Internal Use Only
非公開

TR-H-125

0018

# Effects of Lighting on the Perception of Facial Surfaces

*Harold Hill*    *Vicki Bruce*
*(University of Stirling)*

1995. 1. 31

ATR人間情報通信研究所
〒619-02 京都府相楽郡精華町光台2-2 ☎ 0774-95-1011

**ATR Human Information Processing Research Laboratories**
2-2, Hikaridai, Seika-cho, Soraku-gun, Kyoto 619-02 Japan
Telephone: +81-774-95-1011
Facsimile: +81-774-95-1008

# Effects of Lighting on the Perception of Facial Surfaces

Harold Hill
ATR Human Information Processing Research Laboratories
2-2 Hikaridai, Seika-cho Soraku-gun, Kyoto 619-02, Japan.
e-mail: harold@hip.atr.co.jp fax +81 7749 5 1044

and

Vicki Bruce
Department of Psychology, University of Stirling,
Stirling, FK9 4LA., U. K.
e-mail: vb1@forth.stir.ac.uk, fax: +44 (0)1786 - 467641

Running head: EFFECTS OF LIGHTING ON FACE PERCEPTION

## Abstract

A series of experiments is reported which investigated the effects of variations in lighting and viewpoint on the recognition and matching of facial surfaces. Strong effects of lighting were found which interacted with effects of viewpoint. In matching tasks, changing lighting reduced performance, as did changing view, but changing both did not further reduce performance, suggesting that viewpoint and lighting changes affected a common representational process. There were also differences between top and bottom lighting. Recognizing familiar surfaces, and matching across changes in viewpoint were more accurate when lighting was from above than when it was from below the heads, and matching between different directions of top lighting was more accurate than between different directions of bottom lighting. Control experiments demonstrated that the pattern of effects was not dependent upon the artificial materials and task demands used, and that top-lighting also benefited matching between views of unfamiliar objects (amoebae).

It is argued that edge- or image-based levels of representation are not sufficient to explain the results, particularly the observed differences between top and bottom lighting, but that the results do appear consistent with the use of a light-from-above assumption in the interpretation of facial images.

# Effects of Lighting on the Perception of Facial Surfaces

Faces, like other objects, have to be recognized and processed despite wide variations in viewing conditions. Most current theories of object recognition have concentrated on the problem of variable viewpoint assuming that the use of relatively stable lighting invariant intensity edges will overcome the problem of variable illumination (e.g. Poggio & Edelman, 1990). However, there is evidence that edge-based information may not be sufficient for face recognition (Bruce, Hanna, Dench, Healey, & Burton, 1992) and it has been argued that lighting may therefore be important (Bruce, 1988). This evidence will be reviewed followed by a report of a new series of experiments designed to investigate systematically the effects of lighting and viewpoint on face perception. The results will be discussed in terms of edge-, image- and surface-based theories of object recognition.

Evidence that edge-based information may not be sufficient for face recognition comes from at least three sources. First there have been reports that unelaborated line drawings containing only edge information are poorly recognized (Davies, Ellis, & Shepherd, 1978; Rhodes, Brennan, & Carey, 1987; Bruce, et al., 1992) although other classes of object can be accurately recognized, at least at the level of basic categories, on the basis of similar information (Biederman, 1987).

Photographic negation is also well known to disrupt face recognition (Phillips, 1972; Hayes, Morrone, & Burr, 1986; Hayes, 1988) although it leaves many properties of edges unchanged including their size, position and extent. Hayes (1988) and Hayes et al (1988) showed that negation only disrupts recognition of images containing relatively low spatial frequency information. The recognition of high-pass filtered images of faces, or line drawings lacking pigmented or shaded areas, is unaffected by negation. An edge-based representation could only accommodate such effects if the polarity of contrast across edges formed an integral part of the representation.

Attempts to apply standard edge detection algorithms to face images in order to describe or compress them for machine vision or image transmission have also proved difficult, often resulting in cluttered images (Pearson & Robinson, 1985). One successful machine algorithm for sketching faces in a realistic way incorporates two components, both sensitive to lighting. The first component, resembling an edge-finder in other algorithms, is the "valledge" detector which is sensitive to luminance valleys as well as contrast edges. The second is a threshold component which ensures that the resulting computer-drawn "cartoons" preserve areas of relative lightness and darkness from the original image (Pearson & Robinson, 1985). The cartoons which result from Pearson and Robinson's technique closely resemble sketches produced by a human artist viewing the same image (Pearson, 1992) and are identified almost as well as the photographs from which they were derived (Bruce et al, 1992).

A recent study has set out specifically to examine the effect of changes in lighting direction on face images and image representations (Adini, Moses, & Ullman, 1994). The results showed that changes in lighting, as well as viewpoint, produced greater differences in images and image representations than did changes in identity. This suggests that such descriptions would not be sufficient alone for recognition across changes in viewing conditions - two images of the same face under different viewing conditions would be ranked less similar than two different people's faces shown under the same viewing conditions. Adini et al (1994) obtained these results even when supposedly lighting invariant representations were used including various edge representations, gabor convolutions and first and second derivatives of the image. The conclusion drawn was that such representations are not by themselves sufficient to cope with the problem of variable illumination, at least for faces.

If edge information alone is not sufficient for face recognition some other source or sources of information must be needed. In all three of the above cases it has been suggested that the absence or disruption of shadow and shading information contributes to the effect. For example, the poor recognition of unelaborated line drawings has been attributed in part to the lack of information about shading and shadows (Davies, et al., 1978) - skilled artists add these features to their work in addition to lines. For the automatic cartoons of faces there is evidence that the threshold component is essential for their recognition. Using pictures of famous faces, valledges alone were recognized 67.2% but performance rose to 93.3% when the threshold component was added (against a baseline rate of identification of the original photographs set at 100%) (Bruce, et al., 1992). The threshold component may add necessary low spatial frequency information about shadows and shading as well as about darkly pigmented areas like the hair, eyes and eyebrows. The disruptive effect of photographic negation may also arise in part because it disrupts the interpretation of shading and shadow cues to three-dimensional shape, although the appearance of pigmented areas is also affected (Phillips, 1972). The appearance of photographic negatives is incompatible with any real light source (Johnston, Hill, & Carman, 1992).

If shadows and shading are necessary for accurate face perception, then lighting direction, which is a determinant of shadow and shading, will be important. In the experiments reported here we set out to investigate this by examining systematically the effect of variations in lighting direction on the perception of facial surfaces. Previous work has shown that there are particular problems associated with the perception of faces lit from below (Johnston, et al., 1992) perhaps because the interpretation of shape-from-shadows and shading relies on a light-from-above assumption (Rock, 1973; Ramachandran, 1988a; Ramachandran, 1988b).

Novel stimuli were available for these experiments which are derived from information about shape. These surface representations (see figure 1), which were derived using a laser as a depth ranging device, allowed shape perception to be investigated in isolation from effects of pigmentation and texture. They also allow lighting and viewing direction to be controlled and manipulated more easily than is possible with photographs. The stimuli provide shading, shadow and contour cues to shape which were of particular concern, for the reasons outlined above. However, because these materials are rather unnatural in appearance and difficult to recognize individually (Experiment 1), a replication of the first matching experiment (Experiment 2) using photographic stimuli is also reported (Experiment 7).

Insert Figure 1 about here

The experiments aimed to use these surface images to investigate how the perception of face shape is affected by variations in lighting direction and viewpoint. The first experiment investigated effects of these variables on the recognition of familiar faces. This preliminary study provided the motivation for the next set of experiments (2-6) which used a matching task to investigate effects of these variables in more detail. Finally, three control experiments are reported which checked that the results were not an artifact of the materials or task used.

## Experiment 1

This experiment was designed to test how accurately people could recognize familiar faces from shape information alone, and whether this was affected by the direction of lighting and viewpoint. The experiment followed up a study by Bruce, Healey, Burton, Doyle, Coombes, & Linney (1991) who examined the identification of familiar surface images (Bruce et al, 1991, Experiment 3) and the matching of unfamiliar surface images against photographic line-ups (Bruce et al, 1991, Experiments 1 and 2). In both tasks, performance at recognizing surface images was above chance but well below ceiling. Effects of differences in viewpoint and lighting direction were not investigated systematically in this earlier study. Face recognition is known to be sensitive to viewpoint, with some advantage for the three-quarter view and a decrement for profile (e.g. Bruce, Valentine & Baddeley, 1987). However, lighting direction may also be important for face recognition for the reasons given in the general introduction. In particular we expected that performance might be better with top lighting as a light-from-above assumption is thought to be important for the perception of shape-from-shading and shadow (e.g. Ramachandran, 1988a). Johnston et al (1992) showed that full face photographs of familiar faces were recognized more accurately when light was from above compared with when it was from below the head. Here we used surface images to extend this investigation to include further viewpoints, including the profile view for which information from patterns of shading might be expected to be less useful than for full face and three-quarter view images.

## Method

### Subjects

12 subjects took part in this experiment, all members of the same department as the people used as stimuli.

### Materials

The surface representations used in this and subsequent experiments were produced by using a laser as a depth finding device to measure the shape of consecutive profiles as the subject was rotated in front of the source. Fuller descriptions of the development of the method used to produce the surface representations can be found in Linney (1992), Bruce et al., (1991) and the Appendix. The data obtained was transformed into a database of 16,000 x, y, z coordinate points plotting the shape of the surface of the face which were then joined to produce a wire frame model with approximately 16,000 quadrilateral facets. These were shaded for a particular viewpoint and lighting direction using the Phong shading algorithm (Phong, 1975) which models the effect of a distant point source on a matte (Lambertian) surface.

The three views used were full-face (F), left three-quarter (Q) and profile (P) that is 0°, 45° and 90° rotations around a vertical axis centered on the head with 0° defined as the full-face view. The "Top" (T) and "Bottom" (B) directions of lighting used were from 45° above or below the line of sight respectively. Images were trimmed to remove noise.

The surface representations of eight members of the Department of Psychology at Nottingham University were used in this experiment, four male and four female. Stimuli were presented on a Macintosh computer using a program written in Supercard. Each stimulus was labeled using a randomly assigned stimulus number.

Design

A within subjects 3(View) x 2(Light) x 2(Sex of head) design was used. The levels of view were F, Q and P, of light T(op) and B(ottom) and of sex Male and Female. There were four heads of each sex giving a total of 3 x 2 x 2 x 4 = 48 trials for each subject on the recognition task. Error rates were recorded. It was planned to test for differences between levels of view. In a second stage subjects were also asked to rate each image for likeness on a scale of

1 to 5 giving a second dependent variable. The order of trials was randomized for each subject and each stage of the procedure.

Procedure

The experiment took place in a small windowless room lit by a single fluorescent bulb. Subjects were first shown a list of the names of 14 members of the department whom they might be required to recognize. The list included names of people whose heads were not included in the experiment, in order to reduce the likelihood that people would deduce identities by a process of elimination from a smaller set of names, while still allowing a check that subjects were actually familiar with the people they were being asked to identify. Subjects who reported not being familiar with all the people listed took no further part in the experiment. The list was available during the experiment and subjects were asked to guess from it when uncertain. There were equal numbers of males and females on the list.

The first stage of the actual experiment involved presenting the subjects with all the heads in a random order and asking them to identify each picture. In the second stage the stimuli were presented again in a different random order and the subjects were asked to rate them as likenesses after having been told the actual identity of the person depicted. Rating was 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 2 a) and 2 b) and likeness ratings in figures 3 a) and 3 b). It is clear that male items were much better recognized than female items (overall means 75% and 41%) and rated better likenesses (means 3.6 and 2.3). As a result of this large difference separate analyses were conducted for male and female items.

Insert Figures 2 & 3 about
here

For male items performance was best for top lit three-quarter and full-face views. There was no difference between top and bottom lighting directions for profile views which were both recognized as well as bottom lit full-face and three-quarter views. This pattern of results was confirmed by analysis of variance which showed a significant View x Light interaction, F(2,22)= 7.2, p<<.05. The simple main effect of Light was significant for full-face, F(1,33)= 18.9, p<<.05, and three-quarter, F(1,33)= 6.4, p<.05, but not at profile F(1,33)=.0, n.s. The simple main effect of View was significant for top lit stimuli, F(2,44)= 9.1, p<<.05, but not for bottom lit stimuli, F(2,44)= 1.2, n.s. Planned pairwise comparisons between levels of view (significance level .05/3=.02 to allow for the three possible comparisons) showed that full-face, t(23)=3.9, p<.01, and three-quarter, t(23)=3.5, p<.01, views were recognized significantly better than profiles when lighting was from above.

For female items there were no significant effects of light and view, nor any interaction (all p's > 0.1). The level of performance on the female faces was very low and it is possible that this low level of performance was achieved on the basis of local lighting and viewpoint invariant features.

The pattern of ratings data shown in figures 3 a) and 3 b) was similar to that for accuracy. For male items top lit three-quarter and full-face views were rated the best likenesses. The View x Light interaction was again significant, F(2,22)= 19.3, p<<.05. Again there were simple main effects of light for full-face, F(1,33)= 47.7, p<<.05), and three-quarter views, F(1,33)= 43.9, p<<.05, but not for profiles F(1,33)=.32, n.s. The simple main effect of View was significant for both top lit F(2,44)= 28.3, p<<.01 and bottom lit, F(2,44) = 3.8, p<.05 stimuli. Pairwise comparisons showed that full-face, t(23)=4.9, p<.01, and three-quarter, t(23)=7.1, p<.01, were rated significantly better than profile views when top lit.

For female items there was a main effect of light on likeness ratings, F(1,11)= 11.1, p<<.05, with top lit faces rated better likenesses than bottom lit faces. There was also a main effect of view F(2,22) = 6.4, p<.05; planned comparisons showed that three-quarters views were rated better likenesses than profile, t(23)=3.8, p<.01.

## Discussion

This experiment showed effects of both lighting and viewing direction on the recognition of facial surface representations. It also showed that male representations were better recognized than female representations.

The large difference between identification rates for male and female items replicated a previous finding using surface images (Bruce et al, 1991). This earlier study compared identification of familiar faces from surface images with identification of photographs of the same people taken with hair concealed and eyes closed (i.e. the same conditions under which the laser scans were made). Female faces were recognized as accurately as male faces from the photographs, but female faces were recognized extremely poorly from the surface images. The photographic comparison included by Bruce et al (1991) showed that the very poor performance with female surface images cannot be attributed to the absence of hair and eye features. Other features like skin texture and eyebrows present only in the photographs may be critical, or shape may just be less useful for female faces (which might also explain why their recognition was not sensitive to lighting direction). While differences between male and female items are of interest in their own right, they were not of primary concern here and were not investigated further. Only the better recognized male items were used in the other experiments reported in this paper.

The primary concern of this paper is with the effects of light and viewpoint. On both ratings and accuracy data for male items and ratings data for female items there were advantages for top lighting over bottom lighting. These results with surface images are consistent with those reported by Johnston et al (1992) using full face photographs of real faces (which included a mixture of male and female items, not separately analyzed). Profiles were less well recognized than three-quarter views, but there was no advantage of three-quarter views over full face views for the recognition of these familiar surfaces. The relative performance rates obtained with these different viewpoints are in line with findings using photographs (e.g. Bruce et al, 1987).

The difference between top and bottom lighting provides further evidence that recognition (at least of the male heads used here) is not based on lighting invariant contour information alone - if it was performance would not be sensitive to lighting. The advantage could be because the representations of known individuals incorporate information about familiar lighting direction, or because light-from-above is assumed in the interpretation of face images. For profiles, there was no effect of lighting on the accuracy or ratings for male faces, perhaps because of the salience of lighting independent contour information in this view.

For male items the effects of view and lighting direction appeared to interact. The advantages for full-face and three-quarter views were dependent on lighting being from above, a novel finding emphasizing the importance of lighting direction. Previous experiments reporting advantages for these views over profiles have only used faces lit from above.

The next set of experiments attempted to investigate in more detail the differences between top and bottom lighting found here, and also the relationship between the effect of lighting direction and that of viewpoint.

## Matching Experiments

In experiment 1 interacting effects of lighting and viewpoint were observed on the recognition of surface images of male faces. To investigate the basis of such effects in more detail, we conducted a series of experiments using a matching task. Subjects were presented with pairs of faces and had to decide if they were of the same or of different people, that is whether the faces were the same or different shapes.

The use of a matching task overcame the problem of the relatively low recognition rates reported on experiment 1 and also allowed subjects unfamiliar with the people used as stimuli to be run as subjects. Matching also allows perceptual factors, which were of primary interest, to be investigated independent of any load on memory (Kemp, McManus, & Pigott, 1990). The design of the experiments also allowed us to explore how matching was affected by changes in lighting and viewpoint direction (e.g. matching a profile view shown with top lighting against another shown in bottom lighting compared with matching when both were shown with the same lighting direction) in addition to comparing performance between different directions of lighting and viewpoint (e.g. comparing matching accuracy obtained when lighting of both faces was from above with matching accuracy when lighting of both faces was from below). The general rationale and predictions for the matching experiments will be outlined here followed by a description of the general methods. The individual experiments will then be reported.

It was assumed that matching faces across changes in viewing condition poses similar problems for shape constancy as does identifying faces despite varying viewing conditions. The faces of two different people shown under the same conditions of light and view look more similar in many ways - the same features are visible and receive direct illumination - than two pictures of a single person taken under different conditions, and are more similar according to objective measures (Adini, et al., 1994). In order to match the faces accurately it is necessary to go beyond superficial image similarities to base responses on the shape of the underlying surface. In order to match pairs (or recognize faces) accurately across changes in viewing conditions information is needed that varies more with identity than with viewing condition.

It was expected that changing viewpoint would affect simultaneous matching as this has been shown to affect the matching of unfamiliar faces across brief (Bruce, 1982) as well as longer (Bruce, et al., 1987) delays. Changing viewpoint affects many different sources of information in the image including gray levels, edges and the surfaces visible. The effect of viewpoint is a function of the three-dimensional structure of the object and it may be necessary for a representation to encode this structure in order to cope with changes in view. Alternatively, there is evidence for objects that two-dimensional transformations of the image may be sufficient to match between viewpoints (Poggio & Edelman, 1990; Bülthoff & Edelman, 1992).

The effects of changes in lighting direction are also dependent on the three-dimensional structure of the face. Like changes in viewing direction, changes in lighting direction affect the gray level information in the image, again more than changes in identity (Adini, et al., 1994). It is normally assumed that (at least some components of) edge-based representations will be insensitive to such variations (Biederman, 1987; Ullman, 1989; Poggio & Edelman, 1990) in which case lighting might not be expected to affect matching. An example of edge-based information that would be lighting invariant is the shape of the occluding contour. This changes with changes in view but not with changes in light and might be especially salient for the profile view. In Experiment 1 we found no effect of lighting direction on the recognition of profile views, consistent with the use of lighting-invariant occluding contours. Using a matching task we can ask the further question of whether matching between profile views is affected by a change in lighting direction.

Of course there are also edges that are affected by lighting, for example shadow boundaries, but these would be less useful for representing an object as they are a function of the object casting them as well as the surface on which they are cast. The shape and positions of shadow boundaries are also both viewpoint and lighting dependent. Finally, the effects of a change in lighting may be dependent on whether the view is also changed or not - for example light independent information, such as occluding contours, may be sufficient for matching when view is unchanged but no longer available when view is different.

Experiment 1 and previous work has also shown effects of lighting direction on face processing (Johnston, et al., 1992), with apparent advantages for lighting from above compared with lighting from below. The matching experiments allow further exploration of this effect, and in particular allow us to examine whether lighting from above facilitates matching at particular viewpoints. The results of experiment 1 also suggested that the effects of lighting and viewpoint may not be independent with profiles apparently less affected by lighting direction.

In all the matching experiments, heads were presented in three-quarter and profile views. Two views were sufficient to examine matching across changes in viewpoint. Profile views were used because the results of experiment 1 suggested that these might be interestingly different from the other views, being less sensitive to the effects of lighting. Three-quarter and full-face views behaved in a more similar way. The surface representations used may also be less appropriate for the full-face view because they lack the pigmented features, for example eyes, eyebrows and mouth, normally salient in this view, and so we chose to use three-quarter views in the remaining experiments.

The following experiments differed in the directions of lighting that were used. In brief, experiment 2 used both top and bottom lighting as had experiment 1 while experiments 3-6 investigated possible differences between these by examining effects of changes between different directions of top lighting, and between different directions of bottom lighting. Three control experiments were designed to test that the results were not an artifact of the stimuli or task and are reported as experiments 7- 9. The general methods common to the matching experiments will be described next.

## General Methods For Matching Experiments

### Subjects

Subjects were recruited by advertisement as part of a series of experiments. All subjects had normal or corrected to normal eyesight and were unfamiliar with the format and identities of the surface representations used, with the exception of one group in experiment 2 who knew the people used as stimuli. Each subject took part in only one experiment.

### Materials

Production of the surface images has already been described (Experiment 1, Method and Appendix). Profile and three-quarter views of each head were used throughout. The lighting directions used in each experiment will be described separately. The laser scans of eight males were used for each matching experiment, two for practice and six for the actual experiment. 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 faces side-by-side in pairs. Each image measured 9 cm x 9 cm and was viewed from a distance of approximately 1m, thus subtending a visual angle of about 5° x 5°.

### Design.

The core design for all experiments was a 3(View) x 3(Light) within subjects design. The levels of view were PP (both heads shown in Profile view), QQ (both heads shown in three-Quarter view) or PQ (1 P and 1 Q with left/right positions balanced). The levels of lighting were D1D1 (both heads lit from direction1), D2D2 (both lit from direction2) or D1D2 (one lit from direction1 and one lit from direction2, left/right positions balanced). The two directions of light used, D1 and D2, varied between experiments. Comparisons were planned to look for differences between the different levels of each variable. The full design is shown in Table 1.

As can be seen from Table 1 each cell contained six "Same" trials and six "Different" trials. "Same" trials showed two pictures of the same person while "Different" trials showed pictures of different people. The pairing of identities in "Different" trials was balanced as fully as possible to control for similarity between heads. Whether the trial was a "Same" or "Different" trial determined the correct response which was independent of viewing conditions. For each cell both hits, ("Same" responses to "Same" trials), and false alarms, ("Same" responses to "Different" trials) were recorded.

There were (6+6) x 4 x 3 = 144 trials for each subject. The order was randomized for each subject with the restriction that no two consecutive "Same" trials involved the same person and there were no more than three "Same" or "Different" trials in a row.

Insert Table 1 about here

### 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 images of heads and that their task was to decide whether both images were of the same or different people. It was stressed that the "Same" pairs would not necessarily be identical but would show two pictures of the same person while "Different" trials would show two different people. Subjects were told that pictures might be identical but could also vary in viewpoint or lighting. Subjects were asked to respond with the "Apple" key on the left of the keyboard or the "Shift" key on the right. Subjects were allowed to choose which response they made with which key.

Subjects were given ten practice trials before each experiment using two faces not used in the actual experiment to familiarize them with the stimuli and task. Images during both the practice session and the experiment remained on the screen until a response was made but subjects were encouraged to respond as quickly and as accurately as possible. They were told that error rates and latencies would be recorded. The inter stimulus interval averaged one second. There was a single buffer trial - showing heads not used in the experiment proper - at the start of both practice and experimental sessions. Subjects initiated a session by responding in the normal way to this buffer trial, and they were informed that they could have a break during testing by holding down the key that they had just pressed.

## Treatment of results

In order to reflect sensitivity accurately under 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 these experiments (Norman, 1964; 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 corresponds to chance performance and an A' of 1 to 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. Response latencies were examined for indications of a speed-accuracy trade-off for which there was no evidence, times being slowest in conditions with most errors.

For all experiments, 3(View) x 3(Light) analyses of variance were carried out on transformed data, with levels as described in the design section. Graphs of raw A' s are provided for ease of interpretation. Comparisons were planned to test for differences between the directions of view used, PP and QQ, and between the directions of light, D1D1 and D2D2. We also planned comparisons to test the effect of <u>changing</u> view and light; that is between same view conditions combined, PP and QQ averaged, and the different view condition, PQ, and similarly between same light, D1D1 and D2D2 averaged, and different light, D1D2. Significance level was divided by the number of comparisons, .05/4= .01.

## Experiment 2

This experiment was designed to investigate the effects of lighting direction and viewpoint on a face matching task and also to see how this was affected by prior familiarity with the stimuli.

The pairs of images were presented lit from the top (T) or the bottom (B), as they had been in experiment 1, and in three-quarter (Q) or profile (P) view. It was assumed that matching on the basis of identity would involve many of the same processes as recognition, particularly the visual processing of the images and sensitivity to viewing conditions. As a test of the relationship between matching and recognition two groups of subjects were run, one familiar with the people used as stimuli, as the subjects in experiment 1 had been, and the other unfamiliar, as subjects in the remaining experiments would necessarily be. Familiarity has been reported to facilitate matching in previous work (Young, Hay, McWeeny, Flude, & Ellis, 1985).

General predictions about the effects of light and view were outlined in the introduction to the matching section but specific predictions will be summarized here. First, we expected to find differences between viewpoint and lighting directions as found in experiment 1. Three-quarter views should be better matched than profile views (at least when top lit, see Experiment 1), and top lit faces should be better matched than bottom lit faces (though perhaps not in profile views, see Experiment 1). It was also expected that there would be an effect of changing view, as this is known to affect sequential matching (Bruce, 1982; Bruce et al., 1987), with matching between different viewpoints worse than matching within the same viewpoint. Similarly it was thought that changing lighting direction should also affect performance, at least where lighting invariant information is not sufficient to support matching and recognition. As indicated, theoretical considerations as well as the results of experiment 1 suggested that there might be quite complex interactions between the effects of lighting (direction and change) and view (direction and change).

## Method

Details of the method were as described in the general methods section except for the following. For this experiment two groups of twelve subjects were run. The Familiar group consisted of members of the same department as the people used as stimuli, and the Unfamiliar group were not familiar with any of the people used as stimuli.

The directions of light used in this experiment were Top (T) and Bottom (B) as described for experiment 1.

## Results

Overall performance was quite high in this experiment (mean A'= .87), with subjects familiar with the people used as stimuli performing better (mean A's: familiar = .91, unfamiliar = .83). Familiarity did not affect relative sensitivity to changes in light and view and so data is presented collapsed across familiarity in figure 4.

From Figure 4 we see that, as expected, matching was most accurate when conditions of both view and light were the same, with some overall advantage for top lit (TT) three-quarter (QQ) pairs. Changing view, condition PQ, reduced performance at least when light was the same, TT or BB. When light was different, TB, there was little effect of view. Changing light, condition TB, reduced performance when view was the same. Matching between views, PQ, appeared to be best when light was from above, TT.

---

Insert Figure 4 about here

---

This pattern of effects was confirmed by a 2(Familiarity) x 3(View) x 3(Light) analysis of variance which gave a main effect of familiarity, F(1,22)=8.2, p<.01, and a View x Light interaction F(4,88)=4.1, p<.01. No other effects were significant.

There were simple main effects of light at all levels of view; for QQ pairs, F(2,132)=15.9, p<<.01, for PP pairs, F(2,132)=5.6, p<.01, and for PQ pairs, F(2,132)=5.5, p<.01. There were simple main effects of view for TT pairs F(2,132)=7.3, p<.01 and for BB pairs F(2,132)=17.1, p<.01.

The planned comparisons outlined in the general methods were carried out between the levels of light and view for each of the simple main effects with significance level adjusted to .01. These showed that changing light, (TT+BB)/2 - TB, significantly reduced performance for PP pairs, t(23)=3.1, p<.01, and QQ pairs t(23)=5.2, p<.01. Similarly, changing view, (PP+QQ)/2 - PQ, reduced performance for TT pairs t(23)=2.8, p<.01 and BB pairs, t(23)=5.9, p<.01. Comparisons between the different same light conditions, TT and BB, showed that TT pairs were significantly better matched than BB pairs when view was different, PQ, t(23)=3.0, p<.01. Similarly, QQ pairs were matched marginally better than PP pairs when light was from above, TT, t(23)=2.7, .01< p<.02.

Although A' should measure sensitivity independent of response bias, hits and false alarms were examined to check for possible biases. There was no large bias overall with subjects making 52% "same" response compared to 48% "different" responses. Unfamiliar subjects did respond "Same" more often, 56%, being more likely to make false alarms presumably because they were less able to discriminate between the faces. View had little effect on bias ("Same" responses - PP 53%, QQ 52%, PQ 52%) but light did produce some effect ("Same" responses - TT 54%, BB 58%, TB 45%). Subjects were more likely to confuse faces when they were lit from below but more likely to make misses when lighting was different. This same pattern - a general tendency to respond "different" when light is different - was also found in other experiments in this series, though we do not report these data for remaining experiments. In experiment 9 we ran a control experiment to check whether the effects of lighting observed remained when the procedure was changed to minimize such biases.

## Discussion

The results of Experiment 2 suggest that lighting direction as well as viewing direction is important for matching faces on the basis of shape. Moreover, a change in lighting direction between top and bottom was found to reduce accuracy as much as a change in view. Changed light led to reduced performance for both profile and three-quarter views. The effects of viewing and lighting condition were also found to interact; there was no additional effect of changing view when light was different and the effect of lighting was dependent on view. Importantly there was an advantage for top lighting when subjects had to match between views. It will be argued that these results are not easily explained in terms of edge- or image-based coding but imply the use of a light-from-above assumption.

When both heads were presented under the same conditions of view and light performance was high as expected. Under these conditions images look similar but any difference is evidence that the images are derived from different faces. Matching can therefore be accomplished on the basis of image descriptions alone. The advantage for top lit three-quarter compared with profile views which was found in Experiment 1 was again found in this experiment.

As expected, changing view reduced performance, in line with previous findings (Bruce, 1982; Bruce, et al., 1987), but the effect of lighting change was at least as great. The effect of lighting change even when view was unchanged suggests that lighting invariant but view dependent edge information, for example information about occluding contours, is not alone sufficient to explain face matching. Also there was no additional effect of changing view when light was different, although this would have produced additional changes to lighting invariant edge information. However we must stress that the surface representations used as stimuli did not contain some lighting invariant edge information normally found in images of faces, for example boundaries between areas of different pigmentation. In order to test if this was critical a control experiment using photographs of faces is reported later in the paper (Experiment 8).

Of course changing lighting direction does change some edge information, for example the shape and position of shadow boundaries, which may be important. The particular change in lighting used, from top to bottom, also reverses the contrast of many areas (Johnston, et al., 1992) which could affect a recognition system based on descriptions of the blobs in filtered images (Watt, 1994). These effects on edges or the image could explain the effect of changing lighting reported.

However there are differences between directions of top and bottom lighting which do not seem to be readily explained in terms of edge or image changes. Matching between views was better when light was from above than below. Images such as 1c) and 1d) are not obviously more different than images a) and b). However, because faces are complicated and not top/bottom symmetric it is difficult to predict the magnitude of the effects of view change on the image under top compared with bottom lighting. The next set of experiments controls image differences, while manipulating lighting direction, by assessing effects of inverting the images, which maintains image properties while reversing the apparent direction of lighting.

In summary, while the effect of <u>changing</u> lighting might be explained by many different schemes for visual representation including edge- and image-based schemes, these do not seem readily to explain differences between top and bottom lighting such as the advantage for matching between views when lighting is from above. Instead this result seems to provide evidence for the importance of a light-from-above assumption in face processing. In the next series of experiments we sought further evidence for this proposal.

## Experiments 3-6

These experiments investigated the basis of the differences between <u>directions</u> of top and bottom lighting reported in experiments 1 and 2. They also sought to investigate if the effects of <u>changes</u> in light were different for top and bottom lighting. Investigating these differences should help discriminate between different accounts of the visual representations used.

In all the experiments Profile and three-Quarter views were used but the directions of lighting differed. In Experiment 3 two different directions of top lighting were used. In Experiment 4 the same stimuli were presented figurally inverted, upside-down. This manipulation leaves differences between images unchanged but inverts the apparent direction of lighting. Thus inverting the stimuli also provides a test of differences between top and bottom lighting when image differences were known to be the same. However, inversion also detrimentally affects face processing (e.g. Yin, 1969). If simultaneous matching between images of faces involves some of the same processes as recognizing faces, as has been assumed, inversion would also be expected to detrimentally affect matching. In order to examine differences created by changing the direction of apparent lighting independently of the effects of inverting face images, Experiments 5 and 6 mirrored Experiment 3 and 4 using two directions of bottom rather than top lighting. While the results of each experiment are analyzed separately, comparisons of the pattern of effects between experiments are also important. For example, comparisons of Experiments 3 with 5, and 4 with 6 will reveal effects of lighting direction for upright and inverted faces separately. More crucially, comparisons of the relative effects of inversion between Experiments 3 and 4 (where changing from top to bottom directions may add to the deficit created by inverting face images) and 5 and 6 (where changing from bottom to top directions may reduce the deficit created by inverting face images) allows us to examine the effects of lighting direction, with image differences controlled, over and above any overall decrements due to inversion.

## Method

The design and method for these experiments was as outlined in the general methods section (see also Table 1). Twelve subjects were tested in each experiment, all unfamiliar with the people used as stimuli. The directions of lighting used in experiments 3 and 4 were the original Top direction from experiments 1 and 2 which will be called T1, from 45° above the line of sight, and a new direction of top lighting, T2, which was from 45° above the direction in which the head was facing. Experiments 5 and 6 used the original direction of bottom lighting, from 45° below the line of sight, called B1, together with a new direction of bottom lighting, B2. This was from 45° below the direction in which the head was facing. The choice of these directions of light meant that the effect of changing light involved only a 45° rotation of the direction of light source for three-quarter pairs compared with a 90° rotation for profile pairs.

## Experiment 3

### Heads lit from above and presented upright

Both directions of lighting used in this experiment were from above (T1 and T2). If the effects of lighting found in Experiment 2 were the result of adverse effects of lighting from below, then we would expect to find no effect of lighting direction or change in Experiment 3. However, the different lighting directions used did still result in image changes and changes in illumination dependent edges. If such changes were the basis of the effects of lighting in Experiment 2 then similar effects would be expected here. An effect of changing view was also expected, even though light is from above, as this was found in Experiment 2. Any effects of lighting and view change might again be expected to interact.

### Results

The results for Experiment 3 are summarized in figure 5 a). Overall, performance was better in this experiment than for the equivalent, unfamiliar, group in Experiment 2 (Mean A' s: Experiment 3 = .89 and Experiment 2 = .82). This was probably because all the stimuli were top lit here. As can be seen from figure 5 a) subjects seem able to match between directions of lighting and view with high accuracy when all lighting is from above.

A 3(View) x 3(Light) analysis of variance gave a significant View x Light interaction, $F(4,44)=2.8$, $p<0.05$. Analysis of simple main effects only showed a significant effect of light for PP pairs, $F(2,66)=8.7$, $p<0.05$, although it was close for QQ pairs as well $F(2,66)=2.4$, n.s. This pattern was probably because the difference between lighting directions was greater for profile views than three-quarter views, as explained in the method section. There were simple main effects of view for T1T1 pairs, $F(2,66)=3.3$, $p<0.05$, and T2T2 pairs, $F(2,66)=12.0$, $p<0.05$, but not for T1T2 pairs, $F(2,66)=0.9$, n.s. - as with Experiment 2 there was no additional effect of changing view when light was different.

Planned comparisons showed that PQ pairs were significantly worse than same view pairs, PP and QQ, only when T2T2 lit, $t(11)=3.7$, $p<.01$. The different light condition was also marginally worse for PP pairs, $t(11)=2.8$, $.01<p<.02$. There were no significant differences between the directions of light, T1T1 and T2T2, or between the views, PP and QQ.

Insert Figure 5 about here

## Discussion

Performance in this experiment was better overall than in Experiment 2 and the effects of changing light and view were smaller. This may have been because the images produced by the new levels of light were less different or because all the heads were lit from above in this experiment.

There were still effects of view change when light was the same, as expected, but not when light was different. Comparison with the results of Experiment 2 suggests that the effect of view was less when all light was from above. Again there was no additional effect of view change when light was different although additional lighting invariant contour information would have been changed.

The effect of changing light also appeared less in this experiment, with the simple main effect only significant for PP pairs. Importantly, there was no difference between the levels of light when view was different, PQ. This contrasts with the results of Experiment 2 where top-lit heads were matched across changed viewpoint better than bottom-lit heads or those with changed lighting.

In order to examine the extent to which the results were simply a result of image differences we next presented the same stimuli inverted. This would not affect image differences but would disrupt face processing and, additionally, makes all faces appear lit from below relative to the observer.

## Experiment 4

### Head lit from above presented inverted

Rotating the stimuli 180° in the image plane leaves differences between image or edge level descriptions unchanged. However inversion is well known to disrupt face processing (Yin, 1969; Valentine, 1988) and so might also disrupt matching. In addition, inversion of an image also inverts the lighting direction relative to the observer, making stimuli appear lit from below relative to the observer. Indeed this effect appears to contribute to the usual inversion effect on face perception, since inversion has a greater effect on recognition of faces lit from above than on faces lit from below (Johnston, et al., 1992).

Thus this experiment provided a test of the extent to which matching was based on simple image differences. If matching performance is determined by image differences then the results should be the same as for Experiment 3. However, if knowledge about upright faces and/or a light-from-above assumption is important then performance would be expected to be worse in this experiment.

## Results

As can be seen from figure 5 b) performance was considerably worse in this experiment with strong effects of changing light and view.

The 3(View) by 3(Light) ANOVA showed a View x Light interaction, $F(4,44) = 4.6$, $p<.05$. Analysis of simple main effects showed effects of light for PP pairs, $F(2,66) = 8.0$, $p<<.05$, and QQ pairs, $F(2,66) = 15.9$, $p<<.05$. There were simple main effects of view for T1T1, $F(2,66) = 18.2$, $p<<.05$, and T2T2 pairs, $F(2,66) = 11.0$, $p<<.05$, but not quite for T1T2 pairs, $F(2,66) = 2.6$, $p=.08$.

Planned comparisons ($p<.01$) showed no differences between Profile and Three-quarter views in this experiment or between T1 and T2 lighting. Changing light significantly reduced performance for QQ, $t(11)=3.4$, $p<.01$, and PP pairs, $t(11)=4.5$, $p<.01$. Similarly, changing view significantly reduced performance for T1T1, $t(11)=5.2$, $p<.01$, and T2T2 pairs, $t(11)=3.9$, $p<.01$.

## Discussion

The much lower level of performance on this experiment (A' .77 compared to .89 in Experiment 3) clearly shows that matching performance is not simply determined by image differences. The images and the differences between them were the same in this experiment as in Experiment 3 but the effects of changing lighting or view on performance were much larger. Instead there appears to be an inversion effect similar to that found for the recognition of faces (Yin, 1969). The change in apparent lighting direction so that it was from below may have also contributed to the lower level of performance.

In the next experiment faces were presented upright but lit from below allowing the relative contributions of lighting direction and orientation to be assessed.

## Experiment 5

### Heads lit from below presented upright

This experiment was the same as Experiment 3 except that all lighting was from below. In particular the physical changes in light source position and view were the same magnitudes in this experiment. Their effects on the images may have been different than when lighting was from above, because of the complicated structure of the face, but there seems no a priori reason for expecting changes to be greater with bottom lighting. Thus image- or edge-based accounts would not obviously predict any difference from the pattern of performance found in Experiment 3. However, if bottom lighting does disrupt face processing, as suggested by experiments 1 and 2, then performance across changes in view and/or lighting direction should be poorer in this experiment than Experiment 3.

## Results

As is clear from figure 5 c) the effects of light and view in this experiment were much larger than in Experiment 3. Analysis of variance showed a strong View x Light interaction F(4,44) = 11.0, $p<<0.05$. Analysis of simple main effects showed effects of view for B1B1, F(2,66) = 27.2, $p<<0.05$ and B2B2 pairs, F(2,66) = 25.1, $p<<0.05$. There were simple main effects of light for PP, F(2,66) = 15.0, $p<<.05$, QQ, F(2,66) = 12.2, $p<<.05$, and PQ pairs F(2,66)= 4.4, $p<.05$. However, the effect of light for PQ pairs was not as expected, B1B2 pairs being slightly better matched than pairs whose lighting direction was unchanged.

Planned comparisons (significance level =.01) showed no differences between B1B1 and B2B2 pairs or between PP and QQ pairs. Changing light, significantly reduced performance for PP, t(11)=4.2, $p<.01$, and QQ pairs, t(11)=4.2, $p<.01$. The slight advantage of changing light for PQ pairs was not significant, t(11)=2.4, n.s. Changing view significantly reduced performance for T1T1, t(11)=6.3, $p<.01$, and T2T2 pairs, t(11)=6.1, $p<.01$.

## Discussion

There were clear effects of lighting direction and view in this experiment where heads were presented upright but lit from below. In particular, subjects appeared much less able to match faces across changes in viewing conditions than in Experiment 3 where all lighting was from above. It does not appear that this would be predicted on the basis of changes in low level image properties but it is consistent with the importance of a light-from-above assumption for face perception.

Performance when faces were shown under the same viewing conditions was high in this experiment but, as already discussed in Experiment 2, such matches can be based on image descriptions alone. When light or view were changed performance was significantly worse. The only difference between this and Experiment 3 is that the two directions of lighting used here were from below. The large difference in performance between the experiments is consistent with bottom lighting disrupting face perception but does not seem readily explained in terms of edge or image differences. It appears from this experiment that subjects are worse at matching between directions of lighting, as well as views, when lighting is from below.

Performance was slightly higher when view and light were changed compared with when just view was changed, and this effect is not readily explained. More importantly, however, matching across changes in view was clearly worse in this experiment (Mean A' for PQ trials = .73) compared to Experiment 3 (Mean A' for PQ trials = .86) consistent with bottom lighting disrupting matching across view.

In the last experiment in this series the stimuli from this experiment were presented inverted so that the bottom-lit faces now appeared lit from above.

## Experiment 6

## Heads Lit From Below Presented Inverted

In this experiment the same stimuli as were used in Experiment 5 were presented inverted. The results of Experiment 4 suggest that inverted faces may be poorly matched but this may be offset in this experiment because all lighting would now appear to come from above.

If top lighting is advantageous, performance across changes in light and view might be expected to be better in this experiment than in Experiment 4. However, if differences between top and bottom lighting are because lighting from below the chin results in greater image changes or highlights less useful features performance would be expected to be worse in this experiment. Any decrement in performance in this experiment compared with Experiment 5 would be the result of an inversion effect rather than the change in apparent lighting direction.

## Results

Inspection of figure 5 d) suggests that subjects were very poor at matching across changes in view but better at matching across changes in light compared with Experiment 5. Analysis of variance showed main effects of Light F(2,22) = 7.6, p<<.05 and View F(2,22) = 65.0, p<<.05 but the interaction just missed significance F(4,44) = 2.3, p=.08. The interaction was still plotted to allow comparison with the earlier experiments.

Planned comparisons between the levels of each main effect were carried were carried out (significance level =.01). Changing light significantly reduced performance, t(35)=3.2, p<.01, as did changing view, t(35)=9.3, p<.01. There were no significant differences between PP and QQ or B1B1 and B2B2 pairs.

## Discussion

The results of this experiment suggest that figural orientation is very important for matches across changes in view (PQ)- performance was poor although lighting was from above. However, the reversal of apparent lighting direction did appear to facilitate matching across changes in lighting direction despite the inversion of the faces - changing lighting when view was the same appears to produce less of an effect than in Experiment 5. This possibility will be examined more fully in the discussion of experiments 3-6 below.

## Discussion of Experiments 3-6

Experiments 3-6 were designed as a set and so comparisons between the performance on different experiments are important.

First, the clear effect of inversion shown by comparing experiments 3 with 4 and 5 with 6 rules out an explanation of matching performance in terms of simple image differences. The same images were shown in each pair of experiments but performance was very different. The effect of inversion is also consistent with matching involving some of the same processes as face recognition which shows a similar inversion effect. Knowledge of upright faces seems particularly necessary for matches across changes in view, which were very poor when faces were shown inverted (experiments 4 and 6, PQ trials). Inverted faces may have been being processed as unfamiliar objects, a possibility investigated in Experiment 7 where unfamiliar "amoeba" like objects were used.

Comparing experiments 3 with 5 and 4 with 6 suggests that whether light was from above or below was also critical to performance. When faces were presented upright performance across changes in light and view was generally better when faces were lit from above (Experiment 3) compared with below (Experiment 5). When faces were presented inverted, performance appears somewhat better when lighting appears to come from above (Experiment 6) rather than below (Experiment 4). This general pattern is consistent with the results of experiments 1 and 2, which also showed advantages for top lighting.

Experiments 3-6 combined suggest that effects of lighting change are stronger when light appears to be from above than below, irrespective of the orientation of the heads themselves. To analyze this, we computed for each subject the average reduction in A' from same view, same light conditions to same view but different light conditions. This gives the average effect (cost) of a change in lighting for matching with constant viewpoint for upright top-lit heads (.06), upright bottom-lit heads (.15), inverted top-lit heads (i.e. apparently bottom-lit: .18) and inverted bottom-lit heads (i.e. apparently top-lit: .1). A two-way analysis of variance with factors of head orientation (upright or inverted) and apparent lighting direction (from above or from below with respect to the viewer) yields a highly significant main effect of apparent lighting direction, F(1,44)=9.6, p< .01, and no other significant effects. This pattern cannot be explained in terms of differences between images produced by lighting the face from below the chin or from above the forehead. This advantage for apparent top lighting did not extend to the different view (PQ) condition which was equally poorly matched, regardless of lighting direction, in both the experiments using inverted stimuli. However there was a very large difference in ability to match across views for upright faces between top lighting (PQ trials in Experiment 3) and bottom lighting (PQ trials in

Experiment 5) which is consistent with the advantage of top over bottom lighting for matching between views which was observed in Experiment 2 (and see also Experiments 8 and 9 below).

In the next experiment this apparent advantage of lighting from above for matching between views was further investigated using unfamiliar "amoeba"-like objects instead of faces. This allowed us to examine the effect of inverting the materials, and hence lighting direction, without contamination from more general effects of inversion on familiar objects. Moreover, it allows us to check whether the effects observed with faces generalize to another type of object shape.

## Experiment 7

### Matching "Amoebae"

In this experiment, solid three-dimensional amoeba-like unfamiliar objects were used as stimuli (see figure 6) in order to test whether the effects of lighting and viewpoint reported in Experiments 2-6 were specific to faces. As discussed in the introduction to this paper, lighting may be especially important for faces and if so the effect of variations in lighting may be reduced in this experiment. Faces all share a common configuration while no such restriction was placed on the amoeba-like stimuli used here(see Figure 6 and description in the method section). Lighting invariant features such as occluding contours varied more between amoebae than between faces and thus might be sufficient to match across changes in a view, and reduce the effects of lighting direction.

Previous work using such amoeba-like stimuli has shown that their recognition is sensitive to changes in view (Bülthoff & Edelman, 1992). The images of the amoebae used here, like most three-dimensional objects, vary with viewpoint and so view change would be expected to affect matching (compare left and right columns of Figure 6). Indeed if object knowledge is critical for matching between viewpoints, as the experiments with inverted faces suggested (Experiments 4 and 6), performance would be expected to be poor for the unfamiliar amoebae in the different view condition, although this effect might be offset by the greater variability of the stimuli.

If there is a general advantage associated with top lighting then a difference between top and bottom lighting, similar to that reported for faces (Experiments 2, 3 & 5), would be expected. The stimuli were presented upright and inverted with the expectation that matching would be better when lighting appeared to be from above and that differences in performance with top- and bottom-lit stimuli would reverse when the stimuli were inverted, ruling out explanations in terms of chance differences between top and bottom directions. In this experiment any such effect of orientation would not be complicated by an inversion effect as neither orientation of the amoebae would be expected to be better matched.

To summarize, this experiment was a test of whether the effects of view and lighting condition reported for faces generalized to unfamiliar objects, and also examined whether differences between top and bottom lighting would reverse when the stimuli were inverted.

---

Insert Figure 6 about here

---

## Method

The method for this experiment was the same as for Experiment 2 except for the materials used, and the extra between subjects factor of orientation. The stimuli were made using the ALIAS three-dimensional modeling package on a Silicon Graphics Indigo computer. A pseudo random procedure was used to select control points on a sphere which were then moved normal to the surface by a random amount to produce the "pseudopodia" which could be convex or concave. The modeled surface had the same reflectance properties as the head models used in previous experiments. Eight amoebae were produced in this way, 2 for practice and 6 for the experiment. For each amoeba, a "full-face" upright orientation was randomly assigned and "three-quarter" (Q) and "profile" (P) views were defined by rotating leftwards by 45° and 90° respectively and rendered with Top (T) and Bottom (B) lighting directions as used in Experiments 1 and 2. Inverted stimuli were produced by rotating these images 180° in the image plane.

Subjects were recruited from Doshisya University. Two groups were run differing as to whether they were shown the stimuli "upright" or "inverted". There were 24 subjects in each group.

## Results

The results are plotted in figure 7a) and b). As can be seen changing view had a large effect on matching the amoebae but changing light had much less effect. For matching between views (PQ) there was an advantage for top lighting when the stimuli were upright and "bottom" lighting when inverted. Thus matching was best when lighting appeared to be from above.

Analysis of variance gave a three-way Light x View x Orientation interaction, $F(4,184)=4.2$, $p<<.05$. The simple View x Light interaction was significant for inverted amoebae, $F(4,184)=4.1$, $p<<.05$ but just missed significance for the upright orientation, $F(4,184)=2.3$, $p<0.1$. There was no main effect of orientation - mean A' for Upright = 0.89 and for Inverted = 0.87 - but the pattern of effects was slightly different at the two orientations.

The most noticeable difference is the pattern for different view pairs, PQ. Here matching was best for TT pairs when upright but BB pairs when inverted. Analysis of simple main effects showed a Light x Orientation interaction for the PQ condition as predicted, F(2,276)=6.0. p<<.05. There were simple simple main effects of light for both Upright, F(2,276)=3.1, p<0.05, and Inverted, F(2,276)=6.2, p<<0.05, orientations. Planned comparisons were conducted using a significance level=0.05/3 = 0.02, since differences between PP and QQ trials were of no interest for amoebae, reducing the comparisons potentially of interest from four to three. These showed that TT pairs were marginally better matched than BB pairs for the Upright orientation, t(23)=2.3, .02<p<0.05 but that BB pairs were better for the inverted orientation, t(23)=2.5, p<.02. A 2 (orientation) x 2 (apparent lighting direction - top versus bottom) analysis of variance on the transformed A' scores for TT and BB different-view pairs yielded a significant main effect of lighting direction (F1,46)=5.63, p<0.05) with no other significant effects.

Insert Figure 7 about here

## Discussion

Matching unfamiliar amoeba-like objects appeared to be much more sensitive to changes in view than light. Comparison of figures 7 and 4, experiments 7 and 2, show that the effects of view were greater for amoebae than for faces but the effects of light less.

The strong effect of view reflects the very different images produce by the amoebae when shown from different angles. Also, the experiments reported with inverted faces, experiments 4 and 6, suggest that object knowledge is important for matching across views. The large effect of view was despite the greater variability of the stimuli used here compared to faces.

The greater variability of the stimuli, particularly in their lighting invariant occluding contours, may explain the reduced effect of lighting. These contours would be most useful when the views were the same as they would be changed by a change in view. When view was different such contour information would be different and under these circumstances matching was more accurate when the light was from above.

The advantage for light from above cannot be explained in terms of image properties as which images appeared lit from above depended on orientation. Similarly, if one direction of lighting was better because it highlighted more salient features this direction would still be better when the stimuli were inverted. Nor can the advantage for top lighting in this case be attributed to disruption of a characteristic low frequency pattern of light and dark which might be invoked to explain effects for faces (Watt, 1994) - the objects were unfamiliar and so would not have such a representation available. Instead, the effect appears consistent with the use of a light-from-above assumption.

To conclude, it appears from this experiment that faces may be especially sensitive to the effects of light due to their common configuration and similar arrangements of occluding contours that result. However, some lighting dependent information may be generally useful for matching between views when occluding and other lighting invariant contours differ.

Neither the amoebae nor the face surface representations contained any pigment information, which rules out the use of lighting insensitive information from boundaries between pigment areas. In order to investigate if this was critical for the apparent sensitivity to lighting of face matching we repeated Experiment 2 but using digitized photographs as stimuli.

## Experiment 8

## Matching Photographs of Faces

All the stimuli used in the experiments reported so far have been based on information about shape alone. The purpose of this experiment was to see whether the effects of lighting and viewpoint generalize to more natural stimuli which include pigment and texture information. In order to test this we replicated Experiment 2 using digitized monochrome photographs of real faces under real conditions of illumination.

The surface representations are artificial and, to a certain extent, unfamiliar stimuli though the effect of inversion reported in experiments 4 and 6 does suggest that their matching is mediated by knowledge of upright faces. They do not however contain information available from real faces (or photographs of these) such as information about differently pigmented areas like the lips or differently textured areas like the eyebrows. Both these sources of information contain properties that are lighting invariant, such as the position and shape of their boundaries, and this could facilitate matching across changes in lighting direction. Other properties present in photographs but not surface representations are also viewpoint invariant, such as the reflectance of differently pigmented areas, and this may also affect sensitivity to the effects of viewpoint.

The photographs used in this experiment were taken of people wearing bathing caps to conceal their hair, as it had been for laser scanning, as it was thought that hair would otherwise provide too obvious a cue to identity. Hair has properties which are both lighting and viewpoint invariant, like color and texture. This may explain its importance for the memory of unfamiliar faces (Shepherd, 1981). However, hair is variable in other ways - it can be cut and dyed and it grows - and is less useful for the memory of familiar faces where internal features, which cannot be so easily changed, become more important (Ellis, Shepherd, & Davies, 1979; Young, et al., 1985).

The lighting used in this experiment may also have differed in important ways from the modeled light source used when producing shaded representations of the surface representations, which modeled a directional source effectively at infinity. For this experiment a tungsten bulb situated about a meter from the face was used. This is likely to have created a more complex pattern of shading due to the possibility of highlights and mutual illumination for example.

There were thus a number of ways in which the stimuli used in this experiment may have differed from surface representations and been more similar to heads seen under non-laboratory conditions. These differences may affect relative sensitivity to changes in light and view and allow us to examine the generalizability of the results already reported for surface representations.

## Method

The method for this experiment was the same as for Experiment 2 except for the materials. Twelve subjects unfamiliar with the people used as stimuli were run.

Photographs of eight people were digitized via a video camera connected to a Sun workstation, two for practice and eight for the actual experiment. The stimuli were monochrome. Lighting was from a single 60 Watt bulb mounted in an angle poise lamp 45° above or below the line of sight at a distance of 1 meter. Approximate three-quarter and profile views were again used. Hair was concealed using a black swimming cap.

## Results

The pattern of results for this experiment, as shown in Figure 8, was very similar to that obtained with surface representations in Experiment 2 (figure 4). Again there were effects of changing lighting and viewpoint on matching which appeared to interact. Importantly, matching between views again appears better when light was from above than from below.

A 3(View) x 3(Light) analysis of variance showed that the View x Light interaction was significant, F(4,44)=5.5, p<<.05. It also showed that there were simple main effects of light for all the combinations of view; for PP pairs, F(2,66)=16.5, p<<.05, QQ pairs, F(4,44)=13.1, p<<.05, and PQ pairs, (F4,44)=6.5, p<<.05. The simple main effects of view were only significant for TT, F(4,44)=4.3, p<<.05, and BB pairs, F(4,44)=24.7, p<<.05. Planned comparisons (significance level=.01) showed a significant difference between TT and BB pairs when subjects were required to match across view, condition PQ, t(11)=3.5, p<0.01. There was no significant advantage for QQ over PP pairs (though when light was from below the difference approached significance, t(11)=2.8, p<.02). There were effects of changing view, for BB pairs, t(11)=5.9, but not for TT pairs, t(11)=2.5, n.s. Changing lighting produced a significant drop in performance for PP, t(11)=4.9, p<.01, and QQ pairs, t(11)=4.4, p<.01.

---

Insert Figure 8 about here

---

## Discussion

Interacting effects of lighting direction as well as viewpoint were found in this experiment despite the extra sources of information, including pigment and texture, available from the stimuli. There were also effects of lighting at all levels of view and an advantage for top lighting when subjects were required to match between views. The effects of view were also similar to those reported for surface representations. Therefore, the effects cannot be solely an artifact of using surface representations as stimuli.

The replication of the pattern of effects obtained with surface representations strengthens the conclusion that lighting invariant edges, thought to be sufficient for overcoming the problems of lighting invariance in object recognition (e.g. Poggio & Edelman, 1990) may not be sufficient for face perception. The stimuli used in this experiment contained additional high contrast edges, for example between differently pigmented areas, that could have facilitated matching across lighting conditions particularly when view was the same. The similarity of the results is consistent with the generalizability of findings from experiments using surface representations to real faces. In particular it does not seem that the lack of pigment and texture information can be used to account entirely for the effects of lighting reported. The advantage for top lighting when matching between views is also consistent with a general advantage for top lighting in face processing. The similarity in the effects of lighting (and view) was despite the possible differences discussed in the introduction between the real and modeled light sources.

In the final experiment we tested whether the results were an artifact of response bias due to intermixing trials involving changes in view and light.

## Experiment 9

### Blocked Control

In Experiment 2 we noted a possible response bias associated with changing lighting, where pairs shown with different lighting tended to elicit "different" responses more than pairs shown with the same lighting. In order to investigate if the effects of lighting reported were solely a result of this we re-ran Experiment 2 using blocked groups of trials. In one block pairs of faces were always lit from the same direction, TT or BB, while in the other block pairs always had different lighting, TB.

## Method

The method was the same as for Experiment 2 except that the trials were presented in two different blocks. One block contained all the same light trials, TT and BB, while the other block contained the TB (and BT) trials. There were equal numbers of trials in each block (see Table 1). The order in which the blocks was presented was alternated for the twelve subjects.

## Results

The results of this experiment are illustrated in figure 9. The overall pattern of results is very similar to that found in Experiment 2 (Figure 4).

To test whether blocking significantly affected performance the results of this experiment were combined with those of the equivalent Unfamiliar group in Experiment 2. A 3-way ANOVA was the conducted with Procedure as an additional between subjects factor, the levels being Blocked (Exp. 9) or Mixed (Exp. 2). This ANOVA gave a main effect of Procedure, F(1,22)=7.3, p<.05, but this was independent from the effects of Light and View which interacted, F(4,88)=10.8, p<<.05. No other interactions approached significance, all p's >>.1. Thus subjects were overall more accurate when presented with the trials in different blocks (Mean A' = .90 compared to .83) but this did not interact with the effects of viewpoint or lighting direction.

Each block of this control experiment was also analyzed separately in order to investigate the effects of light and view in this experiment. In the same light block there were interacting effects of View and Light, F(2,22)=5.0, p<.05, as plotted in the filled bars of figure 9. When pictures were identical matching was very accurate but a change in view reduced performance. This reduction was greater when light was from below than when it was from above as shown by a simple main effect of light for PQ pairs, F(1,33)=12.3, p<<.05. The simple main effects of light were not significant for PP, F(1,33)=1.3, n.s., or QQ pairs, F(1,33)=1.9, n.s., showing that there was no difference between TT and BB lighting when view was the same. There were effects of view at both levels of light as shown by analysis of simple main effects; at TT F(2,44)=3.5, p<.05 and at BB pairs F(2,44)=14.6, p<<.05.

In the different light block a one-way ANOVA showed no difference between the conditions of view, F(2,22)=.0, n.s.. This was equivalent to the lack of a simple main effect of view for TB pairs in Experiment 2.

---

Insert Figure 9 about here

---

## Discussion

Although the conditions of presentation in this experiment were very different from Experiment 2 the pattern of results was not. Therefore it seems unlikely that the effects of lighting reported are solely an effect of response bias. Comparing between blocks, trials where only light differed were still worse than identical trials, despite blocking. Also, there was again an advantage for top lighting when subjects had to match between views. In the different light block there was also no effect of view emphasizing that lighting invariant but view dependent features do not seem to determine performance.

Thus Experiment 9 provided a replication of the effects of Experiment 2 under conditions that were expected to change subjects' response criteria. Blocking considerably reduced the tendency to respond "different" when light was changed, though did not eliminate it completely - 48% "different" responses were made in the same light block and 53% in the changed light block. The results of this experiment show that the earlier results were not solely a result of response biases.

## General Discussion

This series of experiments investigated subjects' abilities to recognize and match faces and objects under varying conditions of lighting direction and viewpoint. It appears that for matching faces, changes in lighting direction pose difficulties as great as changes in viewing direction, and that these factors interact. There also seem to be interesting differences between lighting from above or from below which would be consistent with use of a light-from-above assumption in image or shape processing. Effects of inversion suggest that the results cannot be explained in terms of simple pattern matching strategies but that the stimuli were being processed as faces, a conclusion supported by a replication using photographs of real faces. Moreover, the effects of inversion provide further evidence of the advantages of lighting from above. A comparison experiment using unfamiliar three-dimensional objects, "amoeba", suggested that changes in viewpoint may be more disruptive than those of lighting for unfamiliar objects, but that a light-from-above assumption may be employed more generally. In this general discussion we will attempt to tie the results together and consider how they fit with the different theories of object recognition discussed in the introduction.

Effects of changing view were expected and have been reported before in tasks involving face perception (Bruce, 1982; Bruce, et al., 1987) including stimuli of the kind used here (Bruce, et al., 1991). This problem of variable viewpoint has also been widely considered in the object recognition literature (Biederman, 1987; Bülthoff & Edelman, 1992; Poggio & Edelman, 1990; Ullman, 1989). Changing viewpoint affects a lot of the information contained in the image including the shape and position of edges and other image features. The results reported here suggest that some this information, especially that contained in occluding contours affected by view but not light, is not the critical factor at least for matching faces. There were effects of lighting change even when view, and thus occluding contours, remained the same and no additional effect of changing view, and thus occluding contours, when light was different.

The effects of light generally do not seem consistent with theories relying on lighting invariant features such as high contrast edges as an input (Poggio & Edelman, 1990; Ullman, 1989). Not only does face perception appear to be lighting as well as viewpoint dependent but it appears that the effects of viewpoint are dependent on lighting. Matching between views was more accurate when faces were lit from above. The effects of view should not be dependent on the lighting conditions if matching across viewpoint was based on lighting invariant features. The interaction was replicated even when extra lighting invariant features were available, for example boundaries between areas of different pigmentation in the photographs.

Edge information, and in particular shape-from-contour, could still contribute to face matching. Changes in lighting direction as well as changes in viewing direction change the position and shape of possible sources of contour information such as shadow boundaries and isoluminant contours (Koenderink & van Doorn, 1980) and could lead to a reduction in matching accuracy. The effects might not be additive if both manipulations are affecting the same information and do not lead to too great a reduction in the overlap of the available descriptions derived from each image.

However, occluding contours might be expected to be salient in any shape-from-contour scheme (Marr, 1982; Koenderink & van Doorn, 1980) due to their recoverability from the image but they did not seem to be critical to performance. Further, a contour based scheme would not necessarily predict a difference between top and bottom lighting. There is the possibility that changes in view affect contour information more under conditions of bottom lighting for faces but this could not explain the interaction between lighting direction and orientation for the amoebae.

An alternative explanation for the advantage for top lighting is that a light-from-above assumption is critical for a contrast sensitive scheme based on the properties of low spatial frequency oriented filters (Watt & Dakin, 1993; Watt, 1994). Bottom lighting reverses the brightness of many areas, for example the eye sockets, the nostril area and the underside of the chin, in a similar way to the effects of photographic negation (Johnston, et al., 1992) and both manipulations might disrupt such a scheme. Such representations would not be available for unfamiliar objects such as the amoebae and so would not explain the light-from-above advantage found there. Instead the advantages for top lighting for faces and unfamiliar objects must reflect a more general familiarity and facility with light from above.

The explanation that appears most parsimonious for the effects of lighting direction, and its interaction with the effects of view, is in terms of the use of a light-from-above assumption in the derivation of shape-from-shading. The pattern of shading and shadow, the lack of shading, is a function of lighting direction, shape and viewpoint. Lighting direction affects which surfaces receive direct illumination and are thus shaded in the same way as viewpoint affects which surfaces are visible. Areas receiving direct illumination correspond to the surface which would be visible if the face was viewed from the position of the light source, shadowed areas corresponding to occluded surfaces and the boundary between the two corresponding to the occluding contour. In order to process faces seen under different lighting or viewing conditions it may be necessary to recover information about shape, the other determinant of the images projected, and shading as well as contour may be important for this. Indeed these two processes may not be independent - contour can affect the perception of shading directly (Ramachandran, 1988b) and may be used to provide an important constraint for its interpretation (Ikeuchi & Horn, 1981). Different cues to shape have anyway to be integrated to provide a consistent interpretation. To interpret shading patterns in terms of 3D shape requires that the ambiguity of 3D shape-from-shading be resolved in order to decide, for example, whether a particular pattern arises from a top-lit convexity or a bottom-lit concavity. There is a considerable body of evidence which suggests that the visual system uses a light-from-above assumption when interpreting shading patterns (Ramachandran, 1988a, b).

The dependence of shading on both lighting and viewpoint may also explain why these two factors were found to interact - they are affecting the same information. If they were affecting different sources of information their effect might be expected to be additive. Another aspect of the interaction was that matches between views are facilitated by lighting from above for both faces and the unfamiliar amoebae. In order to match across viewpoint it may be necessary to recover cues to three-dimensional shape from each two-dimensional view. Shading may be important for this and depend on a light-from-above assumption, explaining the decrement when light is from below.

In summary, interacting effects of lighting and viewpoint on tasks involving face perception have been reported including an advantage for top lighting for matching between views of faces, which appears consistent with a light-from-above assumption. It has been argued that the results reported are evidence that simple image or lighting invariant edge information is not sufficient for face perception but that a representation of surface shape may also be derived from shading information and used to mediate performance.

# References

Adini, Y., Moses, Y., & Ullman, S. (1994). Face Recognition: the Problem of Compensating for Changes in Illumination Direction (Technical Report No. CS93-21). The Weizmann Institute of Science.

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

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

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

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

Bülthoff, H. H., & Edelman, S. (1992). Psychophysical support for a two-dimensional view interpolation theory of object recognition. Proceedings of the National Acadamy of Science, USA, 89, 60-64.

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.

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.

Hayes, A. (1988). Identification of two-tone images: some implications for high- and low-spatial frequency processes in human vision. Perception, 17, 429-436.

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

Ikeuchi, K., & Horn, B. (1981). Numerical shape-from-shading and occluding boundary. Artificial Intelligence, 17, 141-184.

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

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

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

Linney, A. (1992). The use of 3-D graphics for the simulation and prediction of surgery. In V. Bruce & A. Burton (Eds.), Processing images of faces Norwood: New Jersey: Ablex Publishing Corporation.

Marr, D. (1982). Vision. San Francisco: Freeman.

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

Moss, J. P., Linney, A. D., Grindrod, S. R., Arridge, S. R., & Clifton, J. S. (1987). 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 hit and false-alarm rates. Psychological Review, 71, 243-246.

Pearson, D. (1992). The extraction and use of facial features in low bit-rate visual communication. Philosophical Transactions of the Royal Society of London, B335, 79-85.

Pearson, D., & Robinson, J. (1985). Visual communication at very low data rates. Proc. IEEE, 73(4), 795-811.

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

Phong, B.-T. (1975). Illumination for computer generated pictures. Communications of the ACM, 18(6), 311-317.

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

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

Ramachandran, V. (1988a). Perceiving shape from shading. Scientific American, 259(2), 58-65.

Ramachandran, V. (1988b). Perception of shape from shading. Nature, 331(6152), 163-166.

Rhodes, G., Brennan, S., & Carey, S. (1987). Identification and ratings of caricatures: implications for mental representations of faces. Cognitive Psychology, 190, 473-497.

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

Shepherd, J. W. (1981). Social factors in recognition. In G. Davies, H. Ellis, & J. Shepherd (Eds.), Perceiving and remembering faces London: Academic Press.

Status Report of the Graphics Standards Planning Committee. (1979) Computer Graphics, 13(3).

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 effects of inversion on face recognition. British Journal of Psychology, 61, 471-491.

Valentine, T., & Bruce, V. (1986). The effect of race, inversion, and encoding activity upon face recognition. Acta Psychologica, 61, 259-273.

Watt, R. (1994). A computational examination of image segmentation and the initial stages of human vision. Perception, 23, 383-398.

Watt, R., & Dakin, S. (1993). The integration of information from different scales in face perception. Perception, 22(Supplement), 12.

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

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

# Appendix

All but two of the experiments reported made use of surface representations of faces (see figure 1) and the necessary methods are described here. The representations are based on shape information alone and the techniques allow more accurate control of viewing conditions than is normally possible with photographs.

The surface representations were developed as accurate representations of the geometry of facial surfaces for use in planning facial surgery (Linney, 1992). The subject, wearing a stocking to conceal their hair and with their eyes closed, is rotated 360° in front of a laser source. The laser projects a line onto the surface of the face which follows the contours of a 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 on the line is a function of the shape of the face and, after calibration using known angles, trigonometric principles can be used to derive the relative positions of points on the surface. This is done for successive profiles as the head is rotated in front of the device. 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. The scanner 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 resultant data were transferred to a Sun 3/110 gray 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 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 lighting direction, controlled as described below. Surface were assigned a uniform mid-gray matte reflectance. The shading algorithm used for all the experiments reported in this paper was the Phong diffuse shading algorithm (Phong, 1975). This shading model was chosen as it most accurately models the physics of the illumination of matte surfaces. For Phong shading an intensity is calculated for each facet 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 perpendicular or facing away from the light source are rendered uniformly dark. As the surface modeled was Lambertian surface luminance was independent of viewing direction. 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. No ambient or spectral components were included. Phong shading does not implement ray tracing and so the effects of mutual illumination, which may be important naturally, were also absent. As the surface reflectance was defined to be a uniform gray it contained no information about pigmentation. The position of the light source and viewing direction were specified relative to the center of the head prior to shading as described next.

The available Sun software was adapted to displaying laser heads to allow camera position, and thus viewing direction, to be defined within a head centered world coordinate system. The origin for this coordinate system was the center of the head. The z-axis was defined as running into the face with the x-axis perpendicular and horizontal and the y-axis perpendicular and vertical. Positive directions were into the face, right and up respectively. Three viewpoints were used in the experiments reported here; full-face(F), left three-quarter(Q) and left profile(P). These had viewing directions as defined by camera positions with coordinates (0, 0, -z), (-x, 0, -z) and (-x, 0, 0) respectively. Values of x, y and z were chosen for each view to produce images of uniform size. Light source position, and therefore lighting direction, was defined independently of viewing position.

The alternative light source positions were specified within a normalized device coordinate system again with the z-axis defined as the viewing direction and with orthogonal horizontal and vertical x- and y-axes. Top lighting (T) and bottom lighting (B), used in both recognition and matching experiments, were from 45° above or below the viewing direction, that is as defined by light source positions (0, +/-1, -1). Figure 1 shows examples of a surface representation in all combinations of F, Q and P views with T and B lighting. Lighting directions T2 and B2, used in experiments 3-6, were from 45° above or below the direction in which the head was facing, that is with light source positions (-1,+/-1,-1) for three-quarter views and (-1,+/-1,0) for Profiles. Both light source and viewing position coordinate systems were developed in accordance with ACM core (Status Report of the Graphics Standards Planning Committee, 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 trimmed to remove areas of noise produced by reflection at the top of the head and excess neck.

# Author notes

This research was supported by a SERC research studentship to Harold Hill held at the Universities of Nottingham and Stirling. The work was linked to projects funded by the SERC's Image Interpretation Initiative to Vicki Bruce and Mike Burton, then in Nottingham (GR/F/33698) and Alf Linney at University College London (GR/F/72178). We thank Mike Burton, Pat Healey, Alf Linney and Anne Coombes for the essential contributions they made to this research. Claire O'Malley wrote the Hypercard program used in the reported experiments. Some of the experiments were presented to the European Conference on Visual Perception (ECVP 1993), abstracted as Hill, H. & Bruce, V. (1993) "An investigation into the effects of lighting and viewpoint on face processing with the use of a simultaneous face matching task", *Perception*, 22S, pp22-23. Preparation and revision of the manuscript, and the conduct of Experiment 7, was facilitated by a collaboration between ATR Human Information Processing Research Laboratories and the University of Stirling. We thank Shigeru Akamatsu and other members of HIP Department 2 for their hospitality and facilities. The manuscript was considerably improved as a result of comments by three referees and Joe Lappin. Requests for reprints should be sent to Vicki Bruce at Stirling.

Table 1

The Core Within Subjects Design For Matching Experiments

| | Lighting | | | |
|---|---|---|---|---|
| View | D1D1 | D1D2 | D2D1 | D2D2 |
| PP | 6 "Same"<br>6 "Diff" | 6 "Same"<br>6 "Diff" | 6 "Same"<br>6 "Diff" | 6 "Same"<br>6 "Diff" |
| PQ | 6 "Same"<br>6 "Diff" | 6 "Same"<br>6 "Diff" | 6 "Same"<br>6 "Diff" | 6 "Same"<br>6 "Diff" |
| QQ | 6 "Same"<br>6 "Diff" | 6 "Same"<br>6 "Diff" | 6 "Same"<br>6 "Diff" | 6 "Same"<br>6 "Diff" |

Note: P - Profile view, Q - Three-Quarter view
D1 - Lighting Direction 1, D2 - Lighting Direction 2
"Same" trials showed two pictures of the same person
"Different" trials showed pictures of different people

# Figure Captions

Figure 1: An example of the surface representations used in Experiment 1. The views shown from left to right are profile (P), three-quarter (Q) and full-face (F). The first row shows Top (T) lighting and the second row Bottom (B) lighting.

Figure 2: Percentage correct recognition rates (+/- SE) for Experiment 1.

a) Male items
b)Female items

Figure 3: Likeness ratings (+/- SE) for Experiment 1

a)Male items
b)Female items

Figure 4: Mean A's (+/- SE) for Experiment 2 shown collapsed across familiarity

Figure 5: Mean A's (+/- SE) for Experiment 3 - 6

a) Experiment 3: Upright, top lit heads
b) Experiment 4: Inverted, "Top" lit heads
    - appeared lit from below, relative to the observer
c) Experiment 5: Upright, Bottom lit heads
d) Experiment 6: Inverted, "Bottom" lit heads
    - appeared lit from above, relative to the observer

Figure 6: An example of an "amoeba" used in experiment 7 shown under different viewing conditions. The first column shows "three-quarter" (Q) view and the second "profile" (P). The first row shows top (T) lighting and the second bottom (B) lighting.

Figure 7: Mean A's (+/- SE) for Experiment 7

a) "Upright" amoebae

b)"Inverted" amoebae.

Figure 8: Mean A's (+/- SE) for Experiment 8, photographic stimuli.

Figure 9: Mean A's (+/- SE) for Experiment 9. Results from the same and different light blocks are shown together

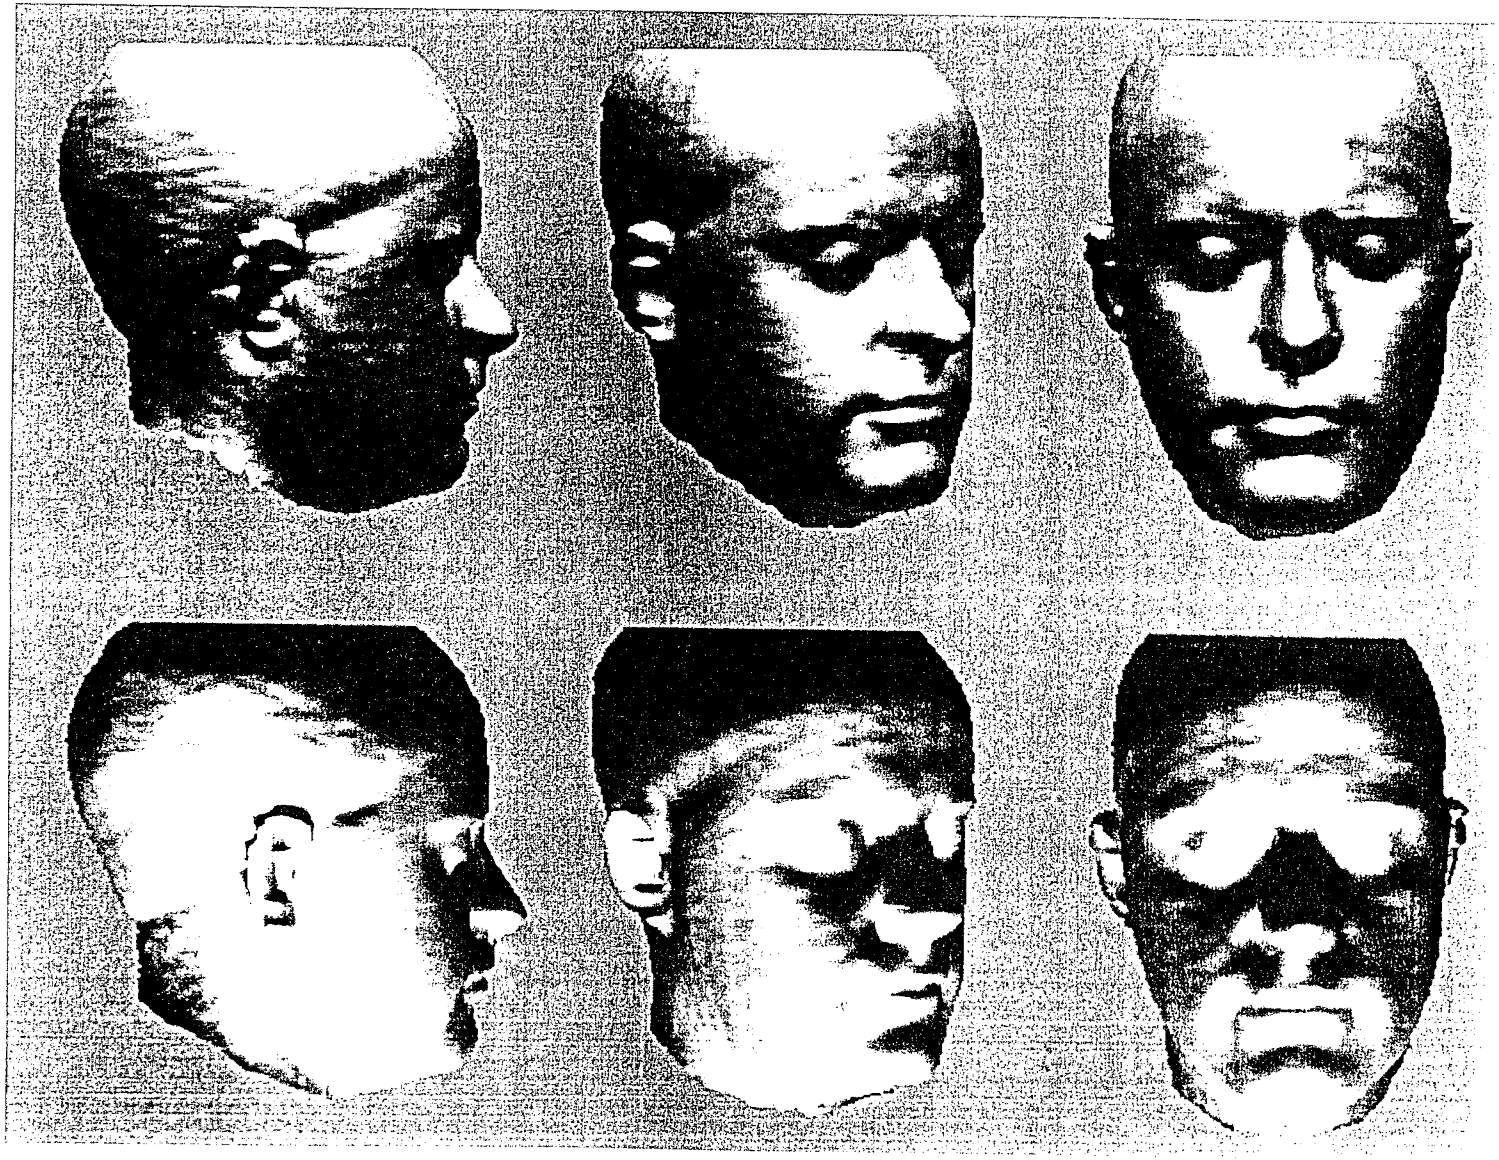

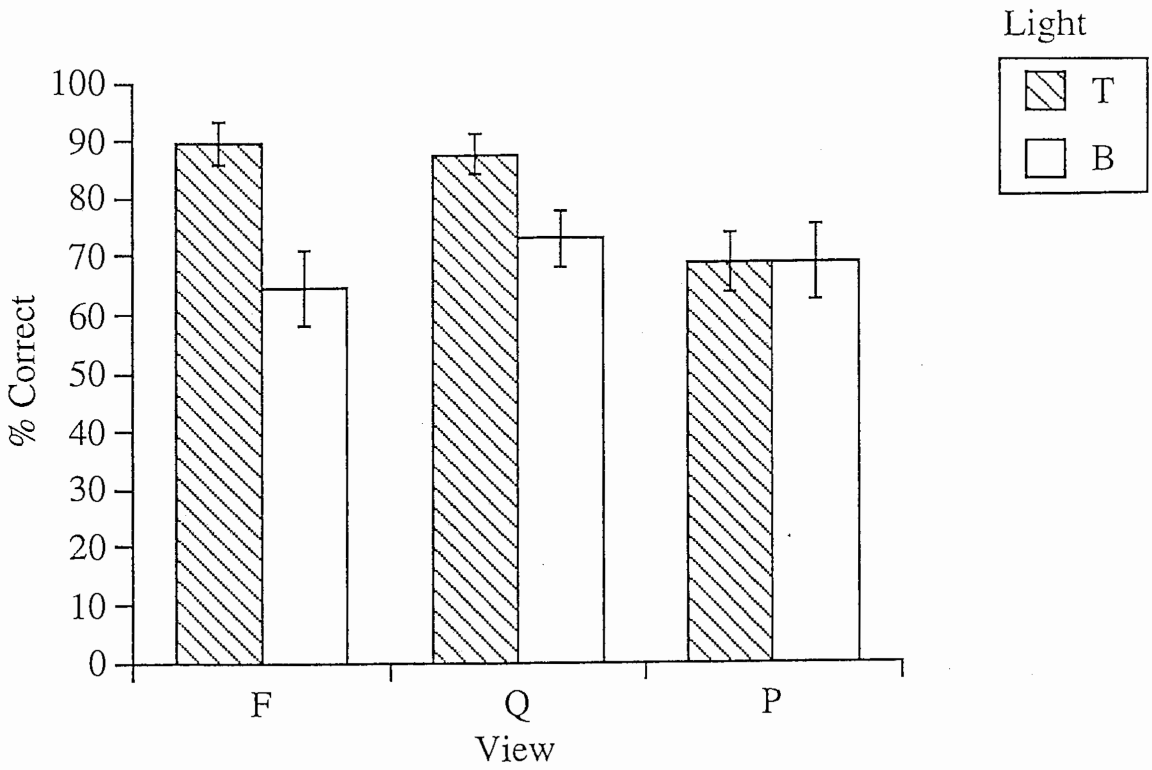

a) Male Items

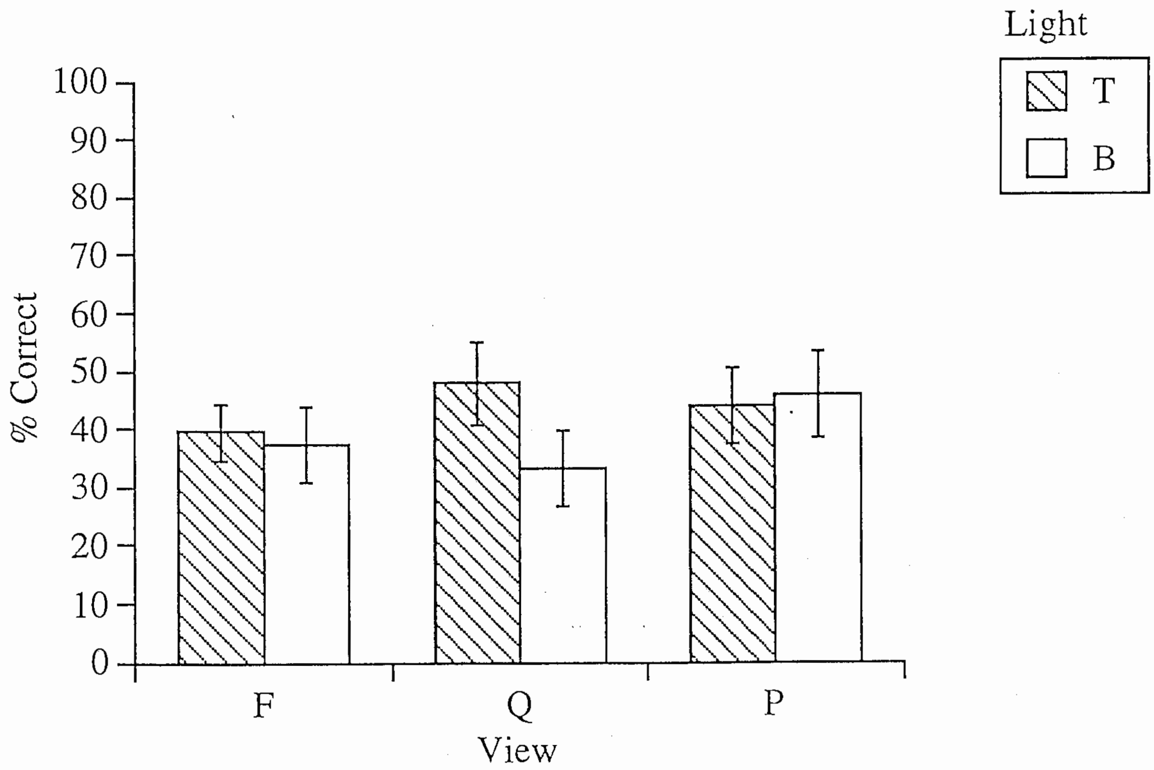

b) Female Items

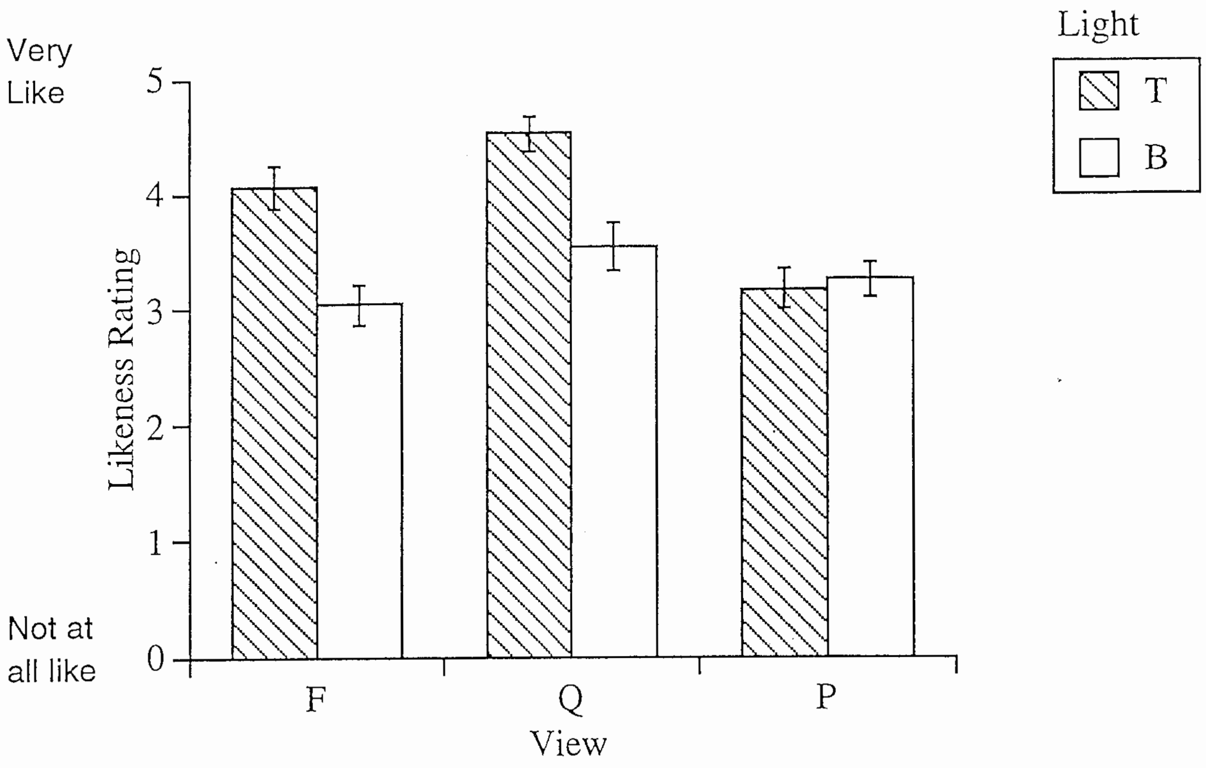

a) Male Items

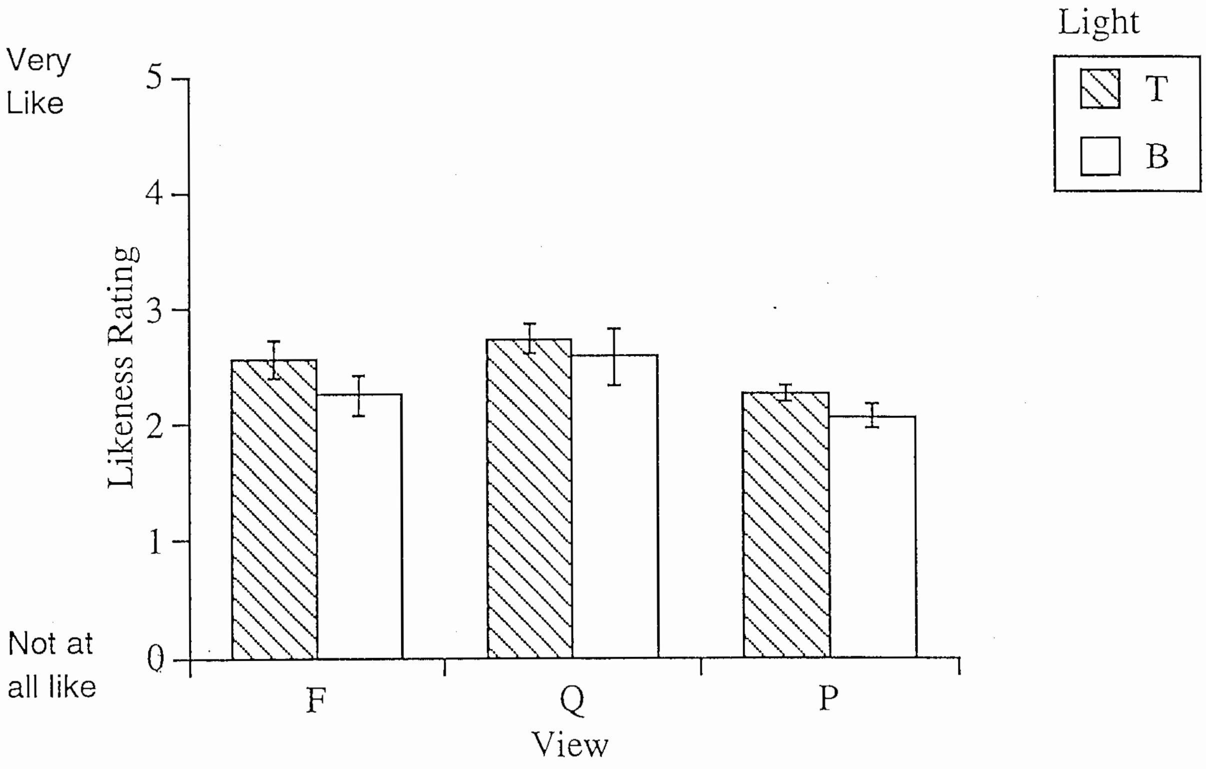

b) Female Items

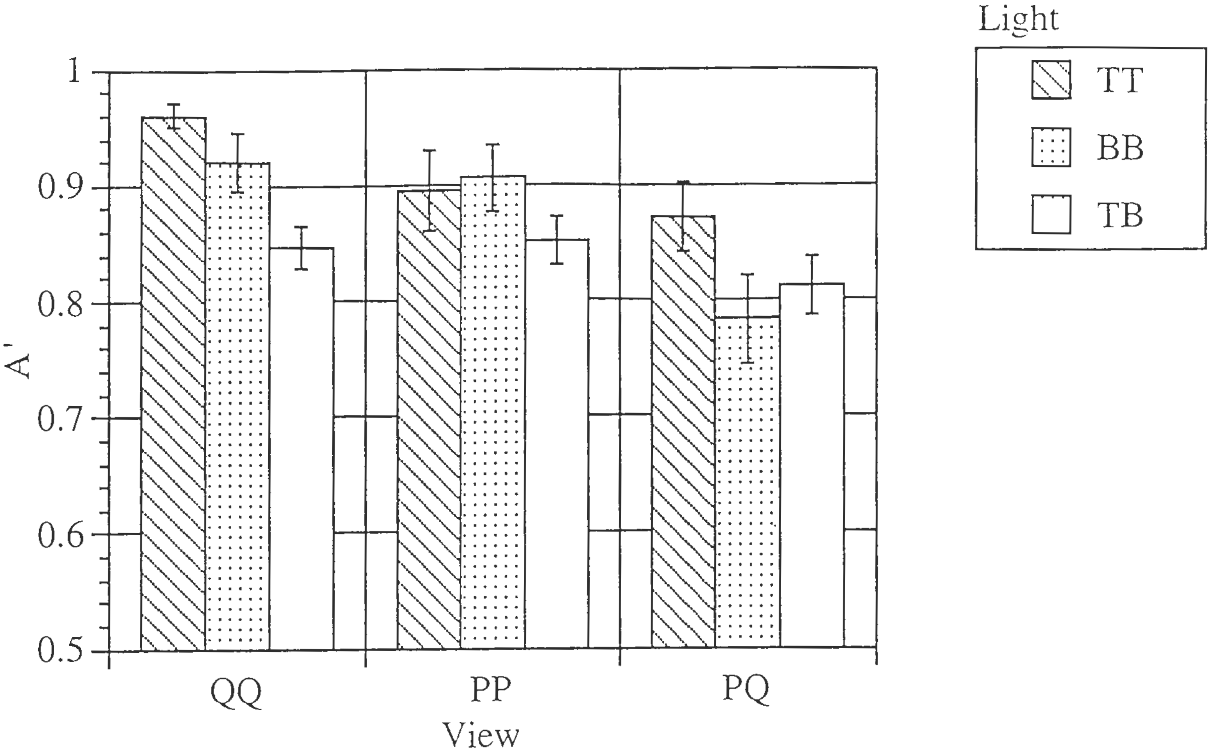
Light
TT
BB
TB
A'
1
0.9
0.8
0.7
0.6
0.5
QQ
PP
PQ
View

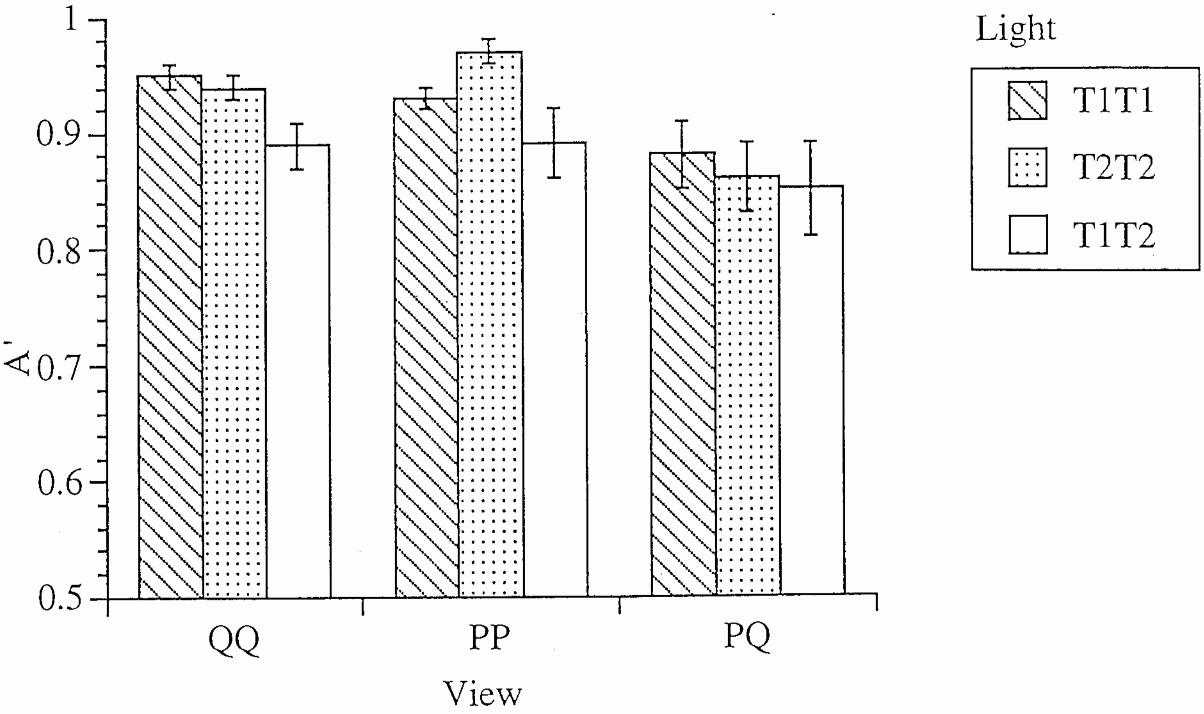

a) Upright, Top Lit Faces

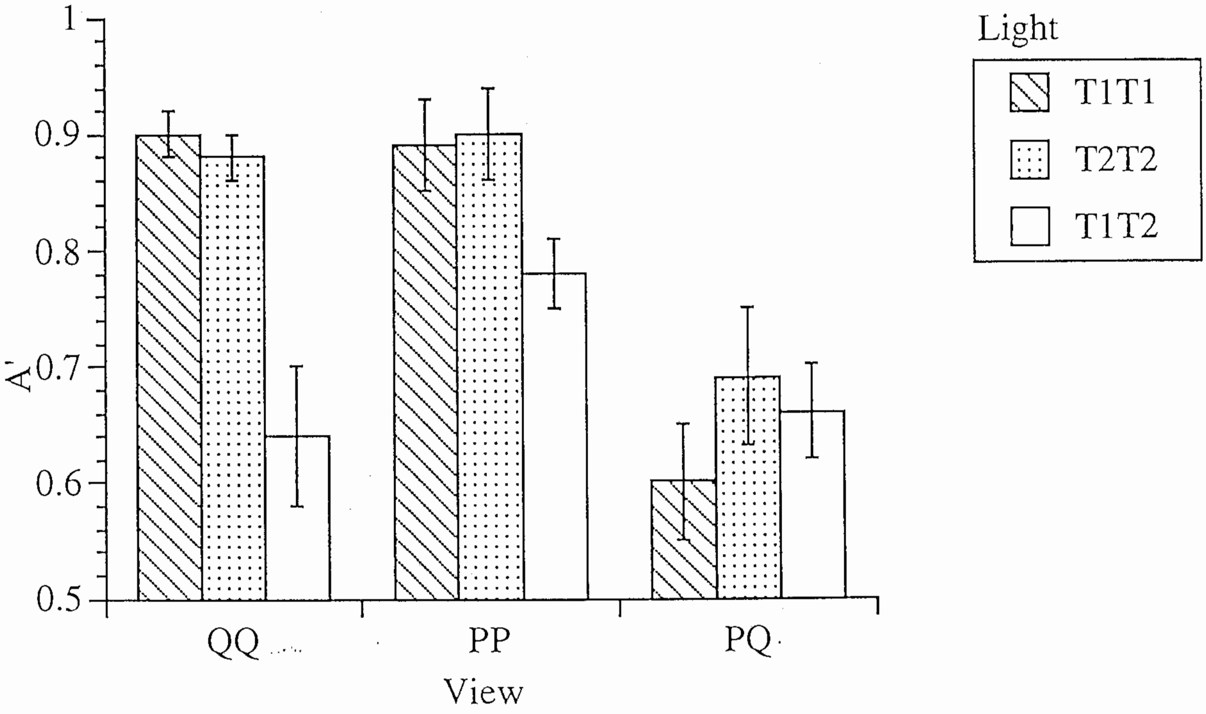

b) Inverted, "Top" Lit Faces

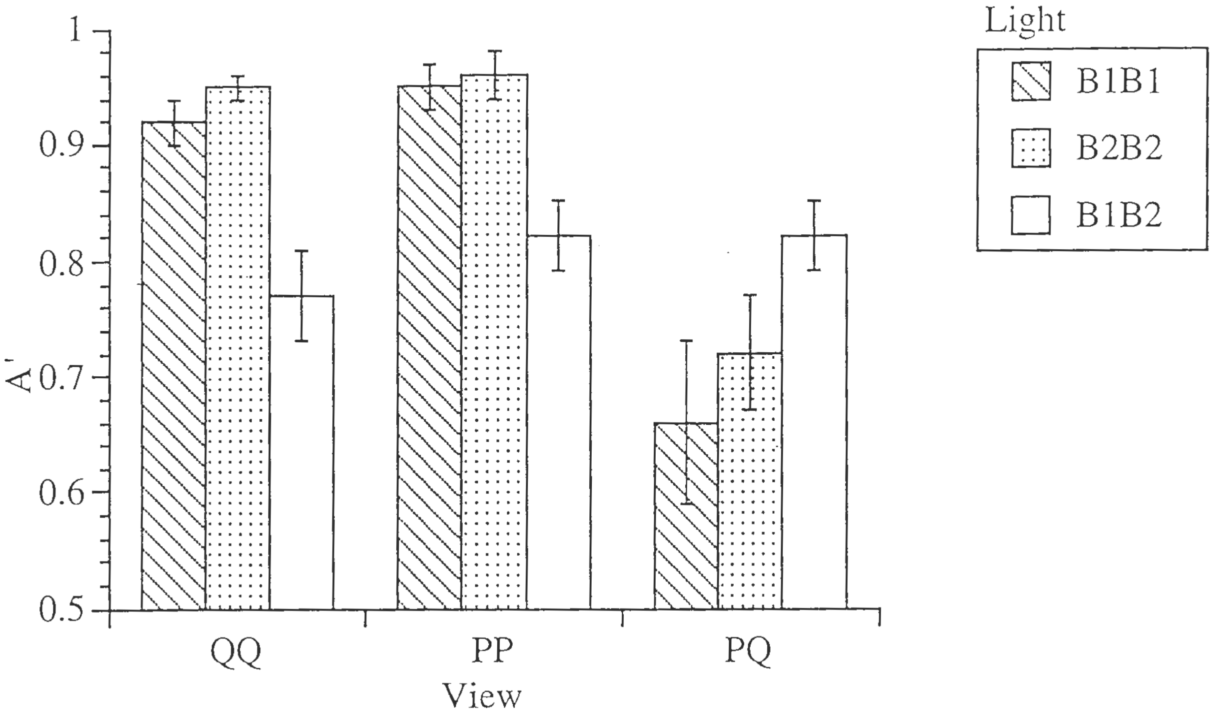

c) Upright, Bottom Lit Faces

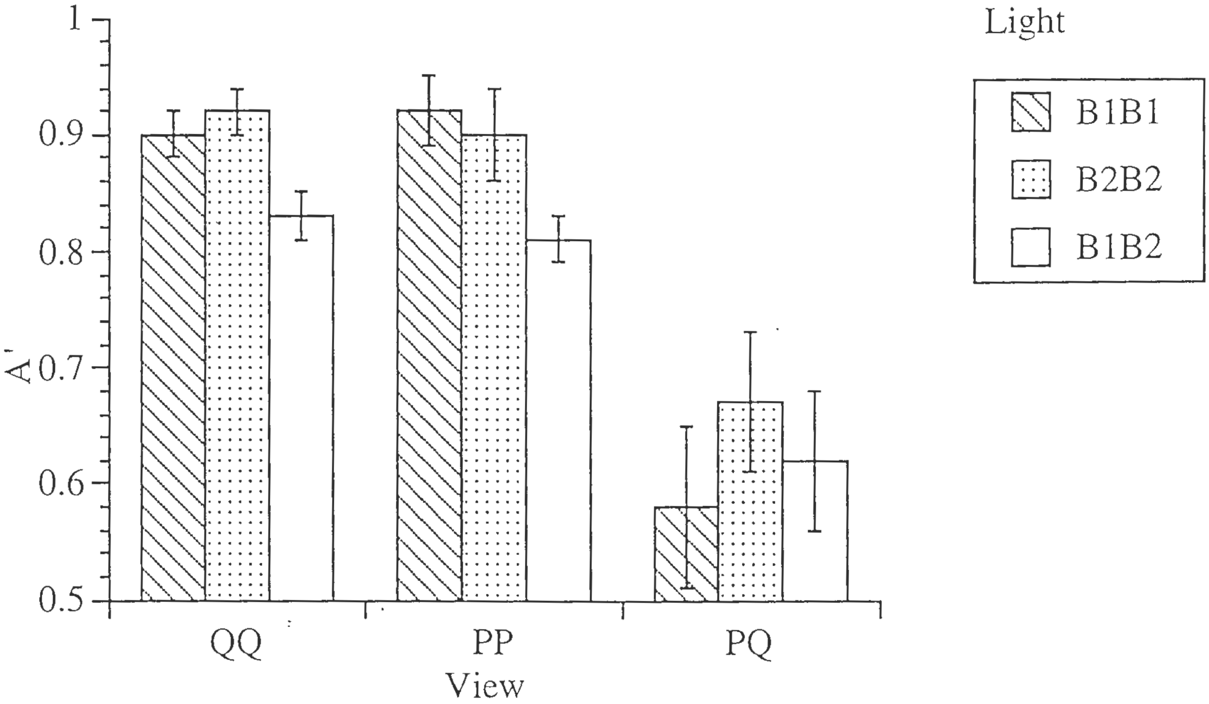

d) Inverted, "Bottom" Lit Faces

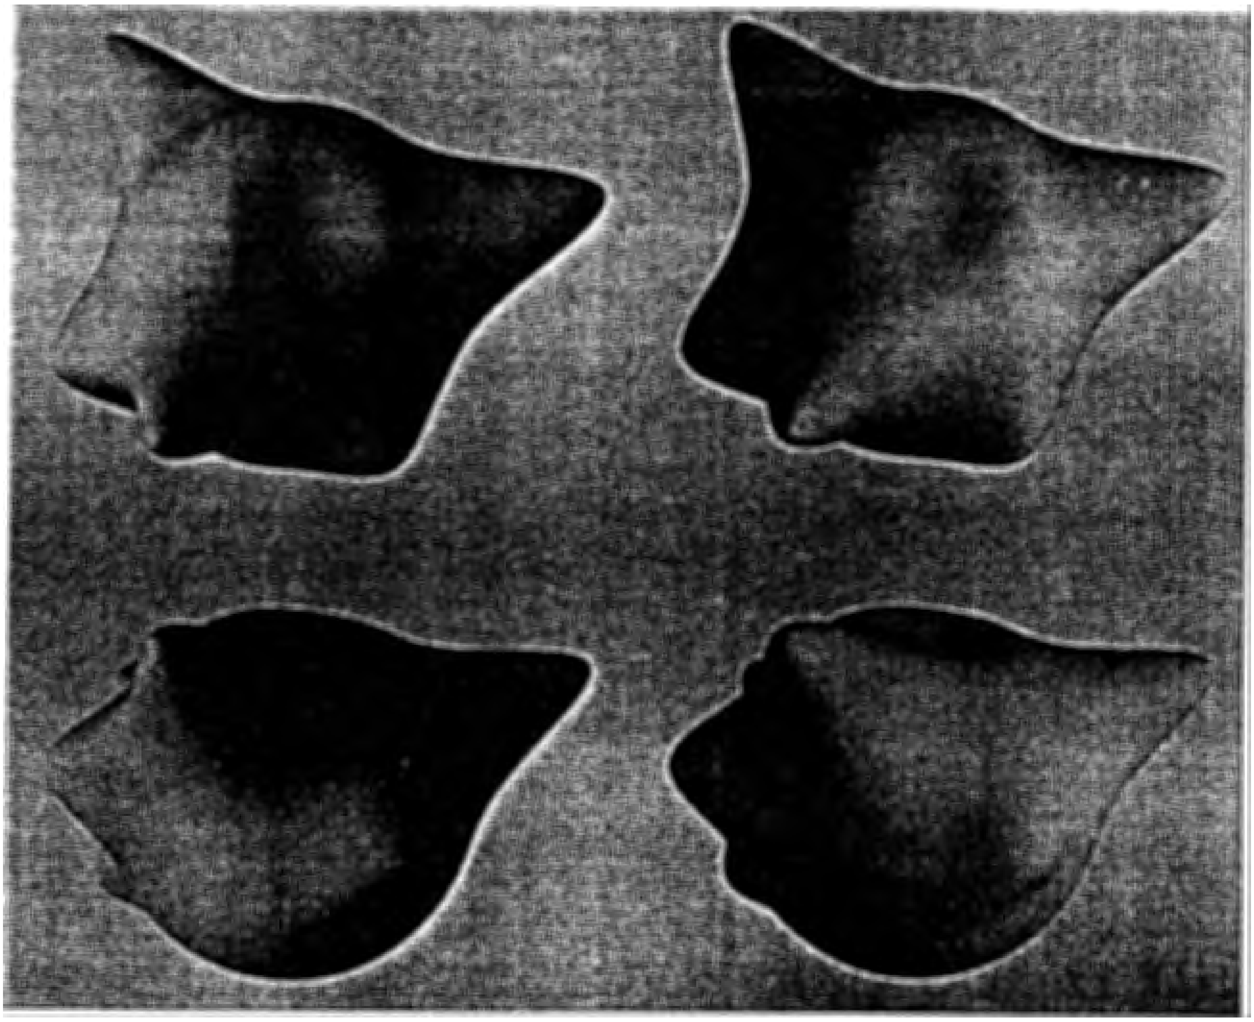

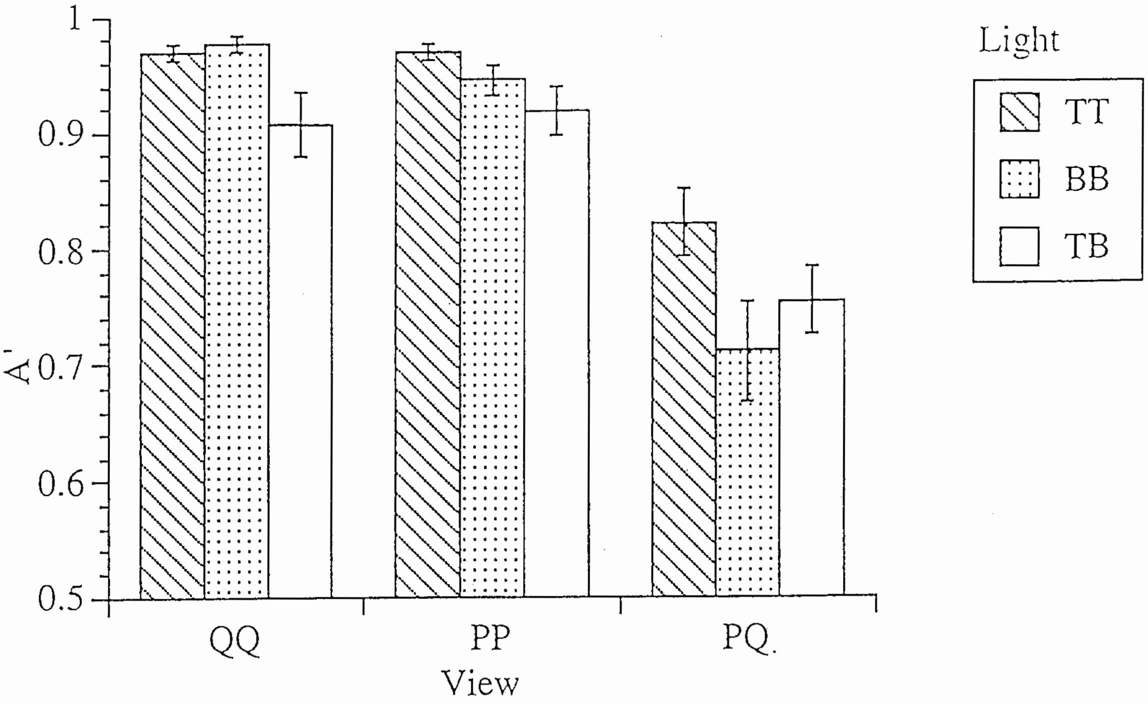

a) Upright Amoebae

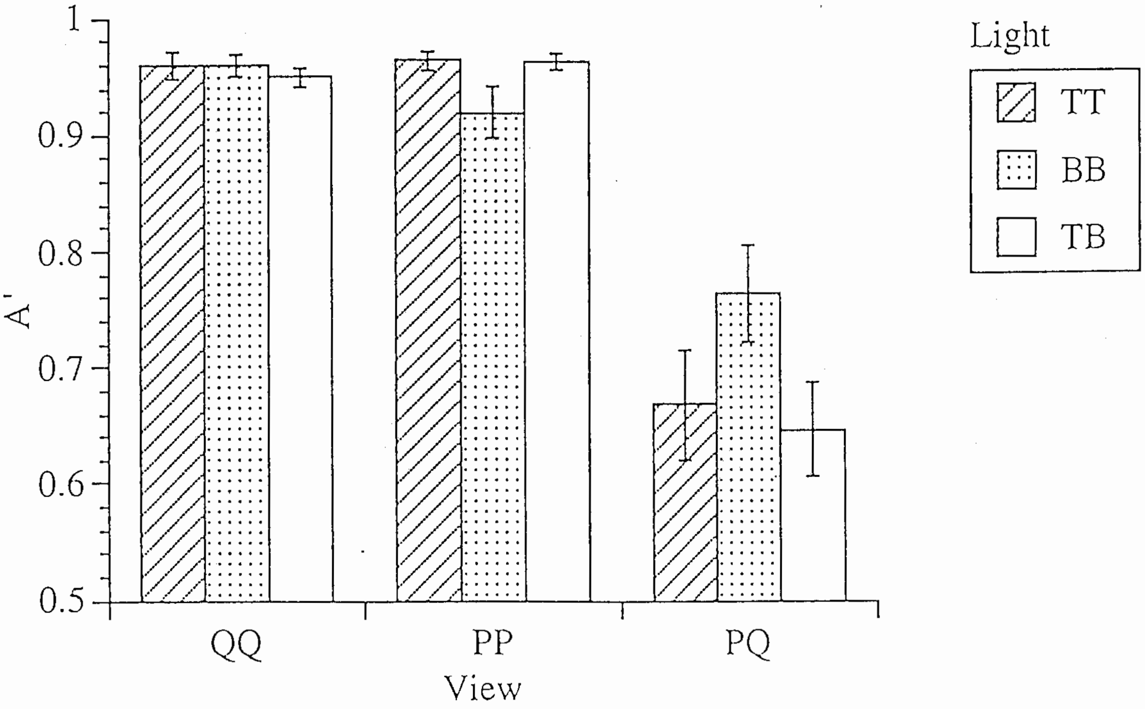

b) Inverted Amoebae

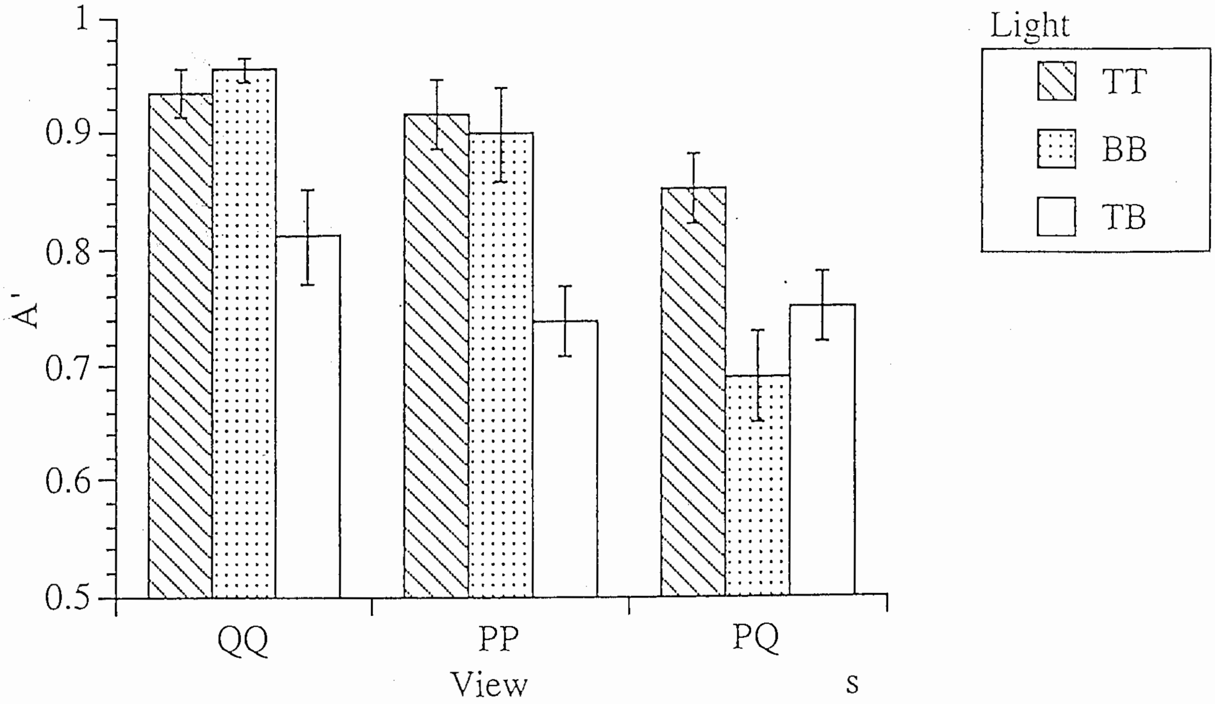
1
0.9
0.8
0.7
0.6
0.5
A'
QQ
PP
PQ
View
s
Light
▧ TT
⬚ BB
☐ TB

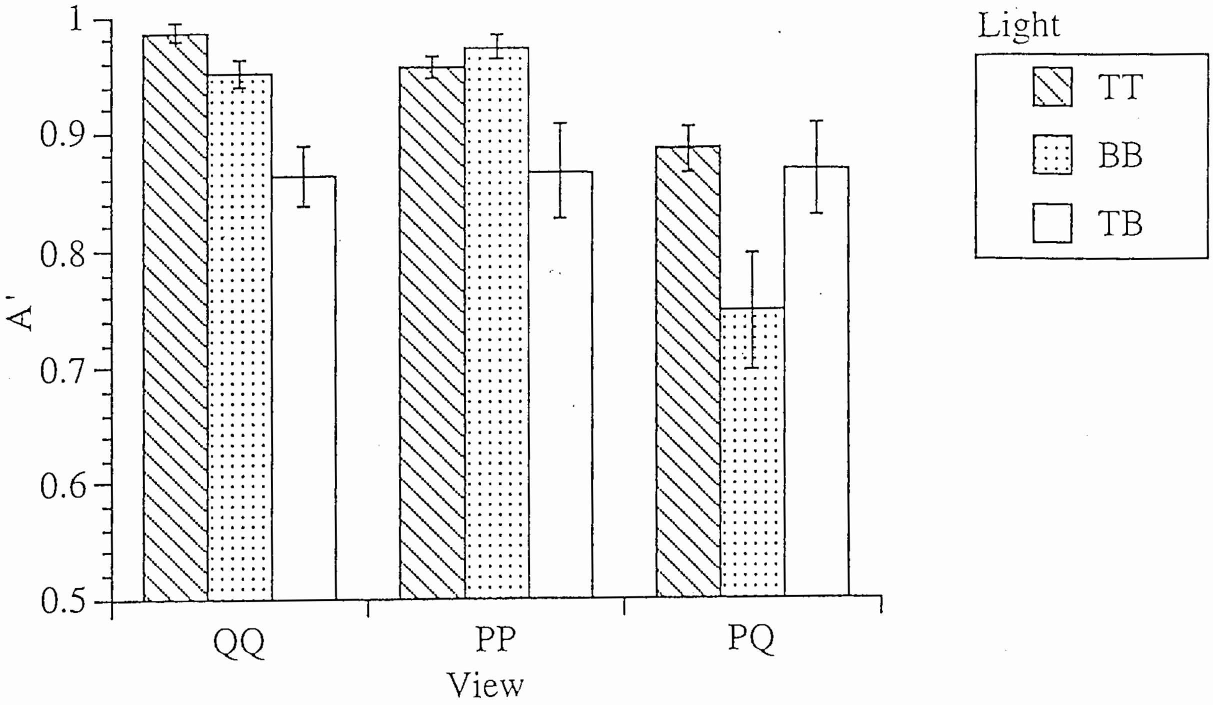

A'
1
0.9
0.8
0.7
0.6
0.5
QQ
PP
PQ
View
Light
TT
BB
TB