COGNITIVE PSYCHOLOGY 25, 281–316 (1993)

Inversion and Configuration of Faces

JAMES C. BARTLETT AND JEAN SEARCY

Program in Human Development and Communication Sciences, The University of Texas at Dallas

If the mouth and eyes of a face are inverted, the altered construction appears grotesque when upright, but not when upside-down. Three studies of this "Thatcher illusion" employed faces that were grotesque when upright because: (a) their eyes and mouths had been *inverted* ("Thatcherized" faces), (b) their eyes and mouths had been *moved* (spatially distorted faces), or (c) they had grotesque posed expressions. Inversion reduced the apparent grotesqueness of *both* Thatcherized *and* spatially distorted faces, but *not* grotesque-expression faces. Moreover, Thatcherized and distorted faces, although not grotesque-expression faces, were judged as more similar to normal, smiling faces when face-pairs were inverted than when they were upright. Similarity ratings to inverted face-pairs were correlated with latencies of response to these pairs in a task that encouraged attention to components (e.g., mouths, eyes) rather than wholistic properties. Similarity ratings to upright face-pairs showed no such correlation, and this and other findings suggested that although similarity ratings to upright faces are based on wholistic information, similarity ratings to inverted faces are based largely on components. The Thatcher illusion reflects a disruption of encoding of wholistic information when faces are inverted.

INVERSION AND CONFIGURATION OF FACES

How does orientation affect perception of a form? Rock (1973) marshalled evidence that two factors are at work: an assignment-of-directions factor and a retinal factor. The former pertains to the fact that many objects look different when rotated or inverted; recognition that such objects have been previously seen is impaired if their orientation is changed from first viewing. Rock's account of this phenomenon is that the internal encoding or "description" of a figure is performed with respect to a reference frame that assigns spatial directions including "top," "bottom," and "sides." A figure's orientation in the physical environment can alter assignment of directions, and, therefore, can alter the figure's encoding.

Rock's (1973) second factor, the retinal factor, is revealed with certain stimuli that are encountered in one orientation much more frequently than

We thank Herve Abdi, Alice O'Toole, Vickie Bruce, Irvin Rock, Jim Tanaka, and Douglas Medin for valuable comments on this research. Correspondence should be addressed to James Bartlett, School of Human Development, Box 830688, Richardson, TX, 75083-0688. This research was supported by NIA Grant RO1-AG07798 to James Bartlett.

0010-0285/93 $9.00

in others. Some such "monooriented" stimuli are encoded much better when presented right-side-up (i.e., in their usual orientation) than when presented upside-down. Words in cursive script and pictures of faces are two examples. Both types of stimuli are poorly encoded if they are inverted—or if the viewer is inverted—such that the assigned directions of top and bottom are reversed from the viewer's egocentric up and down.

A puzzle concerning the retinal factor is the discovery by Yin (1969) that the effects of inversion are much stronger with faces than other monooriented stimuli including pictures of airplanes, houses, stick figures, period costumes, landscapes, and dogs (Yin, 1969; Scapinello & Yarmey, 1970; Diamond & Carey, 1986). This provocative outcome led Yin to propose a "special factor" that is orientation-specific and unique to face stimuli. This idea that faces are "special" (or "unique," see Ellis & Young, 1989) can be questioned based on Rock's (1973) observation that cursive-script processing is disrupted by inversion. It is also opposed to the recent demonstration that effects of inversion on recognition memory can be as strong with dogs as faces, so long as the subjects have expert knowledge of the dogs (Diamond & Carey, 1986).

In light of the preceding results, it appears that Rock's (1973) retinal factor arises in the course of perceptual learning with some (although perhaps not all) monooriented stimuli. Apparently, the skills one acquires in processing such stimuli cannot be recruited, or are poorly recruited, when the stimuli are inverted. This perceptual problem may be solved in some cases by internal correction of a form's orientation (e.g., through mental rotation). However, with complex stimuli including faces such correction is impossible or highly error-prone (Rock, 1973). Hence, the perceptual encoding is impaired in some way.

Accepting this conclusion as a working hypothesis, we are led to certain questions of central importance to theories of perception and recognition memory: What are the effects of perceptual learning and why are they restricted, at least in some cases, to a small range of orientations? What sort of information is poorly encoded when some monooriented stimuli are presented upside-down? To answer such questions, we need to know more about the processing of stimuli that show effects of inversion. What we need to know most is if processing such stimuli is simply impaired by inversion, or if it is altered in a qualitative way. The present research was addressed to this issue.

The question of whether inversion effects are qualitative or quantitative has been addressed with more vigor in the domain of face processing than in other research areas. Valentine (1988) reviewed the face-processing research, and argued that although inversion increases the error of encoding (see Valentine, 1991), there is no persuasive evidence that it causes a shift from one type of processing to another. Studies by Tanaka and Farah (1991) and Valentine and Bruce (1988) lend support to this conclusion.

Despite Valentine's (1988) arguments, several findings suggest that inversion affects the type of information encoded from a face. Sergent (1984) found that latencies for "different" judgments in simultaneous comparisons of faces showed interactive effects of different feature-changes, but only when faces were presented upright. However, Takane and Sergent (1983) found an interactive pattern with inverted faces also. Bruce, Doyle, Dench, and Burton (1991) reported that inversion impaired learning of prototype configurations using computer-drawn faces and houses as stimuli. However, there was no examination of inversion effects with other information besides configurations. Stronger evidence for a qualitative effect of inversion comes from Young, Hellawell, and Hay (1987), who examined latencies for naming parts of famous faces (e.g., top halves of these faces) when presented with complementary parts of other famous faces (e.g., bottom halves). Naming latencies were longer when the parts were aligned than when they were misaligned, suggesting that wholistic or configural processing can interfere with naming one part of a face. Alignment made no difference when the stimuli were inverted, suggesting a qualitative effect of inversion whereby wholistic or configural information is selectively disrupted. However, it is possible to argue, as Valentine (1988) has, that inversion is *generally* disruptive to face processing, adding "error" to configural information as well as feature or component information *without* causing a shift in face processing from one type to another (see also Valentine & Bruce, 1988).

The Margaret Thatcher Illusion

The most dramatic evidence that inversion can cause a shift in face processing is Thompson's (1980) Margaret Thatcher illusion, which is shown in Fig. 1. The illusion is created by inverting the mouth and both eyes of an otherwise upright face (not just Mrs. Thatcher's). The altered face appears highly grotesque, but the grotesqueness is absent when the whole construction is inverted.

The Thatcher illusion confirms prior evidence that inversion interferes with the processing of faces and is a vivid illustration of Rock's (1973) retinal factor. More important, the illusion appears to compel the conclusion that the information encoded from an upright Thatcher stimulus differs in a fundamental way from that encoded from an inverted Thatcher stimulus. Hence, the illusion constitutes a clear demonstration that inversion can qualitatively alter face processing. But what is the nature of this qualitative effect? The answer to this question should have important implications for a proper understanding of Rock's (1973) retinal factor,

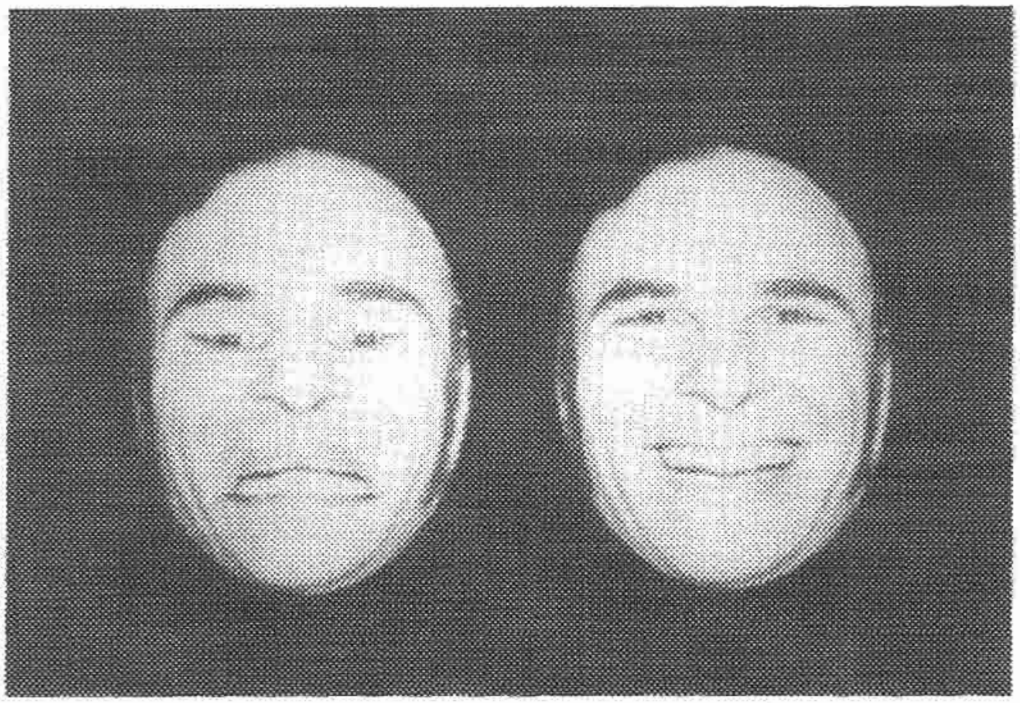

FIG. 1. An example of a normal and Thatcherized face.

and for the related issue of how perceptual learning affects recognition of the visual information we encounter in life.

Theoretical Accounts of the Thatcher Illusion

According to Valentine (1988), the Thatcher illusion demonstrates ". . . that inversion of a face can indeed have a dramatic effect upon the perception of expression . . ." (p. 484). Indeed, a simple account of the Thatcher illusion holds that (a) the grotesque appearance of a Thatcherized face is due to its expression, (b) inversion impairs the encoding of expression, and, therefore, (c) inversion disrupts the perception of grotesqueness of a Thatcherized face. This expression hypothesis is in line with Yin's (1970) suggestion that the type of facial processing most affected by inversion might pertain to facial expression. It also converges with statements by Köhler (1940), as well as many painters (Parks, Coss, & Coss, 1985), that inverted faces lack expression, and it fits Dolezal's (1982) observation that inverting prisms disrupted his perception of facial expressions (albeit only subtle expressions; he had no trouble recognizing overt smiles). Finally, an expression hypothesis is consistent with experimental and neuropsychological evidence that facial-expression encoding is dissociable from facial-identity encoding (see, e.g., Bruce, 1988; Valentine, 1988). Because of this evidence, it is theoretically plausible that a module or system for expression encoding might show stronger inversion effects than modules or systems for recognition of identity. However, a problem for the expression hypothesis is posed by a study of mental rotation of faces by Valentine and Bruce (1988). These authors found that the effects of orientation of faces in an expression-judgment task were weaker than those in familiarity-judgment task, a finding opposed to the expression hypothesis.

A second account of the Thatcher illusion is somewhat more complex. Parks and his colleagues (Parks et al., 1985; Parks, 1983) began with the assumption drawn from Rock's (1973) theory that the effect of inverting a monooriented stimulus reflects the interaction of two frames of reference. One is an object-centered frame conveyed by the structure of the stimulus itself as well as by prior learning. The other is a nonobject-centered frame based on the external environment, or, alternatively, on the viewer's egocentric sense of up versus down based on retinal coordinates. In any case, when a monooriented stimulus is presented upside-down, the assignment of directions based on the object-centered frame is opposite to that based on the nonobject-centered frame, and this weakens the importance of the object-centered frame (Parks, 1983).[1]

Perhaps any account of inversion effects must consider interactions between reference frames. The key claim made by Parks et al. (1985) is that the reference frames involved in perceiving a face affect the appearance of the *features* of this face. Thus, the assignment of top to one part of a mouth (i.e., the upper versus the lower lip) will affect whether the mouth is seen as representing a smile or a grotesque, biting, grimace. A similar argument is made for the eyes which might appear open and upward gazing or strangely squinted and downward gazing depending upon the assignment of top. When Thatcherized faces are presented upright, the locations of tops and bottoms of features are reinforced by both reference frames, and, therefore, the grotesqueness of the eyes and mouth are stable. However, when Thatcherized faces are presented inverted, the locations of tops and bottoms of features vary with the reference frame. Hence, the assignment of directions to features is unstable and so the impression of grotesqueness does not occur.

In a test of the frame-of-reference hypothesis, Parks et al. (1985) had subjects rate the pleasantness of drawings of mouths in upright versus

[1] Although Parks et al. (1985) argued that an environmental reference frame helped in determining the top of facial features, Rock (1988) points out that a retinal reference frame is probably more important (e.g., one can experience the Thatcher illusion by leaving the picture in upright orientation and observing the face normally versus with one's head inverted).

inverted orientation when eyes were placed above them and when eyes were placed below. In agreement with predictions, they found that an upright mouth was judged as more pleasant when the location of the eyes vis-a-vis the mouth supported an assignment of top to the mouth that would indicate a smile (see also Valentine & Bruce, 1985). However, the authors also found that the judged pleasantness of a mouth was affected by the distance between the eyes and mouth, which would appear to be a configural property. Moreover, ratings of pleasantness of the *eyes* paralleled those of the *mouths* that they accompanied, despite the fact that the eyes were vertically symmetric. These findings suggest that, despite task instructions, subjects rated pleasantness of eyes and mouths at least partly on the basis of configural information.[2]

A third account of the Thatcher illusion is based on a family of distinctions between (a) component (or piecemeal feature) and configural information (Bradshaw & Nettleton, 1981; Carey & Diamond, 1977; Ross-Kossak & Turkewitz, 1984; Sergent, 1984), (b) first-order and second-order information (Diamond & Carey, 1986; Rhodes, 1988), (c) relations and attributes (Goldstone, Medin, & Gentner, 1991), (d) global and local information (Navon, 1977), and (f) holistic and analytic information (Kemler, 1983). The terms we will use in the present report are taken from Garner (1978), who distinguished "component" from "wholistic" stimulus characteristics, discussing two different types of components and three different types of wholistic properties. The two types of components—features and dimensions—differ in that the former are either present or absent (e.g., a mole on a face), whereas the latter can exist at two or more mutually exclusive levels (e.g., a nose of a certain shape and size). The three types of wholistic properties are simple wholes, template-based properties and configurations. Simple wholes are just sums of components, and may represent a degenerate concept (Garner, 1978). Template-based properties pertain to the coincidence of stimuli with modal or average forms (e.g., with prototypes or schemas). Configurations are emergent properties that arise from relations among components. The bulk of the face-processing literature has blurred the distinction between features and components, as well as that between template-based properties and configural properties (it is not clear to us that simple wholes have been considered). Hence, the distinction as it is currently drawn can be captured by the terms "component" and "wholis-

[2] To rule out this explanation, Parks et al. (1985) performed a second experiment in which the eyes were removed from the visual field before subjects made their judgments of mouths. They replicated the findings of Experiment 1, but *only* when mouths were in inverted orientation, *not* when they were in upright orientation. Thus, the use of configural information in Experiment 1 was not ruled out definitively.

tic." We subsequently consider whether and how the distinction might be refined.[3]

The Thatcher illusion can be explained by the view that wholistic and component information is processed through different perceptual strategies or modes, and that inversion disrupts the wholistic mode much more than the componential mode (see Carey & Diamond, 1977; Sergent, 1984). Diamond and Carey (1986) put this dual-mode view as follows: "Thompson's (1980) 'Thatcher illusion' provides a striking demonstration that spatial relations among features crucial in the perception of upright faces are not apparent when faces are upside down" (p. 107). Rock (1988) argues similarly that "In viewing an inverted face it is difficult to cope with the relationship of the parts," and that ". . . in Thompson's [inverted] picture, the relationship of the inverted features to the rest of the face is not adequately perceived and thus the grotesqueness is not perceived either" (p. 817).

Just why the wholistic processing mode is affected by inversion has not been made clear. However, Goldstein and Chance (1980) have proposed the existence of a face schema, built up in the course of perceptual learning, that enhances the encoding of upright faces but not that of inverted faces. These authors say little about the nature of the schema, but Rhodes, Brennan, and Carey (1987) have developed the idea that faces are stored in terms of their differences from a "norm face," which captures major contours of a modal or average face in a template-like format (i.e., a two-dimensional array, see also Benson & Perrett, 1991). Similarly, Yuille (1991) has proposed "deformable templates" that have adjustable parameters representing allowable deformations of a schematic face (global templates, see Fischler & Elschlager, 1973), and individual facial features (feature templates). A face is encoded through (a) comparison with the templates, (b) adjustment of the parameters of these templates, and (c) determination of the goodness-of-fit between the adjusted templates and the face.

Neither Yuille (1991) nor Rhodes et al. (1987) discussed the causes of retinal inversion effects. However, both norm faces and deformable templates are visuospatial representations with well defined tops and bottoms. To encode inverted faces using such representations would require a process of mental rotation which preserves spatial relations in a highly precise way. Mental rotation with high resolution might be impossible with inverted faces (Rock, 1973). If so, template-based processing might be precluded with such faces.

[3] We follow Garner (1978) in our spelling of "wholistic," and note that Rhodes' (1988) definition of "second-order information" is closer to that of wholistic information than is Diamond and Carey's (1986) earlier definition.

In summary, the Thatcher illusion can be explained on the assumption that upright Thatcher faces have poor fits to templates for norm faces. These poor fits are the source of the impression of grotesqueness. Since inversion is disruptive to template-based processing, poor fits to templates cannot be detected with inverted Thatcher faces. Hence, inverted Thatcher faces do not appear grotesque.

EXPERIMENT 1

Toward the general goal of improving our knowledge of Rock's (1973) retinal factor, we tested the dual-mode hypothesis against the frame-of-reference and expression hypotheses as theoretical accounts of the Thatcher illusion. We manipulated photographs of 14 smiling faces to produce three types of grotesquely transformed faces: (a) "Thatcherized" images in which the mouths and eyes were inverted, (b) "type-1" distorted images in which the mouths were lowered and the eyes were moved up and closer together, and (c) "type-2" distorted images in which the mouths were moved up and the eyes were moved down and farther apart. We also selected neutral- and grotesque-expression versions of each face stimulus, examples of which are shown in Fig. 2.

Our question was how inverting face stimuli would affect their judged grotesqueness. If inversion disrupts expression encoding (the expression hypothesis), it should reduce the grotesqueness of grotesque-expression faces as well as that of Thatcherized faces. It might *not* reduce grotesqueness of the type-1 and type-2 distorted faces whose grotesqueness is due to the locations of their features and arguably not to facial expression. A different prediction follows from the reference-frame hypothesis. According to this view, the apparent grotesqueness of Thatcherized faces reflects an unusual assignment of directions to their eyes and mouths, and inversion simply weakens this assignment of directions. It follows from this reasoning that: (a) inversion should reduce the grotesqueness of Thatcherized faces, but (b) it should *not* reduce grotesqueness of the other types of faces which are grotesque for other reasons than assignment of directions.

Still another prediction follows if inversion disrupts processing of wholistic information (the dual-mode hypothesis). In this case, inversion should reduce grotesqueness of type-1 and type-2 distorted faces as well as grotesqueness of Thatcherized faces, although it might *not* affect grotesqueness of grotesque-expression faces. The reasoning here is that both distorted faces and Thatcherized faces have been altered with respect to wholistic properties (i.e., configural information and/or deviations from templates). This is not obviously true of the grotesque-expression faces.

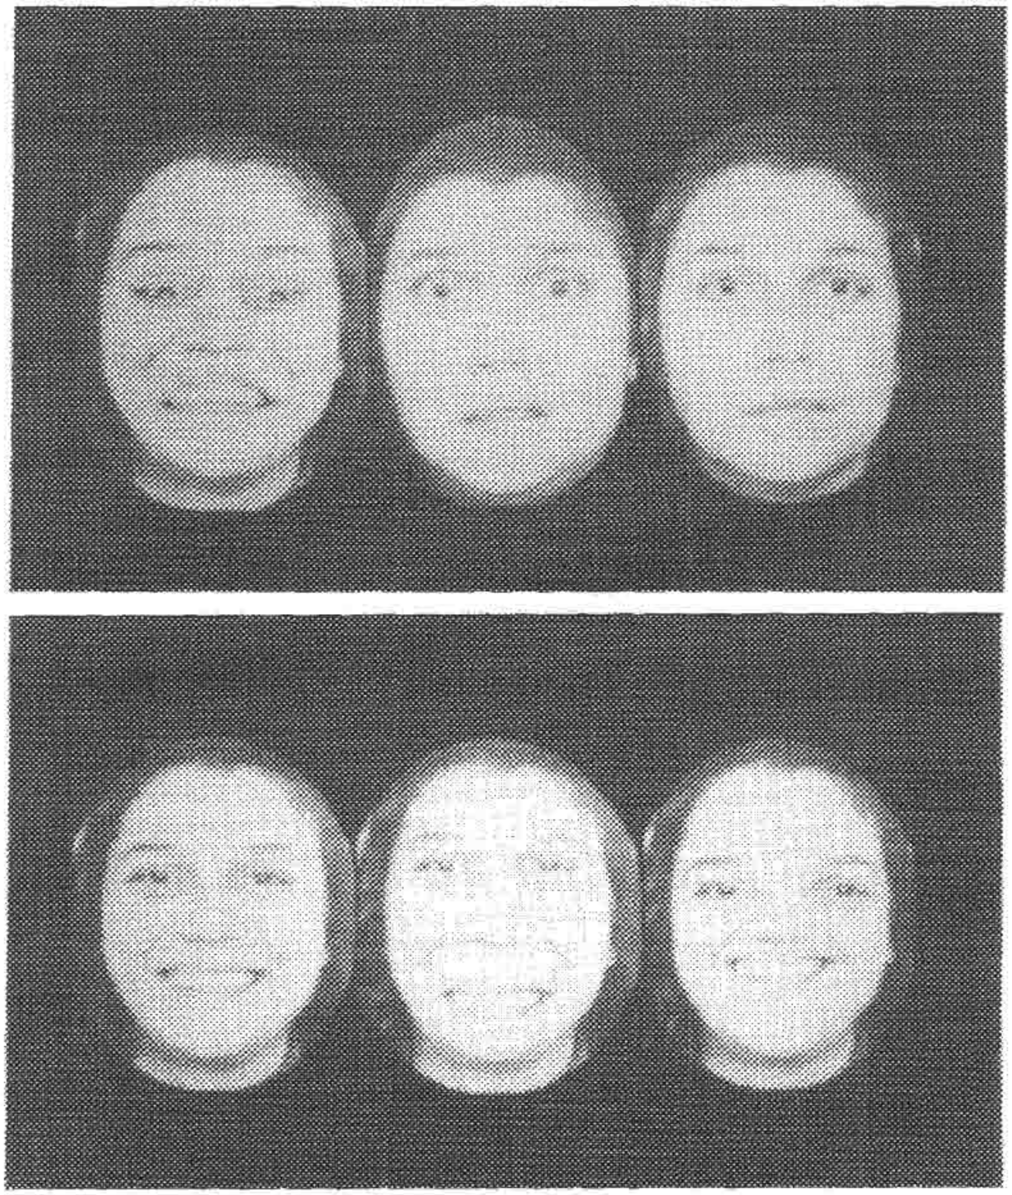

FIG. 2. Examples of a Thatcherized face together with a grotesque-expression and neutral-expression version of this face (top row), and a normal, type-1 distorted and type-2 distorted version (bottom row).

Method

Subjects and procedure. The subjects were 20 undergraduates (70% female; mean age, 27 years) at the University of Texas at Dallas who participated as one means of fulfilling a psychology course requirement. Groups of 1 to 4 subjects were seated approximately 5 feet from the 21-inch monitor on which all stimuli were displayed. They were told that they would see a set of faces some of which had been distorted, and half of which were upside down. The task was to rate each face for "grotesqueness" using a seven-point scale (7, most grotesque). The stimuli were presented in a randomized sequence for 5 s each, with a 2-s interstimulus interval.

Design. The 12 subjects in the expression group viewed a stimulus-set composed of the Thatcherized, smiling-expression, neutral-expression, and grotesque-expression versions of each face. The 8 subjects in the distortion group viewed an overlapping stimulus set composed of the Thatcherized, smiling-expression, type-1-distorted and type-2-distorted versions of each face (see Fig. 1). Since each stimulus-set included four different versions of each of 14 faces, both presentation sequences contained 56 upright face-stimuli and 56 inverted face-stimuli, presented in a randomized sequence of 112 stimuli.

Materials. The materials were derived from Ekman's (1975) pictures of 6 male and 8 female faces of models posing various expressions. For each of the 14 models, we chose a picture showing a normal smiling expression, a picture showing a neutral expression, and a picture showing what the experimenters judged to be the most grotesque expression of all the photographs of that model. The pictures were digitized using a Tamron Fotovix video processor and a 286-type computer equipped with a 16-bit Targa board and Truevision image-processing software (TIPS).

The TIPS software was used to place each face in an oval frame which reduced but did not entirely eliminate hair and background information (see Fig. 2). The pictures were also altered in size to produce an image that was approximately 6 × 5 inches when displayed on a 21-inch monitor. In addition, we transformed the pleasantly smiling picture of each face to produce (a) a Thatcherized version by inverting the mouth and both eyes, (b) a spatially distorted version by moving the mouth downward and the eyes upward and closer together (type 1), and (c) a second spatially distorted version by moving the mouth upwards and the eyes down and further apart (type 2). These transformations were made by inverting or moving rectangular face-regions that closely framed the relevant features and then obscuring the edges of these face-regions through the blending operation in TIPS. We did not precisely control the size of the face-regions that were inverted or moved, or the distances traversed when the regions were moved—our goal in making each transformation was to produce a face that looked grotesque when upright. The images were transferred to optical disc using a Panasonic 2026F optical disc player/recorder upgraded to SVHS quality. The optical disc unit, controlled by the computer, was used for presentation of the faces on a 21-inch SVHS video monitor.

Results

Looking first at the expression condition, Table 1 shows that inversion slightly increased grotesqueness ratings for smiling and neutral-expression faces, while dramatically reducing grotesqueness ratings for Thatcherized faces. Turning to the distortion condition, the effects for smiling and Thatcherized faces were replicated nicely, and the ratings of the spatially distorted faces behaved as predicted by the dual-mode hypothesis: both types of distorted faces were judged as significantly less grotesque when presented upside down.

The ratings were subjected to two analyses of variance (ANOVAs), one for the expression condition and the other for the distortion condition. At the conventional alpha level of .05, both ANOVAs showed reliable effects of orientation, $F(1,11) = 18.1$, $MS_e = .25$ in the expression condition, and $F(1,7) = 49.8$, $MS_e = .43$, in the distortion condition. They also showed

TABLE 1
Mean "Grotesqueness" Ratings (and Standard Deviations) for Upright and Inverted Faces in the Expression and Distortion Conditions of Experiment 1

	Item-type in expression condition			
Orientation	Smiling	Thatcher	Neutral	Grotesque
Upright	2.50 (.85)	6.73 (.27)	3.12 (.73)	4.72 (.56)
Inverted	3.12 (.93)	3.89 (.82)	3.54 (.77)	4.78 (.69)
Difference	−0.62 *	2.84 *	−0.42 *	−0.06
	Item-type in distortion condition			
	Smiling	Thatcher	Distortion 1	Distortion 2
Upright	2.12 (.92)	6.18 (.64)	5.56 (.52)	5.40 (.45)
Inverted	2.46 (.77)	3.42 (.93)	4.63 (.44)	4.14 (.69)
Difference	−0.34 *	2.76 *	0.93 *	1.26 *

Note. *n*s were 12 and 8 in the expression and distortion conditions, respectively. Grotesqueness ratings ranged from 1 (least) to 7 (most).
* The inversion effect gave $p < .05$ by F test.

reliable item-type effects and reliable item-type × orientation interactions, $Fs(3,33) = 65.0$ and 97.5, $MS_e = .508$ and $.161$, respectively, in the expression condition, and $Fs(3,21) = 59.5$ and 55.7, $MS_e = .462$ and $.116$, respectively, in the distortion condition.

The data from the distortion condition suggested that inversion reduced grotesqueness ratings for Thatcherized faces *more* than for spatially distorted faces. Indeed, the item-type × orientation interaction was statistically reliable in a comparison of Thatcherized, type-1, and type-2 distorted faces, $F(2,14) = 35.9$, $MS_e = .106$, and it remained reliable in comparisons of Thatcherized with type-1 distorted faces, $F(1,7) = 40.8$, $MS_e = .164$, and with type-2 distorted faces, $F(1,7) = 51.6$, $MS_e = .087$. Despite a hint that inversion had stronger effects on type-2 distorted faces than type-1 distorted faces, an item-type × orientation interaction was not supported in an ANOVA comparing these two item-types.

To assess the generality of our findings across items, we computed the mean grotesqueness rating for each upright and inverted face presented to each group and plotted the grotesqueness of inverted faces against that of upright faces. As shown in Fig. 3, all 28 Thatcherized faces, and 26 of the 28 (93%) type-1- and type-2-distorted faces, were judged less grotesque when inverted than when upright. In contrast, only 14 of the 56 (25%) unaltered faces (smiling-, neutral-, and grotesque-expression items) were judged less grotesque when viewed upside down. Note that the points for Thatcherized items were generally further below the positive diagonal than were the points for distorted items. This supports the conclusion that

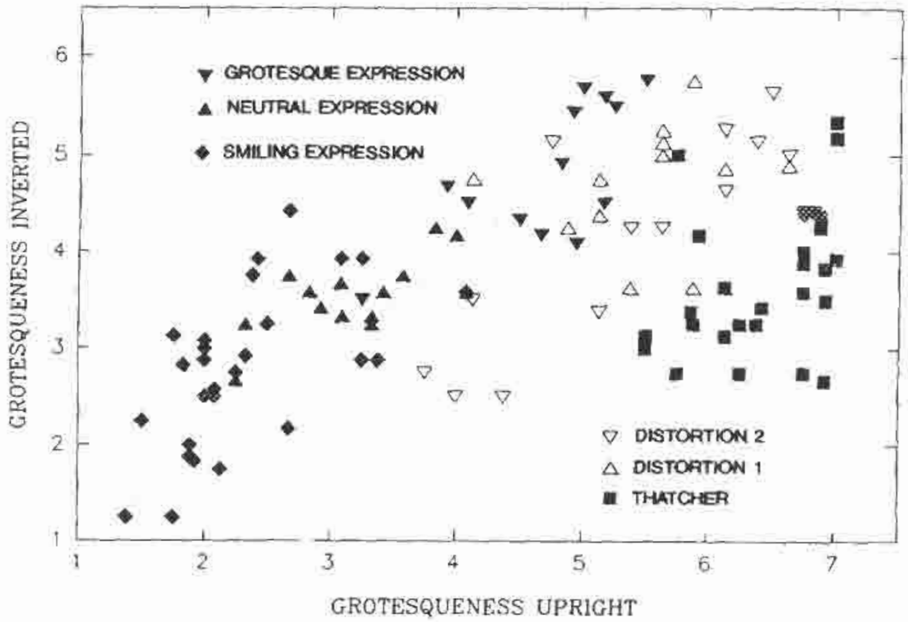

FIG. 3. The effect of inversion on grotesqueness ratings for individual faces in Experiment 1. Each smiling and Thatcherized face is represented by two points, one for the expression condition and one for the distortion condition. Superimpositions occurred at X = 1.88, Y = 2.0 (two smiling faces); X = 4.09, Y = 4.50 (two grotesque faces); X = 6.88, Y = 4.38 (two Thatcherized faces); and X = 6.13, Y = 3.63, and X = 5.75 and Y = 5.00 (a distortion-1 and a Thatcherized face in both cases).

inversion reduced grotesqueness of Thatcherized faces more than that of spatially distorted faces, though it significantly reduced the grotesqueness of both.

The effects of inversion on Thatcherized and spatially distorted faces cannot be attributed to their high level of grotesqueness when presented right-side up. Figure 3 shows that several grotesque-expression faces (inverted, filled triangles) were judged almost as grotesque as some Thatcherized faces (black squares), and more grotesque than some spatially distorted faces (upright- and inverted-empty triangles), when presented right-side up. Yet, these grotesque-expression faces differed from the Thatcherized and spatially distorted faces in their high grotesqueness ratings when presented upside down.

Discussion

Our aim in this study was to test three accounts of the Thatcher illusion. The major claim of the expression hypothesis was that grotesque expressions are less grotesque when inverted. The reference-frame hypothesis held that the grotesqueness of Thatcherized faces reflects the grotesqueness of their individual components (i.e., their inverted mouths and eyes) when "top" is assigned to what is normally the "bottom" of these components (Parks et al., 1985). Such an unusual assignment of directions is supported by two reference-frames when Thatcher faces are upright, but these reference frames conflict—and the assignment of directions is therefore unstable—when Thatcher faces are inverted. According to the dual-mode account, inversion disrupts processing of wholistic information (i.e., configural information or deviations from templates), which is the source of the grotesqueness of Thatcherized faces.

The expression hypothesis was firmly rejected, as the grotesque-expression faces were *not* judged as less grotesque when upside-down than right-side up. Indeed, the (unreliable) trend was in the opposite direction (Table 1, top). The frame-of-reference hypothesis failed to predict the effect of inversion with spatially distorted faces. The dual-mode hypothesis was strongly supported in that two types of spatially distorted faces were similar to Thatcherized faces in being judged as less grotesque when inverted than when upright. Since both types of spatially distorted faces were altered with respect to wholistic information, the finding supports the claim that wholistic encoding is impaired by inversion.

The finding that inversion of grotesque-expression faces did not reduce their judged grotesqueness supports the dual-mode hypothesis, but it is nonetheless surprising in view of painters' observations of the nonexpressive character of upside-down faces (Parks et al., 1985). However, it is possible to argue that inversion interferes with the perception of expressions without making those expressions seem any less grotesque. This possibility is addressed in Experiment 2.

A finding predicted by no hypothesis we considered was that although inversion affected grotesqueness ratings of both Thatcher faces and distorted faces, the effect with Thatcher faces was stronger. One possible account is that inversion disrupts wholistic encoding and *also* weakens subjects' reference-frames for encoding eyes and mouths. Both factors would be operative with Thatcherized faces, but only the factor of wholistic encoding would be operative with spatially distorted faces. An alternative and more parsimonious view is that inversion effects reflect template-based encoding, and not all types of wholistic information are equally dependent on template-based encoding. Specifically, orientation of features (altered in Thatcherized faces) might be more perceptually subtle than spatial-location of features (altered in distorted faces), and encoding of the former might be more dependent on templates for norm faces.

Still a third view draws on Rock's (1973) claim that if a monooriented stimulus is presented upside-down, some sort of perceptual correction takes place. If the stimulus is complex (e.g., if it is a face), this correction occurs separately for individual components. Thatcher faces constitute a special case because when such faces are viewed upside-down, their eyes and mouths are right-side up, and plausibly *not* subject to a correction operation. Since these components are pleasant when viewed right-side up, the inverted Thatcher face appears quite pleasant. In contrast, when distorted faces are viewed upside-down, all of their components (including eyes and mouth) would be subject to correction. They might retain some grotesqueness to the extent that correction preserves spatial relations.[4]

Regardless of the reasons why inversion effects were greater for Thatcherized faces than distorted faces, the main conclusion of this study is clear: Inversion impairs wholistic encoding. This is consistent with the dual-mode hypothesis which distinguishes two face-processing modes, one for wholistic information and the other for component information.

Experiment 1b. An arguable weakness of Experiment 1 is that no grotesque-expression face was judged as grotesque as any Thatcherized face with upright presentation (see Fig. 3). Thus, a critic might argue that grotesque-expression faces that are *truly* grotesque—just as grotesque as upright Thatcher faces—might be judged as less grotesque when presented upside down. To test this possibility, we videotaped each of seven individuals posing four different expressions, one neutral expression, one smiling expression, and two different grotesque expressions (they were coached to make their faces as ugly as they could). We captured images from the videotapes on computer hard disc using TIPS software, and constructed Thatcherized versions of the smiling and neutral-expression pictures. This produced a set of six pictures for each of the seven models (smiling, neutral, Thatcherized-smiling, Thatcherized-neutral, grotesque-expression 1, and grotesque-expression 2). We placed each of the 42 images in an oval frame that masked most of the hair and all clothing as in Experiment 1. These images were transferred to optical disc in upright as well as inverted orientation. The resulting set of 84 images was displayed to subjects via a 21-inch monitor in a randomized sequence at the rate of 7 s per image (stimulus duration, 5 s and interstimulus interval, 2 s). The nine subjects were from the same source used in Experiment 1. Their task was simply to rate each image for grotesqueness using a seven-point scale (7, most grotesque).

The results replicated those of Experiment 1, but, to our initial disappointment, the upright-grotesque-expression faces again received lower grotesqueness ratings than the upright-Thatcherized faces. Fortunately, we were able to identify 4 (of 14) Thatcher stimuli and 4 (of 14) grotesque-expression stimuli whose mean rated grotesqueness in the upright condition was approximately equal (Ms = 3.94 and 4.03, respectively, SDs =

[4] We thank Irvin Rock for suggesting this third account.

.73 and 1.38, respectively). Despite equivalent grotesqueness in the upright condition, the grotesque-expression stimuli were still grotesque when inverted ($M = 4.75$, $SD = 1.25$), whereas the Thatcher stimuli were greatly reduced in grotesqueness ($M = 2.75$, $SD = .71$). Two ANOVAs showed that the stimulus-type × orientation interaction was significant with subjects treated as the random factor, $F(1,8) = 20.5$, $MS_e = .402$, and also with items treated as the random factor, $F(1,6) = 32.5$, $MS_e = .113$ ($p < .002$ in both cases). No other effects reached conventional significance levels.

EXPERIMENT 2

One of our aims for Experiment 2 was to clarify the finding that the rated grotesqueness of grotesque-expression faces was *not* reduced by inversion in Experiment 1. Our preferred interpretation of this counterintuitive finding is that the processing of expression in the present study was impervious to inversion. However, an alternative view is that inversion can impair the processing of expression without making an expression seem any less grotesque (e.g., by making the expression seem less sincere or realistic). To more rigorously examine whether inversion can affect the processing of expression, we had subjects give ratings of the perceptual similarity between faces with grotesque and neutral expressions and these same faces with smiling expressions. If inversion impairs the perception of expression, it should increase subjects' ratings of perceptual similarity.

Another aim of Experiment 2 was to test the implication of the dual-mode hypothesis that the effects of inversion should depend not only on the nature of the stimuli, but also on the task: If a task evokes a wholistic mode of processing when faces are upright, inversion effects should be strong and *qualitative* (i.e., they should be stronger with Thatcher and distorted faces than with grotesque-expression faces). The reason is simply that the wholistic mode of processing usable with upright faces cannot be employed, or is employed with great difficulty, with upside-down faces. However, if a task evokes a componential mode of processing when faces are upright, inversion effects should be small and *quantitative* (i.e., any effects should be no greater for Thatcher and distorted faces than for grotesque-expression faces). The reasoning here is that componential processing is presumed to be possible regardless of orientation.

According to Sergent (1984), tasks of simultaneous paired-comparisons of faces often evoke a componential strategy, especially when subjects are under time pressure. If so, inversion should have minimal effects in such tasks, even if Thatcherized faces are used. To test this line of reasoning, the present experiment examined not only a similarity rating task that allowed for the use of configural processing but also a same-different-face judgment task that was designed to encourage only componential processing.

Both tasks included three types of face-pairs in which one item was always an unaltered, smiling face and the other was: (a) a Thatcherized version of this face, (b) a type-1 or type-2 distortion of this face, and (c) this face with a grotesque or neutral expression. If similarity ratings are based at least partly on configural information, inversion should increase the judged similarity of smiling-Thatcher pairs as well as smiling-distortion pairs. However, if same–different face judgments are based on components, inversion should *not* reduce the latency of correct same-face judgments in response to these same pairs.

Method

Subjects and materials. The 35 subjects were similar to those used in Experiment 1. The face stimuli were also the same as before, except that here they were presented as side-by-side pairs always comprising: (a) a normal, smiling, face next to (b) a changed-expression face, a Thatcherized face, or a spatially distorted face. Only same-face pairs (i.e., two different versions of the same person's face) were presented in the similarity-rating task, but different-face pairs (i.e., pairs in which faces of two persons were shown) were included in the same-different-face classification task. The different-face pairs were produced by pairing each of the 14 models with one of the other models, matching by gender, an attempting to make age and general appearance as similar as possible. Examples are shown in Figs. 4, 5, and 6.

Using the same computer graphics software and optical disc apparatus as in Experiment 1, all face-pairs were constructed such that each of the two faces were in 6 × 5-inch ovals with one appearing on the left side of the (21 inch) screen and the other appearing on the right side of the screen. Each pair of faces was presented four times: (1) in upright orientation with the normal, smiling version on the left side of the screen; (2) in upright orientation with this version on the right; (3) in inverted orientation with the normal, smiling version on the left side of the screen; and (4) in inverted orientation with this version on the right.

Design. Approximately half (18) of the subjects served in the expression condition, in which the presentation sequence included equal numbers of randomly intermixed smiling/neutral-expression pairs, smiling/grotesque-expression pairs, and smiling/Thatcher pairs. The remaining (17) subjects served in the distortion condition, in which the sequence included equal numbers of smiling/type-1 distorted pairs, smiling/type-2 distorted pairs, and smiling/Thatcher pairs.

All of the subjects performed a similarity rating task followed by a same–different classification task. The similarity ratings were made to 168 same-face pairs created through presenting each of the 14 normal, smiling faces with each of three types of comparison face in two spatial arrangements (normal face on the left versus right) and two orientations (upright versus inverted). The same–different classifications were made to these same 168 same-face pairs randomly intermixed with an additional 168 different-face pairs.

Procedure. Prior to the similarity-rating task, subjects were told that they would see a sequence of same-face pairs, and that they should respond to each pair with a judgment of how similar the left-hand stimulus was to the right-hand stimulus, using a seven-point scale on which 1 referred to pairs that were the most similar or hardest to tell apart and 7 referred

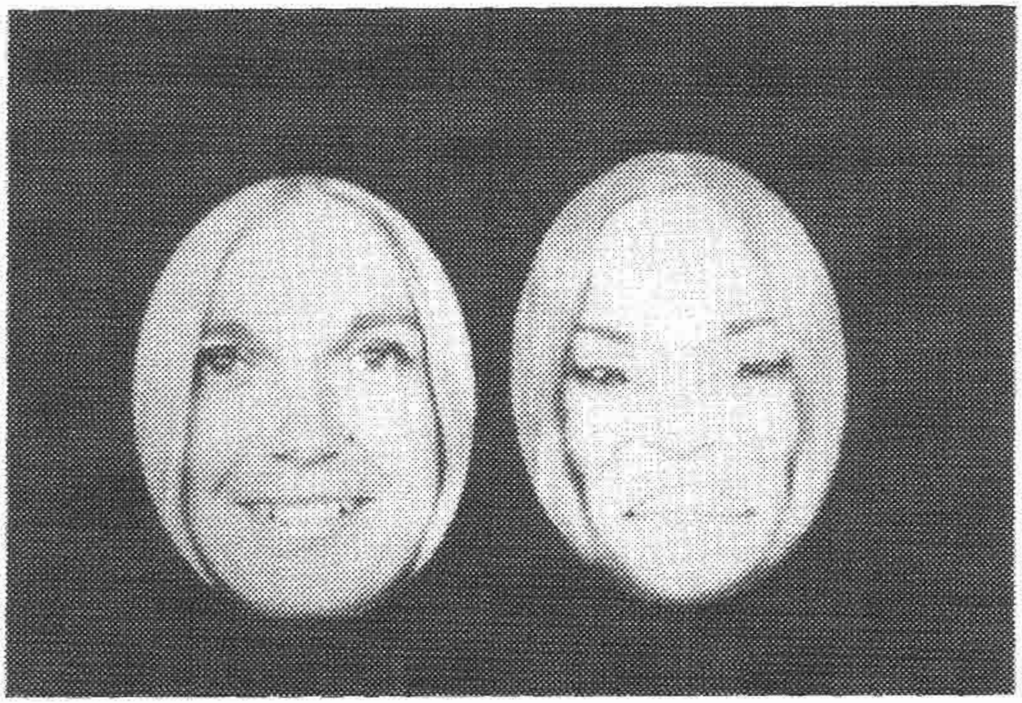

FIG. 4. An example of a Thatcher pair showing a normal, smiling version of one face and a Thatcherized version of a different face.

to pairs that were the least similar or easiest to tell apart. Each stimulus-pair was presented for 5 s, followed by a 2-s interstimulus-interval for subjects to respond by circling numbers on their answer sheets.

Prior to the same–different classification task, subjects were told that they would see another sequence of face-pairs, but that, this time, half of the pairs would show faces of different persons. Their task was to judge, as quickly as possible without sacrificing accuracy, whether each pair comprised pictures of the same person or of two different people. Note that this task involved making "same" judgments to all same-face pairs, including Thatcher and distortion pairs which were different with respect to wholistic information. Hence, attending to and using wholistic information was *not* consistent with task demands, at least not with Thatcher and distortion pairs.

Subjects responded in the classification task by pressing one of two telegraph keys—one labeled "same" and the other labeled "different"—using the index fingers of their left and right hands. The assignment of keys to "same" versus "different" was reversed for half the subjects. The 336-item sequence was preceded by 10 practice trials, with feedback, using pairs of faces drawn from a different set of stimuli.

Results

Similarity ratings. Considering first the expression condition, the top of Table 2 shows that there was a reliable increase in judged similarity when Thatcher pairs were inverted but not when neutral- or grotesque-expression pairs were inverted. Turning to the distortion condition, the bottom of Table 2 shows that the effect of inversion with Thatcher pairs

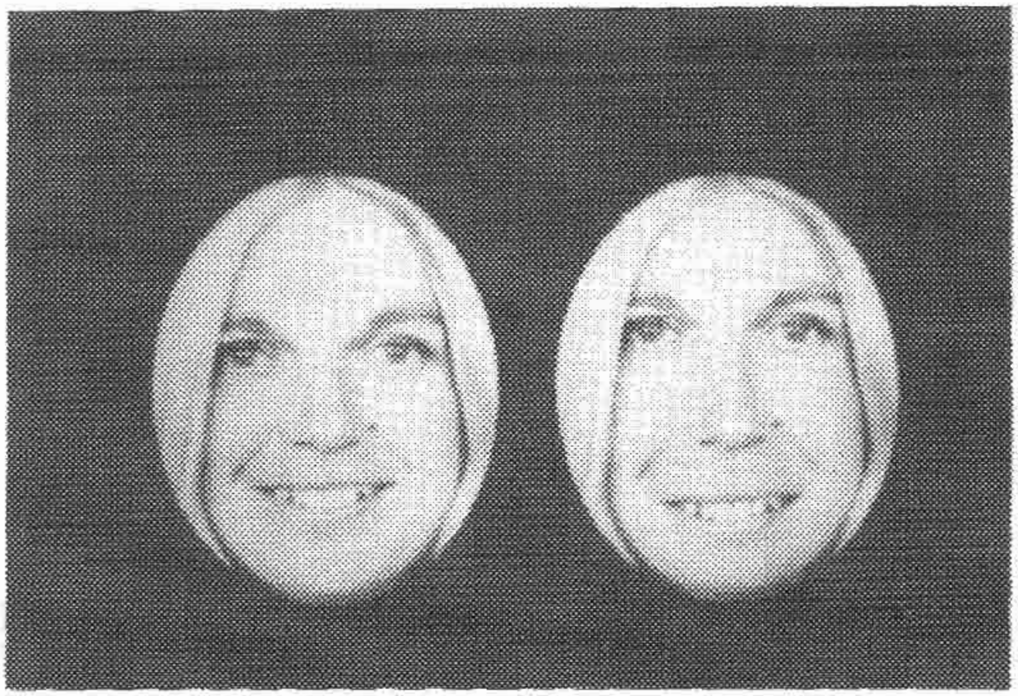

FIG. 5. An example of a type-1 distorted pair showing a normal, smiling version and a type-1 distorted version of the same face.

was once again robust. The effect of inversion with distortion pairs was somewhat less robust, but it was nonetheless reliable and in the same direction. ANOVAs of the expression and distortion conditions supported main effects for pair-type, $F(2,34) = 6.04$, $MS_e = 1.47$ and $F(2,32) = 11.8$, $MS_e = .399$, respectively, main effects for orientation, $F(1,17) = 19.6$, $MS_e = .769$ and $F(1,16) = 122.4$, $MS_e = .448$, respectively, and pair-type × orientation interactions, $F(2,34) = 68.0$, $MS_e = .150$ and $F(2,32) = 27.8$, $MS_e = .150$, respectively. The pattern of data converge with the grotesqueness ratings of Experiment 1 (Table 1): Just as stimulus inversion reduced judged grotesqueness of both Thatcherized faces and spatially distorted faces, but not grotesque- or neutral-expression faces, it also increased judged similarity of both Thatcherized and spatially distorted faces, but not grotesque- or neutral-expression faces, to their unaltered, smiling mates.

Same–different-face-classifications. Subjects found this task to be difficult, and they failed to respond within the allotted 3 s on 12% of the trials in the expression condition, and 6% of the trials in the distortion condition. The conditional probabilities of correct classifications, given that a judgment was made, are shown in Table 3. Two findings stand out. First, correct same judgments were substantially more probable than correct different judgments with the Thatcher and distortion pairs, but not with

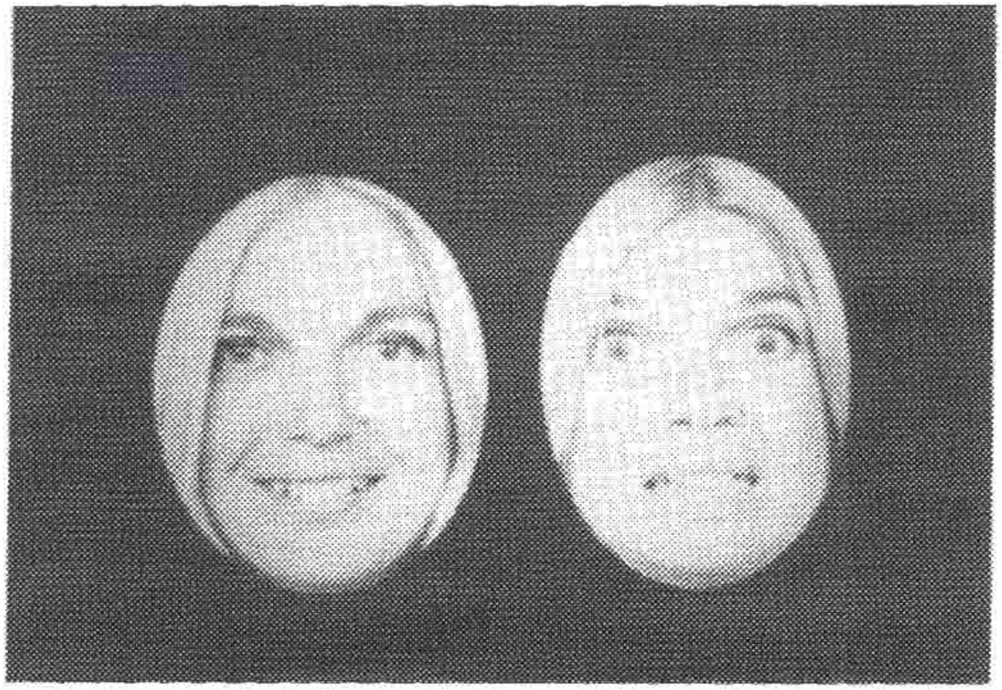

FIG. 6. An example of a grotesque-expression pair showing a normal, smiling version and a grotesque-expression version of the same face.

changed-expression pairs, suggesting a bias to judge the former pairs as "same." Second, inversion effects were generally small, although there was some evidence that inversion reduced accuracy of responses to neutral- and grotesque-expression pairs, and possibly to distortion-2 pairs (on different-face trials).

Accuracy versus bias effects. To distinguish accuracy from bias effects, we computed A' and B'' scores (Grier, 1971) from hit rates (i.e., the values in rows 1, 2, 7, and 8 of Table 3), and false-alarm rates, (i.e., 1.0 minus the values in rows 4, 5, 10, and 11 of Table 3). A's generally vary from .5 (chance discrimination) to 1.0 (perfect discrimination), whereas B'' scores vary from -1.0 (strongest bias for a response) to 0 (no bias) to $+1.0$ (strongest bias against a response).

The A' scores in the expression condition suggested that inversion impaired accuracy of performance with neutral-expression pairs (means = .75 and .67 for the upright and inverted conditions, respectively), and grotesque-expression pairs (*M*s = .72 and .68, respectively), but not with Thatcher pairs (*M*s = .66 and .65, respectively). The A' scores in the distortion condition showed less evidence for inversion effects (*M*s for upright versus inverted presentation were .76 versus .74 for type-1 distortion pairs, .77 versus .73 for type-2 distortion pairs, and .74 versus .76 for Thatcher pairs). An ANOVA of A's in the expression condition

TABLE 2
Mean Similarity Ratings (and Standard Deviations) for Upright and inverted Face-Pairs in the Expression and Distortion Conditions of Experiment 1

Orientation	Pair-type in the expression condition		
	Thatcher	Neutral	Grotesque
Upright	4.89 (1.35)	3.41 (1.07)	4.51 (1.18)
Inverted	2.92 (1.17)	3.39 (0.60)	4.26 (0.78)
Difference	1.97 *	0.02	0.25
	Pair-type in the distortion condition		
	Thatcher	Distortion 1	Distortion 2
Upright	5.13 (0.93)	4.63 (1.00)	3.91 (0.85)
Inverted	2.86 (0.73)	3.53 (1.06)	2.89 (0.70)
Difference	2.27 *	1.10 *	1.02 *

Note. *n*s were 18 and 17 in the expression and distortion conditions, respectively. Ratings ranged from 1 (most similar) to 7 (least similar).

* Inversion effect gave $p < .05$ by F test.

yielded reliable main effects for pair-type, $F(2,34) = 4.94$, $MS_e = .006$, and orientation, $F(1,17) = 10.9$, $MS_e = .005$, as well as a reliable interaction, $F(2,34) = 5.19$, $MS_e = .002$ (all $ps < .02$). The orientation effect was statistically significant with neutral-expression pairs, $F(1,17) = 5.11$, $MS_e = .004$ and grotesque-expression pairs, $F(1,17) = 32.5$, $MS_e = .002$, but not with Thatcher pairs ($F < 1$) in the expression condition. An ANOVA of A's in the distortion condition supported no effects ($ps > .10$).

Bias scores (B''s) appeared unaffected by orientation, and they were not reliably different from 0 with either neutral-expression pairs or grotesque-expression pairs ($Ms = -.10$ and $+.13$, respectively, with $ts\ (17) = 1.00$ and 1.52, respectively, $ps > .10$). However, there was a strong bias that favored "same" judgments with Thatcher pairs in the expression condition ($M = -.75$, $t\ (17) = 13.6$, $p < .001$), and with Thatcher, type-1, and type-2 pairs in the distortion condition ($Ms = -.62$, $-.58$, and $-.63$, respectively, with $ts\ (16) = 6.83$, 6.90, and 7.38, respectively, $ps < .01$). The pattern of bias with Thatcher pairs but not with changed-expression pairs was supported by a reliable pair-type effect in an ANOVA of scores from the expression condition, $F(2,34) = 51.1$, $MS_e = .15$. An ANOVA of scores from the distortion condition supported no effects. This negative outcome is consistent with the finding of bias with type-1 and type-2 distortion pairs as well as Thatcher pairs.

In summary, accuracy scores were reduced by inversion with changed-expression pairs but not with Thatcher pairs and not much (if at all) with

TABLE 3
Conditional Probabilities of Correct Same- and Different-Face Judgments (and Standard Deviations) to Upright and Inverted Pairs in the Expression and Distortion Conditions of Experiment 2

Judgment and orientation	Pair-type in the expression condition		
	Thatcher	Neutral	Grotesque
Same judgments			
Upright	.98 (.04)	.87 (.13)	.76 (.16)
Inverted	.97 (.05)	.82 (.15)	.74 (.15)
Difference	.01	.05	.02
Difference judgments			
Upright	.68 (.16)	.86 (.11)	.87 (.12)
Inverted	.70 (.16)	.77 (.14)	.82 (.12)
Difference	−.02	.09 *	.05 *
	Pair-type in the distortion condition		
	Thatcher	Distortion 1	Distortion 2
Same responses			
Upright	.94 (.08)	.93 (.08)	.95 (.05)
Inverted	.98 (.04)	.94 (.06)	.96 (.06)
Difference	−.04 *	−.01	−.01
Different responses			
Upright	.77 (.16)	.78 (.15)	.79 (.14)
Inverted	.76 (.15)	.76 (.16)	.71 (.14)
Difference	.01	.02	.08 *

* The inversion effect gave $p < .05$ by F test.

distortion pairs. However, Thatcher and distortion pairs evoked a strong bias for "same" judgments.[5]

Latencies of same–different classifications. The bias effects shown by the B'' scores were also reflected in the latencies (Table 4). Mean latencies for correct same judgments were shorter than those for correct different judgments with Thatcher pairs (Ms = 1346 vs 1550, respectively) and distortion pairs (Ms = 1291 vs 1430, respectively), but not with changed-expression pairs (Ms = 1526 vs 1539, respectively). Inversion effects were small and generally unreliable.

In support of the preceding description, an ANOVA of the latencies from the expression condition supported main effects for judgment, $F(1,17) = 15.1$, $MS_e = 28{,}406$, and pair-type, $F(2,34) = 4.41$, $MS_e = 9313$, as well as the judgment × pair-type interaction, $F(2,34) = 38.3$, MS_e

[5] Comparable analyses of Pr and Br scores (Snodgrass & Corwin, 1988) yielded a similar pattern, although in this case there was evidence that inversion impaired discrimination with type-2 distortion pairs as well as changed-expression pairs.

TABLE 4
Latencies (and Standard Deviations) for "Same" and "Different" Judgments to Upright and Inverted Pairs in the Expression and Distortion Conditions of Experiment 2

Judgment and orientation	Pair-type in the expression condition		
	Thatcher	Neutral	Grotesque
Same judgments			
Upright	1385 (211)	1502 (238)	1532 (245)
Inverted	1375 (199)	1512 (204)	1556 (194)
Difference	10	−10	−24
Difference judgments			
Upright	1583 (250)	1507 (276)	1522 (234)
Inverted	1658 (268)	1547 (215)	1580 (231)
Difference	−75 *	−40	−58
	Pair-type in the distortion condition		
	Thatcher	Distortion 1	Distortion 2
Same judgments			
Upright	1334 (216)	1338 (215)	1257 (196)
Inverted	1288 (170)	1284 (201)	1286 (203)
Difference	46	54 *	−29
Difference responses			
Upright	1485 (204)	1398 (214)	1437 (193)
Inverted	1475 (217)	1435 (170)	1449 (218)
Difference	10	−37	−12

* The inversion effect gave $p < .05$ by F test.

= 8048 (there was also a judgment × orientation interaction, $F(1,17) = 5.29$, $MS_e = 6203$, reflecting the fact that, for unknown reasons, a reaction-time advantage for upright presentation was stronger for different- than same-judgments). An ANOVA of latencies from the distortion condition supported only one effect, a main effect of judgment (same vs different), $F(1,16) = 112.3$, $MS_e = 10{,}057$.

In conclusion, the data from the same–different judgments differed from the similarity ratings in showing little or no effect of inversion with Thatcher and distortion pairs. Indeed, the evidence for an inversion effect on same–different judgment accuracy was weaker with Thatcher and distortion pairs than with changed-expression pairs. There was, however, a puzzling bias to judge that Thatcher and distortion pairs were "same."

Item analyses. To clarify the differences between the similarity and same–different-judgment data, we computed a mean similarity rating, and a mean same-judgment reaction-time, for each of the 42 same-face-pairs of each experimental condition when presented upright and when presented inverted. A set of correlations performed on these data showed that the similarity ratings to upright versus inverted face-pairs were only weakly correlated (r = .33), but that the same-judgment latencies to these same pairs were more strongly correlated (r = .68). Both correlations were statistically significant (dfs = 82, ps < .005), but they differed reliably from each other, $z = 3.13, p < .002$. In addition, similarity ratings to upright face-pairs were virtually uncorrelated with same-judgment-latencies (rs < +.10), whereas similarity ratings to *inverted* face-pairs were reliably correlated with same-judgment latencies (rs > +.50, ps < .001). Whether latencies were to upright or inverted face-pairs made no difference to the pattern, which is shown in Fig. 7. We conclude that the information used for same–different classifications overlapped with that used to rate the similarity of inverted face-pairs, but not with that used to rate the similarity of upright face-pairs.

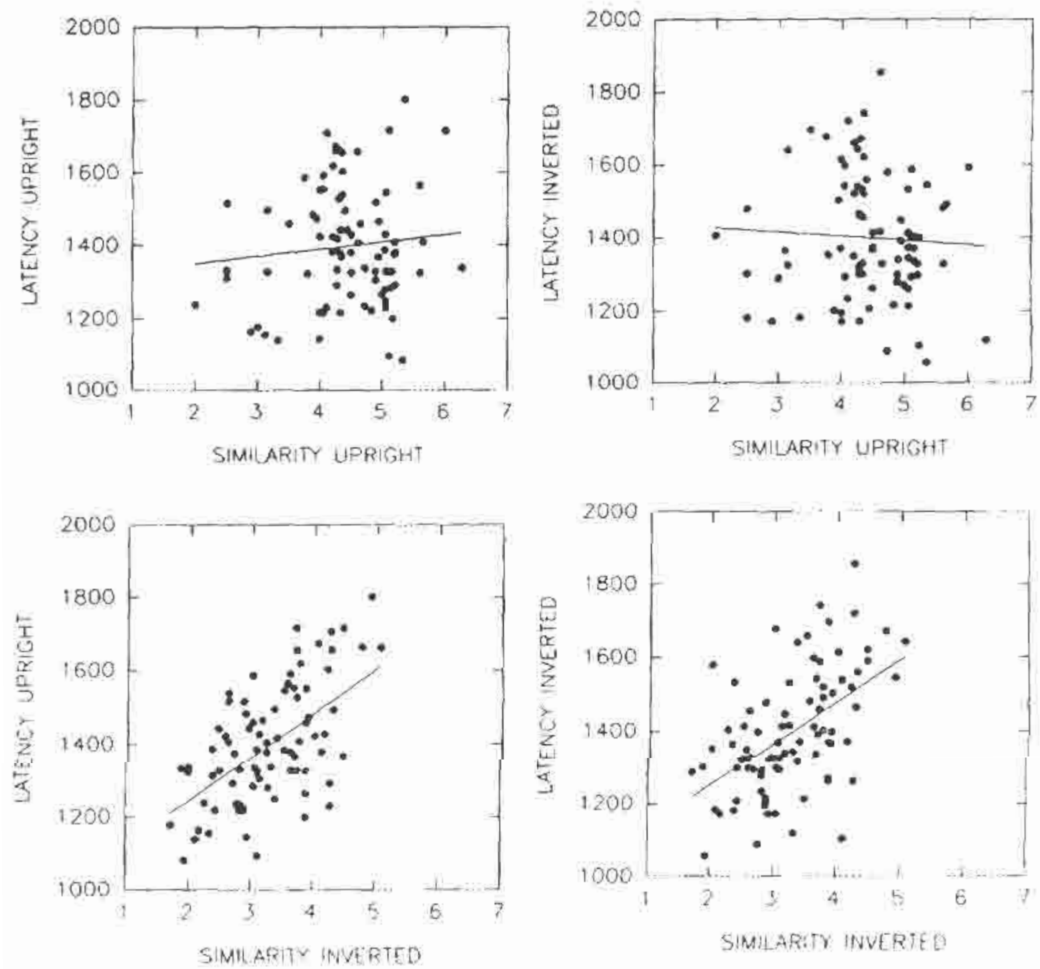

FIG. 7. Scatter plots showing the correlation between similarity ratings (1, greatest similarity) and latencies for same-face judgments to same-face pairs in each orientation (moving clockwise from upper-left plot, rs = +.10, −.06, +.51, and +.56 (dfs = 82).

Discussion

The similarity ratings of Experiment 2 converged with the grotesqueness ratings of Experiment 1. Experiment 1 showed that Thatcher and spatial-distortion faces, but not neutral and grotesque-expression faces, were judged less grotesque when shown upside down. Similarly, the present study showed that Thatcher and spatial-distortion faces, but not neutral- or grotesque-expression faces, were judged as more similar to unaltered, smiling faces when presented upside down. These effects of orientation, which are clearly qualitative, suggest that *both* grotesqueness ratings *and* similarity ratings were based on the processing of wholistic information, and that inversion severely disrupted such processing. There is no suggestion in these data that inversion is disruptive to expression encoding.

In contrast to the similarity ratings, the same–different conditional probabilities and latencies did *not* show strong effects of inversion with Thatcher and distortion pairs. Indeed, an accuracy measure (A') derived from the conditional probabilities was more affected by inversion with changed-expression pairs than with Thatcher and distortion pairs. This outcome suggests that although wholistic information was used as a basis for similarity ratings, another type of information was used for same-face/different-face judgments. An item-analysis further supported this conclusion; it showed that: (a) same-judgment latencies for upright face-pairs and same-judgment latencies for inverted face-pairs were strongly intercorrelated, and (b) both latency measures were significantly correlated with similarity ratings of inverted face-pairs, but not with similarity ratings of upright face-pairs (Fig. 7). The clear implication is that wholistic information was used to judge the similarity of upright pairs, whereas some other type of information was used to judge similarity of inverted pairs. Moreover, this "other type of information" used with inverted pairs overlapped with that used for same-face judgments regardless of orientation. This conclusion is consistent with the dual-mode hypothesis, which claims that the "other type of information" referred to above is component information. Experiment 3 provides new data on this point.

EXPERIMENT 3

An unexpected finding from Experiment 2 was that subjects showed a bias to make same judgments in response to Thatcher and distortion pairs, but not in response to changed-expression pairs. The bias was suggested in the conditional probabilities of correct same and different judgments (Table 3), and was confirmed with signal-detection measures (B''), and in analyses of latencies as well (Table 4).

Why should such a pattern of bias exist? The dual-mode hypothesis suggests the following account: An unaltered version of a face and a Thatcherized or distorted version of this face differ with respect to wholistic information, but they are very well matched with respect to components (e.g., shape and size of mouth, eyes). Hence, encoding Thatcher and distortion pairs with respect to components would produce a good "match," which in turn would tend to influence a subject to make a "same" judgment. Thus, the bias for "same" judgments found with Thatcher and distortion pairs can be simply explained; subjects judged such pairs with respect to their components.

The same line of reasoning can also explain why a same-judgment bias was *not* obtained with changed-expression pairs. Unlike Thatcher and distortion pairs, changed-expression pairs consisted of faces that clearly differed in components (e.g., shape of mouth and eyes, visibility of teeth, wrinkles in the skin, see Fig. 6). Hence, a componential processing mode would produce a relatively poor "match" with such pairs, which itself might explain the absence of bias. Alternatively, since all changed-expression pairs have mismatching components, subjects might tend to judge such pairs primarily on the basis of wholistic information. Indeed, wholistic processing of changed-expression pairs fits the finding that inversion reduced same–different judgment accuracy (A's) only with these pairs.

This dual-process account of the bias effects makes a testable prediction. The major claim of the dual-process view is that inversion of faces precludes wholistic processing, forcing subjects to rely on componential processing. If so, the pattern of bias found in same–different judgments should be reflected in similarity ratings, *so long as the faces are presented upside down*.

To test this prediction, we examined similarity ratings and same-versus-different-face ratings to same-face pairs and different-face pairs in which one face was smiling and the other: (a) was Thatcherized, (b) had a grotesque expression, or (c) had a neutral expression (as in the expression condition of Experiment 2). Both types of ratings were made on a six-point scale on which "1" corresponded to highest similarity (or least difference) as well as highest confidence that a pair was a same-face pair, and "6" corresponded to lowest similarity (greatest difference) as well as highest confidence that a pair was a different-face pair. The pattern of same–different judgments in Experiment 2 suggests that the same–different face ratings should be lower (indicating more bias for a same-face decision) for Thatcher pairs than changed-expression pairs. The new prediction is that the similarity ratings should show this same effect, but *only* when faces are shown upside down. We also expected to replicate the finding that changed-expression pairs, although not Thatcher pairs, would evoke less accurate same–different judgments when presented upside down.

Method

Subjects viewed a list comprising same-face and different-face pairs in which one face was smiling and the other: (a) was Thatcherized, (b) had a grotesque expression, or (c) had a neutral expression. Subjects saw the entire list twice, once while rating similarity on a six-point scale, and once while rating confidence in a same-face versus a different-face classification, again using a six-point scale. The design provided a within-subject comparison of a same–different-face judgment task with a similarity-rating task with control over stimulus materials, presentation conditions, and mode of response.

Subjects and procedure. The 34 subjects were recruited as in prior experiments. They viewed 168 pairs of face-pictures two times; once while performing a same–different-face judgment task, and once while performing either a similarity-rating task or a difference-rating task. The instructions for the same–different-face task were that subjects should judge whether the pictures of each stimulus-pair were made from the same face versus two different faces, and that they should indicate confidence using a scale of 1 (sure same face) to 6 (sure different face). The instructions for the similarity- and difference-rating tasks were modified from those of Experiment 2 to focus on degree of similarity or degree of difference. Subjects made their judgments on a six-point scale on which 1 represented greatest similarity (or least difference), and 6 represented least similarity (or most difference). We found no important differences between the similarity and difference ratings, and we will present the findings collapsed over this variable.

Design. Task order was varied such that half the subjects performed the same-face/different-face task first, whereas the remainder performed the similarity (or difference) task first. Different randomized sequences of the 168 pairs were used for each task, and the sequences used for half of the subjects were reshuffled for the remaining subjects. As in Experiment 2, the 168 face-pairs in each sequence included 84 same-face pairs and 84 different-face pairs, and within each set, one-third of the pairs were Thatcher pairs, one-third were neutral-expression pairs, and one-third were grotesque-expression pairs. Each pair was shown for 5 s with a 3-s interpair interval. Latencies of responses were not measured.

Results

Table 5 displays the average same–different-face ratings and similarity ratings which varied from 1 (most similar or most obviously the same face) to 6 (least similar or most obviously different faces). Inversion increased both types of ratings for same-face-neutral and same-face-grotesque pairs, and it reduced both types of ratings for different-face-neutral and different-face-grotesque pairs. That is, inversion reduced differential responding to same- versus different-face pairs that differed in expression. In contrast, same-face-Thatcher pairs and different-face-Thatcher pairs showed virtually identical effects of inversion. These effects were weakly negative in the same–different-face ratings, but were

TABLE 5
Mean Same–Different Face Judgments and Similarity Ratings (and Standard Deviations) to Thatcher, Neutral, and Grotesque-Expression Pairs in Experiment 3

	Pair-type		
	Thatcher	Neutral	Grotesque
Same–different judgments			
Same-face pairs			
Upright	1.18 (0.33)	1.61 (0.50)	1.86 (0.51)
Inverted	1.30 (0.40)	1.92 (0.58)	2.55 (0.78)
Difference	−.12 *	−.31 **	−.69 **
Different-face pairs			
Upright	4.05 (0.85)	5.09 (0.50)	5.10 (0.54)
Inverted	4.18 (0.68)	4.63 (0.67)	4.57 (0.66)
Difference	−.13	.46 **	.53 **
Similarity ratings			
Same-face pairs			
Upright	2.97 (1.25)	2.50 (0.70)	3.00 (0.81)
Inverted	1.79 (0.51)	2.63 (0.63)	3.22 (0.69)
Difference	1.18 **	−.13	−.22 *
Different-face pairs			
Upright	4.68 (0.71)	4.50 (0.72)	4.79 (0.56)
Inverted	3.65 (0.75)	4.09 (0.69)	4.31 (0.62)
Difference	1.03 **	.41 **	.48 **

* $p < .05$.
** $p < .01$.

strongly positive in the similarity ratings. The pattern reveals that Thatcher pairs were perceived as more similar—but not as more obviously same-face pairs—when they were presented upside down.

The point we wish to emphasize in these data concerns a comparison of Thatcher pairs with grotesque- and neutral-expression pairs in each orientation condition. Figure 8 shows this comparison, collapsing the data from same- and different-face pairs. Note that the average same–different-face ratings (top) were lower for Thatcher pairs than for the changed-expression pairs, as the former showed a bias to make same-face judgments (3.5 was the midpoint of the six-point scale). The average similarity ratings (bottom) showed precisely this same pattern—greater judged similarity for Thatcher pairs than for the changed-expression pairs—but *only* in the inverted condition. When faces were upright, the ratings for the Thatcher pairs were similar to those for the changed-expression pairs. We conclude that there was a bias to judge Thatcher pairs as same-face items and also to judge them as highly similar items. However, the bias to judge such pairs as highly similar was restricted to inverted presentation.

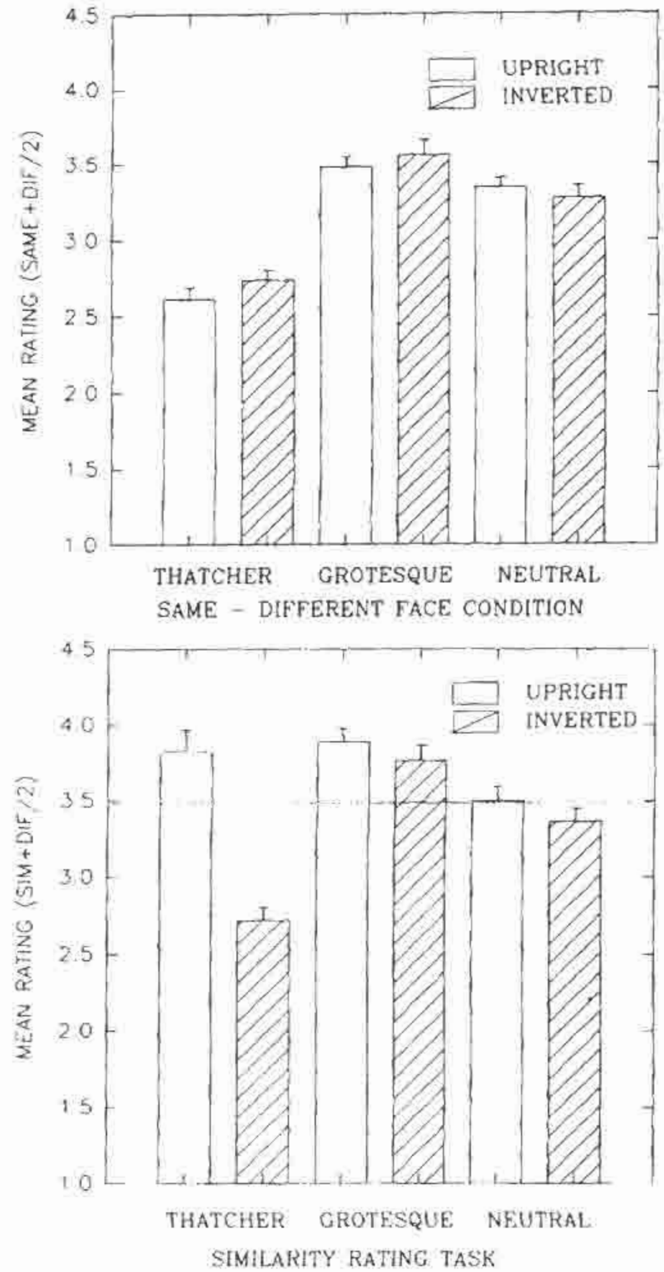

FIG. 8. Means of ratings to same- and different-face pairs in upright and inverted orientation in the same–different-face judgment task and the similarity-rating task of Experiment 3. Higher ratings indicate more dissimilarity or more confidence in a different-face judgment.

The preceding observations were supported by an ANOVA which included the between-subjects variable of task-order (similarity ratings first versus same–different ratings first), and the within-subjects variables of task condition (similarity versus same–different ratings), pair-type

(Thatcher, Neutral, Grotesque), orientation, and same-versus-different-face condition. As shown in Table 6, the four within-subjects variables all had reliable main effects. In addition, there were several interactions involving same–different-face condition and these confirm our observations that: (a) the effects of inversion with changed-expression pairs switched directions between the same-face condition and the different-face condition, whereas (b) the effects of inversion with Thatcher pairs did not switch direction in this way but were affected by task (see Table 5).

Apart from the interactions involving same–different-face condition, the ANOVA revealed reliable interactions between task-condition and orientation, task-condition and pair-type, and orientation and pair-type, all of which were qualified by the higher order interaction among task-condition, pair-type, and orientation. It is the final interaction which is of most importance, supporting the conclusion that subjects showed a bias to judge Thatcher pairs as same-face pairs and also to judge them as high in similarity when they were inverted but not when upright (Fig. 8).

We note that the task-order variable was involved in no reliable effects except for an uninterpretable interaction with task condition and pair-type. Since this interaction involved neither orientation nor same–different-face condition, it does not qualify the findings of most importance here.

TABLE 6
Reliable Effects in the ANOVA of Experiment 3

Effect	*df*	*F*	MS_e	*p*
Main effects				
Same–different face (SD)	1,32	616.5	1.69	.0001
Task-condition	1,32	11.0	2.16	.005
Orientation	1,32	17.9	0.48	.0005
Pair-type	2,64	115.7	0.29	.0001
Simple interactions				
SD × task	1,32	89.4	0.85	.0001
SD × pair	2,64	14.0	0.19	.0001
SD × orientation	1,32	43.1	0.31	.0001
Task × orientation	1,32	66.2	0.19	.0001
Task × pair	2,64	18.9	0.21	.0001
Orientation × pair	2,64	20.3	0.21	.0001
Higher order interactions				
Task-order × task × pair	2,64	9.43	0.21	.0005
SD × task × orientation	1,32	6.5	0.17	.02
SD × pair × orientation	2,64	29.7	0.16	.0001
Task × orientation × pair	2,64	23.5	0.29	.0001

Note. The only significant effect involving task-order was the task-order × task × pair interaction which did not qualify the other higher order interactions.

Discussion

Replicating the findings of Experiment 2, we observed a strong bias favoring "same-face" judgments over "different-face" judgments in response to Thatcher pairs (pairs consisting of normal, smiling faces next to Thatcherized faces), but *not* changed-expression pairs (pairs consisting of smiling faces next to faces with neutral or grotesque posed expressions). The main purpose of Experiment 3 was to test a possible account of this finding suggested by the dual-mode hypothesis.

According to the dual-mode account, subjects' same–different judgments in response to Thatcher pairs are made largely on the basis of component information which is very well matched between the faces of such pairs. Similarity judgments in response to Thatcher pairs can also be based on component information, depending on orientation. Wholistic processing dominates when the face-pairs are upright, but component processing dominates when face-pairs are inverted. It follows that the bias found with same–different judgments should also be reflected in similarity judgments—Thatcher pairs should be judged as more similar than changed-expression pairs—*but only in conditions of inverted presentation*. This was precisely the finding obtained.

Another finding from Experiment 2 that we replicated here concerned discrimination rather than bias effects. Differential responding on same-face trials versus different-face trials was impaired by inversion with changed-expression pairs but not with Thatcher pairs. This effect fits other evidence that inversion impairs encoding of "expression-independent information" which supports recognition of faces of known persons regardless of expression (Bruce & Young, 1986; Valentine & Bruce, 1988). The nature expression-independent information is presently unknown, but an attractive possibility is that such information is essentially wholistic. If this possibility can be confirmed through new research, we will have further support for the dual-mode view that inversion disrupts wholistic encoding.

GENERAL DISCUSSION

The narrow goal of this research was to improve our understanding of the Thatcher illusion. The key findings relevant to this goal were that faces made grotesque through spatial distortion (displacement of the eyes and mouth)—though not faces made grotesque simply through their posed expressions—were similar to Thatcherized faces in that (a) they were judged as less grotesque when inverted than upright (Experiment 1), and (b) they were judged as more similar to unaltered, smiling faces when inverted than upright (Experiments 2 and 3).

These findings have led us to draw two conclusions. First, the Thatcher illusion *cannot* be explained as an effect of inversion on the encoding of expression. Second, a workable account of the Thatcher illusion is that inversion impairs processing of wholistic information, that is, information pertaining to spatial configuration (e.g., symmetry) and/or deviations from templates for norm faces.

Our second conclusion is based on the assumption that the spatial distortion manipulation (Fig. 2) altered wholistic properties of faces as opposed to their component properties. This assumption might be questioned. It is arguable, for example, that the interocular distance between the two eyes is a component property (i.e., a dimension) of faces, and this component was altered in our distorted faces.[6]

We have three replies to this argument. First, the information changed in our distorted faces fits current conceptions of wholistic (i.e., configural and template-based) information, and on this basis we *predicted* that these faces would show strong effects of inversion. Hence, it appears parsimonious in light of current theory to assume that our distorted faces were distorted primarily in their wholistic properties as opposed to their components (although we would not deny that some components may have been altered). Second, we find it impressive that the ratings of the distortion faces converged with the ratings of the Thatcher faces and diverged from the ratings of the changed-expression faces. The pattern suggests that some type of information was altered in the distortion and Thatcher faces, but not in the changed-expression faces. That this type of information was wholistic information is as plausible a notion as we can conceive. Third, a recent study by Rhodes, Tan, and Brake (in press) supports the assumption that the distortion transformation affects wholistic information. Using a recognition-memory paradigm, Rhodes et al. found that inversion impaired discrimination between old and new faces when these differed in eye-and-mouth orientation (Thatcherization), as well as when these differed in eye-and-mouth-spatial-location (spatial distortion). Inversion also impaired discrimination when old and new items differed in individual features, but this latter effect vanished when facial context was removed. The authors interpreted their findings as evidence that inversion impairs "holistic, norm-based coding" of faces which can aid in the encoding of facial components within a facial context (see Carey, 1991). This converges with our own conclusion that inversion impaired processing of wholistic information.

We found no evidence for the frequent observation that inversion impairs processing of facial expression. In fact, the findings of Experiments 2 and 3 converge with prior evidence (e.g., Valentine & Bruce, 1988) that

[6] We thank Vicki Bruce for pointing out this argument.

inversion impairs processing of facial identity *independent* of expression. Having made this point, we also suggest that if inversion impairs processing of wholistic information, it should also impair the processing of expression when wholistic information is the basis of expression. Wholistic information might be especially important as a basis for highly subtle expressions such as mild annoyance or feigned interest (see Dolezal, 1982). In any event, the conclusion to draw from the present research is *not* that inversion does not affect processing of facial expression. Rather, it is that inversion disrupts wholistic encoding more than other types of encoding that can make a facial expression grotesque.

Valentine (1988, 1991) argued that the effects of inversion in face processing tasks are quantitative rather than qualitative, resulting from an increase in the "error" of encoding rather than a switch from one mode to another. This is consistent with some prior findings (e.g., Valentine & Bruce, 1988), and with some of the present findings as well (e.g., Experiment 3, see Fig. 8, top). However, it is clearly inconsistent with the qualitative effects of inversion found in Experiments 1 and 2 and with the correlational analyses of Experiment 2. We find the correlations to be especially provocative due to their counterintuitive nature: Similarity ratings of upright face-pairs were only weakly correlated with those of inverted face-pairs, and the former showed virtually no correlation with latencies for same judgments (Fig. 7, top)—only the ratings of inverted face-pairs showed substantial correlations with same-judgment latencies (Fig. 7, bottom). The pattern suggests that wholistic encoding was used for similarity ratings of upright face-pairs, whereas some other type of encoding—which we have argued is componential encoding—was used for similarity ratings of inverted face-pairs as well as same-face judgments.

In summary, the present data provide strong evidence that inversion disrupts processing of at least one type of wholistic information, but they also suggest that this wholistic information is not always used for paired-comparisons of faces. Indeed, wholistic information is probably ignored in favor of other types of information in a variety of face-processing tasks. Since these other types of information appear relatively insensitive to effects of inversion, the effects of inversion are often merely quantitative, as Valentine (1988) points out. However, qualitative effects of inversion can be found in conditions that evoke a mode or strategy of wholistic encoding when faces are upright.

The broader goal of this research was to increase our understanding of why retinal inversion or monooriented stimuli is sometimes severely disruptive to processing but in other cases not. Most prior discussion of this issue has focused on aspects of the stimulus material, such as whether it consists of face stimuli (Yin, 1969), of stimuli (including but not restricted distinguished based on second-order relations (Diamond & Carey, 1986). Characteristics of subjects have also been discussed, especially the characteristics of age (Carey & Diamond, 1977), right- versus left-hemisphere brain damage (Yin, 1970), and expert knowledge about the stimuli (Diamond & Carey, 1986). In comparison, the factor of task has received but scant attention (but see Corballis, 1988; Sergent, 1984).

From the point of view of general theories of orientation and form, our findings have highlighted the importance of task. The main thrust of the findings from Experiments 1 and 2 is that although inversion dramatically influenced face processing in grotesqueness rating and similarity rating tasks, it did not do so in a same–different-face judgment task. This point was buttressed in Experiment 3, which showed a bias to judge that Thatcher faces, but not neutral- or grotesque-expression faces, represent the same faces as unaltered smiling faces. The similarity ratings showed a similar bias to judge Thatcher pairs as more highly similar than neutral- or grotesque-expression pairs. However, the bias effect in similarity ratings occurred *only* in conditions of inverted presentation. Thus, there was a qualitative effect of inversion on bias in the similarity-rating task, but not in the same–different-judgment task.

We have argued that similarity judgments to upright faces are based on wholistic information, and that same–different judgments regardless of orientation are based on componential information. Even if we are wrong, however, it appears inescapable given our findings that the effects of inversion depend upon the task. That this conclusion was supported using face stimuli—which are well known for showing inversion effects—is particularly impressive in supporting the power of the factor of task. We note that research on mental rotation of facial stimuli (Valentine & Bruce, 1988), as well as nonfacial stimuli (Jolicoeur, 1985, 1988), also supports the conclusion that the effects of orientation of forms vary with the task.

In supporting the importance of the variable of task, the findings presented in this article add support to theories and conceptual frameworks that distinguish different processing modes specialized for encoding different types of information (Carey & Diamond, 1977; Garner, 1978; Goldstone, Medin, & Gentner, 1991; Sergent, 1984). The main question raised by the present research concerns the precise nature of the types of information encoded through these modes. Drawing on Garner's (1978) framework, we have argued that one mode is specialized for the encoding of wholistic information, whereas the others is specialized for the encoding of components. The wholistic information we examined in this paper apparently is specific to orientation of components, as well as the spatial positions of components and/or spatial relations between components.

Beyond this, however, there is little we can say. Whether this information consists of deviations from global templates (Yuille, 1991), or as deviations from "first-order" spatial relations (Diamond & Carey, 1986), or whether these ideas can even be distinguished, has yet to be determined. We also lack knowledge as to whether or how the wholistic/component distinction, as operationalized here, relates to distinctions between relations and attributes (Goldstone et al., 1991), global versus local information (Navon, 1977), and holistic versus analytic information (Kemler, 1983). More research on effects of different facial transformations under different task conditions is obviously required.

REFERENCES

Benson, P. J., & Perrett, D. I. (1991). Perception and recognition of photographic quality facial caricatures: Implications for the recognition of natural images. *European Journal of Cognitive Psychology*, **3**, 105–135.

Bradshaw, J. L., & Nettleton, N. C. (1981). The nature of hemispheric specialization in man. *The Behavioral and Brain Sciences*, **4**, 51–94.

Bruce, V. (1988). *Recognizing faces*. Hillsdale, NJ: Erlbaum.

Bruce, V., Doyle, T., Dench, N., & Burton, M. (1991). Remembering facial configurations. *Cognition*, **38**, 109–144.

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

Carey, S. (1991, November). *Developmental differences in holistic processing of faces?* Paper presented at the 32 annual meeting of the Psychonomic Society, San Francisco, CA.

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

Corballis, M. C. (1988). Recognition of disoriented shapes. *Psychological Review*, **95**, 115–123.

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

Dolezal, H. (1982). *Living in a world transformed: Perceptual and performatory adaptation to visual distortion*. New York: Academic Press.

Ekman, P. (1975). *Pictures of facial affect*. Palo Alto, CA: Consulting Psychologists Press.

Ellis, H. D., & Young, A. W. (1989). Are faces special? In A. W. Young & H. D. Ellis (Eds.), *Handbook of research on face processing*. (pp. 1–26). New York: North-Holland.

Fischler, M. A., & Elschlager, R. A. (1973). The representation and matching of pictorial structures. *IEEE Transactions on Computers*, **22**, 1.

Garner, W. R. (1978). Aspects of a stimulus: Features, dimensions, and configurations. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 99–133). Hillsdale, NJ: Erlbaum.

Gnier, J. B. (1971). Nonparametric indexes for sensitivity and bias: Computing formulas. *Psychological Bulletin*, **75**, 424–429.

Goldstein, A. G., & Chance, J. (1980). Memory for faces and schema theory. *Journal of Psychology*, **105**, 47–59.

Goldstone, R. L., Medin, D. L., & Gentner, D. (1991). Relational similarity and the non-independence of features in similarity judgments. *Cognitive Psychology*, **23**, 222–262.