions were from below rather than above. This is taken as further evidence for a difference between top and bottom lighting as elaborated below.

One difference from experiment 9 was that there was a clear effect of changing lighting direction in both analyses here. The actual change in direction in this experiment was no greater, the only difference being that both directions were from below rather than above the head. When top lit faces were inverted so that they were lit from below relative to the observer (see experiment 10) there was a similar effect of changing between lighting directions. Taken together the results suggest that a light-from-above assumption may be as important as head orientation (upright/inverted) in determining sensitivity to changes in lighting. This was further explored in experiment 12 when bottom lit heads were presented inverted so that they were lit from above relative to the observer.

There were also clear effects of changing viewpoint on performance in this experiment. Changing view alone reduced performance almost as much as for inverted stimuli but this was confounded by an improvement when light was also changed. Taken together it still seems reasonable to conclude that subjects were less able to match across changes in view when stimuli were lit from below. This conclusion is also consistent with the difference between top and bottom lighting when subjects were required to match across changes in view reported in experiment 7. Such a difference is not predicted by the magnitude of change in direction and seems unlikely to be purely explained in terms of feature salience. Instead the result, like the effect of lighting, may reflect a light-from-above assumption which is thought, independently, to be important for shape-from-shading (Ramachandran, 1988a & 1988b). This possibility will be considered more fully in the general discussion. The last experiment reported in this chapter was a repeat of this experiment but with all stimuli presented inverted.

explored in experiment 12 when bottom lit heads were presented inverted so that they were lit from above relative to the observer.

There were also clear effects of changing viewpoint on performance in this experiment. Changing view alone reduced performance almost as much as for inverted stimuli but this was confounded by an improvement when light was also changed. Taken together it still seems reasonable to conclude that subjects were less able to match across changes in view when stimuli were lit from below. This conclusion is also consistent with the difference between top and bottom lighting when subjects were required to match across changes in view reported in experiment 7. Such a difference is not predicted by the magnitude of change in direction and seems unlikely to be purely explained in terms of feature salience. Instead the result, like the effect of lighting, may reflect a light-from-above assumption which is thought, independently, to be important for shape-from-shading (Ramachandran, 1988a & 1988b). This possibility will be considered more fully in the general discussion. The last experiment reported in this chapter was a repeat of this experiment but with all stimuli presented inverted.

4.8 EXPERIMENT 12

INTRODUCTION

In this experiment the same stimuli as were used in experiment 11 were presented inverted. Inversion again entailed a rotation of lighting direction and so all the stimuli in this experiment were lit from above relative to the observer. This was intended to test whether orientation or a light-from-above assumption was more important in determining sensitivity to changes in light and viewpoint. The results of experiments 8 and 10 suggest that inverted heads are not matched well across changes in either light or view but the fact that all stimuli are lit from above in this experiment may offset this. The results of experiments 9 and 11 suggest that sensitivity to changes in light may be particularly affected by a light from above assumption. The results of experiment 8 suggest that inversion reduces matching between views to floor which may suppress any advantage associated with lighting appearing to be from above.

METHOD

The method for this experiment was the same as for experiment 11 except that all stimuli were presented inverted in the same manner as for experiment 8.

RESULTS

Mean A' s for the 2(View) x 2(Light) analysis are presented in figure 4.9 and Table 4.7 shows means of transformed data for the full 3(View) x 3(Light) analysis. As can be seen from the graph changes in view again reduced performance to near chance for inverted stimuli, but changes in light had less of an effect probably because all light was now from above.

A 2(View) by 2(Light) analysis of variance gave a main effect of view ($F_{1,11}=76.1$, $p<<0.05$) but no other significant effects, although the main effect of light came close ($F_{1,11}=3.7$, $p=0.08$) (All other p's > 0.1).

In the 3(View) x 3(Light) analysis there were significant main effects of both view ($F_{2,22}=65.0$, $p<<0.05$) and light ($F_{2,22}=7.6$, $p<<0.05$) and the interaction approached significance ($F_{4,44}=2.3$, $p=0.08$). Different view and light conditions were matched worse than same view and light conditions. The marginal interaction probably reflected the lack of an effect of light when view was different.

There were no significant differences between bottom1 lit pairs and bottom2 lit pairs or between profile pairs and three-quarter pairs in this experiment.

Fig. 4.9: Mean A' s Experiment 12

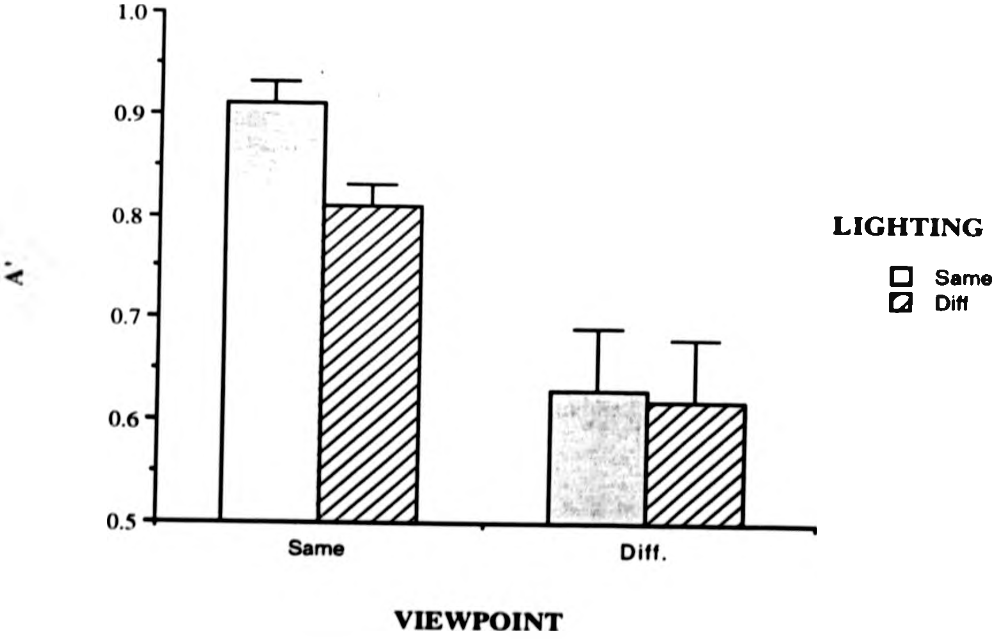

Table 4.7: Means of Transformed A's for the 3(view) x 3(light) Analysis. Experiment 12.

VIEW	LIGHT Bott1/Bott1	 Bott2/Bott2	 Bott1/Bott2
Prf/Prf	2.66	2.71	2.25
T-Q/T-Q	2.66	2.59	2.30
T-Q/Prf	1.94	1.74	1.83

DISCUSSION.

In this experiment, which presented inverted stimuli lit from above relative to the observer, changing view reduced performance, probably to floor, but changing light appeared to have less of an effect.

It was expected that inversion would again disrupt matching between views, although the fact that all light was from above relative to the observer might have offset this. Instead it appears that performance is reduced to floor suppressing any advantage associated with lighting being from above. It appears that information only recoverable from the upright face is critical for matching between views.

There was some evidence in this experiment that lighting from above improved subjects' ability to matching between lighting directions. There was no significant effect of lighting in the 2(View) x 2(Light) analysis and the drop in performance appears less than when the same stimuli were presented upright (see experiment 11). When top lit faces were inverted (experiment 10) there was a definite increase in the effect of changing light, suggesting that a light-from-above assumption may be critical in determining sensitivity. However, subjects were worse at matching across changes in light in this experiment compared to experiment 9 suggesting that head orientation is also important. The results of all the experiments reported in this chapter will be summarised and reported next.

4.9 GENERAL DISCUSSION

The experiments reported in this chapter have shown effects of both lighting and viewpoint on a simultaneous face matching task. The task required a high degree of object constancy as discriminations within the homogeneous class of faces had to be made on the basis of information derived from shape alone. The problem of variable viewpoint for object constancy has been widely studied (see chapter 1) and this manipulation was expected to affect matching performance. However, the problem of variable illumination has received less consideration but the results reported here suggest it may be important for faces.

To summarise the results, an effect of changing between top and bottom lighting, which interacted with the effect of changing between viewpoints, was first reported in experiment 7. There was an effect of changing lighting even when viewpoint remained the same, which was unexpected given the availability of lighting invariant occluding contour information in this condition. Further, there was no additional effect of changing viewpoint once lighting had been changed, although this would have affected the occluding contour unaffected by lighting. This result suggested that information other than occluding contours is critical to matching performance. Lighting does change other sources of contour information, for example shadow boundaries, which could provide clues to shape (Cavanagh & Leclerc, 1989) or be the edges used in edge based theories. However, a difference between top and bottom lighting when subjects were required to match between views was also reported in this experiment. There seems no reason to predict such a difference from edge based theories as there is no reason for expecting bottom lighting to result in greater changes to any edge information. However, a difference between top and bottom lighting would be expected if shape-from-shading, which is thought to incorporate a light-from-above assumption, is important for matching performance. This possibility will be discussed further in the final chapter.

In experiment 8 it was shown that inverting the same stimuli as used in experiment 7 reduced performance across changes in light and, especially, across changes in view. As the magnitudes of image differences were no greater it was concluded that these differences alone did not predict performance. Instead, inversion appears to disrupt the recovery of information useful for matching, perhaps configural information. Subjects also matched poorly across changes in lighting when stimuli were inverted suggesting that this too might be dependent on knowledge about upright faces.

Experiments 9 - 12 further investigated the effects of changes in lighting and viewpoint with both inverted and upright stimuli. In experiment 9 a lighting change was investigated which did not violate any light-from-above assumption as it was between two directions of top lighting. The effects of this change were found to be reduced or absent compared to matching between top and bottom lighting, consistent with the importance of a light-from-above assumption. The only effect of lighting in that experiment seemed to be attributable to an advantage for profiles when they were lit from the direction in which the head was facing. There was still an effect of changing view even when all stimuli were lit from above, although an informal comparison with the equivalent means for experiment 7 (see Tables 4.4 and 4.2b) suggests that the magnitude of the effect was less than when bottom lit stimuli were also included. When these top lit stimuli were inverted (experiment 10), so that they were bottom lit relative to the observer, matching across both lighting and view was less accurate. This was taken as confirmation that matching across viewpoints was substantially affected by inversion, although violation of the light-from-above assumption may also have been important. Matching between lighting directions was also poor when stimuli were inverted although it was not certain if this was because stimuli were inverted or because they were lit from below relative to the observer.

In experiment 11 faces were presented upright but lit from below. Although physical differences between lighting and viewing directions were not greater than in experiment 9 the pattern of results between the two experiment was quite different. When stimuli were lit from below, subjects were less able to match across changes in either lighting or viewpoint. This emphasised the difference between top and bottom lighting and suggested that a light-from-above assumption is important for matching across changes in both view and lighting. This difference does not seem consistent with edge based or image based theories but has a ready explanation in terms of a light-from-above assumption.

In experiment 12 the same stimuli as were used in experiment 11 were presented inverted so that they were lit from above relative to the observer. Inversion still reduced subjects' abilities to match across changes in view, probably to floor. However, changes in lighting had less effect on performance than when the stimuli had been presented upright and much less effect than when "top" lit faces were presented inverted. This advantage for "bottom" lit face over "top" lit faces is inconsistent with the advantage for top lit stimuli presented upright being due to feature salience. Instead, a light-from-above assumption may be critical in determining sensitivity at least to changes in light, although comparison between experiments 9 and 12 suggests that face orientation may also be important.

In summary the results showed effects of light, viewpoint and inversion. The effects of light seem inconsistent with edge based explanations of matching and the effects of inversion with a simple pattern matching account. Differences between top and bottom lighting, like those found in Chapters 2 and 3, suggest that shape-from-shading may be important for the derivation of structural codes which allow matching across changes in viewing conditions. The relationship of the results found using matching to those reported in previous chapters will be elaborated in the final chapter together with their implications for theories of visual aspects of face processing. However, in the next chapter experiments are reported which were designed to test whether these results, particularly the effects of lighting, were artifacts of the task or materials used.

5. CHAPTER 5: CONTROL EXPERIMENTS FOR THE MATCHING TASK

In the last chapter effects of both lighting and view on a simultaneous face matching task were reported. In this chapter three further matching experiments experiments are reported which were intended to test whether the results already reported were an artifact of the task or materials used.

5.1 EXPERIMENT 13

INTRODUCTION

In the last chapter changing view or light was found to disrupt matching performance. It seemed possible that any change from presenting faces under identical viewing conditions might disrupt subjects abilities to match faces, especially if subjects were matching largely on the basis of pictorial codes. This was tested in this experiment where size was varied as well as viewpoint.

Changing the size of a face has profound effects on the image and so if performance is simply image based a change in size would be expected to reduce matching performance. Size constancy of objects despite variable image size is important in the real world as it is related inversely proportionally to viewing distance. The area of the image of a face will quarter when viewing distance is doubled. However, there are other properties of the image that transformations in size leave unchanged, for example proportions and the overall pattern of shading and shadows.

Recent work on the naming of briefly presented objects has been shown to be size invariant, as has an object category decision although "same" responses in an episodic memory task were found to be affected by size (Biederman & Cooper, 1992). The conclusion drawn from these experiments was that the recognition of shape was size invariant but that details of viewing conditions are stored separately. In this experiment subjects were required to match faces on the basis of shape and specifically instructed to ignore details of viewing conditions. If matching involves the recognition of shape and this is size invariant, then no effect of changing size

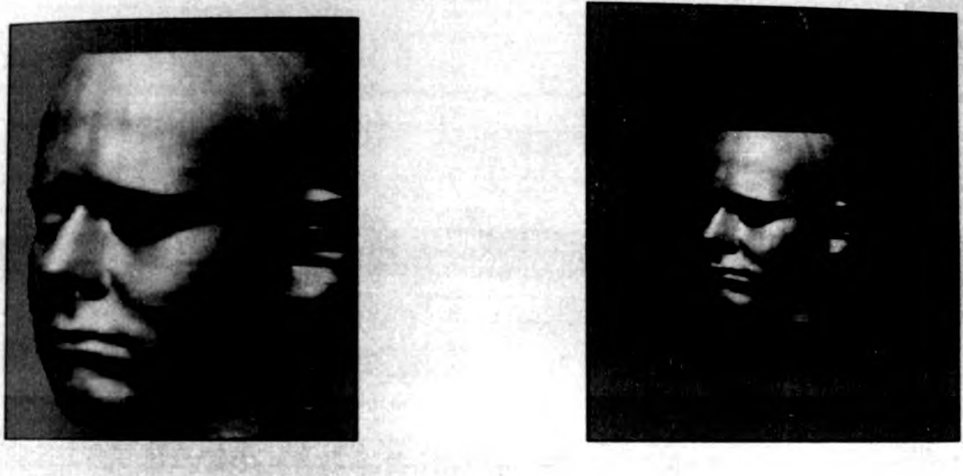

Figure 5.1: Examples of the differently sized stimuli used in experiment 13.

would be expected on matching performance. Previous physiological work has also reported that the responses of "face-cells" are size invariant, although the same work also suggests lighting invariant responses inconsistently with the effects on matching reported (Perrett and co-workers, 1982, 1984, 1989, 1992).

In summary, if any image change affects matching performance effects of size would be expected in this experiment. However, although differently sized images have different image properties, size invariant or derived properties may be sufficient to support matching. A replication of the effect of view was expected with equally sized and differently sized images.

METHODS

The methods for this experiment were the same as for experiment 7 except that only subjects unfamiliar with the people used as stimuli were run and the lighting variable was replaced by a size. Size transformed images were produced by halving the height and width of images using an image processing package so that the area of the image was reduced to 25% of its original size (see figure 5.1). The levels of the size factor which replaced the lighting factor were same or different, with same including both large or both small. Subjects were informed that images would vary in size but that they were to respond "Same" if they thought they were pictures of the same person regardless of size. All lighting was from direction top1, that is from 45° above the viewers line of sight.

RESULTS

The treatment of results was the same as was described for experiment 7. Figure 5.2 shows A' s for the 2(View) x 2(Size) analysis and Table 5.1 gives the means of the transformed data used in the full 3(View) x 3(Size) analysis. As can be seen from figure 5.2 there was a clear effect of changing view, as expected, but no effect of changing size. There did not even appear to be an advantage for identical pictures over differently sized pictures.

A 2(View) x 2(Size) analysis of variance confirmed this pattern of results with a main effect of view ($F_{1,11}=182.4$, $p<<0.05$) but no effect of size or any interaction (p's > 0.1).The full 3(View) x 3(Light) ANOVA also gave a main effect of view ($F_{2,22}=41.6$, $p<<0.05$) with no main effects of size nor any interaction of size with view ($p > 0.1$). Planned comparisons showed no differences between pairs of profiles and pairs of three-quarter views but both same view conditions were matched significantly better than different view pairs.

Fig. 5.2: Mean A' s Experiment 13

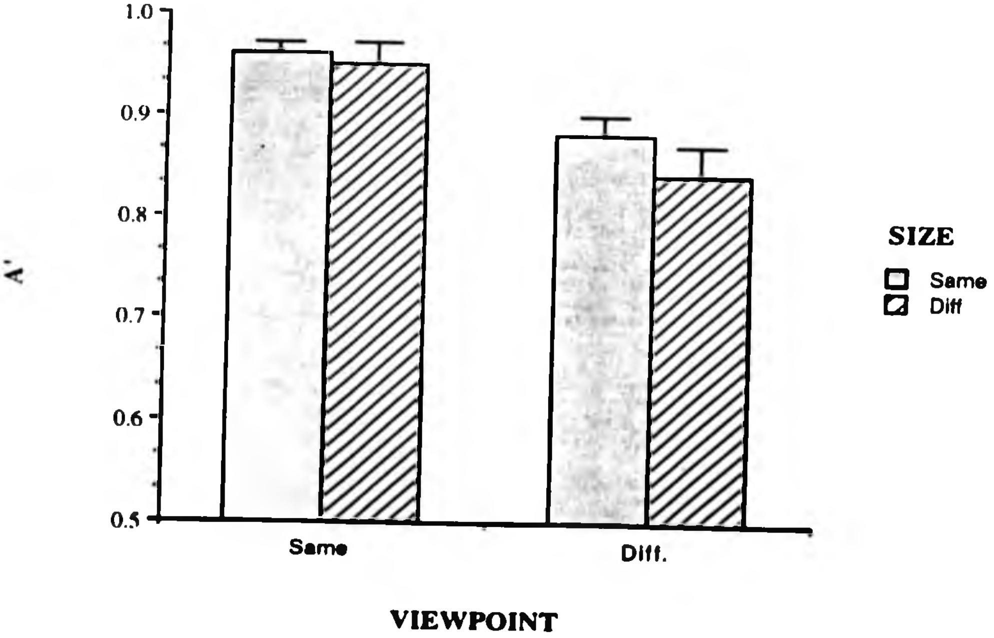

Table 5.1: Means of Transformed A' s for the 3(view) x 3(size) Analysis, Experiment 13.

VIEW	SIZE Large/Large	Small/Small	Large/Small
Prf/Prf	2.79	2.83	2.78
T-Q/T-Q	3.00	2.76	2.76
T-Q/Prf	2.42	2.48	2.33

DISCUSSION.

No effects of changing size were found in this experiment although the effect of changing view was replicated. This showed that matching performance is not reduced by any change in the image; viewpoint and lighting affect performance but size does not. The result is consistent with the conclusion that lighting is important, perhaps because it affects the representation of shape in a way that size does not.

There are many possible explanations for subjects' ability to match across changes in size. As noted in the introduction some properties of the image are size invariant, for example proportions. However, some of these properties are also invariant across changes in light and even view (eg. length of nose: height of head) and so do not predict the pattern of effects on matching reported. An alternative suggestion is that differently sized images produce equivalent outputs in filters tuned to different spatial frequencies which can subsequently be matched. This would not be true for images that were lit differently or that showed heads from different viewpoints.

In conclusion, the absence of an effect of size on a task which has shown effects of viewpoint and lighting shows that not all manipulations that result in changes to the image affect matching performance. The result highlights the apparent importance of lighting and viewpoint for matching the laser head stimuli. In the next experiment the hypothesis that the results might reflect the use of laser heads as stimuli was tested by testing the effects of light and view on digitised photographs.

5.2 EXPERIMENT 14

INTRODUCTION.

The results of experiment 13 showed that not all image changes affect matching performance. This suggests that the unexpected effects of lighting reported in Chapter 4 were evidence for the importance of this variable for the recovery and description of shape not just as a manipulation which affects image properties. However, it could still be argued that the apparent importance of lighting was an artifact of the laser heads used as stimuli which were more than usually lighting dependent. In order to test this the core within subjects design of experiment 7 was repeated but using digitised photographs taken under more natural conditions (see figure 5.3 and Appendix D).

The pictures were taken with the sitters wearing swimming caps, as they had been for laser scanning, as it was thought that hair would provide too obvious a cue to identity. Despite this the stimuli used in this experiment contained information absent from the laser heads, for example information about pigment and texture differences. Some of this information is lighting invariant, in that it can be recovered despite changes in lighting direction, such as which areas of the face are more darkly pigmented and which are hairy. Boundaries between such areas would also produce lighting independent edges in the image which might be used for subsequent processing.

These additional sources of information could make matching performance less sensitive to the effects of lighting. Effects of view have been shown with photographs before (eg. Bruce et al, 1987) and were expected again in this experiment.

METHODS.

The methods for this experiment were the same as for experiment 7 except that only one group of twelve subjects were used, all unfamiliar with the people used as stimuli. Images of eight people were frame-grabbed using a video camera connected to a Sun workstation, two for the practice session and six for the actual experiment. Profile and three-quarter views were used, with both top and bottom lighting. The faces were lit using a single 60 Watt bulb mounted in an angle poise lamp, 45° above or below the viewer's line of sight at a distance of about 1m. The people captured wore a swimming cap to conceal their hair.

a) Top Lit, Three-Quarter

b) Bottom Lit, Three-Quarter

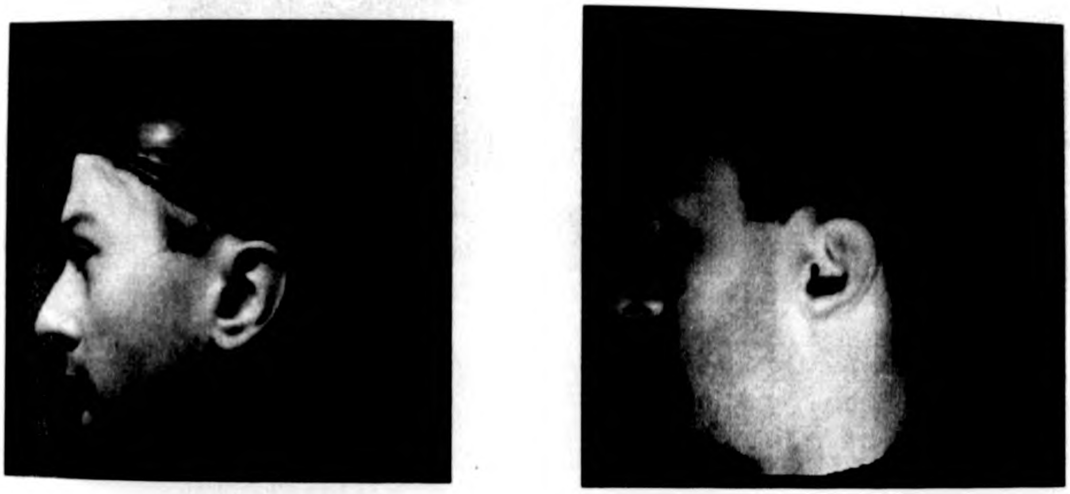

c) Top Lit, Profile

d) Bottom Lit, Profile

Figure 5.3: Examples of the digitised photographs used in experiment 14.

RESULTS

The treatment of results was the same as for previous experiments. Values of A' for the 2(View) x 2(Light) analysis are presented graphically in figure 5.4 and Table 5.2 gives means of transformed data for the 3(View) x 3(Light) analysis. Inspection of figure 5.4 shows a patterns of data similar to that for laser scans as shown in figure 4.2. Changing either light or view reduced performance at least as much as for laser heads but changing both had no additional effect.

A 2(View) x 2(Light) analysis of variance confirmed this pattern of data. There was a significant view x light interaction ($F_{1,11}=11.1$, $p<<0.05$). Analysis of simple main effects showed an effect of light when view was the same ($F_{1,22}=33.1$, $p<<0.05$) but not when view was different ($p>0.1$). There was a simple main effect of view when light was the same ($F_{1,22}=24.7$, $p<<0.05$) but not when light was different ($p>>0.1$).

The 3(View) x 3(Light) ANOVA also gave a significant view x light interaction ($F_{4,44}=5.5$, $p<<0.05$). There were simple main effect of light for pairs of profiles ($F_{2,66}=16.5$, $0<<0.05$), pairs of three-quarters ($F_{2,66}=13.1$, $p<<0.05$) and for different view pairs ($F_{2,66}=6.5$, $p<<0.05$). There were significant simple main effects of view for pairs of top lit face ($F_{2,66}=4.3$, $p<0.05$) and for pairs of bottom lit stimuli ($F_{2,66}=24.7$, $p<<0.05$) but not for differently lit pairs ($p<0.1$). This was the same pattern of effects as reported in experiment 7 for laser heads. Planned comparisons showed no difference between three-quarter and profile views but top lit stimuli were significantly better matched than bottom lit stimuli when view was different, as they were in experiment 7.

Fig. 5.4: Mean A' s From Experiment 14

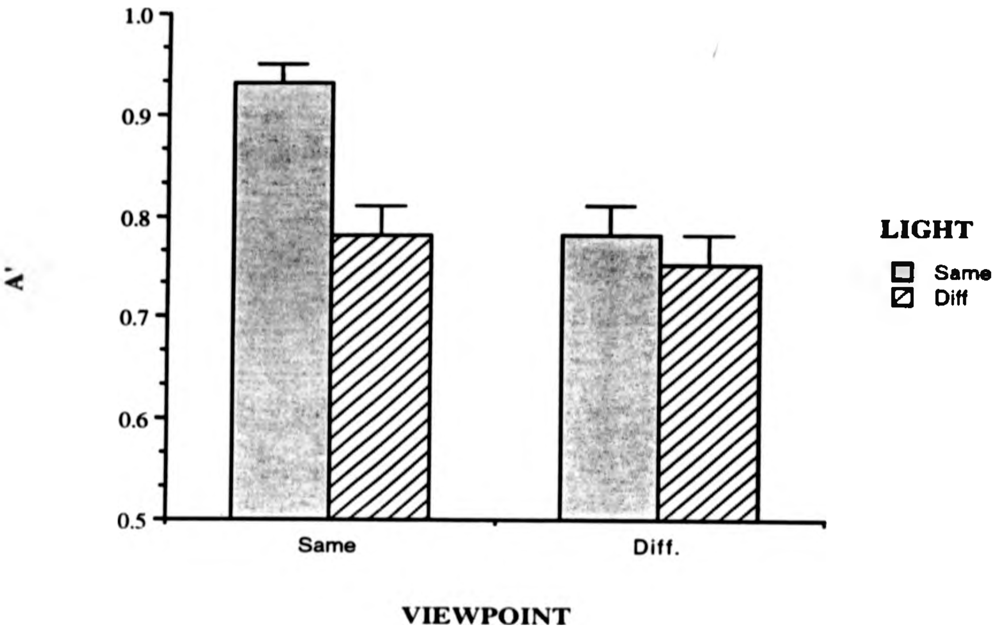

Table 5.2: Means of Transformed A' s for the 3(view) x 3(light) Analysis. Experiment 14.

VIEW		LIGHT	
	Top/Top	Bott/Bott	Top/Bott
Prf/Prf	2.67	2.58	2.08
T-Q/T-Q	2.71	2.81	2.28
T-Q/Prf	2.39	2.01	2.11

DISCUSSION

This experiment showed the same pattern of effects of light and view as were shown in experiment 7 but using digitised photographs instead of laser heads. Performance was at approximately the same level, A' = 0.81, as for the unfamiliar group in experiment 7, A' = 0.82. There was again an interaction between the effects of view and light and the pattern of simple main effects was the same. This suggests that the effects reported on previous experiments, including those of lighting, do not just reflect the use of laser heads as stimuli.

As with laser heads, changing either lighting or view decreased performance. Again there was an effect of lighting even when view remained the same and no additional effect of view when lighting was different, a result that does not seem to be predicted by edge based theories. In this experiment there were edges not present in laser heads which would have been changed by changes in view but not changes in light but this did not importantly affect the pattern of results. Possible sources of lighting and viewpoint invariant information in photographs also did not seem to facilitate matching significantly. The advantage of top over bottom lighting for matching between views was also replicated. The explanations offered in the discussion of experiment 7 seem to also apply here.

The results reported in the previous chapter do not seem to be explained away as just an effect of image changes nor, as shown by this experiment, as a artifact of the materials used. The next experiment reported tested whether the results were due to the simultaneous presentation of stimuli as outlined next.

5.3 EXPERIMENT 15

INTRODUCTION.

This experiment was designed to test another reason why effects of lighting might have been found on the experiments reported in the last chapter and also the experiment with frame-grabbed stimuli just reported. This is that shape-from shading is a global operation which assumes that there is only a single light source illuminating the image. This implies that shape cannot be perceived accurately in two objects illuminated from different angles presented at the same time (Ramachandran, 1988a). When laser scans or photographs of faces were presented under different lighting conditions this would violate a single light source assumption and this might explain the effect of changing light reported. In order to test this stimulus pairs were presented sequentially. As only one image was presented at a time any single light source assumption would be satisfied. Sequential matching of a face is anyway more ecologically valid than simultaneous matching of a single face shown under different lighting conditions!

METHODS.

All details of the design and materials were the same as for experiment 7 except that there was no between subjects familiarity factor and the Supercard program used for simultaneous presentation was adapted for sequential matching. Laser heads were again used for this experiment rather than digitised photographs.

Procedure.

Subjects were informed that they would be presented with pairs of images, one after the other, and were to decide whether the second image was of the same person as the first or of someone different. The procedure was changed from experiment 7 so that what had been the left hand image was presented first for two seconds, then the screen was blank for two seconds, and then there was a beep to cue subjects to respond to what would have been the right hand image, which was presented second for as long as the subject took to make a response. The duration of presentation of the first face was set at two seconds after a pilot study determined that this produced performance at about the same level as was obtained with simultaneous matching. All other aspects of the procedure were the same as for experiment 7.

RESULTS.

The treatment of results was the same as for experiment 7. Figure 5.5 shows A' s for the 2(View) x 2(Light) analysis and Table 5.3 shows means of transformed data for the full 3(View) x 3(Light) analysis. Inspection of figure 5.5 shows that changing either light or view reduced performance but that changing both had no additional effect. There was a clear effect of light even though stimuli were presented sequentially and so the effect of lighting reported in previous experiments cannot solely be attributed to violation of a single light source assumption.

A 2(View) x 2(Light) analysis of variance on transformed data gave a view x light interaction ($F_{1,11}=11.9$, $p<<0.05$). Analysis of simple main effects showed effects of light for same view pairs ($F_{1,22}=21.1$, $p<<0.05$) but not for different view pairs ($p>0.1$). There was a simple main effect of light when view was the same ($F_{1,11}=25.5$, $p<<0.05$) but not when view was different ($p>0.1$). This was similar to the pattern of effects reported for experiment 7 although here the interaction term reached significance.

The 3(View) x 3(Light) ANOVA also gave a view x light interaction ($F_{4,44}=5.7$, $p<<0.05$). Analysis of simple main effects showed an effect of light for pairs of profiles ($F_{2,66}=4.6$, $p<0.05$) and for pairs of three-quarter views ($F_{2,66}=14.1$, $p<<0.05$) but, unlike experiment 7, not for different view pairs ($p>0.1$). There were effects of view for top lit pairs ($F_{2,66}=10.1$, $p<<0.05$) and bottom lit pairs ($F_{2,66}=16.5$, $p<<0.05$) but not for differently lit pairs ($p>0.1$). Planned comparisons showed no significant differences between top/top and bottom/bottom pairs even for the different view condition, another difference from experiment 7 although the difference was in the right direction. Paired three-quarter views did not differ significantly from paired profile views under any condition of light. When light was from above, top/top pairs, only three-quarter view pairs were matched significantly better than profile view pairs. When light was from below different view pairs were matched worse than either level where view remained the same. When light was different there were no differences between levels of view.

Fig. 5.5: Mean A' s For Experiment 15

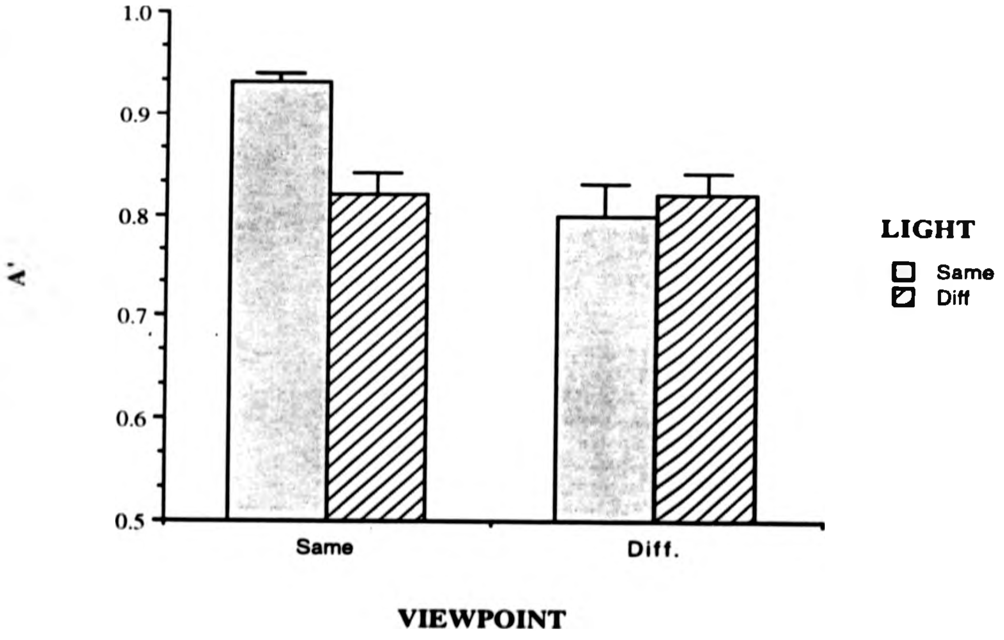

Table 5.3: Means of Transformed A' s for the 3(view) x 3(light) Analysis. Experiment 15

VIEW	LIGHT		
	Top/Top	Bott/Bott	Top/Bott
Prf/Prf	2.57	2.59	2.30
T-Q/T-Q	2.81	2.70	2.27
T-Q/Prf	2.35	2.15	2.29

DISCUSSION.

Overall subjects performed the sequential matching task at about the same level as the simultaneous matching task, as intended (mean A' for this experiment 0.84 and for the unfamiliar group in experiment 7, 0.82). In this experiment there was an effect of light which could not be an artifact of simultaneous matching or, in particular, violation of a single light source assumption. The effect of light again interacted with the effect of view. In the 2(Light) x 2(View) analysis, changing either light or view reduced performance but changing both did not produce any additional effect.

There were differences in the 3(Light) x 3(View) analysis although the main finding was again the same; interacting effects of view and light. However, there was no simple main effect of light when view was different and no significant difference between top and bottom lighting in this condition. Many of the differences were in the right direction but did not reach significance, perhaps because the sequential matching task was less sensitive to such differences. There are many reasons why performance on a sequential matching task might be different from that on a simultaneous matching task including the decay of iconic memory or the effects of a short term visual store (see Humphreys & Bruce, 1989). However, this experiment was not designed to test for such effects and served its purpose in showing that the effect of lighting was not solely due to violation of a single light source assumption. From the similarity of the results reported here it seems likely that when the images were presented simultaneously they did not violate a single light source assumption because, being individually framed, they were treated as separate. Given this, simultaneous matching may be a more useful task as it appears to be more sensitive to differences.

In general the experiments reported reported in this chapter have not provided evidence that the pattern of effects reported for experiments 7 - 12 were artifacts of the method or stimuli used. Experiment 11 showed that not all image changes affect matching performance; a change in size had no effect. That effects of lighting have been reported but not effects of size needs to be reconciled with the physiological evidence, which suggests a pattern of size and lighting invariance together with viewpoint modulation (eg. Perrett et al, 1992). That sensitivity to lighting was not solely an effect of the laser heads used as stimuli was shown by experiment 14, where effects of light were shown when digitised photographs were used as stimuli. The similarity in the pattern of results is consistent with the results derived using laser scans being generalizable to more natural stimuli. Such similarities suggest that laser scans do capture information important for aspects of face processing, a conclusion also supported by the clear effect of inversion on matching laser scans. The last experiment reported, experiment 15, showed that the effect of light reported for both laser scans and digitised photographs was not the effect of violating a single light source assumption. In the final chapter these results will be integrated with those reported in the previous experimental chapters and their implications assessed for the theoretical issues raised in Chapter 1.

6. CHAPER 6: SUMMARY AND GENERAL DISCUSSION

Reliable, although unexpected, effects of lighting direction have been reported on a number of tasks involving face processing. These will be summarised in this chapter, results from different tasks compared and possible theoretical explanations outlined. An account in terms of the description of facial surfaces, thought to be the simplest explanation of the results, will be developed. Possible tests of such a theory will be suggested, which would test between it and alternative, particularly image-based, explanations. Lastly, the relationship of such an account to theories of face processing, object recognition and perception in general will be considered.

6.1 SUMMARY OF RESULTS

Effects of lighting and its relationship to other variables were investigated with three different tasks; perception of a mask as a face, recognition of individuals and matching for identity. In this section results from all the different tasks will be summarised and the relationship between the effects of lighting reported considered.

In Chapter 2 the effect of variations in viewing conditions, including top and bottom lighting, on the strength of the hollow face illusion was investigated. Perception of the illusion was thought to be related to the basic level object categorisation of a face as a face, a between category distinction. The experiments reported showed an effect of lighting direction, with the illusion more compelling when the convex solution was consistent with the light appearing to be from above relative to the observer (experiment 1 & 2). This effect was independent of separate effects of orientation (experiments 1 & 2) and an effect related to the availability of stereoscopic depth information (experiment 4).

The independence of the effect of light was taken as evidence that its effect is on an independent shape-from-shading module (cf. Marr, 1982) providing information important for face perception rather than just on image information. The difference between top and bottom lighting was thought to reflect a light-from-above assumption, thought independently to be used in the recovery of shape-from-shading.

The difference transferred to inverted faces showing that the effect was not the result of more salient features being highlighted. The mask is less likely to be perceived as a face when this would entail it appearing to be lit from below, consistent with a light-from-above assumption. If this explanation is correct the illusion of a "hollow potato" would be expected to be stronger when it appeared lit from above.

The existence of the hollow potato and inverted face illusions supported the notion of a general bias towards the convex solution. Indeed, when this is in conflict with a light-from above assumption it appears a preference for convexity can 'win'. Object specific knowledge was also found to be important for generating the illusion, although there was still an effect of lighting even when this was available (cf. Gregory, 1973). It was argued that object knowledge may act as another independent constraint on the interpretation of shape from different cues. The illusion was also found to be stronger when viewed monocularly rather than binocularly and this effect, like that of lighting, was independent of orientation (experiment 4). This was taken as evidence for stereopsis being important for resolving the illusion at short distances. The independence of the effects of lighting direction, the number of eyes and the availability of object knowledge was argued to be consistent with the processing of shape information by independent "modules", with object knowledge acting as an additional constraint on the eventual interpretation.

In Chapter 3 the difference between top and bottom lighting was found to transfer to another task, the recognition of pigmentless shaded representations of facial surfaces, "laser heads", at least for male faces. This suggested that lighting may be important for within as well as between category distinctions, again perhaps because of its effects on the recovery of shape-from-shading. There appeared to be two possible explanations for this; either top lighting highlights more salient features or shape-from-shading information is useful for distinguishing between faces. A test of this would be to invert the faces; if a light-from-above assumption is critical then faces lit from below should show less of an effect of inversion (Johnston et al, 1992). However, recognition of laser heads was far from perfect and might be reduced to floor by inversion.

The effect of lighting was found to be related to the effect of view with more of an effect for less angled views, especially the full-face. The recognition of profiles was unaffected by lighting and it was argued that this was because occluding contour information is more salient in this view. The effect of lighting was also found to be related to the effect of photographic negation, but that both were dependent on view (experiment 6). There were effects of negation on laser heads showing that the effects of this manipulation are not just on pigment information, although this effect may be more important when photographs are used. It also appeared that negation and bottom lighting may affect image information in complimentary ways, with bottom lit negatives being well recognised. These effects were dependent on view being most pronounced for the full-face view, where lighting and polarity invariant contour information may be least available especially for laser heads.

Due to the relatively poor recognition of laser heads and the limited pool of subjects familiar with the faces used as stimuli the remaining experiments investigated the problem of object constancy across viewing conditions using a simultaneous matching task rather than recognition. This allowed the effects of lighting and view on the perception of information about facial shape to be investigated independently of a load on memory. Again effects of lighting as well as view were found suggesting that lighting dependent information is important for distinguishing between faces across changes in viewing condition.

In the first experiment a change in either light or view was found to disrupt performance but changing both did not produce any additional effect. Thus lighting affected performance even when view, and therefore occluding contours, remained the same and changing these contours did not produce any additional effect when lighting was different. This suggested that edges were not sufficient to explain performance and was consistent with the importance of shading information. It was also found that subjects were better able to match between views when lighting was from above, which did not seem consistent with image or edge based information where the magnitudes of changes would be expected to be the same in either condition. Image analysis could be used to confirm this. Instead the result seemed consistent with lighting from above facilitating the recovery of shape-from-shading which could be used to support matching across views as discussed more fully in the next section. When the stimuli used in the first experiment were presented inverted, performance across changes in lighting and especially view was much poorer. This was not consistent with a simple pattern matching explanation and suggested that information only recoverable from upright faces was important for the task.

In order to investigate if the effect of lighting was due to a difference between top and bottom lighting an experiment was reported which required subjects to match between directions of top lighting. The effects of a change in lighting were found to be reduced or absent. This showed that any change in the image did not affect matching, a result confirmed subsequently when changes in size were found to have no effect. The result was consistent with lighting from above facilitating the derivation of information useful for distinguishing between the shape of faces. The level of performance matching between views was also relatively high. When these stimuli were inverted matching across changes in both lighting and view were disrupted. This may have been due to the violation of a light-from-above assumption as well as because of inversion.

When two directions of bottom lighting were used instead of two directions of top lighting matching across changes in viewing conditions was much poorer. This highlighted the apparent importance of a light-from-above assumption for matches across changes in light as well as view. When these stimuli were inverted matching across changes in light was slightly improved though performance across changes in view were still at floor.

Two experiments were run which confirmed that the results were not a function of the materials used or a result of the simultaneous presentation of images.

To conclude, the results appeared to show that lighting was important for the overall perception of face shape, affected the recognition of familiar individuals and the ability to discriminate between faces on the basis of information derived from shape. In particular a light-from-above assumption appeared to be important, consistent with the importance of shape-from-shading. The relationship of these effects for theories of object recognition will be considered next.

6.2 EFFECTS OF LIGHTING AND THEORIES OF OBJECT CONSTANCY

As has been summarised, the experiments reported in this thesis have reported effects of lighting on tasks ranging from the general, seeing a face as a face, to the specific, discriminating between faces on the basis of information about shape. In this section these effects will be considered within the framework of the theories of object constancy discussed in Chapter 1.

A pure instance-based model of object constancy, along the lines of instance-based models of long term memory (eg. Hintzman, 1986), might explain the effects of lighting simply in terms of the differences lighting introduces to the image of the face. Seeing a the hollow-face illusion as a face would result because it accessed all previous instances of, convex, faces and so it would also be assumed to be convex. Such an account might explain the difference between top and bottom lighting, and between upright and inverted orientations, in terms of the number of similar instances with top lit upright faces eliciting a stronger response.

As regards recognition performance (Chapter 3) the similarity of top lit faces to stored instances might explain their higher recognisability but it is not clear why this should not also be true for profiles (experiment 5). For negatives (experiment 6) there would be few stored instances and recognition would be expected to be poor, though the apparent advantage for bottom lit negatives might be explained in terms of their similarity to top lit positives. Matching performance would also be explained within the context of similarity though it was stressed that simple image differences did not seem to be good predictors of performance. In particular the pattern of effects of top and bottom lighting reported in experiments 9, 10, 11 and 12, where the effects of lighting direction and orientation appeared to interact, would not be predicted by an instance based account (Johnston, 1992).

The main problem with current instance-based models is that notions such as similarity, or even what constitutes an instance, are not independently defined. The matching experiments reported here could be considered as investigations of the important dimensions of similarity as well as tests of an instance-based theory leading to a problem of circularity. Instance-based models may also allow for some processing of the image before storage and this would also affect expected performance and could explain why similarities at the level of images did not correlate with performance. Instance based-models do show how aspects of viewing conditions could be stored (and used to explain context effects) but object constancy still achieved.

Other image-based models learn to associate different images of particular faces and can subsequently recognise novel examples of those faces (eg. Turk & Pentland, 1991). Processing is image based but structural properties of the faces may be abstracted implicitly due to statistical properties of the resulting images. Such models can achieve object constancy despite variations in the input images as such variation is not stored as it does not signal differences between faces. It is not clear what the response of such models to different example of novel faces would be as they are models for recognising particular faces. Also existing implementations of such models often require input images to be normalised for view (see Chapter 1), thus solving many of the problems of object constancy! Such models are claimed to generalize across lighting but this is dependent on their training set and the effects of bottom lighting might be attributed to lack of prior experience. In general image-based theories might be adapted to account for the results reported but this in part reflects their current lack of explicit specification. Possible ways of testing between image-based and other plausible theories will be considered after alternative accounts have been considered.

Another class of solutions to the problem of object constancy considered in the introduction made use of invariant features or properties. In general such features would not predict differences between viewing conditions so long as they were visible, they are invariant. In the experiments reported even when such features were known to be available, for example lighting invariant occluding contours (Chapter 4), they did not appear sufficient to support matching. The effect of orientation on matching cannot be explained in terms of the use of local invariant features, which would have remained the same and been recoverable despite inversion. In general invariant features are rarely found to be sufficient to support object constancy (Ullman, 1989), a conclusion supported by the effects of lighting even when invariant contours were available as reported here. When digitised photographs were used, providing additional possible sources of lighting invariant information including pigment and texture differences, such features still did not appear sufficient to support matching.

The effect of lighting when much edge based information would have remained invariant (ie. when view was the same) does not seem consistent with edge-based theories of object constancy (eg. Biederman, 1987). Edges, and specifically occluding contours (eg. Marr, 1982), are often used to specify the parts in object decomposition models. Faces are not divided into regions by edges, unlike many man-made objects, and even when the boundaries between areas of different pigment or texture were available (experiment 14) they did not appear to support processing. Lighting does affect some edges, for example shadow boundaries and the pattern of isoluminant contours, but the variation in such edges would be as great as the changes in occluding contours due to changes in view (shadow boundaries actually fall where occluding boundaries would fall if the object were viewed from the position of the light source). If such edges were critical, the problems of variable illumination for object constancy would be as great as those of variable viewpoint.

Empirically differences between top and bottom lighting, both of which affect such edges, do not seem to be explained even if lighting sensitive edges are important. Further, negation, which changes only the sign and not the position of such edges, is known to affect performance (eg. Phillips, 1972; Hayes, Morrone & Burr, 1986).

The apparent inadequacy of edge-based accounts for explaining the effect of lighting supports the conclusion that edge based parts such as geons and generalized cones are not appropriate for faces (Bruce, 1988; see also Chapter 1). Even categorisation of a face as a face, the decision level most object decomposition schemes model, was affected by lighting (Chapter 2). However, an object decomposition scheme using another possible set of parts, whose derivation would be expected to be affected by lighting, will be outlined in the next section.

The last approach to object constancy outlined in Chapter 1 was the alignment approach (eg. Ullman, 1989) but this does not seem to explain the effects of lighting reported. In such approaches alignment is thought necessary to overcome the problem of variable viewpoint but in the experiments reported here there were effects of light even when view remained the same. This either suggests that the problem of finding lighting invariant properties for alignment has been underestimated or that how object models would appear under different lighting conditions also would have to be specified. Neither of these possibilities seems to explain why the problem should be greater for bottom lit face than for top lit faces.

To conclude it does not seem that schemes based on edges, invariant features or that use alignment to overcome the problems of variable view easily explain the effects of lighting reported. Lighting certainly does affect the two-dimensional image projection and this may account for its effects on performance. However, image based accounts do not seem to provide explanations of differences between upright and inverted stimuli or top and bottom lighting except in terms of familiarity or similarity to a greater number of stored examples. The apparent interaction between the effects of lighting direction and inversion (experiments 9, 10, 11 & 12) would not be predicted by an image-based scheme and might provide a clear test. In the next section an alternative surface-based scheme that seems to account naturally for the effects of lighting reported will be outlined.

6.3 A SURFACE-BASED ACCOUNT

In this section the essentials of a surface-based object decomposition scheme will be outlined which appears consistent with the pattern of results reported. The limitations and drawbacks of such a scheme will be considered along with possible tests between it and an image-based theory.

Surface-based schemes are characterised by the use of surface descriptions as their building blocks . The parts proposed in one scheme are the classical set of eight surface types defined by Gaussian and mean curvature: Peak, ridge, saddle ridge, minimal, pit, valley, saddle valley and flat (Bruce, Coombes & Richards, 1993). Gaussian curvature is the product of the two principal curvatures and mean curvature half their sum (Watt, 1991). The principal curvatures are the extreme curvatures at any point. Thresholds are used to determine what constitutes zero curvature and decreasing thresholds could be used to pick out areas of maximum curvature first followed by less curved areas.

An alternative way of representing surface curvature is in terms of the local shape index (Koenderink, 1990). In this shape is represented according to the possible combinations of the two principal curvatures normalised by keeping their root means square (RMS) constant. The value of the RMS is used as a second parameter sensitive to scale. In an object recognition scheme surface patches could be characterised using the shape index, though how objects would be segmented has not been specified. It is not known what the "features" for face recognition are but it seems intuitively possible that these could be characterised by different surface types. However it is represented, curvature is a property of the surface and not the viewing conditions and so if it can be recovered from the image it would be naturally useful for representing objects independently of viewing conditions.

There is evidence that people can recover information about curvature from shading information as they can make relative curvature judgements (Johnston, Passmore & Morgan, 1991). It also appears that curvature is recovered directly and not as the derivative of orientation, an important consideration for object constancy across changes in view (Johnston & Passmore, 1992). It is therefore possible that shading information could be used to derive surface parts from the image which could then be used to represent objects. Other sources of information including stereopsis and texture might also contribute to this process of determining surface shape though shading may be especially important for faces (see Chapter 1, section 1.7).

If surface types or parts defined in terms of curvature can be recovered from the image their combination into object descriptions is likely to be important to their usefulness. One scheme currently uses the proportion of surface types at different thresholds in different, interactively defined, areas. This has been used to test a surface-based account to see if it is consistent with the effects of distinctiveness and the detectability of surface changes (Bruce et al, 1993). This however fails to capture relations between parts thought to be especially important for the perception of faces (see Chapter 1, section 1.4). Surface-based parts could be combined in many of the ways suggested for the combination of other parts (eg. Biederman, 1987). Also in this context shading information may provide information about ordinal depth relations (Reichel & Todd, 1989) which may be important for relations between parts as well as relations in the image plane. It also appears from the experiments reported in Chapter 2 that lighting and shading may be important for seeing the face as a face, a whole which is characteristic of "configural processing" (Tanaka & Farah, 1993).

The use of shape-from-shading, thought independently to incorporate a light-from-above assumption, for deriving the parts in a surface-based scheme accounts for many of the effects of lighting reported. The description in terms of surface types would be lighting, and view, independent once derived but lighting could affect the initial derivation of the parts from the image. Bottom lighting, for example, may be disrupting face processing because it violates a light-from-above assumption used for the interpretation of shape from ambiguous patterns of shading. The work on the hollow face illusion illustrates the ambiguity of shape-from-shading and suggests that the direction of lighting may affect categorisation of the stimulus as a face (experiments 1 & 2); the stimulus is more likely to be seen as a face when this is consistent with it appearing to be lit from above. This is so even when the stimulus is inverted showing that the effect is not just due to top lighting highlighting more salient features (different features are highlighting in the inverted and upright face lit from above - see figures 2.1 and 2.2). The illusion highlights the importance of shape-from-shading for face perception (Bruce, 1988) even at the most basic level of categorisation as a face and is consistent with shading specifying the surfaces that comprise a face. The effect of object knowledge, shown by inversion and comparison with a hollow potato (Chapter 2, experiments 1, 2 & 3) suggests that lighting may be important for accessing a stored object model, which would be expected if such a model is represented in terms of surface types accessed by shape-from-shading. It is also possible that bottom lighting is interfering with the characteristic "bar-code" that could be used for categorising faces (Watt & Dakin, 1993), reversing the pattern of light and dark. This is not necessarily contradictory to an explanation in terms of the perception of facial shape as dark bars tend to be associated with concave areas, for example the eye-sockets and the area below the nose, and light bars with convex areas such as the forehead, nose and cheeks and the chin.

The advantage of top lighting for recognising faces (experiments 5 & 6), at least in less angled views, also could be explained in terms of a surface-based theory if information about surface shape can be used to distinguish between faces and is stored for familiar faces. Bottom lighting might prevent the recovery of such information or increase its associated uncertainty because it affects the interpretation of overall surface shape. Negating bottom lit stimuli so that they appear lit from above (see Chapter 3, experiment 6, figure 3.5) offsets the effects of bottom lighting. This effect might also be explained in terms of a "bar-code" if both negation and bottom lighting reverse the polarity of the bars. In general the effect of negation is consistent with an account in terms of shading rather than edges, although the primary effects of negation on recognition are probably because of its effects on pigment information. The effect of negation is also consistent with the use of shape derived from shading rather than just its two dimensional properties as many of these would be unaffected by negation.

The effects of lighting on matching could also be explained if shape-from-shading is used to derive surface-based descriptions. If bottom lighting interferes with or prevents the encoding of surface descriptions this would explain difficulties matching between top and bottom lit faces. It also appears that top lighting facilitates matching between views (experiments 7, 9 & 11), consistent with surface descriptions in terms of curvature being useful for matching across changes in view. Subjects were also better able to match between directions of top lighting than of bottom lighting, consistent with surface descriptions also being used to match across changes in lighting direction. When top lit faces were inverted subjects were worse at matching between lighting directions (experiment 9 & 10) but when bottom lit stimuli were inverted subjects were better (experiments 11 & 12) suggesting that a light-from-above assumption was important possibly because it allows derivation of lighting invariant surface type descriptions even for inverted faces and possibly for unfamiliar objects. A direct test using different levels of both top and bottom lighting with upright and inverted faces might strengthen these conclusion made on the basis of informal between experiment comparisons. Similar tests could also be done using unfamiliar objects with negation used as a control for its effects on image properties.

There was one significant effect of matching between top lit directions, that for profiles (experiment 9). This was explained in terms of one of the directions of lighting highlighting the internal features of the face in this view. This suggests that lighting direction may be important in determining which surfaces are encoded just as view would be expected to be - different surfaces will be visible in different views.

Surface shape is not a function of size and a surface-based scheme is consistent with the lack of an effect of size on matching (experiment 13). The transfer of the effects of light to photographic stimuli, including the differences between top and bottom lighting, is consistent a surface-based scheme being useful for real faces. The surface-types as outlined do not contain information about pigment, reflectance, or texture but it is possible that such properties could be added to the underlying description of shape, though shape itself would remain an important dimension of variation.

In general a surface-based theory seems consistent with the effects of lighting reported. The effect of lighting is explained in terms of its effect on shape-from-shading which is used to derive explicit surface-based descriptions from the image. How surface types are combined into complete object descriptions is not known but the interpretation of three-dimensional shape-from-shading does appear important for perception of the face as a whole. This may affect the recovery of the configural information known to be important for faces and shading may also be important for providing information about the depth relations between features. The experiments reported here have been concerned with between and within category decisions involving faces but surface-based descriptions of parts or the whole of the face might also be useful for expression perception. Expressions often involve characteristic changes in the shape of the face and these could be recovered from a surface-based representation. Some possible tests of a surface-based theory will be considered next, particularly tests of whether surface shape is coded in terms of an explicit structural description or whether it is encoded implicitly from the statistical properties of images.

If bottom lighting does interfere with the ability to derive surface-based representations then differently shaped surfaces should be more difficult to discriminate when lit from below. It is possible to make changes to the surface of laser heads using computer aided design techniques (Bruce et al, 1993) and the just noticeable differences for such changes could be investigated under condition of top and bottom lighting. The noticeability of such changes when they are in areas receiving direct illumination compared to when they are in shadow could also be tested.

As already mentioned an image-based theory would not predict an interaction between the effects of lighting and orientation - there would be no reason for expecting an advantage to be associated with inverting bottom lit stimuli. The pattern of data in experiments 9, 10, 11 and 12 together with a reduction of the effect of inversion for bottom lit heads shown in previous research (Johnston et al, 1992) favours an interpretation in terms of a light-from-above assumption. The discriminability of surface types in upright and inverted top and bottom lit faces could be used to test this as could a matching experiment which included different levels of top and bottom lighting.

A surface-based scheme may be especially important for faces but the differences between top and bottom lighting for inverted faces suggests that the scheme might be generalizable for other smooth surfaces. An effect of lighting on the hollow potato illusion would suggest that a light-from-above assumption generally affects the perception of surface shape. It is not clear why an image-based theory, which encoded shape only implicitly because of its effects on image intensities, would predict an advantage for top lighting except in terms of familiarity.

Photographic negation may also provide a useful way of testing whether subjects are recovering shape explicitly or using implicit information about shape contained in two-dimensional properties of the image. Negation is thought to affect the recovery of shape-from-shading but if image properties can be used directly polarity might be arbitrary.

6.4 CONCLUSIONS

Effects of lighting have been reported on a number of tasks involving the perception of facial surfaces. These effects were not consistent with purely edge-based theories and were taken as evidence for the importance of shading information for the perception of facial surfaces. Differences between top and bottom lighting, consistent with a light-from-above assumption, in particular seem to favour a surface-based rather than an image-based explanation. An account was developed using surface-based primitives derived from shape-from-shading and combined into object descriptions with explicit three-dimensional relational information.

REFERENCES

Barrow, H. G. & Tenenbaum, J. M. (1981). Computational Vision. *Proceeding of the IEEE, 69*, 592-595.

Benton, A. L. & Van Allen, M. W. (1968). Impairment in facial recognition in patients with cerebral disease. *Cortex, 4*, 344-358.

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

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

Biederman, I. & Cooper, E. E. (1992). Size invariance in visual object priming. *Journal of Experimental Psychology: Human Perception and Performance, 18(1)*, 121-133.

Biederman, I. & Ju, G. (1988). Surface vs. edge-based determinants of visual recognition. *Cognitive Psychology, 20*, 38-64.

Binford, T. O. (1972). Visual perception by computer. Paper presented at the IEEE Conference on Systems and Control, December 1971, Miami.

Blake, A. & Bulthoff, H. (1990). Does the brain know the Physics of specular reflection? *Nature, 343*, 165-168.

Bradshaw, J. L. & Wallace, G. (1971). Models for the processing and identification of faces. *Perception & Psychophysics, 9*, 443-448.

Bruce, V. (1982). Changing faces: Visual and nonvisual coding processes in face recognition. *British Journal of Psychology, 73*, 105-116.

Bruce, V. (1988). *Recognising Faces.* Hove: Lawrence Erlbaum Associates Ltd.

Bruce, V., Burton, M., Hanna, E., Healey, P., Mason, O., Coombes, A., Fright, R. & Linney, A. (1993). Sex discrimination: how do we tell the difference between male and female faces? *Perception, 22*, 131-152.

Bruce, V., Coombes, A., & Richards, R. Describing the shapes of faces using surface primitives. *Image and Vision Computing, 11*, 353-363.

Bruce, V., Hanna, E., Dench, N., Healy, P. & Burton, M. (1992). The importance of 'mass' in line drawings of faces. *Applied Cognitive Psychology, 6*, 619-628.

Bruce, V., Healey, P., Burton, M., Doyle, T., Coombes, A. & Linney, A. (1991). Recognising facial surfaces. *Perception, 20*, 755-769.

Bruce, V. & Humphreys, G. W. (1989). Visual cognition: computational, experimental and neuropsychological perspectives. Hove: Lawrence Erlbaum Associates Ltd.

Bruce, V., Valentine, T. & Baddeley, A. (1987). The basis of the 3/4 view advantage in face recognition. *Applied Cognitive Psychology, 1*, 109-120.

Bruce, V. & Young, A. (1986). Understanding face recognition. *British Journal of Psychology*, *77*, 305-327.

Bui-Tuong, P. (1975). Illumination for computer generated pictures. *Communications of the Association of Computing Machinery* (ACM), *18*, 311-317.

Bulthoff, H. H. & Mallot, H. A. (1987). Interaction of different modules in depth perception. *Proceedings of the IEEE International Conference on Computer Vision, London*, 295-305.

Carey, S. & Diamond, R. (1977). From piecemeal to configural representation of faces. *Science*, *195*, 312-314.

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

Coombes, A., Richards, R., Linney, A., Bruce, V. & Fright, R. (1992). Description and recognition of faces from 3-D data. *Proceedings SPIE* San Diego:

Davies, G. M., Ellis, H. D. & Shepherd, J. W. (1978). Face recognition accuracy as a function of mode of representation. *Journal of Applied Psychology*, *63*, 180-187.

Diamond, R. & Carey, S. (1986). Why faces are and are not special: An effect of expertise. *Journal of Experimental Psychology: General*, *115*, 107-117.

Ellis, H. D., Shepherd, J.W. & Davies, G. M. (1979). Identification of familiar and unfamiliar faces from internal and external features: Some implications for theories of face recognition. *Perception*, *8*, 431-439.

Ellis, R. & Allport, D. A. (1985). Multiple representations for visual objects: A behavioural study. *Paper presented at the Easter meeting of the Society for the Study of Artificial Intelligence and Simulation of Behaviour,* Warwick, 10-12 April, 1985.

Ellis, R., Allport, D. A., Humpreys, G. W. & Collis, J. (1989). Varieties of object constancy. *Quarterly Journal of Experimental Psychology, 41A (4),* 775-796.

Endo, M., Takahasi, K., & Maruyama, K. (1984). Effects of observer's attitude on the familiarity of faces: Using the difference in cue value between central and peripheral facial elements as an index of familiarity. *Tohoku Psychologica Folia, 43,* 23-34.

Erens, R. G. F., Kappers, A. M. L. & Koenderink, J. J. (1993). Discrimination of surface shape on the basis of shading. *Perception, 22, supplement,* 101-102.

Fender, D. & Julesz, B. (1967). Extension of Panum's Fusional Area in Binocularly Stabalised Vision. *Journal Optical Society of America, 57,* 819-830.

Fraser, I. H. & Parker, D. M. (1986). Reaction time measures of feature saliency in a perceptual integration task. In H. D. Ellis, M.A. Jeeves, F. Newcombe & A. Young (Eds.). *Aspects of Face Processing.* Dordrecht: Martinus Nijhoff.

Frisby, J. P. (1979). *Seeing: Mind, Brain and Illusion.* Oxford: Oxford University Press.

Galper, R. E. (1970). Recognition of faces in photographic negative. *Psychonomic Science, 19,* 207- 208.

Georgeson, M. A. (1979). Random-dot stereograms of real objects: observations on stereo faces and moulds. *Perception, 8*, 585-588.

Gibson, J. J. (1979). *The Ecological Approach To Visual Perception.* Boston: Houghton Mifflin.

Gregory, R. L. (1973). The confounded eye. In R.L.Gregory & E.H.Gombrich (eds.) *Illusion in Nature & Art.* London: Duckworth.

Haig, N. D. (1986). Investigating face recognition with an image processing computer. In H. D. Ellis, M.A. Jeeves, F. Newcombe & A. Young (Eds.). *Aspects of Face Processing.* Dordrecht: Martinus Nijhoff.

Hay, D. C. & Young, A. W. (1982). The human face. In A. W. Ellis (Ed.), *Normality and pathology in cognitive functions.* London: Academic Press.

Hayes, T., Morrone, M. C. & Burr, D. C. (1986). Recognition of positive and negative bandpass-filtered images. *Perception, 15*, 595-602.

Hietanen, J., Perrett, D., Oram, M., Benson, P. & Dittrich, W. (1992). The effects of lighting conditions on responses of cells selective for face views in the macaque temporal cortex. *Experimantal Brain Research, 89*, 157-171.

Hintzman, D. L. (1986). "Schema abstraction" in multiple-trace memory model. *Psychological Review, 93*, 411-428.

Homa, D. , Haver, B. & Schwartz, T. (1976). Perceptibility of schematic face stimuli: Evidence for a perceptual Gestalt. *Memory and Cognition, 4*, 176-185.

Horn, B. K. P. (1977). Understanding Image Intensities. *Artificial Intelligence, 8,* 201-231.

Ikeuchi, K. & Horn, B. K. P. (1981). Numerical shape from shading and occluding boundaries. *Artificial Intelligence, 17,* 144-184.

Johnston, A. (1992). Object constancy in face processing: intermediate representations and object forms. *Irish Journal of Psychology, 13(4),* 426-439.

Johnston, A., Hill, H. & Carman, N. (1992). Recognising faces: effects of lighting direction, inversion, and brightness reversal. *Perception, 21,* 365-375.

Johnston, A. & Passmore, P. (1992). Shape from shading: psychophysical evidence for the primacy of surface curvature perception. *Perception, 21,* 54-55.

Johnston, A., Passmore, P. & Morgan, M. J. (1991). Curvature discrimination thresholds for shaded and textured surfaces. *Invest. Opthal. Vis. Sci., 32,* 1179.

Kemp, R., McManus, I. C. & Pigott, T. (1990). Sensitivity to the displacement of facial features in negative and inverted images" *Perception, 19,* 531-543.

Koenderink, J. J. (1990). *Solid Shape.* Cambridge, Mass.: MIT Press.

Koenderink, J. J. & van Doorn, A. J. (1980). Photometric invariants related to solid shape. *Opt. Acta, 27,* 981-996.

Linney, A. D. (1992). The use of 3-D computer graphics for the simulation and prediction of facial surgery. In Bruce, V. & Burton, A.M. (Eds.) *Processing Images of Faces*. (pp. 149-178). Norwood, New Jersey: Ablex Publishing Corporation.

Lowe, D. G. (1985). *Perceptual organization and visual recognition.* Boston: Kluwer.

Lowe, D. G. (1987). The viewpoint consistency constraint. *The International Journal of Computer Vision, 1*, 57-72.

Luria, S. M. (1978). Comparison of eye-movements over faces in photographic positives and negatives. *Perception, 7,* 349-358.

Marr, D. (1982). *Vision*. San Fransisco: Freeman.

Marr, D. & Nishihara, H. K. (1978). Representation and recognition of the spatial organisation of the spatial organization of thre-dimensional shapes. *Proceedings of the Royal Society of London, B204,* 301-328.

McNicol, D. (1972). *A Primer in Signal Detection Theory.* London: Allen & Unwin.

Mingola, E. & Todd, J. T. (1986). Perception of solid shape from shading. *Biological Cybernetics, 53,* 137-151.

Moss, J. P., Linney, A. D., Grindrod, S. R., Arridge, S. R. and Clifton, J. S. (1987). The three-dimensional visualisation of the face and skull using computerised tomography and laser scanning techniques. *European Journal of Orthodontics ,9,* 247-253.

Norman, D. A. (1964). A comparison of data obtained with different false-alarm rates. *Psychological Review, 71*, 243-246.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115*, 39-57.

Pearson, D. E. (1983). Evaluation of feature-extracted images for deaf communication, *Electronics Letters, 19*, no. 16, 629-631.

Pearson, D. E., Hanna, E. & Martinez, K. (1990). Computer-generated cartoons. In H. Barlow, C. Blakemore & M. Weston-Smith (Eds.), *Images and Understanding: thoughts about images, ideas about understanding* (pp. 46-60). Cambridge: Cambridge University Press.

Pearson, D. E. & Robinson, J. A. (1985). Visual Communication at Very Low Data Rates, *Proc. IEEE, 73*, no. 4, 795-811.

Pearson, D. E. & Six, H. (1982). Low data-rate moving-image transmission for deaf communication, *Proc. IEE Conference on Electronic Image Processing*, York, IEE Conference Publication no. 214, 204-208.

Penry, J. (1971). *Looking at faces and remembering them.* London: Elek Books.

Pentland, A. P. (1982). Finding the direction of illumination. *Journal of the Optical Society of America, 72*, 448-455.

Pentland, A. P. (1986). Perceptual organisation and the representation of natural form. *Artificial Intelligence, 28*, 293-331.

Perrett, D. I., Harries, M. H., Bevan, R., Thomas, S., Benson, P. J., Mistlin, A. J., Chitty, A. J., Hietenan, J. K. & Ortega, J. E. (1989). Frameworks of analysis for the neural representation of animate objects and actions. *Journal of Experimental Biology, 146*, 87-114.

Perrett, D. I., Hietanen, J. K., Oram, M. W. & Benson, P. J. (1992) Organization and functions of cells responsive to faces in the temporal cortex. *Philosophical Transactions of the Royal Society of London, 335*, 23-30.

Perrett, D. I., Mistlin, A. J., Chitty, A. J., Smith, P. A., Potter, D. D., Broennimann, R. & Harries, M. (1988). Specialised face processing and hemispheric asymmretry in man and monkey: evidence from single unit reaction time studies. *Behaviour and Brain Research, 29*, 245-258.

Perrett, D. I., Rolls, E. T. & Caan, W. (1982). Visual neurons responsive to faces in the monkey temporal cortex. *Experimental Brain Research, 47*, 329-342.

Perrett, D. I., Smith, P. A. J., Potter, D. D., Mistlin, A. J., Head, A. S., Milner A. D. & Jeeves, M. A. (1984) Neurons responsive to faces in the temporal cortex: studies of functional organization, sensitivity to identity and relation to perception. *Human Neurobiology, 3*, 197-208.

Phillips, R. J. (1972). Why are faces hard to recognise in the photographic negative? *Perception & Psychophysics, 12*, 425-426.

Phillips, W. A. (1974). On the distiction between sensory storage and short term visual memory. *Perception & Psychophysics, 16*, 283-290.

Poggio, T. & Edelman, S. (1990). A network that learns to recognise three-dimensional objects. *Nature, 343*, 263-266.

Posner, M. I. & Keele, S. W. (1967). Decay of visual information from a single letter. *Science, 158*, 137-139.

Rae, G. (1976). A non-parametric measure of recognition performance. *Perceptual & Motor Skills, 42*, 98.

Ramachandran, V. S. (1988a). Perception of shape from shading. *Nature 331*, No. 6152, 133-166.

Ramachandran, V. S. (1988b). Perceiving Shape form Shading. *Scientific American, 259*, 58-65.

Reichel, F. & Todd, J. (1990). Perceived depth inversion of smoothly curved surfaces due to image orientation. *Journal of Experimental Psychology, Human Perception & Performance, 16, No.3*, 653-664.

Rhodes,G., Brake, S & Atkinson, A. P. (1993). What's lost in inverted faces? *Cognition, 47*, 25-57.

Roberts, A. & Bruce, V. (1988). Feature saliency in judging the sex and and familiarity of faces. *Perception, 17*, 475-481.

Rock, I. (1973). *Orientation and form*. New York: Academic Press.

Rosch, E. (1978). Principals of categorisation. In E. Rosch & B. Lloyd (Eds), *Cognition and Categorisation*. Hillsdale, N.J.: Lawrence Erlbaum Associates Inc.

Sakurai, K., Endo, M. & Maruyama, K. (1985). Depth reversal in the aperture viewing of facial shading patterns. *Tohoku Psychologica Folia, 44 (1-4)*, 70-80.

Sergent, J. (1984). An investigation into component and configural processes underlying face recognition. *British Journal of Psychology, 75*, 221-142.

Sergent, J. (1986). Microgenesis of face perception. In H. D. Ellis, M.A. Jeeves, F. Newcombe & A. Young (Eds.). *Aspects of Face Processing*. Dordrecht: Martinus Nijhoff.

Shepherd, J.W., Davies, G. M., & Ellis, H. D. (1981). Studies of cue saliency. In G. Davies, M. A. Jeeves, F. Newcombe, & A. Young (Eds.), *Perceiving and Remembering Faces*. London: Academic Press.

Stevens, K. A. (1981). The visual interpretation of surface contours. *Artificial Intelligence, 17*, 47-73.

Tanaka, J. W. & Farah, M. J. (1993). Parts and wholes in face recognition. *The Quarterly Journal of Experimental Psychology, 46a (2)*, 225-245.

Todd, J.,T. & Mingolla, E. (1983). Perception of Surface Curvature and Direction of Illumination from Patterns of Shading. *Journal Experimental Psychology, Human Perception & Performance , 9, No. 4*, 583-595.

Todd, J. T. & Reichel, F. D. (1989). Ordinal structure in the visual perception and cognition of smoothly curved surfaces. *Psych. Rev., 4*, 643-657.

Turk, M. & Pentland, A. (1991). Eigenfaces fo recognition. *Journal of Cognitive Neuroscience, 3*, 71-86.

Ullman, S. (1989). Aligning pictorial descriptions: An approach to object recognition. *Cognition, 32*, 193-254.

Valentine, T. (1988). Upside-down faces: a review of the effect of inversion upon face recognition. *British Journal of Psychology, 79,* 471-491.]

Valentine, T. & Bruce, V. (1986). The effect of race, inversion and encoding activity upon face recognition. *British Journal of Psychology, 61*, 471-491.

Watt, R. J. (1991). *Understanding Vision.* London: Academic Press.

Watt, R. J. & Dakin, S. C. (1993). The integration of information from different spatial scales in face perception. *Perception, 22(supplement),* 12.

Witkin, A. P. (1981). Recovering surface shape and orientation from texture. *Artificial Intelligence, 17*, 17-45.

Witkin, A. P. & Tenenbaum, J. M. (1983). On the role of structure in vision. In J. Beck, B. Hope & A. Rosenfeld (Eds.) *Human Machine Vision* . New York: Academic Press.

Yellott, J. I., & Kaiwa, J. L. (1979). Depth inversion despite stereopsis: the appearance of random-dot stereograms on surfaces seen in reverse perspective. *Perception, 8*, 135-142.

Yin, R. K. (1969). Loooking at upside-down faces. *Journal of Experimental Psychology, 81*, 141-145.

Young, A. W. & Bruce, V. (1991). Perceptual categories and the computation of "grandmother". *European Journal of Cognitive Psychology, 3(1)*, 5-49.

Young, A. W., Hay, D.C., McWeeny, K. H., Flude, B. M., & Ellis, A. W. (1985). Matching familiar and unfamiliar faces on internal and external features. *Perception, 14*, 737-746.

Young, A. W., Hellawell, D. & Hay, D. C. (1987). Configural information in face perception. *Perception, 16*, 747-759.

APPENDIX

ANOVA SUMMARY TABLES

CHAPTER 2.

EXPERIMENT 1.

```
FACTOR:          sub      orient       light        dir        dist
LEVELS:           24           2           2          2         192
TYPE  :       RANDOM      WITHIN      WITHIN     WITHIN        DATA

SOURCE                SS      df                 MS          F       p
=====================================================================
orient       272857.5208       1        272857.5208     57.286   0.000 ***
os/          109550.9792      23          4763.0861

light         29106.7500       1         29106.7500     10.988   0.003 **
ls/           60924.2500      23          2648.8804

ol             5271.0208       1          5271.0208      1.359   0.256
ols/          89180.9792      23          3877.4339

dir          260043.5208       1        260043.5208     56.908   0.000 ***
ds/          105099.9792      23          4569.5643

od             3008.3333       1          3008.3333      1.711   0.204
ods/          40429.1667      23          1757.7899

ld               17.5208       1            17.5208      0.008   0.929
lds/          49730.4792      23          2162.1947

old            4181.3333       1          4181.3333      2.092   0.162
olds/         45980.6667      23          1999.1594
```

EXPERIMENT 2.

```
FACTOR:         sub     orient      light        dir      dist
LEVELS:          24          2          2          2       192
TYPE  :      RANDOM     WITHIN     WITHIN     WITHIN      DATA

SOURCE                    SS      df                MS          F         p
===============================================================================
orient       1009635.0469       1   1009635.0469     92.051   0.000 ***
os/           252267.5781      23     10968.1556

light          35779.3802       1     35779.3802      4.752   0.040 *
ls/           173165.7448      23      7528.9454

ol               845.8802       1       845.8802      0.127   0.725
ols/          153254.7448      23      6663.2498

dir           386732.7552       1    386732.7552     73.141   0.000 ***
ds/           121611.8698      23      5287.4726

od             10428.2552       1     10428.2552      3.479   0.075
ods/           68951.8698      23      2997.9074

ld             25507.1302       1     25507.1302      8.284   0.008 **
lds/           70816.4948      23      3078.9780

old             1116.5052       1      1116.5052      0.292   0.594
olds/          88032.6198      23      3827.5052
```

EXPERIMENT 3.

```
FACTOR:         sub        typ      orient       move         dist
LEVELS:          12          2           2          2           96
TYPE  :      RANDOM     WITHIN      WITHIN     WITHIN         DATA

SOURCE                  SS      df                 MS          F         p
==============================================================
typ          208134.3750       1        208134.3750     45.291     0.000 ***
ts/           50550.8750      11          4595.5341

orient       140990.0104       1        140990.0104     33.177     0.000 ***
os/           46745.9896      11          4249.6354

to           160966.2604       1        160966.2604     60.049     0.000 ***
tos/          29486.6146      11          2680.6013

move         319935.0417       1        319935.0417     67.976     0.000 ***
ms/           51772.3333      11          4706.5758

tm            30888.3750       1         30888.3750     15.037     0.003 **
tms/          22595.6250      11          2054.1477

om            10901.3438       1         10901.3438      6.387     0.028 *
oms/          18775.6562      11          1706.8778

tom           18398.3437       1         18398.3437     12.036     0.005 **
toms/         16814.0313      11          1528.5483
```

EXPERIMENT 4.

```
FACTOR:          sub        eyes      orient        dir        dist
LEVELS:           12           2           2          2          96
TYPE  :       RANDOM      WITHIN      WITHIN     WITHIN        DATA

SOURCE                SS      df                MS          F        p
=======================================================================
eyes          86400.0000       1        86400.0000     14.676    0.003 **
es/           64760.7500      11         5887.3409

orient        91760.6667       1        91760.6667     19.075    0.001 **
os/           52915.0833      11         4810.4621

eo             1261.5000       1         1261.5000      0.539    0.478
eos/          25756.2500      11         2341.4773

dir           71068.1667       1        71068.1667      7.842    0.017 *
ds/           99682.5833      11         9062.0530

ed              322.6667       1          322.6667      0.150    0.706
eds/          23670.0833      11         2151.8258

od             5828.1667       1         5828.1667      4.124    0.067
ods/          15546.5833      11         1413.3258

eod             337.5000       1          337.5000      0.143    0.712
eods/         25940.2500      11         2358.2045
```

CHAPTER 3

EXPERIMENT 5

PERCENTAGE CORRECT

Male Items

```
FACTOR:          sub         view        light        perc
LEVELS:           12            3            2          72
TYPE  :       RANDOM       WITHIN       WITHIN        DATA

SOURCE                     SS     df              MS          F        p
=======================================================================
view                1684.0278      2        842.0139      3.880    0.036 *
vs/                 4774.3056     22        217.0139

light               3133.6806      1       3133.6806      9.477    0.011 *
ls/                 3637.1528     11        330.6503

vl                  1892.3611      2        946.1806      7.180    0.004 **
vls/                2899.3056     22        131.7866
```

Female Items

```
FACTOR:          sub         view        light        perc
LEVELS:           12            3            2          72
TYPE  :       RANDOM       WITHIN       WITHIN        DATA

SOURCE                     SS     df              MS          F        p
=======================================================================
view                 486.1111      2        243.0556      0.664    0.525
vs/                 8055.5556     22        366.1616

light                425.3472      1        425.3472      2.374    0.152
ls/                 1970.4861     11        179.1351

vl                   902.7778      2        451.3889      1.300    0.293
vls/                7638.8889     22        347.2222
```

LIKENESS RATINGS.

MALE ITEMS.

```
FACTOR:          sub        view       light         rat
LEVELS:           12           3           2          72
TYPE  :       RANDOM      WITHIN      WITHIN        DATA

SOURCE                  SS     df                 MS          F       p
===========================================================================
view                8.1997      2             4.0998     14.606   0.000 ***
vs/                 6.1753     22             0.2807

light               7.3472      1             7.3472     48.899   0.000 ***
ls/                 1.6528     11             0.1503

vl                  4.6997      2             2.3498     19.323   0.000 ***
vls/                2.6753     22             0.1216
```

FEMALE ITEMS.

```
FACTOR:          sub        view       light         rat
LEVELS:           12           3           2          72
TYPE  :       RANDOM      WITHIN      WITHIN        DATA

SOURCE                  SS     df                 MS          F       p
===========================================================================
view                2.4601      2             1.2300      5.674   0.010 *
vs/                 4.7691     22             0.2168

light               1.0634      1             1.0634     10.310   0.008 **
ls/                 1.1345     11             0.1031

vl                  0.1580      2             0.0790      0.537   0.592
vls/                3.2378     22             0.1472
```

EXPERIMENT 6

Transformed A'

```
FACTOR:         sub       polar      light        view         asab
LEVELS:          16           2          2           3          192
TYPE   :     RANDOM      WITHIN     WITHIN      WITHIN         DATA

SOURCE                  SS       df               MS           F        p
==========================================================================
polar               0.3991        1           0.3991       2.545    0.131
ps/                 2.3520       15           0.1568

light               0.0169        1           0.0169       0.127    0.726
ls/                 1.9923       15           0.1328

pl                  1.3307        1           1.3307      15.285    0.001 **
pls/                1.3059       15           0.0871

view                0.8754        2           0.4377       1.831    0.178
vs/                 7.1699       30           0.2390

pv                  0.3539        2           0.1770       1.555    0.228
pvs/                3.4147       30           0.1138

lv                  1.1875        2           0.5937       4.310    0.023 *
lvs/                4.1323       30           0.1377

plv                 1.5708        2           0.7854       8.892    0.001 ***
plvs/               2.6498       30           0.0883
```

CHAPTER 4

TRANSFORMED A'

EXPERIMENT 7

```
2(FAMILIARITY) x 2(VIEW) x 2(LIGHT).

FACTOR:         sub      known       view      light       asab
LEVELS:          24          2          2          2         96
TYPE  :      RANDOM    BETWEEN     WITHIN     WITHIN       DATA

SOURCE                   SS      df              MS          F       p
===============================================================
known                1.4438       1          1.4438      8.674   0.007 **
s/k                  3.6617      22          0.1664

view                 0.7528       1          0.7528     15.715   0.001 ***
vs/k                 1.0539      22          0.0479

kv                   0.0339       1          0.0339      0.708   0.409
vs/k                 1.0539      22          0.0479

light                0.7461       1          0.7461     27.272   0.000 ***
ls/k                 0.6018      22          0.0274

kl                   0.0000       1          0.0000      0.000   0.992
ls/k                 0.6018      22          0.0274

vl                   0.2405       1          0.2405      4.208   0.052
vls/k                1.2575      22          0.0572

kvl                  0.0094       1          0.0094      0.164   0.689
vls/k                1.2575      22          0.0572
```

```
2(KNOWN) x 3(VIEW) x 3(LIGHT)

FACTOR:          sub       known        view       light        asab
LEVELS:           24           2           3           3         216
TYPE  :       RANDOM     BETWEEN      WITHIN      WITHIN        DATA

SOURCE                    SS      df              MS          F        p
===================================================================
known                 3.5106       1          3.5106      8.221    0.009 **
s/k                   9.3949      22          0.4270

view                  3.6389       2          1.8194     15.773    0.000 ***
vs/k                  5.0754      44          0.1153

kv                    0.1910       2          0.0955      0.828    0.444
vs/k                  5.0754      44          0.1153

light                 3.2862       2          1.6431     22.572    0.000 ***
ls/k                  3.2030      44          0.0728

kl                    0.0192       2          0.0096      0.132    0.877
ls/k                  3.2030      44          0.0728

vl                    1.6292       4          0.4073      4.050    0.005 **
vls/k                 8.8509      88          0.1006

kvl                   0.7403       4          0.1851      1.840    0.128
vls/k                 8.8509      88          0.1006
```

EXPERIMENT 8.

```
2(VIEW) x 2(LIGHT).

FACTOR:        sub       view      light        asab
LEVELS:         12          2          2          48
TYPE   :    RANDOM     WITHIN     WITHIN        DATA

SOURCE                SS     df              MS         F      p
================================================================
view              2.4752      1          2.4752    43.194  0.000 ***
vs/               0.6303     11          0.0573

light             0.3536      1          0.3536     5.919  0.033 *
ls/               0.6572     11          0.0597

vl                0.2745      1          0.2745     4.326  0.062
vls/              0.6980     11          0.0635
```

```
3(VIEW) x 3(LIGHT)

FACTOR:        sub       view      light        asab
LEVELS:         12          3          3         108
TYPE   :    RANDOM     WITHIN     WITHIN        DATA

SOURCE                SS     df              MS         F      p
================================================================
view              6.2315      2          3.1157    36.210  0.000 ***
vs/               1.8930     22          0.0860

light             1.1503      2          0.5752     4.745  0.019 *
ls/               2.6668     22          0.1212

vl                1.7025      4          0.4256     3.380  0.017 *
vls/              5.5415     44          0.1259
```

EXPERIMENT 9.

2(VIEW) x 2(LIGHT)

```
FACTOR:          sub        view       light        asab
LEVELS:           12           2           2          48
TYPE  :       RANDOM      WITHIN      WITHIN        DATA

SOURCE                    SS      df              MS          F        p
========================================================================
view                  0.3893       1          0.3893     13.301    0.004 **
vs/                   0.3220      11          0.0293

light                 0.1975       1          0.1975      2.366    0.152
ls/                   0.9183      11          0.0835

vl                    0.0887       1          0.0887      1.474    0.250
vls/                  0.6616      11          0.0601
```

3(VIEW) x 3(LIGHT).

```
FACTOR:          sub        view       light        asab
LEVELS:           12           3           3         108
TYPE  :       RANDOM      WITHIN      WITHIN        DATA

SOURCE                    SS      df              MS          F        p
========================================================================
view                  1.3051       2          0.6526     10.226    0.001 ***
vs/                   1.4040      22          0.0638

light                 0.9302       2          0.4651      5.190    0.014 *
ls/                   1.9715      22          0.0896

vl                    0.6594       4          0.1649      2.800    0.037 *
vls/                  2.5906      44          0.0589
```

EXPERIMENT 10

TRANSFORMED A'

2(VIEW) x 2(LIGHT)

```
FACTOR:         sub        view       light        asab
LEVELS:          12           2           2          48
TYPE  :      RANDOM      WITHIN      WITHIN        DATA

SOURCE                   SS      df                  MS          F       p
===============================================================================
view                 1.7976       1              1.7976     20.051   0.001 ***
vs/                  0.9862      11              0.0897

light                0.5988       1              0.5988     17.482   0.002 **
ls/                  0.3768      11              0.0343

vl                   0.7795       1              0.7795     15.871   0.002 **
vls/                 0.5403      11              0.0491
```

3(VIEW) x 3(LIGHT)

```
FACTOR:         sub        view       light        asab
LEVELS:          12           3           3         108
TYPE  :      RANDOM      WITHIN      WITHIN        DATA

SOURCE                   SS      df                  MS          F       p
===============================================================================
view                 6.3737       2              3.1868     18.620   0.000 ***
vs/                  3.7654      22              0.1712

light                3.2452       2              1.6226     18.506   0.000 ***
ls/                  1.9289      22              0.0877

vl                   2.0854       4              0.5213      4.571   0.004 **
vls/                 5.0185      44              0.1141
```

EXPERIMENT 11

TRANSFORMED A'

```
2(VIEW) x 2(LIGHT)

FACTOR:         sub        view       light        asab
LEVELS:          12           2           2          48
TYPE  :      RANDOM      WITHIN      WITHIN        DATA

SOURCE                   SS      df                 MS           F        p
===========================================================================
view                 1.2449       1             1.2449      32.446    0.000 ***
vs/                  0.4220      11             0.0384

light                0.0774       1             0.0774       2.263    0.161
ls/                  0.3765      11             0.0342

vl                   1.7918       1             1.7918      27.367    0.000 ***
vls/                 0.7202      11             0.0655
```

```
3(VIEW) x 3(LIGHT)

FACTOR:         sub        view       light        asab
LEVELS:          12           3           3         108
TYPE  :      RANDOM      WITHIN      WITHIN        DATA

SOURCE                   SS      df                 MS           F        p
===========================================================================
view                 6.1852       2             3.0926      27.711    0.000 ***
vs/                  2.4552      22             0.1116

light                1.6883       2             0.8442       7.209    0.004 **
ls/                  2.5761      22             0.1171

vl                   3.7799       4             0.9450      11.020    0.000 ***
vls/                 3.7732      44             0.0858
```

EXPERIMENT 12

TRANSFORMED A'

```
2(VIEW) x 2(LIGHT)

FACTOR:         sub        view       light        DATA
LEVELS:          12           2           2          48
TYPE  :      RANDOM      WITHIN      WITHIN        DATA

SOURCE                  SS     df                  MS          F        p
===========================================================================
view                4.1436      1              4.1436     76.071    0.000 ***
vs/                 0.5992     11              0.0545

light               0.3719      1              0.3719      3.739    0.079
ls/                 1.0941     11              0.0995

vl                  0.3122      1              0.3122      2.930    0.115
vls/                1.1721     11              0.1066
```

```
3(VIEW) x 3(LIGHT)

FACTOR:         sub        view       light     astrans
LEVELS:          12           3           3         108
TYPE  :      RANDOM      WITHIN      WITHIN        DATA

SOURCE                  SS     df                  MS          F        p
===========================================================================
view               11.4330      2              5.7165     65.022    0.000 ***
vs/                 1.9341     22              0.0879

light               1.6479      2              0.8240      7.599    0.003 **
ls/                 2.3855     22              0.1084

vl                  0.9573      4              0.2393      2.252    0.079
vls/                4.6759     44              0.1063
```

CHAPTER 5

EXPERIMENT 13

TRANSFORMED A'

```
2(VIEW) x 2(SIZE).

FACTOR:          sub        View        Size        asab
LEVELS:           12           2           2          48
TYPE  :       RANDOM      WITHIN      WITHIN        DATA

SOURCE                    SS      df                MS          F       p
===============================================================================
View                  1.5980       1            1.5980    138.309   0.000 ***
Vs/                   0.1271      11            0.0116

Size                  0.0675       1            0.0675      1.808   0.206
Ss/                   0.4107      11            0.0373

VS                    0.0161       1            0.0161      0.685   0.425
VSs/                  0.2579      11            0.0234
```

```
3(VIEW) x 3(SIZE)

FACTOR:          sub        View        Size        asab
LEVELS:           12           3           3         108
TYPE  :       RANDOM      WITHIN      WITHIN        DATA

SOURCE                    SS      df                MS          F       p
===============================================================================
View                  4.0074       2            2.0037     41.573   0.000 ***
Vs/                   1.0603      22            0.0482

Size                  0.2337       2            0.1169      1.947   0.166
Ss/                   1.3202      22            0.0600

VS                    0.3740       4            0.0935      2.133   0.093
VSs/                  1.9287      44            0.0438
```

EXPERIMENT 14

TRANSFORMED A'.

```
2(VIEW) x 2(LIGHT).

FACTOR:        sub        view       light        asab
LEVELS:         12           2           2          48
TYPE  :     RANDOM      WITHIN      WITHIN        DATA

SOURCE                  SS      df               MS          F        p
=========================================================================
view                0.7778       1           0.7778     13.772    0.003 **
vs/                 0.6212      11           0.0565

light               0.8576       1           0.8576     22.743    0.001 ***
ls/                 0.4148      11           0.0377

vl                  0.3788       1           0.3788     11.086    0.007 **
vls/                0.3758      11           0.0342
```

```
3(VIEW) x 3(LIGHT)

FACTOR:       subj        view       light        asab
LEVELS:         12           3           3         108
TYPE  :     RANDOM      WITHIN      WITHIN        DATA

SOURCE                  SS      df               MS          F        p
=========================================================================
view                3.4368       2           1.7184     17.567    0.000 ***
vs/                 2.1521      22           0.0978

light               3.5645       2           1.7822     26.871    0.000 ***
ls/                 1.4592      22           0.0663

vl                  1.6848       4           0.4212      5.543    0.001 **
vls/                3.3431      44           0.0760
```

EXPERIMENT 15

TRANSFORMED A'

```
2(VIEW) x 2(LIGHT).

FACTOR:        sub        view       light        asab
LEVELS:         12           2           2          48
TYPE  :     RANDOM      WITHIN      WITHIN        DATA

SOURCE                  SS      df                  MS          F       p
==========================================================================
view                0.3838       1              0.3838     13.821   0.003 **
vs/                 0.3054      11              0.0278

light               0.2377       1              0.2377      9.167   0.011 *
ls/                 0.2853      11              0.0259

vl                  0.4208       1              0.4208     11.896   0.005 **
vls/                0.3891      11              0.0354

3(VIEW) x 3(LIGHT).

FACTOR:        sub        view       light        asab
LEVELS:         12           3           3         108
TYPE  :     RANDOM      WITHIN      WITHIN        DATA

SOURCE                  SS      df                  MS          F       p
==========================================================================
view                2.0623       2              1.0312     13.524   0.000 ***
vs/                 1.6774      22              0.0762

light               1.5991       2              0.7996      8.180   0.002 **
ls/                 2.1505      22              0.0977

vl                  1.2549       4              0.3137      5.692   0.001 ***
vls/                2.4251      44              0.0551
```

APPENDIX B:

Examples of six of the male laser heads used in the experiments reported in Chapters 3, 4 and 5.

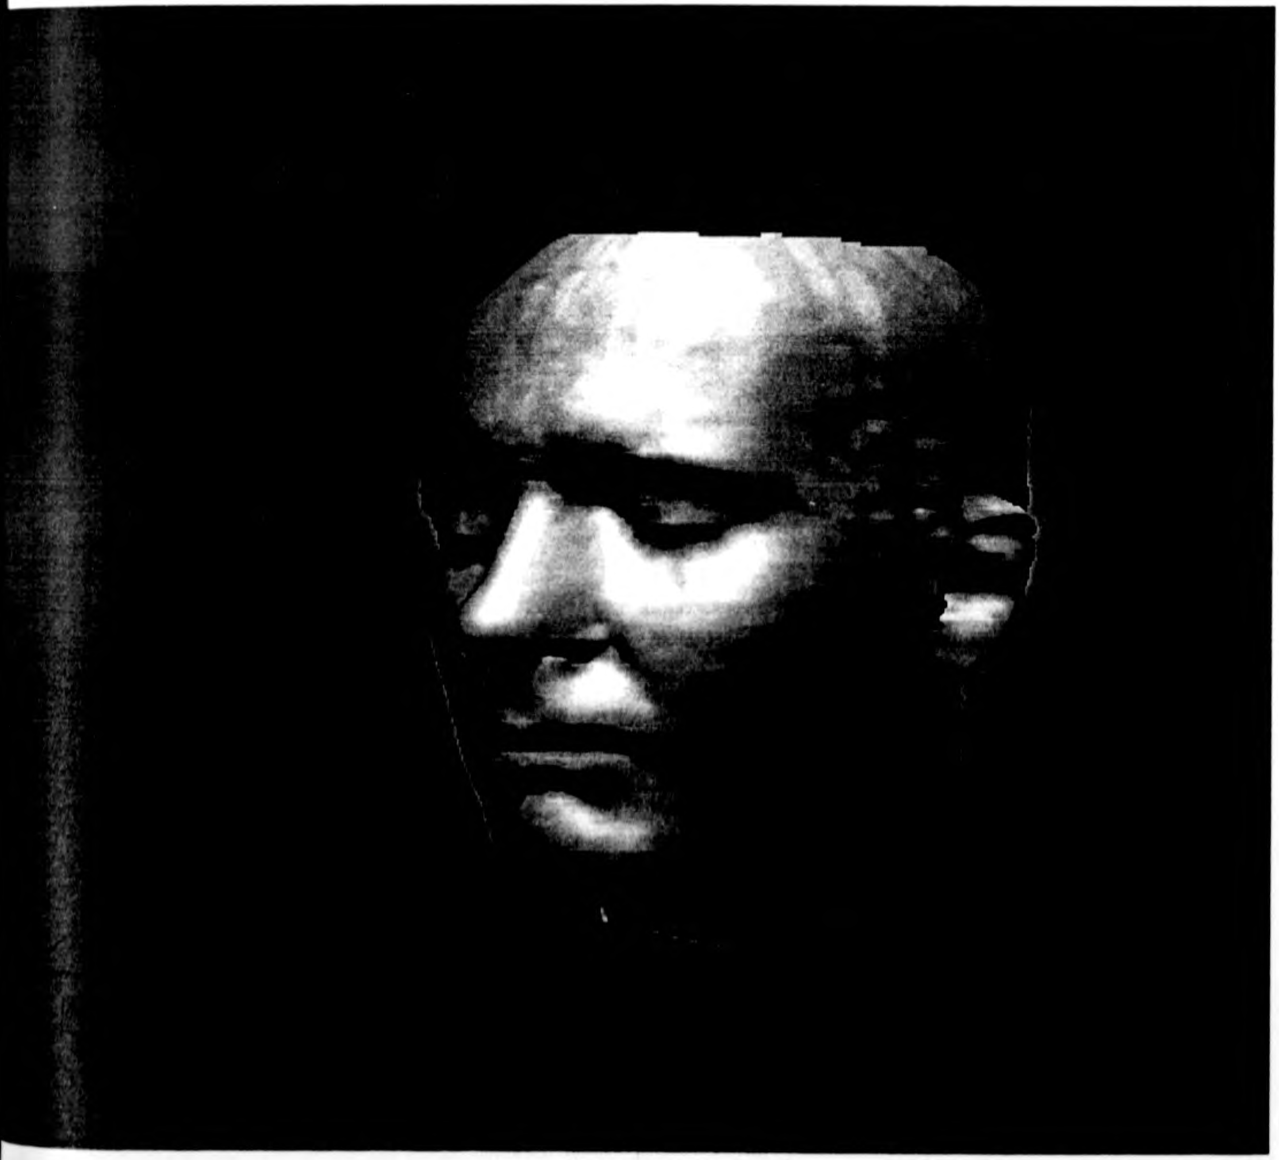

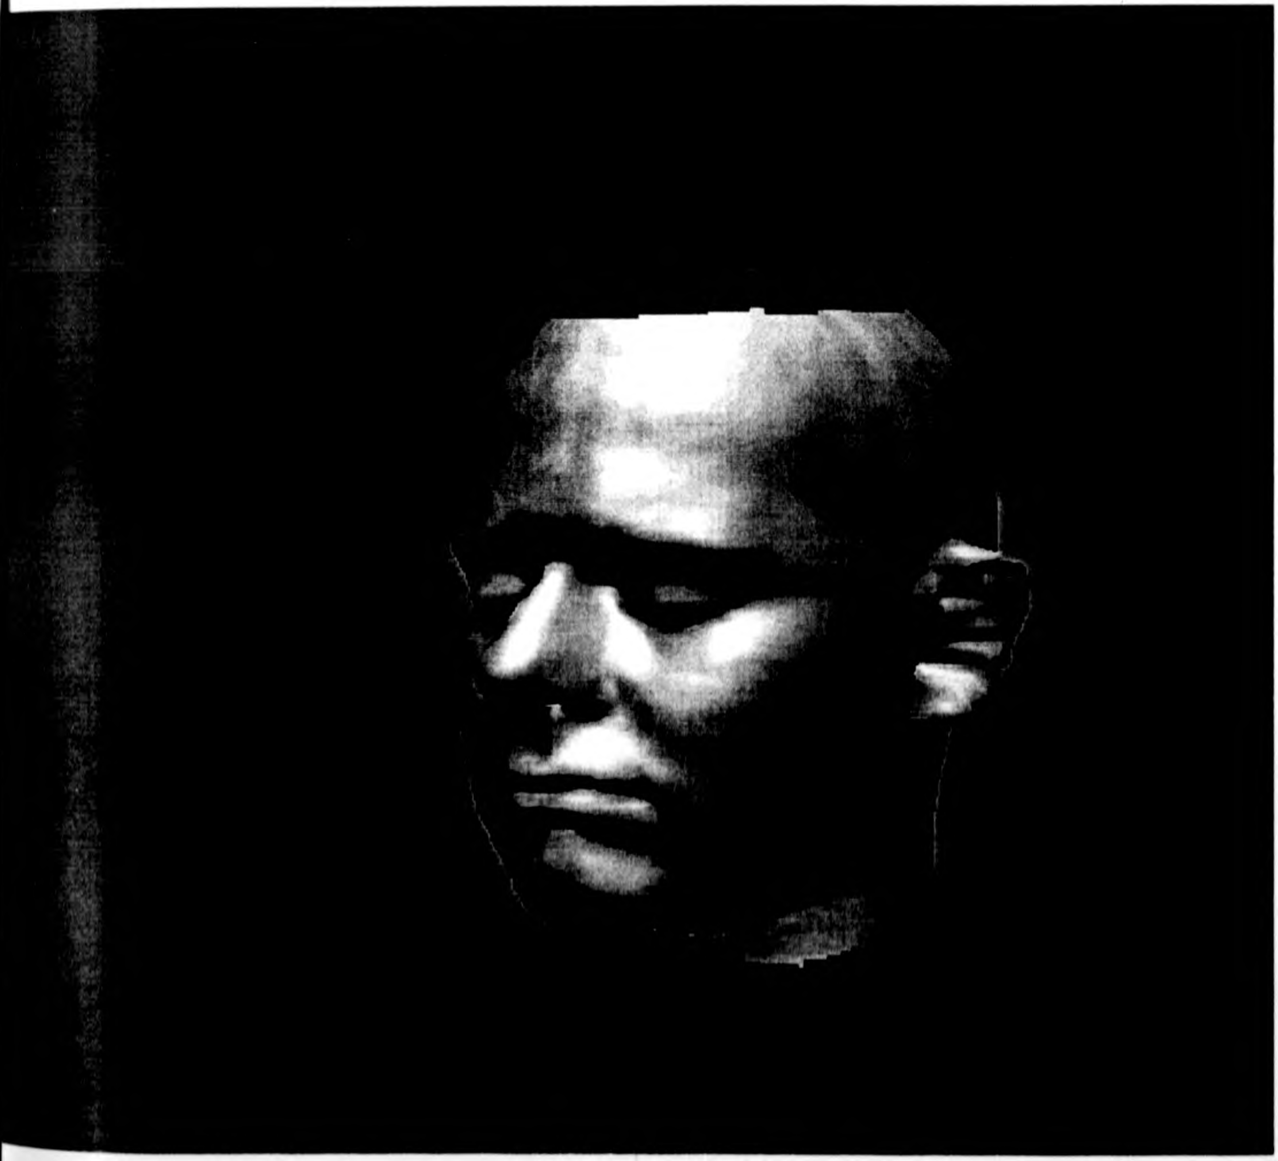

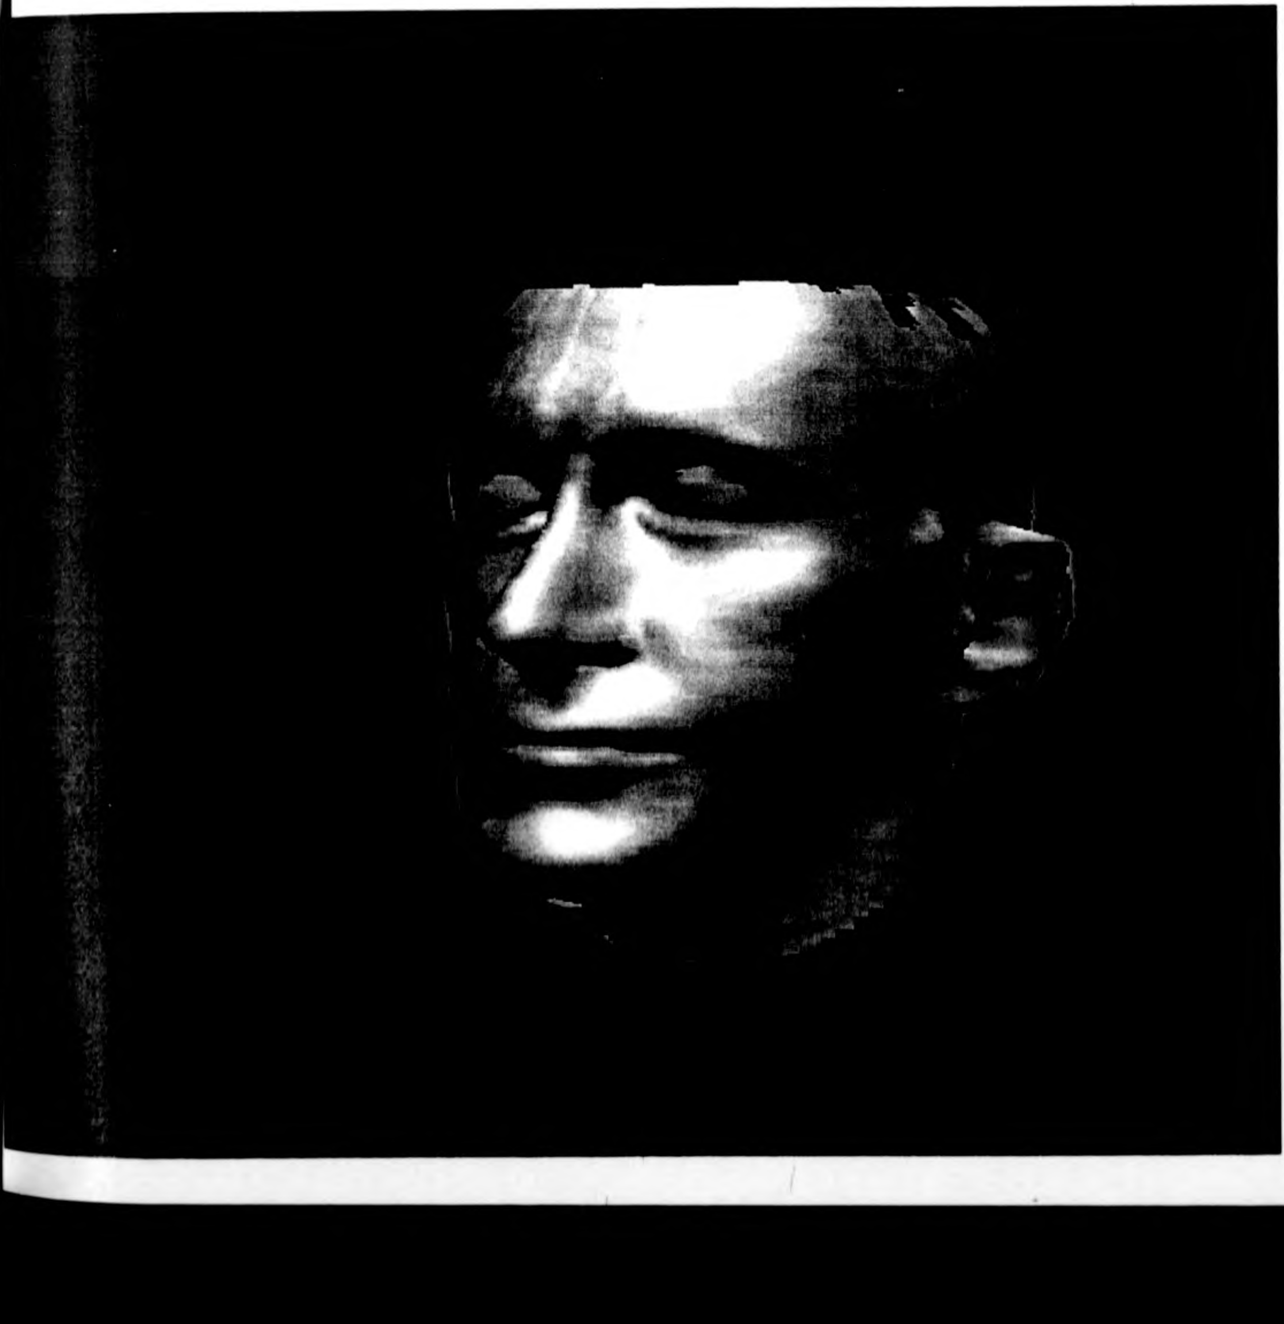

APPENDIX C:

Examples of the four female laser heads used in experiment 5, Chapter 3.

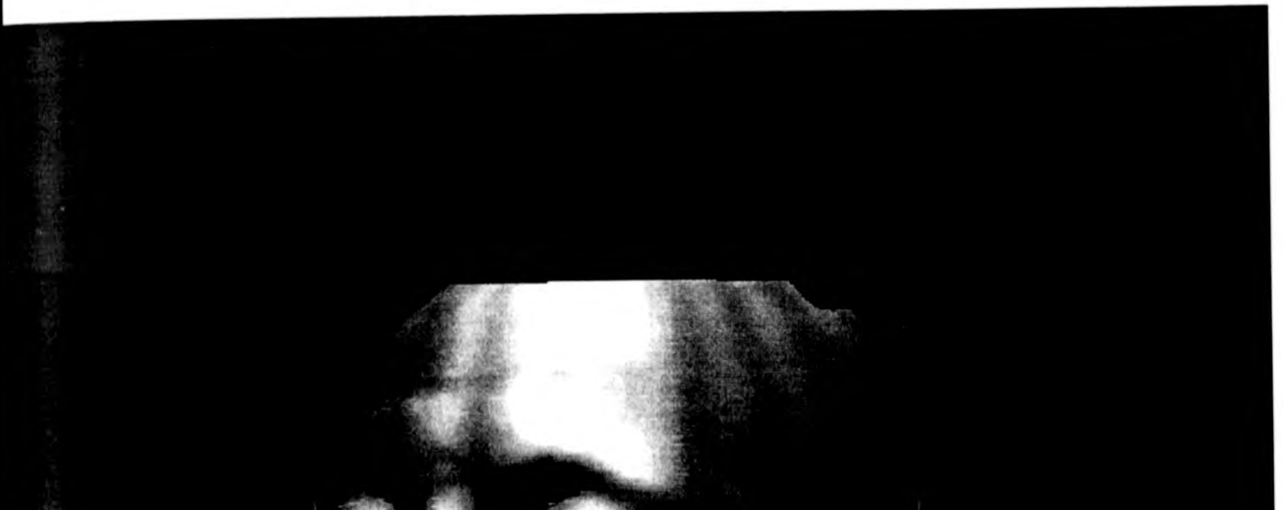

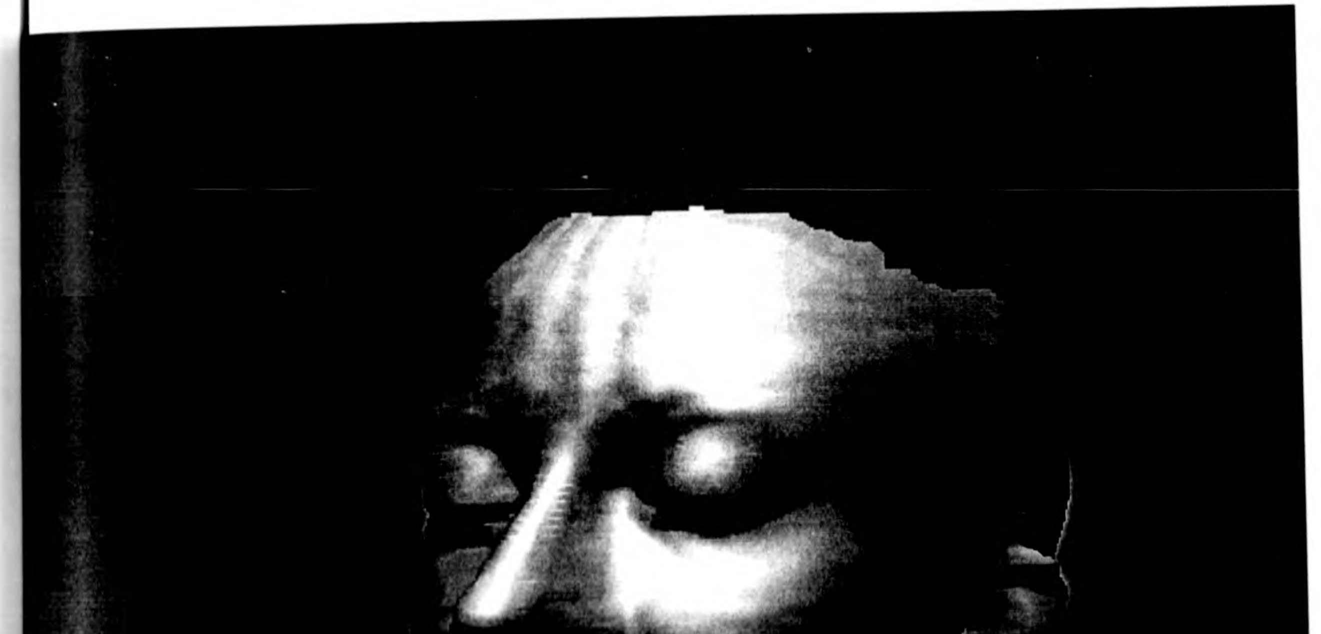

APPENDIX D:

A male item shown in photographic and laser head format for comparison.