

# Detection of Mooney faces is robust to image asymmetries produced by illumination

Lindsay M. Peterson

School of Psychology,  
University of New South Wales, Sydney, Australia



Colin W. G. Clifford

School of Psychology,  
University of New South Wales, Sydney, Australia



Colin J. Palmer

Department of Psychology,  
National University of Singapore, Singapore  
School of Psychology,  
University of New South Wales, Sydney, Australia



Face detection relies on the visual features that are shared across different faces. An important component of the basic spatial configuration of a face is symmetry around the vertical midline. Although human faces are *structurally* symmetrical, they can be asymmetrical in an image due to the direction of lighting or the position of the face. In the experiments presented here, we examined how face detection from simple contrast patterns that occur across the face is affected by the image asymmetries associated with variations in the horizontal lighting direction. We presented observers with two-tone images of faces (Mooney faces) that isolated the unique pattern of contrast in the shading and shadows on a face, illuminated from a wide range of horizontal directions. In two experiments, we found that face detection is surprisingly robust to these lighting changes, with sensitivity in discriminating between face and non-face patterns reduced only at the most extreme lighting directions. This tolerance to changes in the horizontal lighting direction depended partly on the orientation of the face, vertical lighting direction, and contrast polarity. Our results provide insight into how contrast cues produced by shading and shadows occurring across the facial surface are utilized by the visual system to detect human faces.

and a mouth; an example can be seen [Figure 1A](#). The visual system is sensitive to this schematic of a face, with both adults and newborns displaying attentional biases toward these face-like patterns ([Farroni et al., 2005](#); [Tomalski, Csibra, & Johnson, 2009](#)). Single-cell electrophysiology has also identified cells in the macaque temporal cortex with tuning to simple face-like contrast patterns (e.g., [Kobatake & Tanaka, 1994](#); [Ohayon, Freiwald, & Tsao, 2012](#)). This suggests that the visual system may employ a template-matching approach to detect faces, where incoming visual signals are compared against a template that captures the basic spatial configuration of a face in a coarse pattern of contrast ([Tsao & Livingstone, 2008](#)). The existence of a basic face template is further supported by evidence that distorted faces can still be detected as long as the general facial shape is not disrupted ([Hershler & Hochstein, 2005](#); [Pongakkasira & Bindemann, 2015](#)). A key challenge to face detection is to deal with natural variability in the appearance of faces across contexts (e.g., due to changes in viewing angle, lighting conditions, face identity, or facial expression). Any template of visual features that underlies face detection is likely to be shaped by real-world exposure to faces, such that any biases in exposure (e.g., in the spatial orientation or lighting conditions in which faces are most commonly viewed) may constrain how invariant face detection is to such contextual factors. The current study explores the level of invariance that occurs in face detection performance when the presence of a face is conveyed only in a coarse pattern of contrast, which may have implications for understanding the visual features that contribute to face detection in human vision.

## Introduction

The ability of the human visual system to detect faces relies on extracting visual features that are common to different faces. These common features can be expressed in part by a simple stimulus such as a light-colored oval containing dark blobs that represent two eyes, a nose,

Citation: Peterson, L. M., Clifford, C. W. G., & Palmer, C. J. (2023). Detection of Mooney faces is robust to image asymmetries produced by illumination. *Journal of Vision*, 23(12):9, 1–20, <https://doi.org/10.1167/jov.23.12.9>.

<https://doi.org/10.1167/jov.23.12.9>

Received May 30, 2023; published October 26, 2023

ISSN 1534-7362 Copyright 2023 The Authors



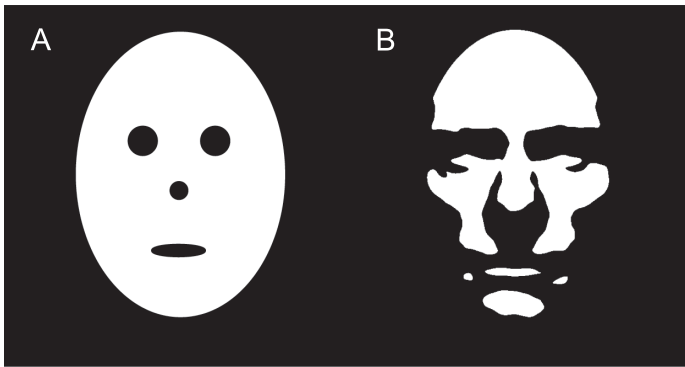


Figure 1. Simple stimuli that capture the basic structure of a human face can give a strong impression of a face. (A) Features that are shared across different faces—eyes, nose, and mouth—can be depicted by dark blobs within a light-colored oval. (B) Impoverished stimuli, such as a two-tone or Mooney face, can also produce a strong percept of a human face.

One underlying component of the visual structure of the human face is the pattern of shading and shadows that occurs across the facial surface. For example, shadows tend to form below the brows when a face is lit from above. This effect of illumination on the appearance of the face arises from the interaction between face geometry and the direction of lighting. Importantly, these patterns of shading and shadows across the face can provide cues that are sufficient to drive face detection. Consider the image of a face shown in Figure 1B. This two-tone image, or Mooney face (Mooney, 1957), was created by thresholding an image of a face with a uniform gray reflectance to isolate the coarse pattern of contrast produced by shading and shadows on the face (described further in Methods). This method of generating two-tone images ensures that the coarse luminance pattern in the images is produced by interactions between lighting direction and facial shape, rather than variations in surface reflectance. Although lots of information about the face is absent in the two-tone image, such as skin color and texture, as well as some facial features, the face is still readily perceived in the image.

Simple contrast cues produced by the pattern of illumination across a face, like those depicted in Figure 1B, can play an important role in face detection. In a recent study, Palmer, Goddard, and Clifford (2022) reported that two-tone faces can be discriminated from two-tone non-face objects when the two-tone stimuli were generated using the method described above that isolates contrast cues produced by the pattern of shading and shadows across the facial surface. The ability of observers to detect faces in these visual patterns was dependent on the vertical lighting direction, however, such that performance was facilitated for two-tone images that were consistent

with light arriving from above the face rather than from below the face. This result is consistent with the advantage of overhead lighting in face processing described in previous research. For example, newborns prefer to look at faces that are lit from above compared with below (Farroni et al., 2005), and faces are more easily recognized when illuminated from above (Enns & Shore, 1997; Hill & Bruce, 1996; Johnston, Hill, & Carman, 1992; Liu, Collin, Burton, & Chaudhuri, 1999). The results of Palmer et al. (2022) also support the notion of the visual system being somewhat tuned to the statistics of illumination in natural environments; as light tends to arrive from above our heads (e.g., light provided by the sun), the visual system may have developed a “prior” for this type of lighting that influences the interpretation of shape from shading (Mamassian & Goutcher, 2001; Ramachandran, 1988; Sun & Perona, 1998).

In sum, the interaction between facial shape and lighting creates a unique pattern of shading that captures the basic spatial configuration of a face, and these shading cues can be exploited for face detection. A key aspect of the spatial configuration of a face is symmetry around the vertical midline. However, although the *structure* of a face may be vertically symmetrical, an image of a face is often not symmetrical. This asymmetry can be due to a range of factors, such as the direction of lighting, facial expression, or the position of the face relative to the viewer (Adini, Moses, & Ullman, 1997; Favelle, Hill, & Claes, 2017). Face processing is sensitive to these factors; for example, the ability to match the identity of two sequentially presented faces is impaired when the faces are illuminated from different horizontal directions (Braje, 2003; Braje, Kersten, Tarr, & Troje, 1998). The effect of horizontal lighting direction on the symmetry of visual features is illustrated in Figure 2 and is particularly striking. We generated two-tone images that isolated the pattern of contrast across the face produced by shading and shadows, and we then calculated the symmetry of these images across a range of horizontal lighting directions (or light-source azimuths) and head rotations. Symmetry was given by the proportion of corresponding pixels that were white in both halves of the image. The images of faces rotated directly toward the observer and illuminated by a central light source are almost completely symmetrical. For all three of the head rotations, changing the azimuth of the light source introduced considerable asymmetries into the image. The portion of a face visible in the image reduces drastically as the light-source azimuth becomes more extreme (i.e., further away from 0°), with only a sliver of the face visible at the extreme azimuths.

To further elucidate how minimal visual cues are used to detect faces, in the current paper we examine how face detection based on the broad patterns of contrast on a face is affected by changes in the horizontal lighting

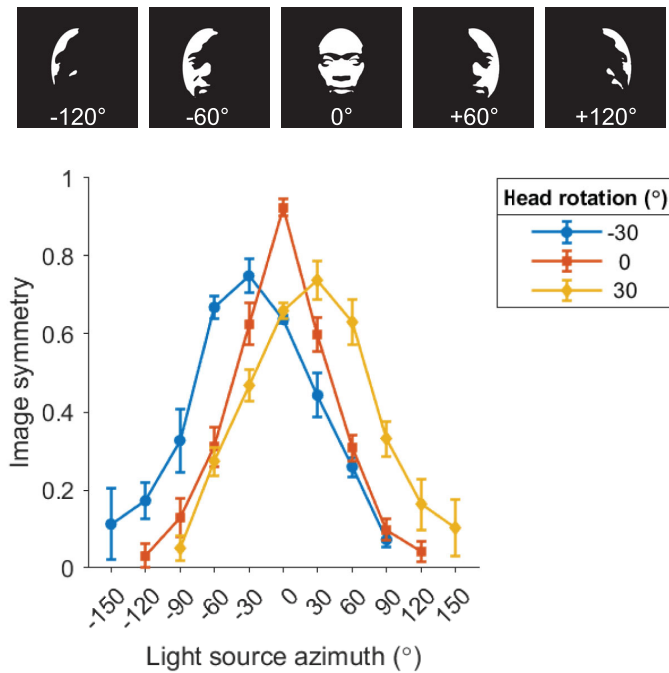


Figure 2. Image symmetry of two-tone faces across variations in horizontal lighting direction and head rotation. The horizontal axis is the azimuth of the light source illuminating the faces, where the azimuth is relative to the observer's perspective (rather than relative to the face in the image). The vertical axis is the symmetry of the two-tone images. Image symmetry was given by the proportion of white pixels shared across the left and right halves of the image (within an elliptical mask around the face, which is described further in the Stimuli section). Each marker represents the image symmetry averaged across face identities (six in total), and the error bars represent the  $\pm 1$  standard deviation. Some examples of the two-tone images for a face with a head rotation of  $0^\circ$  are shown at the top of the figure. Note that the two-tone images from which image symmetry was calculated were also the stimuli in Experiment 1, and the method used to create these stimuli is described below in the text.

direction, across the large image asymmetries associated with these changes. We used three-dimensional (3D) rendering to isolate visual cues produced by the pattern of shading and shadows on the face across a range of lighting directions. Participants were presented with an image of a two-tone face or non-face object and indicated whether they saw a human face in the image. In both experiments, we tested whether sensitivity at discriminating faces from non-faces is tuned to the horizontal lighting direction. One hypothesis is that face detection performance from simple contrast cues will be closely tied to the changes in symmetry that occur in the image across horizontal lighting directions. That is, discrimination sensitivity may peak for frontal lighting that is relative to the face in the image and reduce as the lighting direction becomes more averted and more

asymmetries are introduced into the image (following the narrow tuning of image symmetry across horizontal lighting direction depicted in Figure 2). This would be consistent with a template-matching approach to face detection that exploits the structural symmetry of the human face but suffers when extraneous factors (here, lighting direction) introduce horizontal asymmetries into the visual appearance of the face. Alternatively, discrimination sensitivity may be *broadly tuned* to the horizontal lighting direction, indicating that the visual system can tolerate the large image asymmetries caused by variations in horizontal illumination. This would be consistent with an approach to face detection that employs multiple templates to capture variation in the appearance of the human face under different lighting directions, for example, or a generic template that captures commonalities in facial appearance across lighting directions despite the considerable changes in the image that occur. A third hypothesis is that face detection performance may be best for *non-central* horizontal lighting directions, if non-central lighting aids face processing by creating shadowed areas on the face that facilitate shape from shading (e.g., see Chen, Chen, & Tyler, 2013).

In the first experiment, we manipulated the horizontal lighting direction, rotation of the head, and image orientation. By manipulating the rotation of the head, we could test whether discrimination sensitivity peaks for faces that are illuminated front-on relative to the *observer's perspective* or relative to the *orientation of the face in the image*. If the visual system has a prior expectation of central lighting (similar to the light-from-above prior), we would expect discrimination sensitivity to be best for frontal lighting that is relative to the observer's perspective. Conversely, sensitivity peaking for frontal lighting that is relative to the face in the image would be consistent with image symmetry operating as a cue for face detection, as the two-tone images are most symmetrical for this type of lighting. We also manipulated the image orientation such that the two-tone images were presented upright and upside down. Previous studies have indicated that the detection of two-tone faces is facilitated by the typical upright configuration of the face (Kanwisher, Tong, & Nakayama, 1998; Palmer et al., 2022), which suggests that faces are processed holistically by the visual system (Rossion, 2008; Rossion, 2009; Taubert, Apthorp, Aagten-Murphy, & Alais, 2011). As such, the inclusion of the image orientation manipulation allowed us to test whether tolerance to variations in the horizontal lighting direction when discriminating two-tone faces from non-faces is associated with holistic processing.

In the second experiment, we manipulate the vertical lighting direction such that faces could be illuminated from above or from below. Previous work has suggested that the detection of two-tone faces is facilitated when the image is consistent with light from above

(Brodski, Paasch, Helbling, & Wibrall, 2015; Palmer et al., 2022), consistent with the idea that the visual system has adapted to some extent for conditions in which light arrives more strongly from overhead. The inclusion of the vertical lighting manipulation allowed us to test whether the ability of observers to tolerate variations in the horizontal lighting direction depends on sensory patterns that are familiar to us (i.e., facial shading patterns that are generated by overhead lighting). We also tested whether tolerance to variations in the horizontal lighting direction depends on natural contrast polarity. That is, when discriminating faces from non-faces based only on the shading information available in the two-tone images, is it critical that the parts of the face we expect to receive stronger illumination are white and those we expect to be in shadow are black? Similar to the vertical lighting manipulation, we would expect to see an advantage for natural contrast polarity images (compared with reversed polarity images) if detection is facilitated by familiar sensory patterns.

## Experiment 1

### Methods

#### Participants

Thirty-nine participants (32 female, six male, one preferred not to say; median age, 19 years) were recruited from a database of undergraduate students enrolled in a first-year psychology course at the University of New South Wales (UNSW), Sydney. One further participant completed the experiment but was excluded, as described in the Analysis section. Participants were required to have normal or corrected-to-normal vision to complete the experiment. Recruitment and experimental procedures were approved by the Human Research Ethics Advisory Panel C at the School of Psychology, UNSW Sydney, and in adherence with the tenets of the Declaration of Helsinki.

The key effect of interest was that of the horizontal lighting direction (or light-source azimuth) on discrimination sensitivity. As this specific effect has not been examined previously, we relied on the results from Palmer et al. (2022), who reported a large effect of the vertical lighting direction on the discrimination of two-tone faces and non-faces ( $\eta_p^2 = 0.85$  in Experiment 1) to inform the estimated effect size for the current experiment. A power analysis (G\*Power 3.1.9.7) (Faul, Erdfelder, Lang, & Buchner, 2007) indicated that fewer than five participants would be necessary to detect a main effect of this magnitude in a repeated-measures analysis of variance (ANOVA) with 95% power and  $\alpha$

$= 0.05$ . However, we aimed to collect approximately 35 to 40 participants such that there was sufficient power to detect other, potentially smaller effects. An additional power analysis indicated that a sample size of 40 participants would be sufficient to detect an effect of  $\eta_p^2 = 0.07$  with 95% power.

#### Apparatus

This was an online experiment that was implemented using jsPsych 6.3.1 (De Leeuw, 2015) and JATOS 3.3.6 (Lange, Kühn, & Filevich, 2015). Participants were required to complete the experiment on a computer desktop or laptop with a minimum screen resolution of 800 pixels  $\times$  600 pixels. The stimuli were created using Blender 2.93.1 (The Blender Foundation, Amsterdam, The Netherlands) and MATLAB R2021a (MathWorks, Natick, MA). MATLAB was also used to conduct the data analysis.

#### Stimuli

The stimuli were black-and-white images of faces and non-face objects that were created using 3D models of human heads and objects. To create images of realistic faces, we used high-resolution 3D models of six human heads that were generated from scans of real individuals. These 3D models were acquired from the Ten24 3D scan store (<https://www.3dscanstore.com/>). Each model was placed into a rendering environment in Blender; this allowed us to produce images of each model with different head positions and under different lighting conditions. The rendering environment contained the 3D models, a single light source, and a camera. The light source was a 60 cm square plane that emitted light in the direction of the models. The light source was positioned 1.5 m from the model and had an elevation of  $+45^\circ$  (i.e., all of the models were lit from above). The light source could have an azimuth of  $-120^\circ$ ,  $-90^\circ$ ,  $-60^\circ$ ,  $-30^\circ$ ,  $0^\circ$ ,  $+30^\circ$ ,  $+60^\circ$ ,  $+90^\circ$ , or  $+120^\circ$  relative to the rotation of the head, with the model being illuminated from front-on at  $0^\circ$ . Possible head rotations were  $-30^\circ$ ,  $0^\circ$ , and  $+30^\circ$  around the vertical axis, where at  $0^\circ$  the face is oriented directly toward the camera. For example, a model that had a head rotation of  $-30^\circ$  could be illuminated by the light source with an azimuth of  $-150^\circ$ ,  $-90^\circ$ ,  $-60^\circ$ ,  $-30^\circ$ ,  $0^\circ$ ,  $+30^\circ$ ,  $+60^\circ$ , or  $+90^\circ$  (in absolute terms, or relative to the observer's perspective). Examples of the human models can be seen in the top row of Figure 3.

We also created six non-face objects in Blender. We aimed to create objects that shared some structural characteristics with human faces (including their outline, which is described further below) such that participants would need to rely on the internal structure of the models to discriminate the faces and non-faces. The non-face objects were ellipsoids whose dimensions

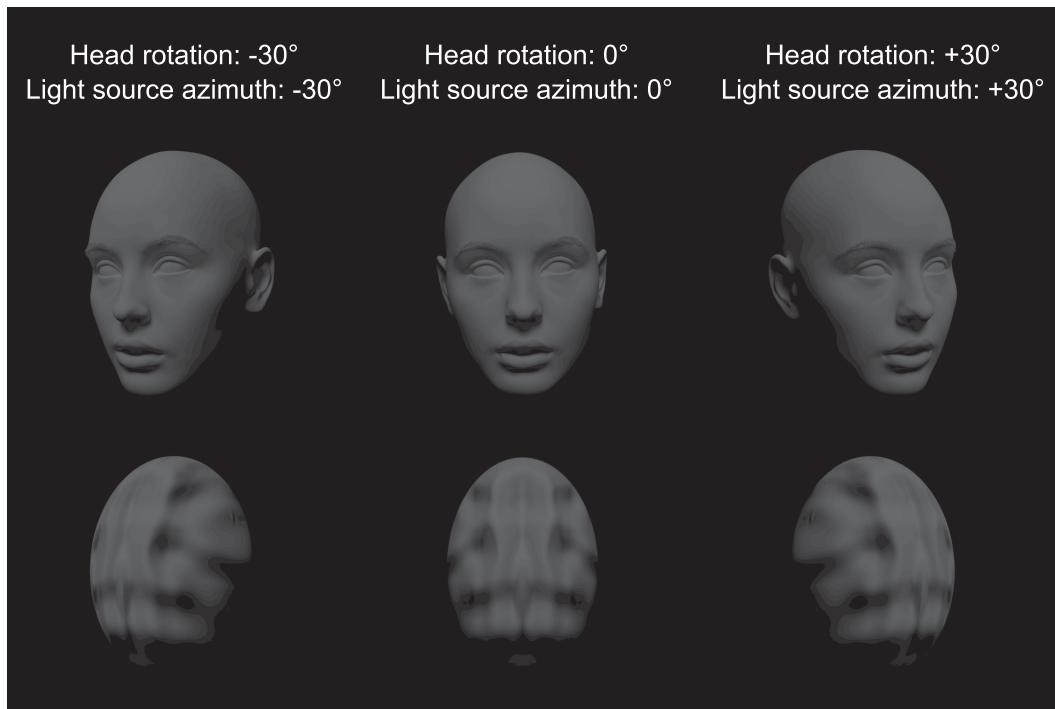


Figure 3. Examples of the grayscale images of human and non-human models that were produced by the rendering process described in the [Experiment 1 Stimuli](#) section. Each model is illuminated from front-on relative to the rotation of the model. Note that the light-source azimuth is given in absolute terms.

were selected to be similar to those of the human models. The surface normals were displaced using a Perlin noise texture that created smooth curvature across the surface of the objects, and this texture was mirrored around the vertical axis. Examples of the non-human models can be seen in the bottom row of [Figure 3](#).

The rendering process was controlled by custom scripts in Python 3.9.2 and performed with Cycles, a physically based rendering engine that simulates the path of light rays in 3D space to produce realistic images. All of the models were rendered with a uniform gray Lambertian reflectance. The camera was positioned 2 m from the models and had a resolution of 945 pixels  $\times$  945 pixels. The rendering process produced 324 grayscale images in total: 12 identities (six face, six non-face)  $\times$  9 light source azimuths  $\times$  3 rotations. The images were encoded in a linear color space with a 32-bit depth per color channel.

The rendered images were then processed in MATLAB to produce the two-tone images that were presented to participants during the experiment. First, the images were low-pass filtered with a two-dimensional (2D) Gaussian kernel ( $\sigma = 5.0$  pixels). Each image was then cropped with an elliptical mask (width of 204 pixels and height of 306 pixels) to remove the external features of the faces and non-faces, such that discrimination between faces and non-faces had

to be made on the basis of internal surface curvature. The cropped grayscale images were then converted to two-tone images using the method from [Otsu \(1979\)](#), which determines the optimal threshold for each image based on the histogram of gray levels in the image. The optimal threshold is selected such that the within-class variance of the pixels above and below the threshold is minimized (where the two classes are above-threshold and below-threshold pixels). For each image, pixels that had an intensity value above the selected threshold were changed to white, and pixels that had an intensity value below the threshold were changed to black. Finally, image orientation was manipulated by flipping each two-tone image upside down. This image processing resulted in a stimulus database of 648 two-tone images: 12 identities (six face, six non-face)  $\times$  9 light-source azimuths  $\times$  3 rotations  $\times$  2 image orientations. Examples of the two-tone faces and non-faces can be seen in [Figures 4](#) and [5](#), respectively.

### Design and procedure

Participants completed a face detection task in which they were asked to discriminate human faces from non-faces. The experiment had a within-subjects design with factors of head rotation ( $-30^\circ$ ,  $0^\circ$ , and  $+30^\circ$ ), light-source azimuth ( $-120^\circ$ ,  $-90^\circ$ ,  $-60^\circ$ ,  $-30^\circ$ ,  $0^\circ$ ,  $+30^\circ$ ,  $+60^\circ$ ,  $+90^\circ$ , and  $+120^\circ$ ), and image orientation

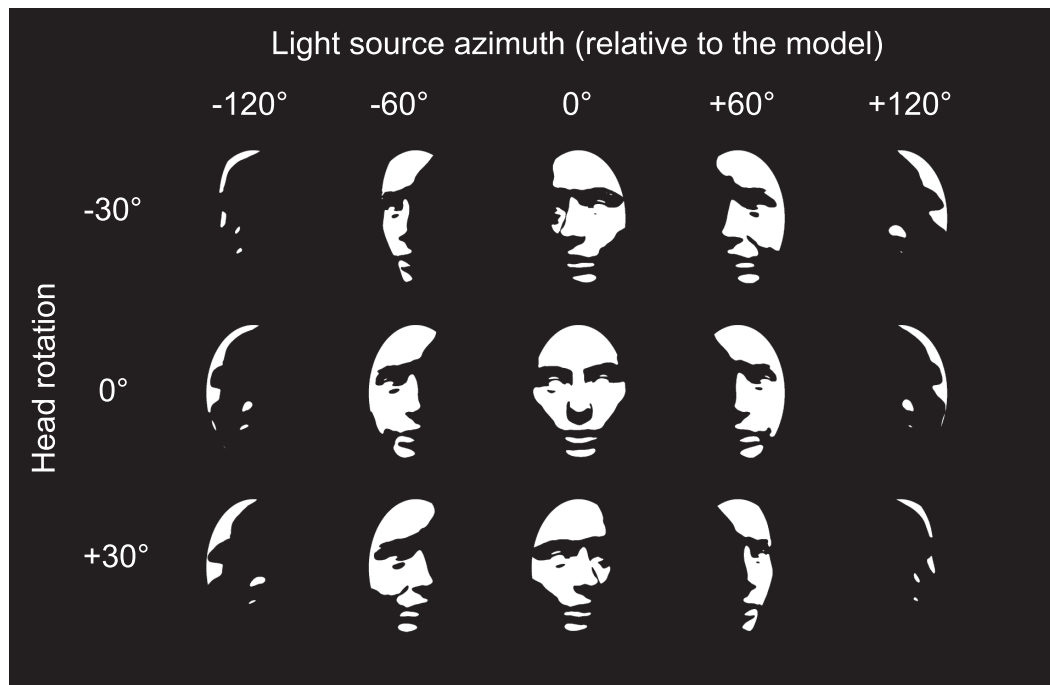


Figure 4. Examples of the two-tone images for one of the models from [Experiment 1](#). The light-source azimuths are relative to the rotation of the head; for example, the absolute azimuths for the  $-30^\circ$  rotation condition are  $-150^\circ$ ,  $-90^\circ$ ,  $-30^\circ$ ,  $+30^\circ$ , and  $+90^\circ$ . The grayscale images used to create the two-tone images shown in the middle column can be seen in the top row of [Figure 3](#).




Figure 5. Examples of the two-tone images for one of the non-face objects from [Experiment 1](#). The light-source azimuths are relative to the rotation of the object; for example, the absolute azimuths for the  $-30^\circ$  rotation condition are  $-150^\circ$ ,  $-90^\circ$ ,  $-30^\circ$ ,  $+30^\circ$ , and  $+90^\circ$ . The grayscale images used to create the two-tone images shown in the middle column can be seen in the bottom row of [Figure 3](#).

(whether the image was presented upright or flipped upside-down). Here, the light-source azimuth is expressed relative to the rotation of the head, where an azimuth of  $0^\circ$  corresponds to a face (or non-face) being lit from directly front-on. The experiment consisted of 648 trials across four runs, with 162 trials per run. There was a self-paced rest break between each run. In addition to these trials, participants also completed 20 practice trials prior to beginning the experiment.

A calibration was performed at the beginning of the experiment to ensure that the stimuli were presented at approximately the same size for all participants. This calibration process involved participants holding a credit card (or a card with the dimensions  $8.6 \text{ cm} \times 5.4 \text{ cm}$ ) up to their computer screen and adjusting the size of a box on the screen until it was the same size as the card. On each trial, participants were first presented with a fixation cross for 1000 ms. A randomly selected stimulus was then presented for 100 ms. The image size was randomly selected to be 30%, 40%, or 50% of the actual image size on each trial, meaning that the size of the face varied between 3.2 and 5.4 cm in width and 4.8 and 8.1 cm in height throughout the experiment. The size of the image was varied to increase task difficulty and encourage participant concentration throughout the experiment. Following the presentation of the image, participants viewed the response screen that contained the prompt: “Was it an image of a human face? Press K for YES. Press L for NO.” The prompt remained on the screen until the participant responded, with the key correspondence randomized across participants.

### Analysis

Upon completion of the experiment, there were 648 data points in total per participant, with 12 data points (six face trials and six non-face trials) for each experimental condition. If a participant indicated that they saw an image of a face on a face trial, this was considered to be a hit. If a participant indicated that they saw an image of a face on a non-face trial, this was considered to be a false alarm. The hit rate and false alarm rate were then used to calculate each participant’s sensitivity in discriminating faces from non-faces:

$$d' = z(P_H) - z(P_{FA})$$

where the  $z$ -score for the false alarm rate is subtracted from the  $z$ -score for the hit rate (Kingdom & Prins, 2010). Note that 0.5 was added to hit and false alarm counts of 0 and subtracted when the rate was 100%. We also calculated the response bias or decision criterion ( $C$ ):

$$C = -0.5(z(P_H) + z(P_{FA}))$$

where an unbiased participant would have a criterion of zero, and negative values correspond to a participant

having a bias toward indicating that they saw an image of a face on a given trial (Stanislaw & Todorov, 1999). One participant was excluded from the analysis because their overall  $d'$  (i.e., their sensitivity across all trials) was below zero. The analysis described here was conducted with  $n = 39$ .

For each participant, we calculated the centroid of the distribution of  $d'$  values across light-source azimuths. The centroid is a weighted mean that represents the center point of the distribution and can be interpreted as the light-source azimuth at which the discrimination of faces from non-faces is best. As such, the centroid of the distribution of  $d'$  for each head rotation condition would be informative of whether frontal lighting that is relative to the face in the image or the observer is important for face detection. The centroid for each head rotation and spatial inversion condition was calculated from

$$\text{Centroid} = \frac{\sum_{i=1}^n (d_i \cdot a_i)}{\sum |d_i|}$$

where  $a_i$  is a given light-source azimuth condition, and  $d_i$  is the discrimination sensitivity for that azimuth condition. Absolute  $d'$  values were used in the denominator, as a few participants had negative  $d'$  for some of the experimental conditions. We calculated the centroids for both the relative and absolute light-source azimuths. For the centroid calculation for the absolute azimuths, the range of azimuths was restricted to  $-90^\circ$  to  $+90^\circ$  such that the centroids for each head rotation condition were calculated from the same range of azimuths.

### Results

Overall, participants were able to discriminate faces from non-faces using only the shading cues that were available in the two-tone images. As can be seen in Figure 6, participants had close to ceiling performance for images that were spatially upright. This ability to detect faces in the two-tone images was largely robust to changes in the horizontal lighting direction. We consider this to be a surprising result given the drastic image changes associated with variations in light-source azimuth (see Figures 2 and 4, for example).

The mean discrimination sensitivity for each experimental condition is shown in Figure 6. For the upright images, participants’ discrimination sensitivity only reduced for the extreme light-source azimuths where only a small portion of the face was visible. A repeated-measures ANOVA indicated that there was a significant main effect of light-source azimuth on discrimination sensitivity,  $F(8, 304) = 138.55, p < 0.001, \eta_p^2 = 0.79$ , and this main effect seems to be driven by the reduced discrimination sensitivity at the extreme azimuths. However, varying the light-source azimuth

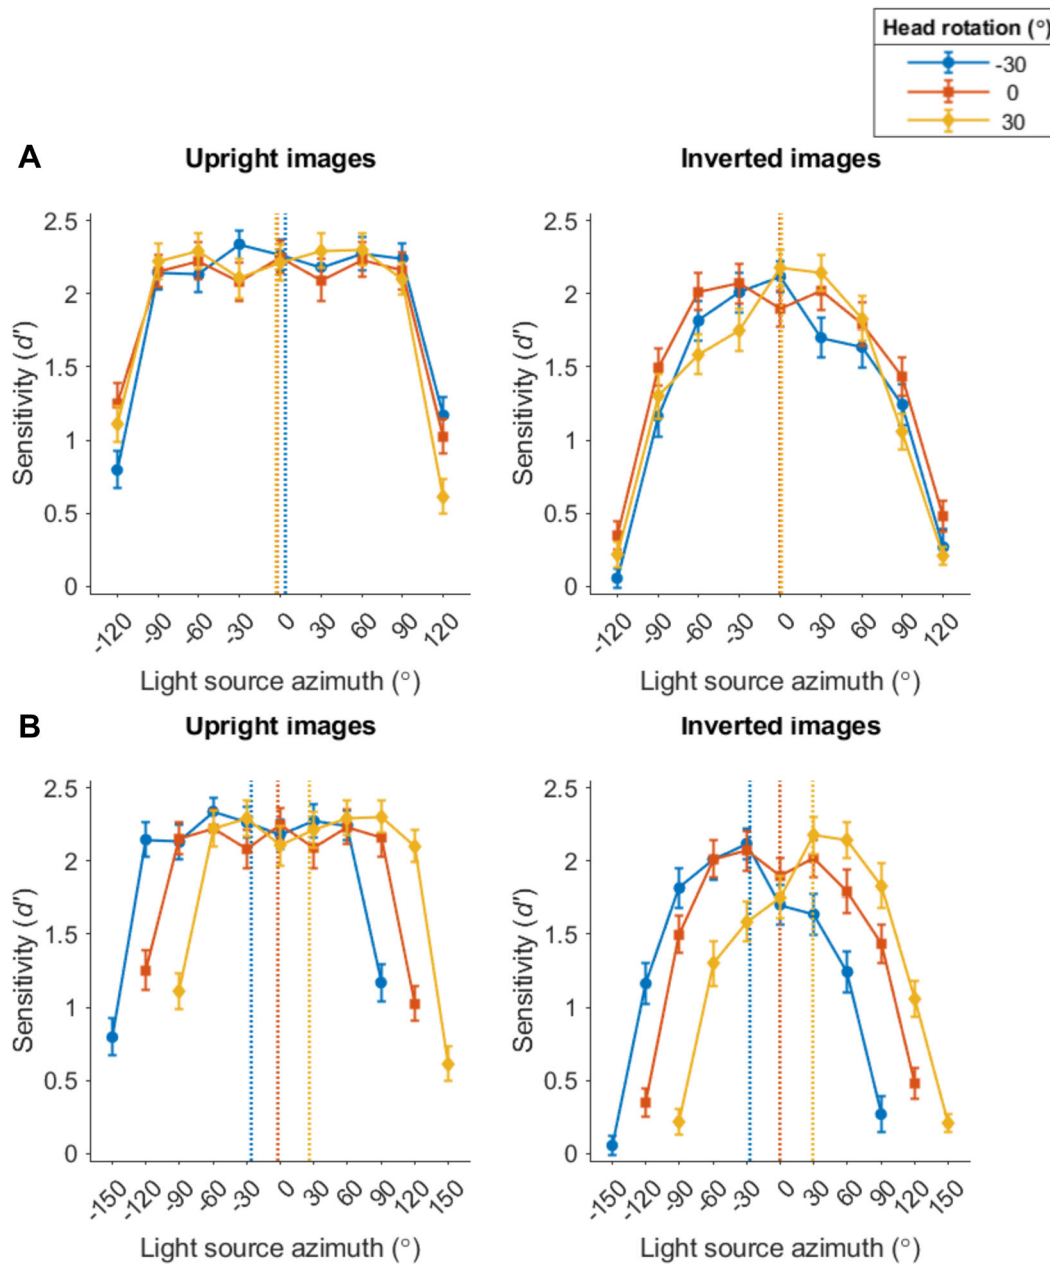


Figure 6. Mean discrimination sensitivity for each condition in [Experiment 1](#). The horizontal axis is the azimuth of the light source and the vertical axis is the mean  $d'$  across participants. **(A)** Discrimination sensitivity is plotted with the light-source azimuth expressed relative to head rotation. **(B)** Discrimination sensitivity is plotted with the absolute light-source azimuths. The dashed lines represent the mean centroid for each head rotation condition. The error bars attached to each marker represent the  $\pm 1$  standard error of the mean. Note that the centroids for the absolute azimuths depicted in panel B were calculated from the complete distribution of  $d'$  values for each condition rather than the  $-90^\circ$  to  $+90^\circ$  range used in the analysis. The proportion of “face” responses for the face and non-face trials for each condition is depicted in Supplementary Figure S1 in the Supplementary Materials.

had a greater effect on discrimination sensitivity for inverted images compared with upright images. This was reflected in a significant interaction between light-source azimuth and spatial inversion,  $F(8, 304) = 14.49$ ,  $p < 0.001$ ,  $\eta_p^2 = 0.28$ , and suggests that tolerance to changes in the horizontal illumination of the face relies partly on viewing the face in its most

typical upright configuration. Unsurprisingly, spatially inverting the images also led to worse discrimination sensitivity overall, reflected in a main effect of spatial inversion,  $F(1, 38) = 132.18$ ,  $p < 0.001$ ,  $\eta_p^2 = 0.77$ .

In addition to examining how face detection is affected by changes in the horizontal lighting direction, we were interested in whether discrimination sensitivity



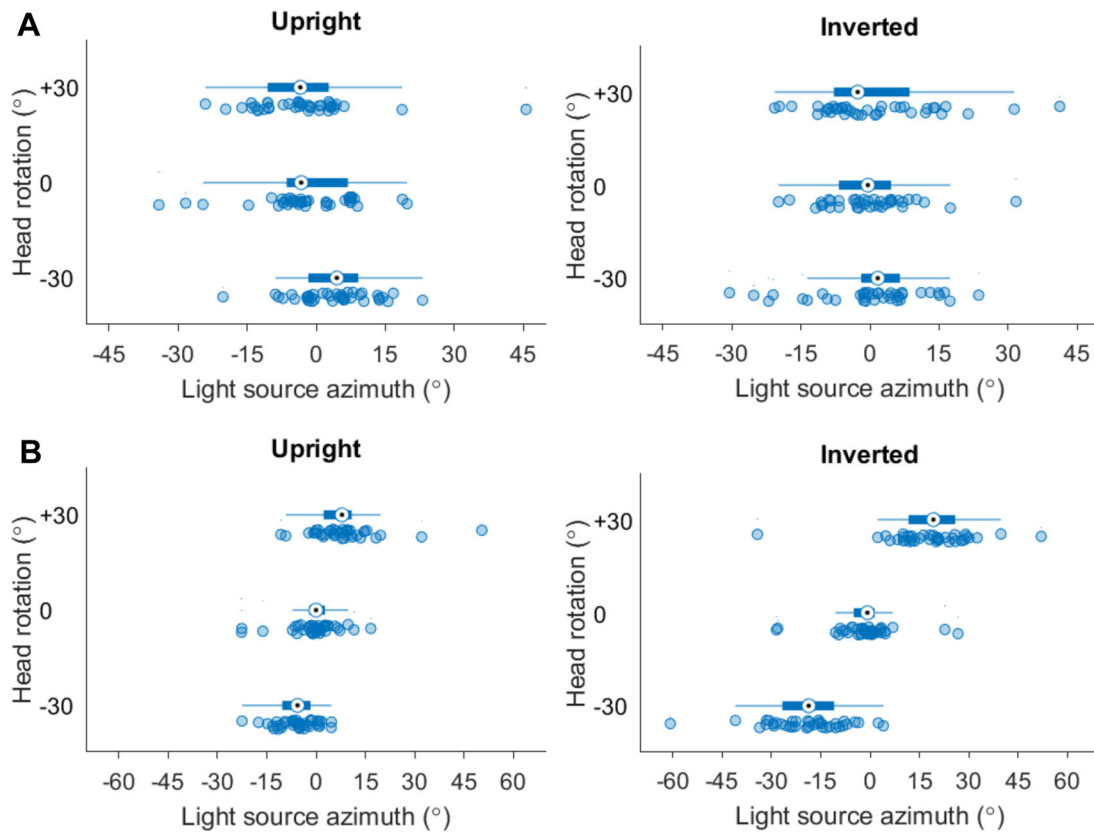


Figure 7. Summary of the centroids for each experimental condition in Experiment 1. (A) The centroids for relative light-source azimuths. (B) The centroids for the absolute light-source azimuths. The marker within each box plot represents the median, the left and right edges of the box represent the 25th and 75th percentiles, and the whiskers depict the range of the centroids (ignoring outliers). The individual markers below each boxplot represent the centroids for each participant. Note that the centroids for the absolute light-source azimuths depicted in this figure (panel B) were calculated from the  $-90^{\circ}$  to  $+90^{\circ}$  range of azimuths (see text for details).

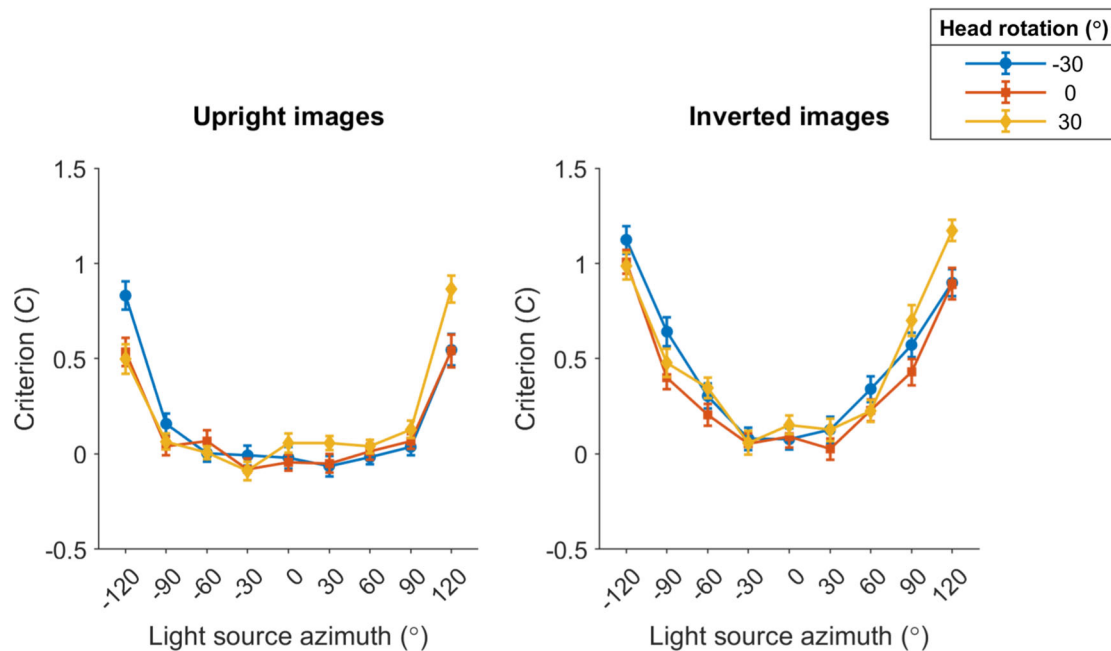


Figure 8. Mean response bias for each condition in Experiment 1. This horizontal axis is the light-source azimuth and the vertical axis is the mean criterion (C). The error bars attached to each marker represents the  $\pm 1$  standard error of the mean across participants.

is best for front-on illumination that is relative to the face in the image or relative to the participant's perspective. As can be seen in Figure 6, the distributions of  $d'$  for each head rotation condition overlap within one another for the relative light-source azimuths (Figure 6A) and are separated for the absolute azimuths (Figure 6B), suggesting that sensitivity is best for front-on illumination that is relative to the face. To examine this statistically, we compared the centroids for each head rotation condition. If front-on illumination relative to the face leads to better sensitivity, we would expect the centroids to be close to zero for the relative light-source azimuths. We would also expect the centroids to shift away from zero for the absolute azimuths: a leftward shift for the  $-30^\circ$  head rotation condition and a rightward shift for the  $+30^\circ$  condition. Conversely, if sensitivity was best for front-on illumination relative to the participant's perspective, we would expect the centroids to be close to zero for the absolute light-source azimuths and shifted for the relative azimuths (rightward for the  $-30^\circ$  head rotation condition and leftward for the  $+30^\circ$  condition).

The centroids for both the relative (Figure 6A) and absolute (Figure 6B) light-source azimuths suggest that participants were best at discriminating faces from non-faces for front-on illumination *relative to the face in an image*. When the light-source azimuth was coded relative to the face, the half-difference between the centroids for the  $-30^\circ$  and  $+30^\circ$  head rotation condition was not significantly different from zero for the upright condition,  $t(39) = 1.10$ ,  $p = 0.28$ , Cohen's  $d = 0.17$ , and the inverted condition,  $t(39) = -0.09$ ,  $p = 0.93$ , Cohen's  $d = -0.01$ . For the absolute light-source azimuths, the centroid half-difference was significantly different from zero for both the upright condition,  $t(39) = -5.92$ ,  $p < 0.001$ , Cohen's  $d = -0.94$ , and the inverted condition,  $t(39) = -12.48$ ,  $p < 0.001$ , Cohen's  $d = -1.97$ . The centroid half-differences for the relative and absolute light-source azimuths are significantly different from each other for both spatially upright images,  $t(39) = -10.40$ ,  $p < 0.001$ , Cohen's  $d = -1.64$ , and inverted images,  $t(39) = -14.30$ ,  $p < 0.001$ , Cohen's  $d = -2.26$ . Thus, the horizontal direction of lighting relative to the face appeared more important for face detection than the direction of lighting relative to the viewer. The centroids for all participants can be seen in Figure 7.

To ensure the robustness of our conclusions, we repeated the centroid analysis on the pooled data from all participants, with the centroids calculated without taking absolute  $d'$  values in the denominator. Bootstrapped 95% confidence intervals (CIs) for the centroid half-differences were calculated from the pooled data by resampling the subject pool with replacement 10,000 times and repeating the centroid calculation for each iteration. The bootstrapped 95% CIs indicate that the centroid half-differences were significantly different from zero for the absolute

light-source azimuths (upright images: 95% CI,  $-9.43$  to  $-4.96$ ; inverted images: 95% CI,  $-21.56$  to  $-1.57$ ) and not significantly different from zero for the relative light-source azimuths (upright images: 95% CI,  $-0.97$  to  $4.43$ ; inverted images: 95% CI,  $-6.73$  to  $1.70$ ). This is consistent with the results of the main centroid analysis.

There was also a significant main effect of head rotation,  $F(2, 76) = 6.14$ ,  $p = 0.003$ ,  $\eta_p^2 = 0.14$ , where sensitivity was slightly better for the  $0^\circ$  head rotation condition, and this effect interacted with light-source azimuth,  $F(16, 608) = 4.41$ ,  $p < 0.001$ ,  $\eta_p^2 = 0.10$ . There was a small interaction between head rotation and spatial inversion,  $F(2, 76) = 5.69$ ,  $p = 0.005$ ,  $\eta_p^2 = 0.13$ , where sensitivity for the  $0^\circ$  head rotation condition was slightly less affected by spatial inversion compared with the other two rotations. There was also a significant three-way interaction among light-source azimuth, head rotation, and spatial inversion,  $F(16, 608) = 2.27$ ,  $p = 0.003$ ,  $\eta_p^2 = 0.06$ .

The mean criterion or response bias across participants for each condition is shown in Figure 8. A repeated-measures ANOVA indicated that there was a significant effect of head rotation on response bias,  $F(2, 78) = 20.96$ ,  $p < 0.001$ ,  $\eta_p^2 = 0.35$ , where participants were slightly less biased toward a non-face response for the  $0^\circ$  rotation condition. Response bias was also affected by the horizontal lighting direction,  $F(8, 312) = 136.11$ ,  $p < 0.001$ ,  $\eta_p^2 = 0.78$ , with participants having a greater bias toward responding "non-face" at the extreme light-source azimuths. There was a significant interaction between head rotation and light-source azimuth,  $F(16, 624) = 7.38$ ,  $p < 0.001$ ,  $\eta_p^2 = 0.16$ ; the increase in criterion for the extreme light-source azimuth conditions depended on the rotation condition, where the bias toward "non-face" responses was greater for extreme light-source azimuths. Participants were also more biased toward responding "non-face" for spatially inverted faces compared with spatially upright faces,  $F(1, 39) = 120.51$ ,  $p < 0.001$ ,  $\eta_p^2 = 0.77$ . The interaction between light-source azimuth and image orientation was significant,  $F(8, 312) = 12.95$ ,  $p < 0.001$ ,  $\eta_p^2 = 0.25$ , as well as the interaction among all three factors,  $F(16, 624) = 1.67$ ,  $p = 0.047$ ,  $\eta_p^2 = 0.04$ .

## Experiment 2

In Experiment 1, we found that face detection in contrast patterns is robust to considerable changes in the horizontal illumination of the face, particularly when viewing faces in their typical upright configuration, although participants did show a greater bias toward responding that they did not see a face in the image as the horizontal lighting became more extreme. In Experiment 2, we measured the effects of light-source

azimuth, light-source elevation, and contrast polarity on the discrimination of faces from non-faces. We were interested in how the vertical lighting direction would influence tolerance to horizontal lighting changes, given that previous studies have reported that two-tone faces illuminated from above are more easily detected than those illuminated from below (Brodski et al., 2015; Palmer et al., 2022). This advantage for faces that are lit from above is consistent with the greater familiarity of the visual system with overhead lighting (Mamassian & Goutcher, 2001; Ramachandran, 1988; Sun & Perona, 1998). In [Experiment 2](#), we tested whether tolerance to changes in the horizontal lighting direction depend on familiar lighting conditions by presenting participants with two-tone faces and non-faces that were illuminated from above and below. Because natural contrast polarity is important for the detection of two-tone faces (Farroni et al., 2005; Palmer et al., 2022; Tomalski et al., 2009), we were also interested in whether the tolerance to horizontal lighting changes we observed in [Experiment 1](#) is dependent on the contrast polarity of the two-tone face. To test this, participants in [Experiment 2](#) were presented with two-tone faces and non-faces that had natural contrast polarity (as in [Experiment 1](#)) and reversed contrast polarity.

## Methods

### Participants

Forty-one participants (27 female, 13 male, one non-binary; median age, 19 years) completed the experiment. An additional participant completed the experiment but their data were excluded from the analysis (as described below). The recruitment procedures were as described for the previous experiment. As with [Experiment 1](#), the horizontal lighting direction was the key effect of interest, and only a small sample size ( $n < 5$ ) would be necessary to detect this effect in a repeated-measures ANOVA with 95% power ( $\alpha = 0.05$  and  $\eta_p^2 = 0.79$ , based on the results of the previous experiment). However, our target sample size was 35 to 40 participants to facilitate comparisons across experiments and allow us to detect smaller effects. As reported for the previous experiment, a power analysis indicated that a sample size of 40 participants would be sufficient to detect an effect of  $\eta_p^2 = 0.07$  with 95% power.

### Apparatus

The apparatus was as reported for the previous experiment.

### Stimuli

The stimuli consisted of two-tone images of the six faces and six non-faces that were used in the previous experiment. The key differences in the stimuli for the current experiment was the removal of the head rotation and spatial inversion manipulations and the addition of the light-source elevation and contrast polarity manipulations. All of the models were oriented at  $0^\circ$  on the vertical axis (i.e., facing toward the camera) and could be illuminated from  $+45^\circ$  (i.e., lit from above) as well as  $-45^\circ$  (i.e., lit from below). The range of light-source azimuths was as described for the previous experiment. The rendering of the grayscale images in Blender and processing of those images in MATLAB to create the two-tone images were as described for the previous experiment. The contrast polarity of the two-tone images was reversed by switching the white regions of the image to black and vice versa. There were 432 two-tone images in total: 9 light-source azimuths  $\times$  2 light-source elevations  $\times$  2 contrast polarities  $\times$  12 identities (six face, six non-face). Examples of the two-tone images presented during the experiment can be seen in [Figure 9](#).

### Design and procedure

The experiment had a within-subjects design with factors of light-source azimuth ( $-120^\circ$ ,  $-90^\circ$ ,  $-60^\circ$ ,  $-30^\circ$ ,  $0^\circ$ ,  $+30^\circ$ ,  $+60^\circ$ ,  $+90^\circ$ ,  $+120^\circ$ ), light-source elevation ( $-45^\circ$ ,  $+45^\circ$ ), and contrast polarity (natural polarity, reversed polarity). The experiment had 432 trials across four runs, with a self-paced rest break between each run. To keep the background constant within a run, two of the runs consisted of only the natural contrast polarity images and the other two runs contained the reversed polarity images. The order of the runs was randomized across participants. All other experimental procedures were as described for the previous experiment.

### Analysis

There were 432 data points for each participant, with 12 data points (six face trials and six non-face trials) for each combination of light-source azimuth, light-source elevation, and contrast polarity. Similar to the previous experiment, we applied signal detection theory to calculate each participant's discrimination sensitivity. One participant was excluded from the data analysis because their overall discrimination sensitivity was less than zero; the analysis described here was conducted with  $n = 41$ .




Figure 9. In [Experiment 2](#), we manipulated the light-source azimuth and elevation and the contrast polarity of the images. In this figure, one of the human faces is shown illuminated by a light source with an elevation of  $+45^\circ$  (i.e., lit from above) and  $-45^\circ$  (i.e., lit from below) and an azimuth of  $-60^\circ$ ,  $0^\circ$ , and  $+60^\circ$  for both natural and reversed contrast polarity.

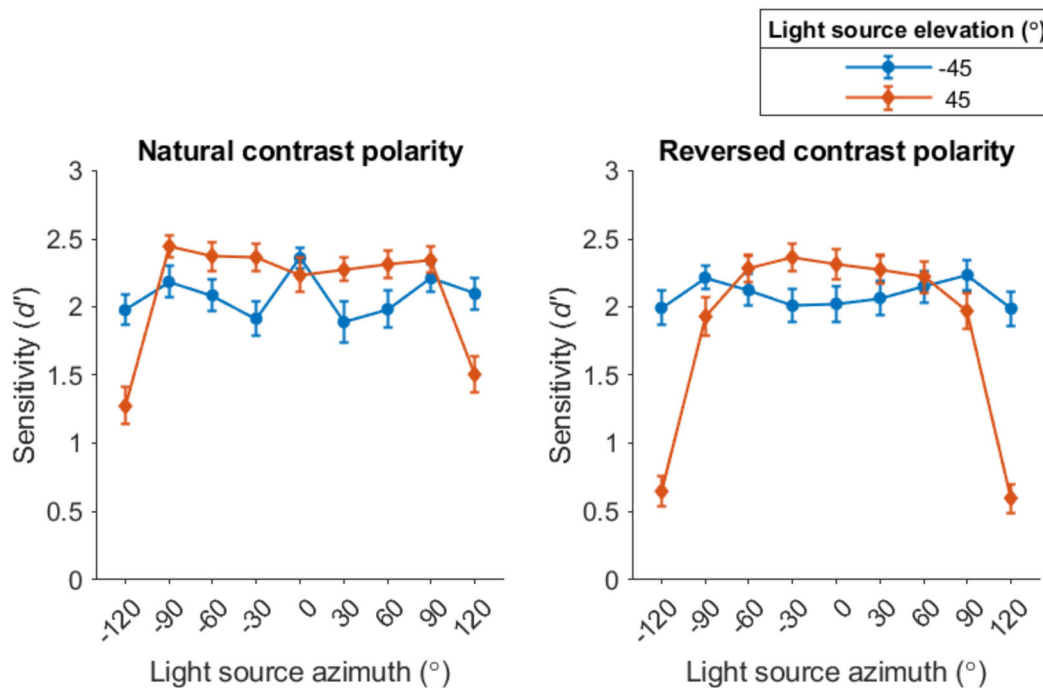


Figure 10. Mean discrimination sensitivity for each condition in [Experiment 2](#). The horizontal axis is the light-source azimuth, and the vertical axis is the mean discrimination sensitivity across participants. The error bars attached to each marker represent the  $\pm 1$  standard error of the mean. The proportion of “face” responses for the face and non-face trials for each condition can be seen in [Supplementary Figure S2](#) in the [Supplementary Materials](#).

## Results

Consistent with the results from the previous experiment, participants were able to utilize the broad contrast patterns on a face to discriminate faces from non-faces across most of the horizontal lighting directions. As can be seen in [Figure 10](#), discrimination

sensitivity remained quite high despite changes in the horizontal lighting direction, although sensitivity was reduced for the extreme lighting directions. A repeated-measures ANOVA indicated that there was a significant main effect of light-source azimuth,  $F(8, 320) = 43.28$ ,  $p < 0.001$ ,  $\eta_p^2 = 0.52$ , and light-source elevation,  $F(1, 40) = 7.18$ ,  $p = 0.01$ ,  $\eta_p^2 = 0.15$ . There




Figure 11. Mean response bias for each condition in Experiment 2. This horizontal axis is the light-source azimuth and the vertical axis is the mean criterion ( $C$ ). The error bars attached to each marker represents the  $\pm 1$  standard error of the mean across participants.

was also a significant interaction between light-source azimuth and light-source elevation,  $F(8, 320) = 44.25$ ,  $p < 0.001$ ,  $\eta_p^2 = 0.53$ , and this interaction varied across the contrast polarity conditions: three-way interaction  $F(8, 320) = 5.61$ ,  $p < 0.001$ ,  $\eta_p^2 = 0.12$  (see Figure 10).

There are two main features of the results that appear to explain the interaction between light-source azimuth and light-source elevation. First, sensitivity was reduced for the extreme horizontal lighting directions (i.e.,  $\pm 120^\circ$ ) for the faces lit from above but not for faces lit from below (Figure 10). In other words, there was greater tolerance for extreme angles of horizontal illumination for faces lit from below. Second, in the natural contrast polarity images, sensitivity was comparable for faces illuminated from above and below for the  $0^\circ$  azimuth condition (central lighting), but sensitivity was reduced for the faces lit from below as the light-source azimuth moved away from  $0^\circ$ . Post hoc tests indicated that sensitivity was not significantly different for faces illuminated from above and below with an azimuth of  $0^\circ$ ,  $t(40) = 0.93$ ,  $p = 0.36$ , Cohen's  $d = 0.15$ , although there is an advantage for faces illuminated from above as the light-source azimuth moves away from  $0^\circ$ : for  $\pm 30^\circ$ ,  $t(81) = -5.41$ ,  $p < 0.001$ , Cohen's  $d = -0.60$ ; for  $\pm 60^\circ$ ,  $t(81) = -4.12$ ,  $p < 0.001$ , Cohen's  $d = -0.46$ ; for  $\pm 90^\circ$ ,  $t(81) = -3.35$ ,  $p = 0.001$ , Cohen's  $d = -0.37$ . This advantage of lighting from above persists until the horizontal lighting direction

reaches  $\pm 120^\circ$ ,  $t(81) = 7.60$ ,  $p < 0.001$ , Cohen's  $d = 0.84$ , as depicted in the left panel of Figure 10. Additionally, there was a significant reduction in sensitivity between the  $0^\circ$  azimuth and the mean of the  $\pm 30^\circ$  azimuth conditions for the faces illuminated from below,  $t(40) = 3.95$ ,  $p < 0.001$ , Cohen's  $d = 0.62$ , and this change in sensitivity was not evident for faces lit from above: the mean difference in  $d'$  for faces lit from above was  $-0.08$  and for faces lit from below was  $0.46$ ,  $t(40) = 3.46$ ,  $p = 0.001$ , Cohen's  $d = 0.54$ . In sum, faces lit from above and below showed distinct tuning across the horizontal lighting direction.

Overall, discrimination sensitivity was not greatly affected by changes in contrast polarity. Sensitivity across the horizontal and vertical lighting directions generally follows the same pattern for both natural and reversed contrast polarity images. However, sensitivity for the  $\pm 120^\circ$  light-source azimuths was reduced to a greater extent for the faces lit from above in the reversed polarity condition compared with the natural polarity condition: for light-source azimuth and contrast polarity interaction,  $F(8, 320) = 6.27$ ,  $p < 0.001$ ,  $\eta_p^2 = 0.14$ ; it seems that the reduced sensitivity in these conditions was driving the significant main effect of contrast polarity,  $F(1, 40) = 12.47$ ,  $p = 0.001$ ,  $\eta_p^2 = 0.24$ , as well as the interaction between contrast polarity and light-source elevation,  $F(1, 40) = 28.03$ ,  $p < 0.001$ ,  $\eta_p^2 = 0.41$ . The similarity in discrimination sensitivity for the natural polarity and reversed polarity

images across the light-source azimuths suggests that it is the pattern of contrast across a face that facilitates detection rather than the polarity.

The mean response bias across participants for each condition is depicted in [Figure 11](#). A repeated-measures ANOVA showed a significant effect of light-source azimuth,  $F(8, 320) = 43.72, p < 0.001, \eta_p^2 = 0.52$ , and elevation,  $F(1, 40) = 23.12, p < 0.001, \eta_p^2 = 0.37$ , on response bias, where participants were more biased to respond “non-face” for lighting from above (on average) and for extreme horizontal lighting directions. There was also a significant interaction between these two factors,  $F(8, 320) = 33.71, p < 0.001, \eta_p^2 = 0.46$ , where the tendency to respond “non-face” was greater at the extreme horizontal directions for lighting from above. The effect of light-source azimuth also depended on the image polarity, where the bias toward responding “non-face” at the extreme horizontal lighting directions was greater for reversed polarity faces,  $F(8, 320) = 11.89, p < 0.001, \eta_p^2 = 0.23$ . The interaction between light-source elevation and image polarity,  $F(1, 40) = 36.58, p < 0.001, \eta_p^2 = 0.48$ , indicates that the tendency to respond “non-face” for lighting from above was greater for the reversed polarity images compared with the natural polarity images.

## Discussion

The pattern of shading and shadows that falls across a face provides visual cues that can enable face detection. The experiments presented here measured sensitivity at discriminating faces from non-faces using simple contrast patterns produced by shading across the face under a range of horizontal lighting directions. In [Experiment 1](#), we found that sensitivity in detecting spatially upright faces is surprisingly robust to variations in light-source azimuth, despite the considerable change in the contrast pattern occurring across the face as the light-source azimuth changes. This tolerance depended partly on the upright configuration of the face, as sensitivity was more narrowly tuned to light-source azimuth for spatially inverted faces. We expanded on this result in [Experiment 2](#), finding that faces that lit from below show a different pattern of tuning across horizontal lighting directions compared with faces lit from above, including greater robustness to extreme horizontal angles of illumination. The results of [Experiment 2](#) also indicate relatively little difference in detection performance between natural and reversed contrast polarity images, suggesting that it is the *pattern* of the contrast that is critical for face detection across a range of horizontal illuminations more so than the *polarity* of those contrast differences.

## Tolerance to image asymmetries produced by horizontal lighting direction changes

The key finding of our experiments is that the visual system is able to tolerate large variations in horizontal lighting direction when detecting faces from the broad patterns of contrast present in the two-tone images. This tolerance is quite impressive given the significant image asymmetries associated with these lighting variations (quantified in [Figure 2](#)), which contrasts notably with the consistently high sensitivity in discriminating faces from non-faces across most horizontal lighting conditions shown in [Figure 6](#). The relationship between image symmetry and sensitivity is plotted in [Figure 12](#) and suggests that face detection performance in this task was largely unaffected by image asymmetry, particularly for upright images. We did not find much tuning to the horizontal lighting direction for upright faces, although there was more pronounced tuning for spatially inverted faces (which is discussed further below). As indicated by the centroid analysis in [Experiment 1](#), detection performance was facilitated by illumination that was front-on relative to the face in the image, even when that face was oriented away from the observer. This effect was present for both image orientations. It could be argued that image symmetry explains this effect, given that symmetry reaches its maximum for central lighting that is relative to the face (as depicted in [Figure 2](#)). However, the mismatch between the broad tuning of discrimination sensitivity and narrow tuning of image symmetry to the horizontal lighting direction suggests that detection is not dependent on image symmetry. A more likely explanation is that central lighting that is relative to the face represents the lighting angle at which most of the face is visible in the image, and, as such, the contrast pattern created by this lighting angle is most informative of a face. The broad tuning that we observed in both experiments indicates that detection performance is at its worst when most of the face is no longer visible in the image due to cast shadows falling across the face, and the lighting angle at which this occurs depends on the rotation of the head.

## A template-matching approach to face detection

Face detection may be achieved by comparing incoming visual input with some form of internal representation or template of a face, potentially incorporating a lighting model (e.g., as described by [Moore & Cavanagh, 1998](#)). It has been argued that the visual system may use a template-matching approach to perform face detection ([Tsao & Livingstone, 2008](#)). As discussed previously, some behavioral findings are

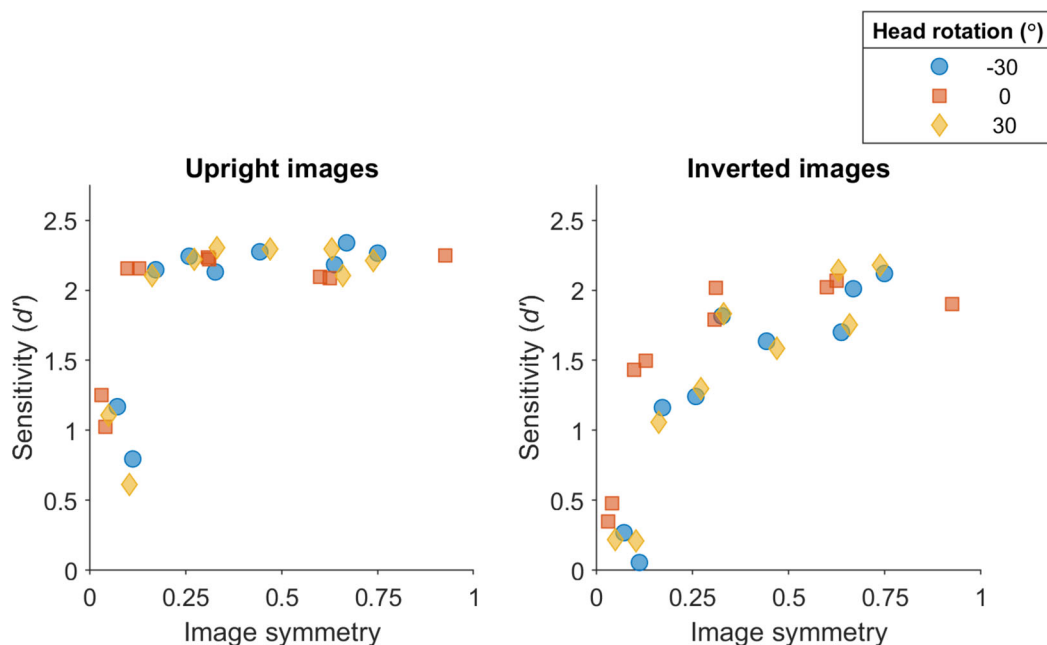


Figure 12. Discrimination sensitivity as a function of image symmetry. The horizontal axis is the symmetry of the two-tone faces for a given condition, averaged across facial identity (as depicted in Figure 2). The vertical axis is the discrimination sensitivity averaged across participants for each condition in Experiment 1. Markers represent the different horizontal lighting conditions, which varied in image symmetry as depicted in Figure 2.

suggestive of the use of face templates that capture the configuration of the human face in the form of a simple contrast pattern (Farroni et al., 2005; Tomalski et al., 2009), and there is evidence for “face cells” in the macaque temporal cortex that are activated by simple face-like contrast patterns (e.g., Kobatake & Tanaka, 1994; Ohayon et al., 2012). A template-matching mechanism is also a potential explanation for face pareidolia, where a broadly tuned template that is sensitive to a basic configuration of facial features may result in illusory face percepts but ensure that a genuine face stimulus is not missed (Caruana & Seymour, 2022; Omer, Sapir, Hatuka, & Yovel, 2019; Paras & Webster, 2013).

How might a template-matching approach deal with the considerable change in the contrast pattern that can occur across the face under horizontal variations in lighting? One possibility is a single non-specific template that captures features that tend to be present in the appearance of faces across a wide range of viewing conditions. Another possibility is the use of multiple templates that are each tuned to the appearance of faces under specific lighting conditions, specific head orientations, or interactions between these. To visualize what form such a template (or templates) may take to exploit the cues provided by shading across the face, we used the two-tone faces from our study to create a series of face templates that capture broad contrast patterns tending to occur under different horizontal lighting and head orientation conditions (Figure 13). The templates were generated by averaging over different

groups of two-tone images, and these average faces (see the left image in each panel in Figure 13) were then thresholded to create two-tone templates (see the right image in each panel in Figure 13). Although the lighting-specific templates share similarities with the two-tone stimuli, the templates that capture features present *across* horizontal lighting conditions (i.e., the non-specific and rotation specific templates) consist of a horizontal contrast pattern across the forehead and nose. It is interesting to note the similarities between these templates and the vertical “bar code” structure of a face described by Dakin and Watt (2009), who argued that this structure is unique to faces and potentially beneficial for face processing. These templates also capture some of the contrast relationships between different regions of the face (e.g., the forehead is brighter than the eye socket) that are central to the “ratio template” introduced by Sinha (2002), which is a computer-vision approach to face detection that is robust to certain lighting changes. As such, the templates in Figure 13 provide some insight into the types of templates that could be useful for detecting human faces under lighting variations.

### Sensitivity is tuned to the horizontal lighting direction for spatially inverted faces

In Experiment 1, detection performance was worse for images that had been spatially inverted, indicating

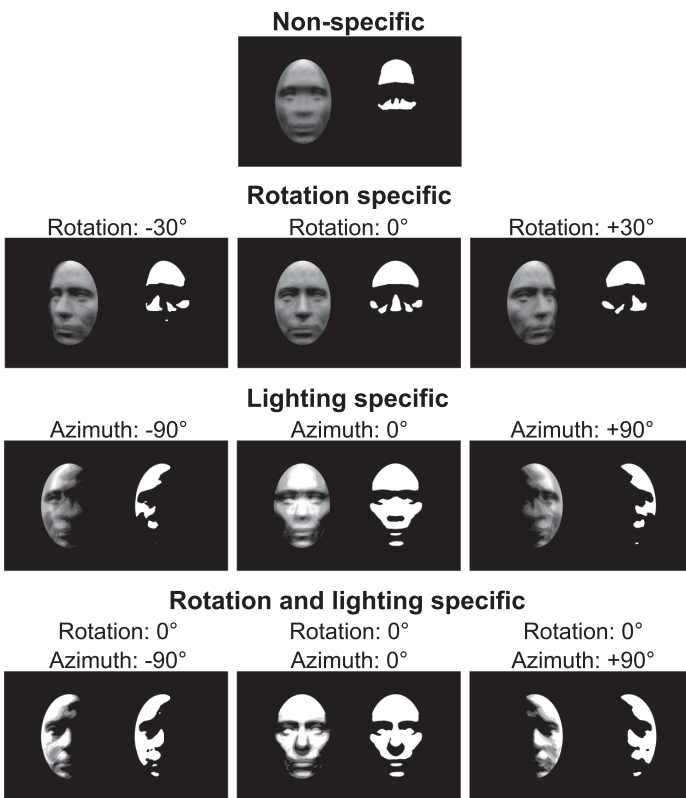


Figure 13. Examples of face templates that were generated from the two-tone images used in the current study. The non-specific template was created by averaging the two-tone faces for all face identities, light-source azimuth, and head rotation conditions (from Experiment 1). The rotation-specific templates were created by averaging over facial identity and light-source azimuth. The lighting-specific templates were created by averaging over facial identity and head rotation. The rotation- and lighting-specific templates were created by averaging over only facial identity. For each panel, the average two-tone image used to create the template is shown on the left and the corresponding template is shown on the right.

that the detection of faces based on the simple contrast patterns present in our stimuli is facilitated by the typical upright configuration of the face. This spatial inversion effect is consistent with previous studies that have examined face detection with Mooney faces (Kanwisher et al., 1998; Palmer et al., 2022), as well as grayscale faces (Garrido, Duchaine, & Nakayama, 2008; Lewis & Edmonds, 2003). The spatial inversion effect is often interpreted as evidence of holistic processing: Spatial inversion disrupts the integration of facial features, causing inverted faces to be more difficult to identify (Rossion, 2008; Rossion, 2009; Taubert et al., 2011). As such, the narrower tuning we observed for spatially inverted faces compared with upright faces (see Figure 6) suggests that our ability to tolerate horizontal lighting changes may depend on

holistic processing of the available features. That is, the large image asymmetries in two-tone images that are associated with horizontal lighting changes can be effectively handled with holistic processing. Although it seems that there is evidence of holistic processing of the two-tone faces, it is interesting to note the contribution of individual facial features to detection. Here, consider our results from Experiment 2 that are shown in Figure 10: Participants were able to detect faces even when only a small part of the nose and mouth was visible in the image (see the lit from below faces in Figure 14). This suggests that some isolated local features are sufficient for detection, pointing toward the joint contribution of part-based and holistic processing to face detection (Canas-Bajo & Whitney, 2020).

### Detection depends on the vertical lighting direction

Interpretation by the visual system of shading information is influenced by the assumption that light typically arrives from above our heads (Mamassian & Goutcher, 2001; Ramachandran, 1988; Sun & Perona, 1998), particularly when there is uncertainty regarding the direction of lighting (Morgenstern, Murray, & Harris, 2011). The influence of prior experience with light from above is evident in face processing. For example, Palmer et al. (2022) found strong tuning in face detection performance to the vertical lighting direction, where human observers were better able to detect faces in two-tone images that were consistent with light arriving from above the face. Overhead lighting also facilitates the recognition of human faces (Enns & Shore, 1997; Hill & Bruce, 1996; Johnston et al., 1992; Liu et al., 1999), including two-tone faces (Peterson, Susilo, Clifford, & Palmer, 2023). In the current study, we found that the discrimination of faces from non-faces in two-tone images was generally better for overhead lighting compared with lighting arriving from below the face. As can be seen in Figure 10, detection performance was slightly worse for faces lit from below compared with above for the non-extreme horizontal lighting conditions. A notable exception to this is the natural contrast polarity faces that were lit from below with a light-source azimuth of 0°, where performance was comparable to that for faces illuminated from above. Taken together with the findings of previous studies, it seems that the detection of human faces from simple contrast patterns is facilitated by sensory patterns that are familiar to us, such as those patterns associated with light arriving from above a face. Conversely, we did not find evidence of a prior for central lighting along the horizontal dimension, as indicated by the centroid analysis in





Figure 14. At the extreme horizontal lighting directions, the pattern of contrast for faces illuminated from below seems to be more easily detectable compared with faces illuminated from above. The faces and non-faces shown in this figure are illuminated by a light source with an azimuth of  $+120^\circ$ .

**Experiment 1**, in which performance appeared to depend more closely on the horizontal lighting direction relative to the rotation of the face rather than central lighting relative to the observer.

For faces that are lit from above, the tolerance to horizontal lighting variations ended at the highly averted lighting directions (i.e., the  $\pm 120^\circ$  light-source azimuth conditions). This drop in performance is reasonable; for these conditions, only a small portion of the face was visible in the image and appeared quite similar to the non-faces (compare the  $\pm 120^\circ$  condition stimuli shown in [Figures 4](#) and [5](#)). In contrast, the detection of faces illuminated from below was more robust to extreme angles of horizontal lighting. As can be seen [Figure 9](#), there was a drop in detection performance for the faces lit from above when the face was illuminated from  $\pm 120^\circ$ , but we did not see this change in performance for the faces illuminated from below. This pattern is also present in the response bias (see [Figure 10](#)). An explanation for this is the type of information present in the two-tone images for the two light-source elevation conditions. Consider the two-tone images in [Figure 14](#); the image of a face illuminated from below contains recognizable facial features, with parts of the nose, mouth, and chin visible in the image. In comparison, the pattern of contrast for the face illuminated from above is less informative of face (with only a sliver of the forehead visible). For the highly averted lighting directions, it is possible that participants were more likely to mistake a face for a non-face when the lighting was from above compared with below due to these image differences.

### Effect of reversing contrast polarity on detection

Although there was a significant effect of contrast polarity on detection in [Experiment 2](#), we were struck by the similarity in performance across the contrast polarity conditions. Compare the two plots in [Figure 10](#); apart from the  $\pm 120^\circ$  light-source azimuth conditions, there is not much difference in discrimination sensitivity among the contrast polarity conditions. This suggests that the pattern of contrast facilitates face detection, rather than the polarity of the contrast.

We were interested in the interaction between contrast polarity reversal and light-source elevation, as previous research has suggested that reversing the polarity of faces that are lit from below can reduce the adverse effect of bottom lighting on recognition performance ([Johnston et al, 1992](#); [Liu et al., 1999](#)). This is likely due to the contrast polarity reversal causing the faces that are lit from below to appear to be lit from above (see the bottom right panel in [Figure 9](#), for example) and is consistent with the preference of the visual system for patterns of contrast associated with overhead lighting. [Palmer et al. \(2022\)](#) did report a small interaction between contrast polarity and lighting elevation in which detection slightly improved for faces lit from below when contrast polarity was reversed, although the overall advantage for faces lit from above remained. Our results are somewhat consistent with those of [Palmer et al. \(2022\)](#); we report that there was an overall advantage for faces that were illuminated from above for most of the non-extreme lighting directions, although there was little difference in detection

performance for faces illuminated from above and below when contrast polarity was reversed (as can be seen in Figure 10).

## Conclusions

The aim of the experiments presented here was to examine how face detection based only on broad patterns of contrast on a face is affected by changes in the horizontal lighting direction. In two experiments, we showed that the discrimination of faces from non-faces based on these cues is remarkably robust to variations in horizontal lighting despite the large image asymmetries associated with these variations. This tolerance to variations in horizontal lighting appears to rely partly on the upright configuration of the face (potentially implicating holistic processing) and relates to the pattern of luminance occurring across the face independent of its contrast polarity. Our results also extend on those of Palmer et al. (2022) by demonstrating that, although the advantage of lighting from above does persist across horizontal lighting directions, there are instances in which detection is better for faces illuminated from below. Overall, our results demonstrate that the visual system can utilize the unique patterns of contrast produced by shading and shadows across the internal features of the face and that these cues are beneficial for face detection across a range of lighting directions. The ability of observers to accommodate considerable changes in the pattern of contrast across the face produced by different horizontal lighting directions has implications for understanding how a template-matching approach may be implemented in human vision.

*Keywords:* face detection, Mooney faces, illumination, shading, shadows

## Acknowledgments

Supported by a grant from the Australian Research Council Discovery Project (DP200100003). CJP was also supported by an Australian Research Council Discovery Early Career Researcher Award (DE190100459).

Commercial relationships: none.  
Corresponding author: Lindsay M. Peterson.  
Email: lindsay.peterson@unsw.edu.au.  
Address: School of Psychology, University of New South Wales, Sydney, Australia.

## References

- Adini, Y., Moses, Y., & Ullman, S. (1997). Face recognition: The problem of compensating for changes in illumination direction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7), 721–732, <https://doi.org/10.1109/34.598229>.
- Braje, W. L. (2003). Illumination encoding in face recognition: Effect of position shift. *Journal of Vision*, 3(2):4, 161–170, <https://doi.org/10.1167/3.2.4>.
- Braje, W. L., Kersten, D. J., Tarr, M. J., & Troje, N. F. (1998). Illumination effects in face recognition. *Psychobiology*, 26(4), 371–380, <https://doi.org/10.3758/BF03330623>.
- Brodski, A., Paasch, G. F., Helbling, S., & Wibrall, M. (2015). The faces of predictive coding. *The Journal of Neuroscience*, 35(24), 8997–9006, <https://doi.org/10.1523/JNEUROSCI.1529-14.2015>.
- Canas-Bajo, T., & Whitney, D. (2020). Stimulus-specific individual differences in holistic perception of Mooney faces. *Frontiers in Psychology*, 11, 1–10, <https://doi.org/10.3389/fpsyg.2020.585921>.
- Caruana, N., & Seymour, K. (2022). Objects that induce face pareidolia are prioritized by the visual system. *British Journal of Psychology*, 113(2), 496–507, <https://doi.org/10.1111/bjop.12546>.
- Chen, C.-C., Chen, C.-M., & Tyler, C. W. (2013). Depth structure from asymmetric shading supports face discrimination. *PLoS One*, 8(2), e55865, <https://doi.org/10.1371/journal.pone.0055865>.
- Dakin, S. C., & Watt, R. J. (2009). Biological “bar codes” in human faces. *Journal of Vision*, 9(4):2, 1–10, <https://doi.org/10.1167/9.4.2>.
- De Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior Research Methods*, 47(1), 1–12, <https://doi.org/10.3758/s13428-014-0458-y>.
- Enns, J. T., & Shore, D. I. (1997). Separate influences of orientation and lighting in the inverted-face effect. *Perception & Psychophysics*, 59(1), 23–31, <https://doi.org/10.3758/BF03206844>.
- Farroni, T., Johnson, M. H., Menon, E., Zulian, L., Faraguna, D., & Csibra, G. (2005). Newborns’ preference for face-relevant stimuli: Effects of contrast polarity. *Proceedings of the National Academy of Sciences, USA*, 102(47), 17245–17250, <https://doi.org/10.1073/pnas.0502205102>.
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G\*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior*

- Research Methods*, 39(2), 175–191, <https://doi.org/10.3758/BF03193146>.
- Favelle, S., Hill, H., & Claes, P. (2017). About face: Matching unfamiliar faces across rotations of view and lighting. *i-Perception*, 8(6), 1–22, <https://doi.org/10.1177/2041669517744221>.
- Garrido, L., Duchaine, B., & Nakayama, K. (2008). Face detection in normal and prosopagnosic individuals. *Journal of Neuropsychology*, 2(1), 119–140, <https://doi.org/10.1348/174866407X246843>.
- Hershler, O., & Hochstein, S. (2005). At first sight: A high-level pop out effect for faces. *Vision Research*, 45(13), 1707–1724, <https://doi.org/10.1016/j.visres.2004.12.021>.
- Hill, H., & Bruce, V. (1996). The effects of lighting on the perception of facial surfaces. *Journal of Experimental Psychology: Human Perception and Performance*, 22(4), 986–1004, <https://doi.org/10.1037/0096-1523.22.4.986>.
- Johnston, A., Hill, H., & Carman, N. (1992). Recognising faces: Effects of lighting direction, inversion, and brightness reversal. *Perception*, 42(11), 1227–1237, <https://doi.org/10.1068/p210365n>.
- Kanwisher, N., Tong, F., & Nakayama, K. (1998). The effect of face inversion on the human fusiform face area. *Cognition*, 68(1), 1–11, [https://doi.org/10.1016/S0010-0277\(98\)00035-3](https://doi.org/10.1016/S0010-0277(98)00035-3).
- Kingdom, F. A. A., & Prins, N. (2010). *Psychophysics: A practical introduction*. London: Academic Press.
- Kobatake, E., & Tanaka, K. (1994). Neuronal selectivities to complex object features in the ventral visual pathway of the macaque cerebral cortex. *Journal of Neurophysiology*, 71(3), 856–867, <https://doi.org/10.1152/jn.1994.71.3.856>.
- Lange, K., Kühn, S., & Filevich, E. (2015). “Just Another Tool for Online Studies” (JATOS): An easy solution for setup and management of web servers supporting online studies. *PLoS One*, 10(6), e0130834, <https://doi.org/10.1371/journal.pone.0130834>.
- Lewis, M. B., & Edmonds, A. J. (2003). Face detection: Mapping human performance. *Perception*, 32(8), 903–920, <https://doi.org/10.1068/p5007>.
- Liu, C. H., Collin, C. A., Burton, A. M., & Chaudhuri, A. (1999). Lighting direction affects recognition of untextured faces in photographic positive and negative. *Vision Research*, 39(24), 4003–4009, [https://doi.org/10.1016/S0042-6989\(99\)00109-1](https://doi.org/10.1016/S0042-6989(99)00109-1).
- Mamassian, P., & Goutcher, R. (2001). Prior knowledge on the illumination position. *Cognition*, 81(1), B1–B9, [https://doi.org/10.1016/S0010-0277\(01\)00116-0](https://doi.org/10.1016/S0010-0277(01)00116-0).
- Mooney, C. M. (1957). Age in the development of closure ability in children. *Canadian Journal of Psychology*, 11(4), 219–226, <https://doi.org/10.1037/h0083717>.
- Moore, C., & Cavanagh, P. (1998). Recovery of 3D volume from 2-tone images of novel objects. *Cognition*, 67(1-2), 45–71, [https://doi.org/10.1016/s0010-0277\(98\)00014-6](https://doi.org/10.1016/s0010-0277(98)00014-6).
- Morgenstern, Y., Murray, R. F., & Harris, L. R. (2011). The human visual system’s assumption that light comes from above is weak. *Proceedings of the National Academy of Sciences, USA*, 108(30), 12551–12553, <https://doi.org/10.1073/pnas.1100794108>.
- Ohayon, S., Freiwald, W. A., & Tsao, D. Y. (2012). What makes a cell face selective? The importance of contrast. *Neuron*, 74(3), 567–581, <https://doi.org/10.1016/j.neuron.2012.03.024>.
- Omer, Y., Sapir, R., Hatuka, Y., & Yovel, G. (2019). What is a face? Critical features for face detection. *Perception*, 48(5), 437–446, <https://doi.org/10.1177/0301006619838734>.
- Otsu, N. (1979). A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, 9(1), 62–66.
- Palmer, C. J., Goddard, E., & Clifford, C. W. (2022). Face detection from patterns of shading and shadows: The role of overhead illumination in generating the familiar appearance of the human face. *Cognition*, 225, 105172, <https://doi.org/10.1016/j.cognition.2022.105172>.
- Paras, C. L., & Webster, M. A. (2013). Stimulus requirements for face perception: An analysis based on “totem poles.” *Frontiers in Psychology*, 4, 18, <https://doi.org/10.3389/fpsyg.2013.00018>.
- Peterson, L. M., Susilo, T., Clifford, C. W. G., & Palmer, C. J. (2023) Discrimination of facial identity based on simple contrast patterns generated by shading and shadows. *Vision Research*, 212, 108307, <https://doi.org/10.1016/j.visres.2023.108307>.
- Pongakkasira, K., & Bindemann, M. (2015). The shape of the face template: Geometric distortions of faces and their detection in natural scenes. *Vision Research*, 109, 99–106, <https://doi.org/10.1016/j.visres.2015.02.008>.
- Ramachandran, V. S. (1988). Perception of shape from shading. *Nature*, 331(6152), 163–166, <https://doi.org/10.1038/331163a0>.
- Rossion, R. (2008). Picture-plane inversion leads to qualitative changes of face perception. *Acta Psychologica*, 128, 274–289, <https://doi.org/10.1016/j.actpsy.2008.02.003>.
- Rossion, R. (2009). Distinguishing the cause and consequence of face inversion: The perceptual field hypothesis. *Acta Psychologica*, 132, 300–312, <https://doi.org/10.1016/j.actpsy.2009.08.002>.

- Sinha, P. (2002). Qualitative representations for recognition. In H. H. Bülthoff, C. Wallraven, S. W. Lee, & T. A. Poggio (Eds.), *Lecture Notes in Computer Science: Vol. 2525. Biologically Motivated Computer Vision* (pp. 249–262). Berlin: Springer, [https://doi.org/10.1007/3-540-36181-2\\_25](https://doi.org/10.1007/3-540-36181-2_25).
- Stanislaw, H., & Todorov, N. (1999). Calculation of signal detection theory measures. *Behavior Research Methods, Instruments, & Computers*, *31*(1), 137–149, <https://doi.org/10.3758/BF03207704>.
- Sun, J., & Perona, P. (1998). Where is the sun? *Nature Neuroscience*, *1*(3), 183–184, <https://doi.org/10.1038/630>.
- Taubert, J., Apthorp, D., Aagten-Murphy, D., & Alais, D. (2011). The role of holistic processing in face perception: Evidence from the face inversion effect. *Vision Research*, *51*(11), 1273–1278, <https://doi.org/10.1016/j.visres.2011.04.002>.
- Tomalski, P., Csibra, G., & Johnson, M. H. (2009). Rapid orienting toward face-like stimuli with gaze-relevant contrast information. *Perception*, *38*(4), 569–578, <https://doi.org/10.1068/p6137>.
- Tsao, D. Y., & Livingstone, M. S. (2008). Mechanisms of face perception. *Annual Review of Neuroscience*, *31*, 411–437, <https://doi.org/10.1146/annurev.neuro.30.051606.094238>.