Research Article

REVISITING THE PERCEPTION OF UPSIDE-DOWN FACES

Janice E. Murray,[1] Esther Yong,[1] and Gillian Rhodes[2]

[1]University of Otago, Dunedin, New Zealand, and [2]University of Western Australia, Nedlands, Australia

Abstract—*In two experiments, the effect of orientation on face perception was assessed. Using a scale from 1 (*normal*) to 7 (*bizarre*), participants rated normal, unaltered faces and faces in which changes had been made to spatial-relational properties (eyes and mouth inverted or relative position of the eyes and mouth altered) or to component properties (eyes whitened and teeth blackened). For unaltered and component-distortion faces, bizarreness ratings increased linearly as orientation increased from 0° to 180°. For spatial-distortion faces, a discontinuity in the function relating orientation and bizarreness was in evidence between 90° and 120°. The results provide support for the view that there is a qualitative difference in the processing of upright and inverted faces due to the disproportionate effect of inversion on the encoding of spatial-relational information.*

What encoding mechanisms underlie humans' remarkable ability to discriminate faces? A possible answer is suggested by the observation that face encoding is more vulnerable to inversion than is encoding of other classes of stimuli (for reviews, see Searcy & Bartlett, 1996; Valentine, 1988). There have been suggestions that spatial relations among components (e.g., relative position of the eyes and nose) are more important than the individual components themselves in face perception, and that it is the perceptual encoding of spatial-relational information that is disproportionately affected by inversion (e.g., Carey & Diamond, 1977; Diamond & Carey, 1986; Rhodes, Brake, & Atkinson, 1993; Searcy & Bartlett, 1996).[1] Alternatively, Valentine (1988, 1991) has argued that inversion adds noise to the encoding process, affecting spatial-relational and component information equally.

The two views differ in their predictions for how encoding of a face changes as the face is rotated from the upright. The dual-mode view (Carey & Diamond, 1977) suggests a qualitative difference in what is encoded; whereas both spatial-relational and component information can be encoded in upright faces, predominantly component information is encoded in inverted faces. This implies a discontinuity in the function relating departure from upright and success at face encoding at the point where coding shifts from reliance on spatial-relational to component information. No such discontinuity is expected according to the noise view, because it is argued that the encoding of spatial-relational information is not disproportionately impaired when the face is inverted (Valentine, 1988, 1991). Rather, this view suggests a quantitative effect on encoding; face encoding simply becomes increasingly difficult and error prone with increasing departures from upright. Valentine and Bruce (1988) sought evidence for the discontinuity that might be expected if there is a shift from spatial-relational to component information processing as a face departs from upright orientation. In sequential face-matching and recognition tasks, they found that response time increased linearly with increases in orientation, with no deviation from linearity. Valentine and Bruce viewed these results as contradictory to the dual-mode view, and argued that the effect of change in orientation on face processing is to increase the difficulty in encoding spatial-relational information. However, they acknowledged that the procedures they used may not have been sensitive enough to detect the suggested shift in processing strategy.

In the experiments we report here, we revisited this question of a discontinuity in the function relating orientation and face encoding. Spatial-relational and component information were manipulated separately to determine whether a discontinuity occurs in the perception of spatial-relational information, as expected according to the dual-mode view. In the experiments, unaltered faces and faces made to look grotesque through altered spatial-relational or component properties were rated for bizarreness over 24 orientations. We chose to use bizarreness ratings rather than other measures because we were specifically interested in exploring the possibility that the orientation effect occurs in the perceptual encoding of faces, rather than in the retention and retrieval of memory representations. Other tasks, such as recognition and sequential matching, contain a memory component that is irrelevant to the question of perceptual encoding.

Past studies have shown that the inversion effect occurs in the perceptual encoding of faces. Bartlett and Searcy (1993; Searcy & Bartlett, 1996) found that whereas inversion reduced the perceived grotesqueness of faces in which spatial information had been altered, inversion produced no reliable change in the perceived grotesqueness of faces that had been componentially distorted through whitening of the eyes and blackening of the teeth. Although consistent with the dual-mode hypothesis, Searcy and Bartlett's results do not provide a direct test of the notion that there is a qualitative shift in processing strategy at some intermediate orientation between 0° and 180°. We sought to do this in our studies. The critical condition involved spatial-relational distortion, with component-distortion and unaltered conditions serving as controls. If there is a switch from processing spatial-relational information to processing of predominantly component information with increasing departures from upright, then a discontinuity in the function relating bizarreness and orientation should be observed for faces in which bizarreness is explicitly created by manipulating spatial-relational properties. No such discontinuity would be expected for unaltered faces, or faces made bizarre through a manipulation of component properties. This outcome would support the dual-mode view. Alternatively, if the inversion effect is derived from an increased difficulty in face encoding rather than a switch from one mode to another, then a continuous reduction in bizarreness ratings would be expected for spatial-relationally altered faces. This finding would support the noise view.

Address correspondence to Janice E. Murray, Department of Psychology, University of Otago, Box 56, Dunedin, New Zealand; e-mail: jmur@psy.otago.ac.nz.

1. Various terms, such as configural, holistic, second-order relational, and spatial-relational, have been applied to describe the information derived from the encoding of the spatial relations among facial components. In this article, we adopt the term spatial-relational information to refer to this type of information.

EXPERIMENT 1

In Experiment 1, to create a perception of bizarreness derived from changes in spatial-relational information, we exploited a perceptual effect that is highly sensitive to inversion, namely, the Thatcher illusion (Thompson, 1980). In Thatcherized faces, spatial-relational information is altered by inverting the eyes and mouth in a face that is normal in all other respects. The resulting face is perceived as grotesque or bizarre looking when viewed upright. What is notable about this effect is that it is dramatically reduced when the entire altered face is inverted. To create a perception of bizarreness derived from changes to component information, we whitened the eyes and blackened portions of the teeth in otherwise normal faces. The Thatcherized, unaltered, and component-distortion faces were presented at 24 orientations and rated for bizarreness.

Method

Participants

Twenty-four students (12 females) from the University of Otago participated for payment or credit toward a course requirement. All had normal or corrected-to-normal vision. Each participant was tested in a single session lasting approximately 50 min.

Materials

The stimuli were black-and-white photographs of four male and four female faces. The models were undergraduate students of the University of Otago who agreed to have their photos used in face perception experiments. None of the models had facial hair or wore glasses, all had short to shoulder-length hair, and all were photographed smiling, with visible teeth. Each photo was scanned and edited to produce a face that was 55 mm by 75 mm in size, subtending 4° of visual angle against a light gray background. Three versions of each face were produced. For the Thatcherized version, the eyes and mouth were inverted in the upright face. For the component-distortion version, the pupils of the eyes were whitened and portions of the teeth were blackened. The third version consisted of the unaltered face. An example of each version is shown in Figure 1. Each face in each version was shown in 24 orientations, the upright view and 23 rotated views in 15° steps, for a total of 576 stimuli. Presentation of stimuli and collection of responses were controlled by Micro Experimental Laboratory software (Schneider, 1988) on an IBM-compatible computer.

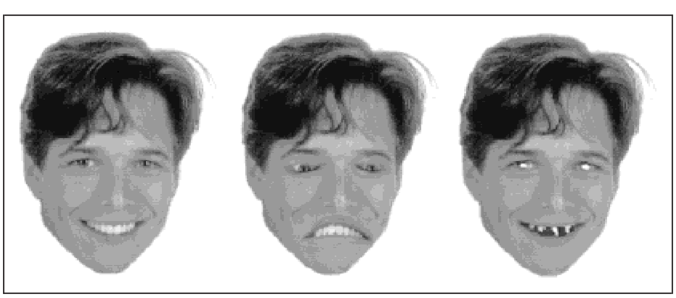

Fig. 1. Examples of the stimuli used in Experiment 1: an unaltered face (left), a Thatcherized face (middle), and a component-distortion face (right).

Procedure

Each face was presented in the center of the computer screen for 3 s. Participants were instructed to rate the bizarreness of each face on a scale from 1 to 7, with 1 representing the most normal-looking face and 7 representing the most bizarre. Participants made their response by pressing keys marked 1 through 7 on the keyboard. The next trial began after the participant's response.

The experimental trials were divided into three blocks of 192 trials. In each block, each version of each of the eight faces was presented at eight different orientations, so that equal numbers of faces, versions, and orientations were presented in each block. The order in which trials were presented in each block was determined randomly for each participant, and the order of the blocks was counterbalanced across participants. Ten practice trials using nonexperimental faces preceded the 576 experimental trials.

Results and Discussion

Ratings for clockwise and counterclockwise rotations of equal magnitude were averaged. The mean ratings for unaltered, component-distortion, and Thatcherized faces at each of the 13 orientations are shown in Figure 2. As is evident in the figure, with increasing

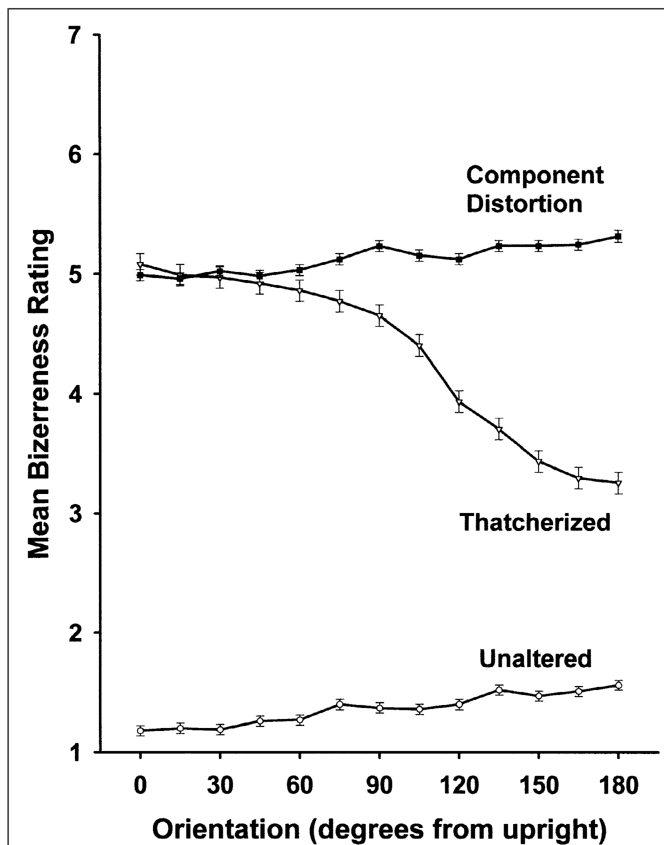

Fig. 2. Mean bizarreness ratings (1 = *most normal*, 7 = *most bizarre*) and standard errors for unaltered, Thatcherized, and component-distortion faces as a function of orientation (degrees from upright) in Experiment 1. Error bars are within-subjects errors (Loftus & Masson, 1994) calculated separately for each function.

departures from the upright, bizarreness ratings gradually increased for unaltered and component-distortion faces, and decreased for Thatcherized faces. This interaction between orientation and face version was found to be significant, $F(24, 552) = 56.17$, $p < .001$, in an analysis of variance (ANOVA) in which the main effects of version, $F(2, 46) = 137.59$, $p < .001$, and orientation, $F(12, 276) = 23.6$, $p < .001$, were also significant. All three functions suggest a linear change in bizarreness ratings with increases in orientation. This apparent linear effect of orientation was significant for unaltered faces, $F(1, 23) = 9.18$, $p < .001$; component-distortion faces, $F(1, 23) = 16.62$, $p < .001$; and Thatcherized faces, $F(1, 23) = 111.11$, $p < .001$.[2] What is equally apparent in Figure 2 is the striking departure from linearity shown for Thatcherized faces, $F(11, 253) = 11.89$, $p < .001$. For the unaltered and component-distortion faces, deviation from linearity was not significant.

The results are clear-cut, and a number of points can be made. Thatcherized faces, in which spatial-relational information was manipulated for the explicit purpose of producing bizarreness, were perceived as less bizarre with increasing departures from upright. Most important, the linear function relating orientation and bizarreness had a significant discontinuity that was not present in the functions for component-distortion and unaltered faces. These results provide strong support for the view that there is a qualitative difference in the processing of upright and inverted faces (Carey & Diamond, 1977). This view suggests that at some point as a face's orientation departs further from upright and approaches 180°, a shift in processing strategy must take place. The present results suggest that this shift occurs between 90° and 120° (see Fig. 2).

The absence of a discontinuity in the functions for component-distortion and unaltered faces suggests that spatial-relational information did not play a role in any perceived bizarreness in these two conditions. However, there was an effect of orientation, with both component-distortion faces and unaltered faces being perceived as more bizarre with increasing departures from upright. This suggests that perceived changes in bizarreness derived from component properties also occur following changes in orientation. Parks, Coss, and Coss (1985) reported that when viewed in isolation, an inverted upturned mouth was judged more unpleasant than its upright counterpart. The present results suggest that when individual facial components are presented in the context of a face, they are also perceived as more bizarre as they approach upside down (see also Rakover & Teucher, 1997, for results of a recognition task). The fact that this small linear effect was found for both unaltered and component-distortion faces indicates that this effect of orientation occurs independently of the overall level of perceived bizarreness. It appears that regardless of the initial perception of bizarreness when a face is upright, perception of bizarreness based on component properties increases as the components are rotated further from the upright view.

Finally, it is apparent in Figure 2 that even after the shift in the function, bizarreness ratings continued to decrease linearly for Thatcherized faces. This was confirmed in an analysis of the ratings at orientations 120° and greater, $F(1, 23) = 48.70$, $p < .001$. If the postdiscontinuity portion of the function reflects processing of predominantly component properties (Diamond & Carey, 1986), then this decrease in bizarreness ratings can be considered in terms of the explanation for the linear effect of orientation for component-distortion and unaltered faces. That is, bizarreness based on perception of component properties is affected by inversion, so that upright components are viewed as normal, whereas inverted components are perceived as relatively bizarre. In the case of Thatcherized faces, the eyes and mouth become upright with rotation of the face to 180°, and any perception of bizarreness based on these components would decrease as the postdiscontinuity face is rotated closer to 180°. The one difficulty with this account is that the magnitude of the decrease in ratings for Thatcherized faces was greater than the corresponding increase for unaltered and component-distortion faces. We reserve discussion of this point for the General Discussion.

EXPERIMENT 2

In Experiment 1, the spatial-relational manipulation used not only resulted in a change to the relations among components, it also changed the topological structure of the face. In Experiment 2, we chose a spatial-relational manipulation that did not alter the top-bottom structure of components in the face. Spatial-relational information was altered by changing the internal spacing between the eyes as well as moving the eyes and mouth up or down (Bartlett & Searcy, 1993).

This manipulation (which we refer to as the spatial-distortion condition) allowed us to look at potential changes in the perception of component properties following manipulation of spatial-relational information. In addition to changing spatial-relational properties, moving the mouth downward might be expected to alter the perception of the component property "chin." Also, as well as being encoded in terms of spatial-relational information, the distance between the eyes might be encoded as a component property. If both spatial-relational and component properties are affected by changes in the relative position of facial components, then both properties could contribute to the perception of bizarreness in the spatial-distortion condition, and be affected differently by changes in orientation.

On the basis of the results of the previous experiment, we expected to see a discontinuity in the function relating orientation and bizarreness in the spatial-distortion condition. We argue that this discontinuity signals the shift from spatial-relational processing to largely component processing. In addition, if indeed the spatial distortions used had the effect of creating component properties that also appeared bizarre, then perception of bizarreness would be maintained to some degree following the shift to component processing. Therefore, we predicted that the overall postshift ratings in the spatial-distortion condition would be closer to the ratings observed for component-distortion faces than for unaltered faces.

Finally, because the orientation of all facial components was the same in all three conditions, the postdiscontinuity effect of orientation on bizarreness ratings for spatial-distortion faces was expected to mirror the effect found for component-distortion and unaltered faces, both in magnitude and in direction of effect. Based on the findings of Experiment 1, we predicted a relatively small increase in bizarreness ratings from approximately 120° onward. Such an effect would be in contrast to that observed for Thatcherized faces in Experiment 1, and would provide additional support for the claim that postdiscontinuity ratings reflect processing of predominantly component properties.

2. For each face version, contrast weights of −6, −5, −4, −3, −2, −1, 0, 1, 2, 3, 4, 5, and 6 were applied to the results for the 13 orientations in order to assess the linear effect of orientation.

Method

Twenty-four students (11 females) from the University of Otago participated for credit toward a course requirement. The stimuli and procedure were identical to those of Experiment 1 with one exception. The Thatcherized faces were replaced with two spatial-distortion conditions. The first type of distortion involved moving the eyes further apart and downward, and the mouth upward. For the second distortion, the eyes were moved upward and closer together, and the mouth was moved downward. Examples of the two types of distortion are shown in Figure 3. All participants received the unaltered and component-distortion versions of the faces, and one type of the spatial-distortion versions. Equal numbers of participants were randomly assigned to the two groups defined by type of spatial distortion.

Results and Discussion

Ratings for clockwise and counterclockwise rotations of equal magnitude were averaged. An initial analysis of the two types of spatial distortion did not reveal any significant effect of type or any interaction of type with orientation. Therefore, in subsequent analyses, the data from the two spatial-distortion conditions were not differentiated.

The mean ratings for unaltered, component-distortion, and spatial-distortion faces are shown in Figure 4. The ANOVA revealed a significant main effect of version, $F(2, 46) = 71.87$, $p < .001$, and an interaction between version and orientation, $F(24, 552) = 17.00$, $p < .001$. As is evident in Figure 4, bizarreness ratings increased linearly for both unaltered faces, $F(1, 23) = 31.73$, $p < .001$, and component-distortion faces, $F(1, 23) = 11.16$, $p < .01$. There was no significant deviation from linearity in either condition. For spatial-distortion faces, the linear effect of orientation was also significant, $F(1, 23) = 18.32$, $p < .01$, as was the deviation from linearity, $F(11, 253) = 20.96$, $p < .001$.

Important aspects of Experiment 1 were replicated using a different manipulation of spatial-relational information. First, the results for spatial-distortion faces showed a systematic decline of bizarreness ratings with increasing departures of the face from upright, with a significant discontinuity in the function occurring between 90° and 120°. And second, the results for unaltered and component-distortion faces showed a small linear increase in bizarreness ratings with increases in orientation.

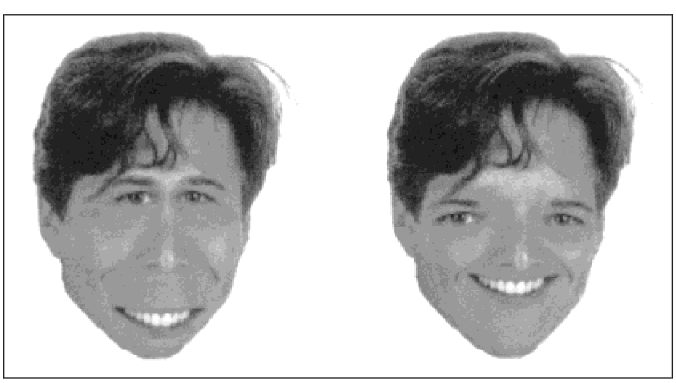

Fig. 3. Examples of the spatial-distortion faces used in Experiment 2. In the face on the left, the eyes were moved upward and closer together and the mouth moved downward. In the face on the right, the movement of the eyes and mouth was in the opposite direction. The original face can be seen in Figure 1.

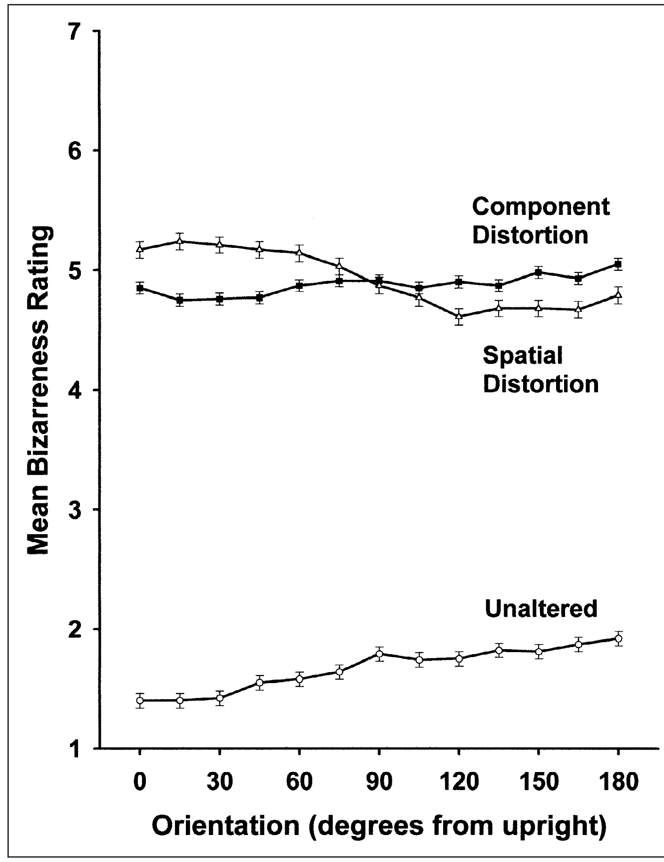

Fig. 4. Mean bizarreness ratings (1 = *most normal*, 7 = *most bizarre*) and standard errors for unaltered, spatial-distortion, and component-distortion faces as a function of orientation (degrees from upright) in Experiment 2. Error bars are within-subjects errors (Loftus & Masson, 1994) calculated separately for each function.

The condition in which spatial-relational information was manipulated revealed two additional notable findings. First, it is clear from Figure 4 that the perception of bizarreness was still maintained to a large extent following the shift in processing from spatial-relational information to component information: Bizarreness ratings were not much different from the ratings for component-distortion faces and markedly different from the ratings for unaltered faces over the same range of orientations.

Second, as predicted, the pattern of decreasing bizarreness ratings shown for spatial-distortion faces for orientations between 0° and 120° was reversed for orientations at 120° and above. For these post-discontinuity orientations, bizarreness ratings increased with increases in orientation, and the magnitude and direction of the effect did not appear to differ from that shown for unaltered and component-distortion faces over the same range of orientations. This pattern was confirmed in an analysis of orientations at 120° and above, which found a significant linear effect of orientation, $F(1, 23) = 4.46$, $p < .05$, that did not differ across the three conditions.

GENERAL DISCUSSION

Dissociations between face and object processing (even when objects are homogeneous) indicate different neural systems for faces and objects (see Farah, 1996, for a review). Moscovitch, Winocur, and Behrmann (1997) recently reported the case of C.K., a patient who had an impoverished component-based object recognition system but an intact face recognition system. Thus, how face- and object-processing systems may differ is in their reliance on the encoding of spatial-relational information.

The main support for this view is the disproportionate inversion effect found for faces, with inversion appearing to selectively disrupt the coding of spatial-relational properties (Bartlett & Searcy, 1993; Rhodes et al., 1993; Searcy & Bartlett, 1996). The present work adds significantly to that evidence. When spatial-relational information was altered by inverting the eyes and nose in a face or by changing the relative positions of these components, a discontinuity in the function relating orientation and bizarreness was found. For both types of spatial-relational change, the apparent shift in processing mode occurred between 90° and 120°. This is the first recorded estimate of the range of sensitivity of a spatial-relational encoding mechanism in normal face perception. It is of interest to compare our estimate with the results of some preliminary testing with patient C.K. Moscovitch et al. (1997) reported that C.K. could not recognize famous faces at orientations beyond 45° to 90° from upright. Although our estimate is not entirely coincident with the results reported by Moscovitch et al., the similarity between the two measures is encouraging if one takes into account the likelihood of individual differences.

The results also suggest that the effect of changes to spatial-relational properties is not necessarily restricted to the encoding of spatial-relational information. In Experiment 2, bizarreness ratings for spatial-distortion faces in the postdiscontinuity portion of the curve were maintained at a level comparable to that obtained for component-distortion faces. From this we conclude that both component and spatial-relational properties are affected when changes are made to the relative positions of components. This conclusion hinges on acceptance of the assumption that postshift ratings predominantly reflect encoding of component properties. We suggest that there is support for this assumption. In both experiments, orientation effects continued to be observed following the discontinuity in the function. Whether or not bizarreness ratings decreased (Experiment 1) or increased (Experiment 2) was predicted precisely by the orientation of the components rather than the face, and was independent of the direction of the prediscontinuity orientation effect, which occurred at orientations at which both spatial-relational and component information could be easily encoded. Additionally, the size of the postdiscontinuity orientation effect for spatial-distortion faces paralleled the orientation effect found for unaltered and component-distortion faces.

The results for Thatcherized faces suggest that encoding of component properties in the postdiscontinuity portion of the curve may not be the whole story. With Thatcherized faces, unlike the spatial-distortion faces in Experiment 2, the magnitude of the postdiscontinuity orientation effect was greater than would be predicted by the functions for component-distortion and unaltered faces over the same range of orientations. This raises the possibility that some form of non-component-based information is available for processing following the observed discontinuity, at least for Thatcherized faces. The fact that this magnitude difference was found only for Thatcherized faces suggests that the processing of faces in which the topological structure of the face is violated may represent a special case of face perception. Further work is required to resolve this issue. Regardless of the reasons for the larger postdiscontinuity effect of orientation for Thatcherized faces, the main finding is clear: The encoding of spatial-relational information is disproportionately impaired as the orientation of a face deviates from upright, reflecting a qualitative difference in the processing of upright and inverted faces.

Acknowledgments—This research was supported by an Otago Research Grant. We are grateful to David La Rooy for his assistance with stimulus preparation. We thank Jim Enns, Paddy McMullen, and an anonymous reviewer for their comments.

REFERENCES

Bartlett, J.C., & Searcy, J. (1993). Inversion and configuration of faces. *Cognitive Psychology*, *25*, 281–316.
Carey, S., & Diamond, R. (1977). From piecemeal to configurational representation of faces. *Science*, *195*, 312–314.
Diamond, R., & Carey, S. (1986). Why faces are and are not special: An effect of expertise. *Journal of Experimental Psychology: General*, *115*, 107–117.
Farah, M.J. (1996). Is face recognition special: Evidence from neuropsychology. *Behavioral Brain Research*, *76*, 181–189.
Loftus, G.R., & Masson, M.E.J. (1994). Using confidence intervals in within-subjects designs. *Psychonomic Bulletin & Review*, *1*, 476–490.
Moscovitch, M., Winocur, G., & Behrmann, M. (1997). What is special about face recognition? Nineteen experiments on a person with visual object agnosia and dyslexia but normal face recognition. *Journal of Cognitive Neuroscience*, *9*, 555–604.
Parks, T.E., Coss, R.G., & Coss, C.S. (1985). Thatcher and the Cheshire cat: Context and the processing of facial features. *Perception*, *14*, 747–754.
Rakover, S.S., & Teucher, B. (1997). Facial inversion effects: Parts and whole relationships. *Perception & Psychophysics*, *59*, 752–761.
Rhodes, G., Brake, S., & Atkinson, A.P. (1993). What's lost in inverted faces? *Cognition*, *47*, 25–57.
Schneider, W. (1988). Micro Experimental Psychology: An integrated system for IBM PC compatibles. *Behavior Research Methods, Instruments, & Computers*, *20*, 206–217.
Searcy, J.H., & Bartlett, J.C. (1996). Inversion and processing of component and spatial-relational information in faces. *Journal of Experimental Psychology: Human Perception and Performance*, *22*, 904–915.
Thompson, P. (1980). Margaret Thatcher: A new illusion. *Perception*, *9*, 483–484.
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. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. *The Quarterly Journal of Experimental Psychology*, *43*, 161–204.
Valentine, T., & Bruce, V. (1988). Mental rotation of faces. *Memory & Cognition*, *16*, 556–566.

(Received 9/24/99; Revision accepted 4/27/00)