

Separate influences of orientation and lighting in the inverted-face effect

JAMES T. ENNS and DAVID I. SHORE

University of British Columbia, Vancouver, British Columbia, Canada

Studies of the inverted-face effect typically use photos as stimuli. Inverting photos not only misorients the face but also reverses important shading and shadow cues. We decoupled the influence of spatial orientation and the direction of lighting in three experiments and found that the relation between these factors varied with the task given to observers. When the task required identification of faces (Experiments 1 and 3), the factors were additive, consistent with a strategy of mental rotation of the face prior to an interpretation of the shading cues. When faces were assigned to coarse categories (Experiments 2 and 3), these factors interacted, consistent with a more piecemeal approach to face processing. We propose that the identification of a specific individual depends on configurational information, which is preserved if the image of an inverted face is mentally rotated before the identification process is begun.

It is difficult to recognize and identify individuals from inverted photos of their faces (Diamond & Carey, 1986; Hochberg & Galper, 1967; Valentine, 1988; Yin, 1969, 1970). This paper examines the relations between two factors that have been confounded in most research on the inverted-face effect (IFE): the spatial orientation of the face and the direction of lighting in the photo. In the course of this investigation, we discovered a revealing interaction between the task given to observers and the apparent strategies employed in the perception of inverted faces.

Background

The present study was not specifically designed to address the underlying causes of the IFE. Nonetheless, because the questions we asked were relevant to several prevailing theories, we will briefly outline these views in order to provide a context for the present study.

By one account (hereafter, the *faces-are-special* account), views of upright human faces are represented in neural substrates different from those involved in the perception of other objects. Support for this position comes from studies of hemispheric specialization (Hillger & Koenig, 1991; Sergent, 1984a), the clinical condition of prosopagnosia (Damasio, 1985), and single-cell recordings from awake monkeys (Hietanen, Perrett, Oram, Benson, & Dittrich, 1992). By this account, inverted faces are no more “difficult” to process than are any other objects; upright faces receive *privileged* processing by the brain.

In another view (hereafter, the *expertise* account), inverted faces are difficult to identify because faces are members of a special class of stimuli in which each member can be considered a variant of an underlying proto-

typical configuration, members can only be distinguished from one another by the relations among a set of parts that they all have in common, and the defining relations for a member of a class are acquired with extensive learning. Support for this position comes from developmental studies (Carey & Diamond, 1977; Diamond & Carey, 1986; Goldstein & Chance, 1980), as well as from studies using caricatures (Rhodes, 1994). Inverted faces, therefore, obscure the very properties that are essential for face identification—namely, the perceptual distance and direction of any given face from the common prototype.

A third view (hereafter, the *faces-are-complex* account) sees the IFE as reflective of general principles of object perception. It assumes that (1) inverted objects are identified only after they have first been “normalized” into a standard orientation (e.g., by mental rotation; see Rock, 1973), (2) faces are highly complex three-dimensional objects (Bruce et al., 1991), and (3) mental rotation is more difficult for complex objects than for simple objects, either because rotation of complex patterns can only be done piecemeal (Rock, 1973) or because rotation time increases with pattern complexity (Hochberg & Gellman, 1977; Valentine & Bruce, 1988). In either case, inverted faces are difficult to process because they are among the most complex objects we are asked to mentally rotate before identification processes can be applied.

Despite the obvious differences between theories, a theme that is common to each is that face recognition involves more than perceptual sensitivity to particular visual features; it requires acute sensitivity to the relations among features (i.e., to a configuration or *gestalt*). For the *faces-are-special* account, it is asserted that a face-like configuration is the very trigger-feature for privileged processing. For the *expertise* view, differences in configuration are the prime attributes by which one face can be discriminated from another. For the *faces-are-complex*

Correspondence should be addressed to J. T. Enns, Department of Psychology, University of British Columbia, 2136 West Mall, Vancouver, BC, Canada V6T 1Z4 (e-mail: jenns@cortex.psych.ubc.ca).

account, it is the complex configuration of the face that makes it difficult to mentally rotate.

Independent of these theoretical positions, there is now also an abundance of evidence that the configurational properties of faces play a critical role in the IFE. When feature relations in face stimuli have been varied independently of the component features themselves, larger effects of stimulus inversion have consistently been associated with the relational manipulation (Farah, Tanaka, & Drain, 1995; Hillger & Koenig, 1991; McKelvie, 1991; Rhodes, 1994; Rhodes, Brake, & Atkinson, 1993; Sergent, 1984b; Tanaka & Farah, 1993).

Rationale for the Present Study

When an image of a face is inverted, two separable factors make the configuration of the face difficult to perceive. The most obvious factor is, of course, that of spatial orientation; it is difficult to compare the relations among features in patterns that do not share the same orientation (e.g., Farah et al., 1995). However, a second factor is that of direction of lighting. Visual systems have evolved to process shading and shadows in an environment in which the light source is typically overhead and objects are generally convex (Ramachandran, 1988). Behaving under this assumption, humans typically interpret a shading gradient that runs from light-at-top to dark-at-bottom as a surface slanted away from the viewer; shading that runs in the opposite direction is seen as a surface slanted toward the viewer (Bruce et al., 1991; Gregory, 1973; Kleffner & Ramachandran, 1992; Ramachandran, 1988). From this perspective, it is surprising that most IFE research has confounded spatial orientation with the direction of lighting. Stimuli have typically been culled from yearbook photos (Faw, 1992); turning such pictures upside down not only misorients the face but it also creates an image in which the assumed direction of lighting is not the one preferred by the visual system. We therefore

began with an IFE experiment in which the direction of lighting in the stimuli was varied independently from face orientation, as shown in Figure 1.

There are several possible outcomes for such a study that are of theoretical importance. One prediction arises from the hypothesis that the mental operations involved in normalizing the orientation of the face are independent from, and perhaps even completed *before*, the operations involved in interpreting the shading cues. This would make Figure 1C the easier of the two inverted faces to recognize, since once this image has been rotated, its features are appropriately lit from above, thereby facilitating the task of interpreting the shape-from-shading cues. In contrast, the features in Figure 1D would be lit from below after mental rotation, and it should therefore be more difficult to interpret the shading cues in this picture. This hypothesis leads directly to the strong prediction that the difference in performance between Figures 1C and 1D should be exactly the same as the difference in performance between Figures 1A and 1B. That is, although there is some cost associated with the requirement to normalize an inverted face, once normalized, recognition will be influenced in a similar way by the shading cues for shape. This outcome is illustrated in Figure 2A.

Another possibility is that the factors of orientation and direction of lighting interact. For instance, abnormal lighting (i.e., from below the face) may itself slow down the speed of mental rotation. In that case, Figure 1C might actually be more difficult to recognize than Figure 1D. Alternatively, direction of lighting and spatial orientation might each contribute to the ease/difficulty of face recognition, such that the benefits of lighting from above in Figure 1C are offset by the costs of inverted orientation in that condition. In that case, we would expect Figures 1B, 1C, and 1D to each be comparably difficult in comparison to the standard photo in Figure 1A. This possibility is shown in Figure 2B.

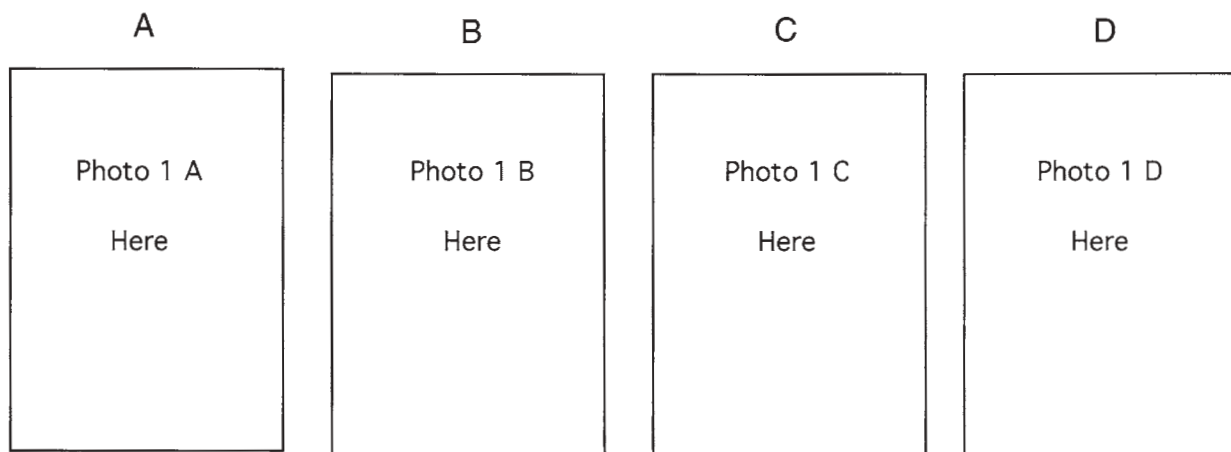


Figure 1. Examples of poses for a single model in each of the four face types in the study. (A) A standard photo that is upright and brow-lit. (B) Upright and chin-lit. (C) Inverted and brow-lit. (D) Inverted and chin-lit. Note that the direction of lighting is given with respect to the face of the model.

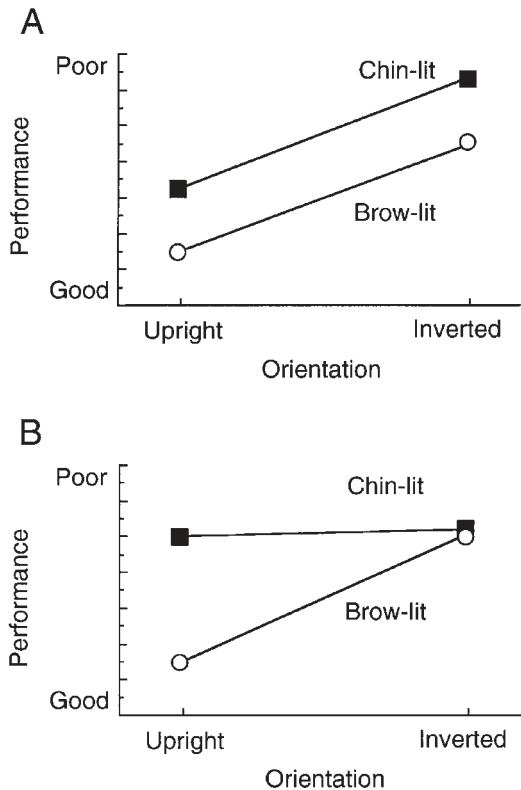


Figure 2. Two possible outcomes of the orthogonal manipulation of face orientation and direction of lighting. (A) The two factors are independent, with the orientation correction process occurring before the interpretation of shape from shading. (B) The two factors interact, with the three nonstandard faces being comparably difficult to process.

We note here that, following the completion of two of the three experiments reported here, a related study was published by Johnston, Hill, and Carman (1992). This provided the opportunity to compare our data with those obtained under interestingly different conditions. Their study consisted of a new-old recognition task in which faces of former classmates were presented in two orientations (upright, inverted) and two lighting conditions (brow-lit, chin-lit). The dependent measure was the number of faces correctly recognized as belonging to a participant's graduating class. This yielded results corresponding to Figure 2B, in that there was a strong interaction between orientation and lighting, with the three nonstandard faces all being comparably more difficult to recognize than the standard face.

General Stimulus Materials

Eight female students were recruited as models from the University of British Columbia, as shown in Figure 3. Each gave written consent to have their photos used in the study and published. Model faces were lit with a single 100-W incandescent bulb, 100 cm from the face, at a 45° angle directly above or below. Other controls included: (1) external facial features (e.g., hair, ears) were eliminated by draping a black cloth around each face; (2) several poses of each face

were used to reduce responses based on specific features of the picture and to favor responses based on invariant features of an individual's face; (3) models were instructed to maintain a neutral emotional expression; and (4) models were photographed lying down to eliminate differences in the effects of gravity for upright and inverted faces.

The black-and-white photos (Illford Delta 400 film, 1/60 sec, F-stop 1.4) were developed as 3.5 × 5 in. matte prints (Illford paper) and then scanned digitally (256 gray levels) so that they occupied 185 × 275 pixels on a Macintosh computer screen. The overall contrast in each image was set by equating the brightest point. A total of six poses were used from each model, three different poses in each of the two lighting conditions (brow- and chin-lit). None of the observers in any of the experiments recognized the models as familiar people.

EXPERIMENT 1 Memory for Recently Learned Faces

Experiment 1 consisted of two phases. In a study phase, observers learned the names of eight novel upright and brow-lit faces. In a subsequent test phase, they attempted to name each face in each of the four face types shown in Figure 1.

Method

Observers. Thirty-eight students (24 women, 14 men) in a 3rd-year university course in perception served as observers in return for extra course credit.

Materials. Each observer received a covered booklet of photocopied faces (8.5 × 11 in.). The booklet was stapled along one edge (assigned as the "top") to maintain a standard orientation for viewing all photos. The learning phase consisted of 10 pages, each containing two rows of four photos (65% reduction), with sufficient space below each photo for the observers to write in a name. These photos were always shown in the standard brow-lit and upright orientation, with one of the three poses for each model chosen randomly. Each page had a different random arrangement of the eight faces, but each observer received the same arrangements in the same order. On the first page, each of the photos were shown with a name appearing directly below it (see Figure 3).

The test phase consisted of 32 pages, with one 3.5 × 5 in. photo per page. Poses representing each of the model (1–8), orientation (upright, inverted), and lighting (brow-lit, chin-lit) conditions were presented in a different random order for each observer.

Procedure. The experiment was run in a single group session lasting about 30 min. The observers were informed that they had the opportunity to participate in an experiment concerning memory for faces. After informed consent was obtained, the observers were instructed to study the names of the eight faces on the first page of the booklet for 30 sec. During this time, the eight names were written on the blackboard, so that the task essentially became one of matching known names to memorized faces. The observers were then instructed to turn the page and attempt to write the correct name below each face from memory. They were given 30 sec to complete this task, after which time the correct names were announced. This process continued until all subjects had correctly matched the names to all eight faces twice in a row. The data for 8 observers were excluded, because they failed to reach this criterion within the learning phase. In the test phase, the observers were given 10 sec to name each face before being told to turn the page and continue. At the conclusion, the experimenter engaged the observers in a discussion of the role of lighting factors and spatial orientation in the perception of faces.

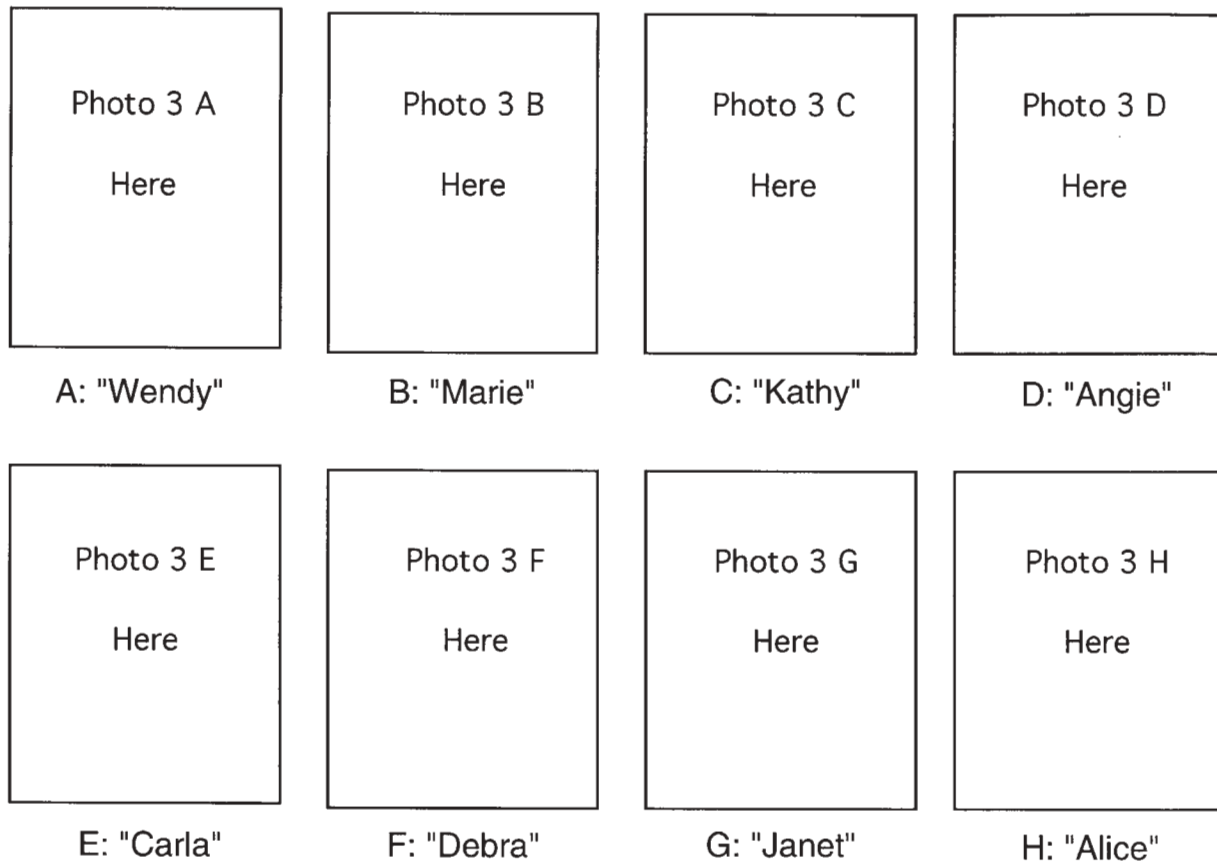


Figure 3. Examples of standard poses (upright, brow-lit) and their assigned names for each of the eight models used in the study.

Results

Mean percentage error is shown in Figure 4. A repeated measures analysis of variance (ANOVA) revealed significant main effects of orientation [$F(1,29) = 22.04, MS_e = 222.99, p < .001$] and lighting [$F(1,29) = 15.91, MS_e =$

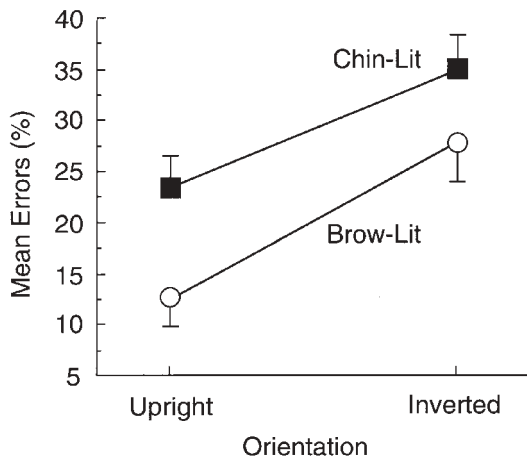


Figure 4. Mean percentage errors in Experiment 1: Memory for recently learned faces.

$177.39, p < .001$]. These two effects were additive, as shown by the nonsignificant interaction [$F(1,29) = 0.49, MS_e = 157.70$].

The data therefore conformed to the prediction shown in Figure 2A. This is consistent with observers first normalizing the orientation of the image (perhaps through mental rotation) before interpreting the shading cues in the photo. Recall that the one previous study to examine these factors (Johnston et al., 1992) found the pattern shown in Figure 2B, which is inconsistent with separate influences of orientation and lighting. We therefore began to investigate which of the many procedural differences might be responsible for this discrepancy.

EXPERIMENT 2
Same-Different Discrimination

There were many differences between the methods of Experiment 1 and those of Johnston et al. (1992). Among them was that our task required observers to specifically identify each face by name; their task tested only for familiarity or a “feeling of knowing” (Mandler, 1980). The stimuli were also very different: ours consisted of faces from a relatively small set of individuals for whom we had eliminated external features; theirs consisted of a much

larger set, and external facial features were not as carefully controlled.

We began by simplifying the observer's task. We opted for a simultaneous *same-different* task, both to minimize any memorial aspects and because it permitted us to examine both response time (RT) and accuracy. If this simple task replicated the interactive data pattern of Johnston et al. (1992; see our Figure 2B), it would suggest that there was nothing about the long-term memory requirements or the new-old recognition response that was responsible for this pattern. Indeed, such a finding would open up the possibility of a fine-grained examination of the interaction of orientation and lighting factors in face perception (as opposed to memory). Alternatively, if this task replicated the additive pattern found in Experiment 1 (Figure 2B), we would have to consider the possibility that these factors play different roles in face perception versus memory, as has been suggested (Valentine, 1988).

Method

Observers. Ten students (4 women, 6 men) from the undergraduate subject pool at the University of British Columbia served as observers. All had normal or corrected-to-normal vision.

Stimuli. On each trial, two faces were presented, one on the right and one on the left of a fixation point (0.43°) at the center of the screen. One face was always the standard (upright and brow-lit), and it appeared randomly on the left or right. Each image was $9.9^\circ \times 14.1^\circ$, centered 7.8° from fixation. There were five faces in the set (A, B, C, D, and E in Figure 3).

Procedure. The observer's task was to press one of two keys to indicate whether the two faces were of the same person or two different people. The observers were instructed that half of the trials would actually show the same people, that some of the faces would appear upside down, and that faces could be lit either from above or from below.

After being introduced to the task and shown a typical display, the observers practiced until they made four correct consecutive responses. Testing consisted of a total of 750 trials (15 blocks of 50 trials) in a 1-h session. Between blocks, the screen displayed percentage accuracy from the previous block, and the observers were permitted to begin the next block at their own pace. Each of the 120 unique conditions (5 models \times 3 poses \times 2 trial types \times 2 orientations \times 2 lightings) was presented randomly and approximately equally often.

Each trial began with a blank screen for 716 msec, followed by the display, which remained on until a response had been made or until 3 sec had elapsed. Feedback was presented at the center of the screen for 700 msec (+ for correct, - for incorrect, 0.5° visual angle).

Results

Mean correct RT and percentage errors are shown in Figure 5. The RT data revealed significant effects of orientation [$F(1,9) = 70.91$, $MS_e = 3,476.11$, $p < .001$] and lighting [$F(1,9) = 96.30$, $MS_e = 1,347.21$, $p < .001$] and a significant interaction [$F(1,9) = 24.94$, $MS_e = 1,626.79$, $p < .001$]. The IFE was significant for both the brow-lit faces [$t(9) = 12.24$, $p < .01$] and the chin-lit faces [$t(9) = 5.17$, $p < .01$].

The error data showed a similar pattern [orientation, $F(1,9) = 21.61$, $MS_e = 5.66$, $p < .01$; lighting, $F(1,9) = 30.59$, $MS_e = 9.89$, $p < .001$; lighting \times orientation, $F(1,9) = 7.50$, $MS_e = 8.33$, $p < .05$]. The IFE was signif-




Figure 5. Mean correct response time (upper panel) and mean percentage errors (lower panel) in Experiment 2: *Same-different* discrimination.

icant for the brow-lit faces [$t(9) = 4.64$, $p < .01$], but not for chin-lit faces [$t(9) < 1$].

These results therefore replicated the interactive pattern reported by Johnston et al. (1992) and shown in Figure 2B. This indicates that there is no necessary link between this data pattern and face memory. It is just as easily obtained in a task in which observers are merely asked to make a perceptual *same-different* judgment.

EXPERIMENT 3

Identification Versus Categorization

So far, we have observed three instances in which face orientation and direction of lighting interact (Johnston et al., 1992, Experiments 1 and 2; present Experiment 2) and one instance in which they are additive (present Experiment 1). Looking across these experiments, it is obvious that face perception versus memory is not at issue in these differences (see Valentine, 1988). However, what is unique about the one case of additivity is that it involves a task in which faces must be identified as belonging to one or another individual. In the *new-old* recognition task,

face identification needs to proceed only to some criterion level of familiarity (Mandler, 1980); in the *same-different* task, individuals need not be identified at all. The faces must simply be judged as belonging to the same (or different) individual.

Why might a task in which faces must be identified result in additivity for the factors of orientation and lighting? A clue may be found in the extent to which observers are required to process the configurational properties of the face rather than merely its unique features. Tasks that require observers to identify individual faces would seem naturally to demand more configurational processing than tasks that ask observers either to inspect for differences (Experiment 2) or to group faces into coarse categories of relative familiarity (Hillger & Koenig, 1991; Johnston et al., 1992; Sergent, 1984b). Furthermore, to the extent that configurational processing is necessary in a task, observers may be more likely to attempt a mental rotation of the entire image (Farah et al., 1995). If so, then a mental rotation of the image would permit inverted brow-lit faces (Figure 1C) to be identified more readily than inverted chin-lit faces (Figure 1D), since the direction of lighting would be appropriately from above the face once the image had been rotated.

To test this hypothesis, we used the same stimuli in two different perceptual tasks. In an identification task, observers were shown a series of faces and asked to identify each one by matching it to one of five photos that were constantly displayed below the computer screen. In the categorization task, observers were shown a similar series of faces and asked to assign each face to either an "in group" of three photos constantly displayed or an "out group" of three photos that were not displayed. Our prediction was that additivity of orientation and lighting would be observed in the identification task, where faces had to be assigned individual labels, but not in the classification task, where faces could be categorized more coarsely into one of two groups.

Method

Observers. Ten students (4 women, 6 men) served as observers in the identification task; 16 different students (10 women, 6 men) participated in the classification task.

Apparatus and Stimuli. In both tasks, a single pose was presented in the center of the screen on each trial, subtending approximately $9.9^\circ \times 14.1^\circ$ visual angle. In the identification task, five numbered photos (A, B, C, D, and E in Figure 3) were always visible directly below the computer screen. In the classification task, only three photos were shown below the screen and designated as the "in group." These three faces were chosen randomly from the set A, B, C, D, E, and F in Figure 3, with five different sets being tested in all.

Procedure. The observers in the identification task pressed one of five keys as rapidly as possible in response to each face. Those in the categorization task pressed one of two keys, depending on whether the face was part of the "in group" or "out group." In both tasks, the observers first practiced until they made four correct consecutive responses. Testing consisted of a total of 500 trials (10 blocks of 50 trials) in a 1-h session. Each trial began with a blank screen for 1,520 msec. A fixation point then appeared for 155 msec, followed by a blank screen for 16 msec, and then the face for 110 msec. Other details were identical to those of Experiment 2.

Results

Identification task. Mean correct RT and percentage errors are shown in Figure 6. The RT data revealed significant main effects of orientation [$F(1,9) = 38.55$, $MS_e = 2,234.01$, $p < .001$] and lighting [$F(1,9) = 23.33$, $MS_e = 483.88$, $p < .001$]. These two factors were additive, as shown by their nonsignificant interaction [$F(1,9) < 1$, $MS_e = 1,067.50$].

The error data revealed a similar pattern [orientation, $F(1,9) = 24.50$, $MS_e = 48.06$, $p < .001$; lighting, $F(1,9) = 5.49$, $MS_e = 10.06$, $p < .05$; lighting \times orientation, $F(1,9) = 2.28$, $MS_e = 2.45$].

Categorization task. Mean correct RT and percentage errors are shown in Figure 7. There were significant RT effects of orientation [$F(1,15) = 218.85$, $MS_e = 712.38$, $p < .001$] and lighting [$F(1,15) = 9.81$, $MS_e = 1,652.42$, $p < .01$] and a significant interaction [$F(1,15) = 6.23$, $MS_e = 959.59$, $p < .05$]. The IFE was significant for the brow-lit faces [$t(15) = 4.41$, $p < .01$], but not for the chin-lit faces [$t(15) < 1$].

The pattern for errors was similar, although only the effect of lighting was significant [orientation, $F(1,15) = 3.02$, $MS_e = 75.11$; lighting, $F(1,15) = 8.10$, $MS_e = 25.20$, $p < .05$; orientation \times lighting, $F(1,15) < 1$, $MS_e = 19.86$]. Again, the IFE was significant for the brow-lit

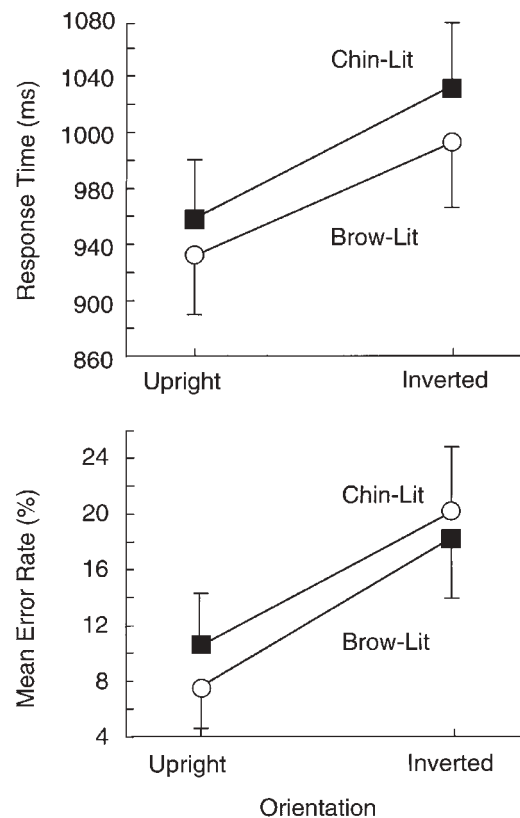


Figure 6. Mean correct response time (upper panel) and mean percentage errors (lower panel) in Experiment 3: identification task.

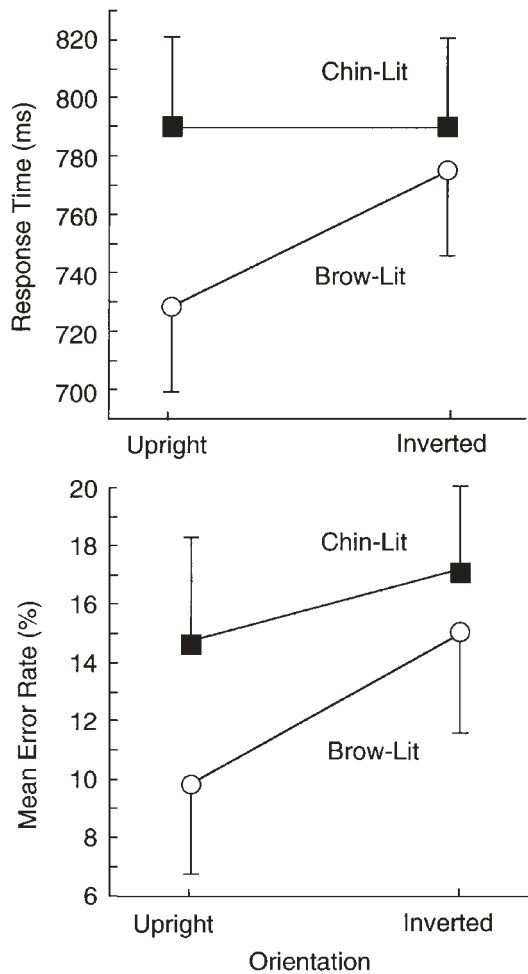


Figure 7. Mean correct response time (upper panel) and mean percentage errors (lower panel) in Experiment 3: categorization task.

faces [$t(15) = 2.91, p < .025$], but not for the chin-lit faces [$t(15) = 1.88$].

GENERAL DISCUSSION

The common practice of studying the IFE by using photos of faces that have merely been inverted introduces a potential confound to the study of face perception. Not only is the spatial structure of the face inverted but so is the apparent direction of lighting in the photo and, therefore, the cues to facial shape that can be derived from shading and shadows. When we separated the effects of facial orientation from those of lighting direction in the present study, we found a pattern of results that was more complex than we had anticipated. That is, the relation between orientation and lighting direction varied with the task given to observers. However, this finding in itself appears to be quite telling with regard to the problem of face perception.

When the observers were required to identify a specific individual whose face had been inverted, either from memory in Experiment 1 or in a match-to-sample task in Experiment 3, the factors of orientation and lighting were additive, according to the model of task performance illustrated in Figure 2A. That is, the observers appeared to solve the spatial orientation problem separately (perhaps by first mentally rotating the image) from the problem of shape interpretation based on the shading cues in the face. On the other hand, when the observers were merely required to discriminate the faces as members of one or another coarse category, as in Experiments 2 and 3 (categorization task), these factors interacted according to the pattern in Figure 2B. Here, all the faces that were non-standard, because of abnormal orientation, abnormal lighting, or both, were comparably difficult to perceive. This form of interaction may indicate an interdependence between the operations involved in mental rotation and the perception of shape from shading (e.g., bottom-lit photos are especially difficult to mentally rotate), or it may reflect the privileged status afforded the processing of standard faces, as predicted by the faces-are-special view summarized in the introduction. More research will be needed to resolve this specific question.

It is nonetheless important to consider what might be behind the apparent connection between tasks requiring the identification of specific individuals and the additivity of orientation and lighting factors. One explanation concerns the specificity of information required. In order to identify a particular individual from among a set of similar individuals, it may be necessary to perform an analysis involving configurational properties of the face. Feature-based comparisons may be insufficient, as is indicated by the expertise account of face perception summarized in the introduction and by the empirical evidence on the importance of relational properties in face perception. If so, then only a transformation of an inverted face that preserved the important configurational properties would be sufficient to identify an inverted face, and mental rotation of the entire face is such a transformation. In order to merely discriminate one face from another, or to categorize faces into groups based on common features, such a full analysis may not be required. Observers may then opt not to mentally rotate in the inverted, especially when one considers the mental effort involved (Hochberg & Gellman, 1977; Rock, 1973).

Several other aspects of the data are consistent with this interpretation. For example, it is noteworthy that the additive pattern of results seems to coincide with tasks that are generally more difficult for observers to perform. A comparison of the average RTs and errors in Experiment 3 shows that the identification task was generally slower (by 200 msec) and slightly more error prone (by 2%). Second, the average size of the IFE was larger in those tasks where additivity was observed. This is consistent with a mental rotation operation being used selectively in those tasks.

The possibility that the observers were using a mental rotation strategy in all tasks, but that different aspects of the image were being rotated in different tasks, should also be considered. For example, when the task is face discrimination or categorization, it may not be necessary to rotate the entire image. Instead, it may be sufficient to rotate only a part of the face, such as the nose or the eyes. This would certainly speed up the orientation normalization process (Hochberg & Gellman, 1977). Alternatively, observers may rotate facial features one at a time, until a threshold of discrimination or familiarity has been met. Either of these possibilities could produce the interaction of orientation and lighting seen in these simpler tasks.

What about the possibility that the observers were using a feature-based strategy in all tasks? One attempt at such a parsimonious account of the present results might begin by proposing that some threshold number of features must be identified to perform each task. The factors of orientation and lighting each contribute to the difficulty of this feature extraction process for observers. Furthermore, some tasks require a larger number of features to differentiate the correct responses (e.g., naming or matching individuals) than do others (e.g., discriminating one face from another, or deciding whether a face is familiar). The main finding that such a view has difficulty predicting is the specific pattern of additivity found for the factors of lighting and orientation in the more difficult tasks. Why should the relative difficulty of feature extraction be exactly the same for brow- and chin-lit upright faces (Figures 1A and 1B) as for brow- and chin-lit inverted faces (Figures 1C and 1D)? This line of thought leads to speculation that there must be orientation-specific features that are completely independent from shading-specific features. This seems implausible to us, in large part because of the inherent relations between the direction of shading (i.e., orientation) and the assumed lighting direction in the interpretation of any shading features (Ramachandran, 1988). On the other hand, such additivity is quite naturally consistent with the existence of a configuration-preserving operation, such as mental rotation.

THEORETICAL IMPLICATIONS

As mentioned in the introduction, the IFE itself is sometimes taken as support for the view that face perception involves specialized processes (Carey & Diamond, 1977; Rhodes et al., 1993; Sergent, 1984a; Yin, 1969). The finding that orientation and lighting interact, as they do in some tasks, is completely consistent with this position. In the interaction that was observed (see Figure 2B), only the standard upright face receives privileged processing. However, the finding that orientation and lighting are sometimes additive in other tasks is a little more difficult to incorporate. Why should the relations between these factors be so easily changed by the task? One way for the faces-are-special view to accommodate this finding is to propose that face stimuli can sometimes be submitted to the specialized face processors after they

have been normalized for orientation. But this begs the question, what determines that a stimulus is a face in need of such an operation?

It is easier to accommodate the finding of task dependency within the expertise and faces-are-complex views. Tasks that require the identification of specific individuals would naturally depend more heavily on configurations, since configurations are thought to be the defining features of individual faces. Observers would therefore be prudent to seek out a normalizing transformation for an inverted face that would preserve the configurational qualities. Mental rotation is one such transformation, although there may be others.

Finally, it is worth noting that neither Johnson et al. (1992) nor we have observed one particular form of interaction between orientation and lighting—namely, a complete crossover interaction superimposed on a main effect of orientation. In such an outcome, Figure 1D would be easier to perceive than Figure 1C by an amount equal to the benefit of Figure 1A over Figure 1B. This outcome would be consistent with a model of task performance in which observers interpreted the shading cues in the image *before* mentally rotating the face. That is, if the interpretation based on shading was performed first, then it should be easier to do so in Figure 1D because that image is already lit from above. The subsequent mental rotation would thus begin before that in Figure 1C, where the image is bottom lit, making shading interpretation more difficult. However, since this outcome was never observed, we conclude that the mental rotation of faces in our Experiments 1 and 3 (identification task) was an operation that was applied to a relatively early representation in the visual system, one that had not yet performed a three-dimensional surface interpretation of the shading patterns in the image.

REFERENCES

- BRUCE, V., HEALEY, P., BURTON, M., DOYLE, T., COOMBES, A., & LINNEY, A. (1991). Recognizing facial surfaces. *Perception*, **20**, 755-769.
- CAREY, S., & DIAMOND, R. (1977, January 21). From piecemeal to configurational representation of faces. *Science*, **195**, 312-314.
- DAMASIO, A. R. (1985). Prosopagnosia. *Trends in Neurosciences*, **8**, 132-135.
- DIAMOND, R., & CAREY, S. (1986). Why faces are and are not special: An effect of expertise. *Journal of Experimental Psychology: General*, **105**, 107-117.
- FARAH, M. J., TANAKA, J. W., & DRAIN, H. M. (1995). What causes the face inversion effect? *Journal of Experimental Psychology: Human Perception & Performance*, **21**, 628-634.
- FAW, H. W. (1992). Recognition of unfamiliar faces: Procedural and methodological considerations. *British Journal of Psychology*, **83**, 25-37.
- GOLDSTEIN, A. G., & CHANCE, J. E. (1980). Memory for faces and schema theory. *Journal of Psychology*, **105**, 47-59.
- GREGORY, R. L. (1973). The confounded eye. In R. L. Gregory & E. H. Gombrich (Eds.), *Illusion in nature and art* (pp. 49-96). London: Duckworth.
- HIETANEN, J. K., PERRETT, D. I., ORAM, M. W., BENSON, P. J., & DITTRICH, W. H. (1992). The effects of lighting conditions on responses of cells selective for face views in the macaque temporal cortex. *Experimental Brain Research*, **89**, 157-171.
- HILLGER, L. A., & KOENIG, O. (1991). Separable mechanisms in face

- processing: Evidence from hemispheric specialization. *Journal of Cognitive Neuroscience*, **3**, 42-58.
- HOCHBERG, J., & GALPER, R. E. (1967). Recognition of faces: I. An exploratory study. *Psychonomic Science*, **9**, 619-620.
- HOCHBERG, J., & GELLMAN, L. (1977). The effect of landmark features on mental rotation times. *Memory & Cognition*, **5**, 23-26.
- JOHNSTON, A., HILL, H., & CARMAN, N. (1992). Recognizing faces: Effects of lighting direction, inversion, and brightness reversal. *Perception*, **21**, 365-375.
- KLEFFNER, D. A., & RAMACHANDRAN, V. S. (1992). On the perception of shape from shading. *Perception & Psychophysics*, **52**, 18-36.
- MANDLER, G. (1980). Recognizing: The judgment of previous occurrence. *Psychological Review*, **87**, 252-271.
- MCKELVIE, S. J. (1991). Effects of processing strategy and transformation on recognition memory for photographs of faces. *Bulletin of the Psychonomic Society*, **29**, 98-100.
- RAMACHANDRAN, V. S. (1988). Perceiving shape from shading. *Scientific American*, **259**, 76-83.
- RHODES, G. (1994). Secrets of the face. *New Zealand Journal of Psychology*, **23**, 3-17.
- RHODES, G., BRAKE, S., & ATKINSON, A. P. (1993). What's lost in inverted faces? *Cognition*, **47**, 25-57.
- ROCK, I. (1973). *Orientation and form*. New York: Academic Press.
- SERGEANT, J. (1984a). Configural processing of faces in the left and right cerebral hemispheres. *Journal of Experimental Psychology: Human Perception & Performance*, **10**, 554-572.
- SERGEANT, J. (1984b). An investigation into component and configural processes underlying face perception. *British Journal of Psychology*, **75**, 221-242.
- TANAKA, J. W., & FARAH, M. J. (1993). Parts and wholes in face recognition. *Quarterly Journal of Experimental Psychology*, **46A**, 225-245.
- VALENTINE, T. (1988). Upside-down faces: A review of the effect of inversion upon face recognition. *British Journal of Psychology*, **79**, 471-491.
- VALENTINE, T., & BRUCE, V. (1988). Mental rotation of faces. *Memory & Cognition*, **16**, 556-566.
- YIN, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, **81**, 141-145.
- YIN, R. K. (1970). Face recognition: A dissociable ability? *Neuropsychologia*, **23**, 395-402.

(Manuscript received July 28, 1995;
revision accepted for publication February 15, 1996.)