The Quarterly Journal of Experimental Psychology Section A: Human Experimental Psychology
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/pqja20

When inverted faces are recognized: The role of configural information in face recognition

Helmut Leder & Vicki Bruce
Published online: 22 Oct 2010.

To cite this article: Helmut Leder & Vicki Bruce (2000) When inverted faces are recognized: The role of configural information in face recognition, The Quarterly Journal of Experimental Psychology Section A: Human Experimental Psychology, 53:2, 513-536

To link to this article: http://dx.doi.org/10.1080/713755889

THE QUARTERLY JOURNAL OF EXPERIMENTAL PSYCHOLOGY, 2000, *53A* (2), 513–536

When inverted faces are recognized: The role of configural information in face recognition

Helmut Leder
Freie Universität Berlin, Germany
Vicki Bruce
University of Stirling, Scotland, U.K.

The identification of upright faces seems to involve a special sensitivity to "configural" information, the processing of which is less effective when the face is inverted. However the precise meaning of "configural" remains unclear. Five experiments are presented, which showed that the disruption of the processing of relational, rather than holistic, information largely determines the occurrence as well as the size of the face-inversion effect. In Experiment 1, faces could be identified either by unique combinations of local information (e.g. a specific eye colour plus hair colour) or by unique relational information (e.g. nose–mouth distance). The former showed no inversion effect, whereas the latter did. A combination of local and relational information (Experiment 2) again produced an inversion effect, although this effect was smaller than that found when only relational information was used. The results were replicated in Experiment 3 when differences in the brightness of local features were used instead of specific colour combinations. Experiment 4 used different retrieval conditions to distinguish relational from holistic processing, and demonstrated again that spatial relations between single features appeared to provide crucial information for face recognition. In Experiment 5, the importance of relational information was confirmed using faces that also varied in the shapes of local features.

In everyday life the face is the main source of information about other humans: Faces tell us who people are, their age and sex, and at least occasionally how they feel. We are very skilled in "reading" these types of information. But when faces are presented upside-down then this information extraction is severely disrupted, and it is particularly difficult to recognize and identify faces when they are inverted. Although inversion also affects recognition of other mono-orientated objects, the effect seems to be profound for faces (Yin, 1969) and it seems justified to speak of a special face-inversion effect (FIE).

Requests for reprints should be sent to Helmut Leder, Department of Psychology, Freie Universität Berlin, Habelschwerdter Allee 45, 14195 Berlin, Germany. Email: lederh@zedat.fu-berlin.de

This research was supported by a research fellowship to Leder from the Swiss National Research Council, and by a grant from the Economic and Social Research Council to Bruce and Burton in Stirling.

We want to thank James Bartlett and an anonymous reviewer for helpful comments on earlier versions of this manuscript.

http://www.tandf.co.uk/journals/pp/02724987.html

This effect has been demonstrated in a variety of tasks and settings (see Valentine, 1988, for a review) and has often been used as an indicator of "special" face processing (see, for example, Enns & Shore, 1997). However, the cause of the FIE is still under debate. Many theorists agree that the effect is related to the expertise involved in upright face processing (Goldstein & Chance, 1980). More specifically, it has been proposed that in addition to coding of "local" features, upright face processing relies in some special way on processing of the "configuration" of face features. It is the processing of this latter kind of information that may be disrupted when faces are turned upside-down.

However, what precisely is meant by "configural processing"? Different researchers at different times have discussed the coding of spatial relationships ("relational" coding), "second-order" configural processing, and "holistic" (template-like) processing. The definitions of these terms have not always been precise, and the existing literature does not allow us to draw firm conclusions about the nature of the "special" processing of upright faces that is impeded in inversion. In this paper we will use the face-inversion effect to explore more precisely what it is that is coded from upright faces and thus what may underlie our particular expertise with these patterns. First, however, we define the terms to be used and explored in our studies.

Local Feature Information

Local features are the independent constituent elements of an object. In Biederman's (1987) influential model of basic-level object recognition, an object is composed of a number of component shapes ("geons"), and the class of object is determined by both the geon type and the spatial relationships between different geons. In this theory, each geon can be considered a local feature of the object. The eyes, nose, and mouth are obvious candidates to be the local features of faces. Importantly, these features differ from each other in various dimensions, such as size, colour, texture, and shape.

Configural Information

The term configural information covers a broad range of different definitions, which are sometimes hard to distinguish conceptually or empirically from each other. A thoughtful theoretical approach stems from Diamond and Carey (1986), who proposed that faces contain local information, and configural information of two types: first-order relational and second-order relational information. First-order information consists of spatial relations between constituent parts of an object, such as the placement of the eyes above the nose. It is this first-order relationship that defines a set of face features as a face, in the same way that first-order relationships between geons define any other class of object. Second-order relational information describes the relative size of these spatial relations with respect to an underlying prototype. As all faces share a common first-order arrangement, the critical information in which faces differ from each other is this second-order information. Diamond and Carey's framework places special emphasis on the precise spatial relationships between different local features. This is a version of the configural processing view that we now refer to as "relational" processing, and such a view has been adopted by a number of other researchers too, though sometimes only implicitly. According to this position face processing is special in that it is particularly sensitive to this kind of relational information—a sensitivity that could also be disrupted by changing the orientation of the face and therefore cause the FIE.

This is not the only interpretation of what is meant by configural processing, however. The most extreme alternative view is "holistic" processing, which derives from earlier views that faces were processed as Gestalt patterns or templates. Recently, for example, Tanaka and Farah (1993) and Farah, Tanaka, and Drain (1995) defined a "holistic" representation as one in which the parts are not represented explicitly—or, in a less radical version, in which parts are explicitly represented to a lesser degree than in non-holistic processing. Thus the extreme holistic position is one where a face pattern is not decomposed into local features at all, and thus spatial relationships between these features will not be explicitly represented.

Configural Information and the Face-inversion Effect

The face-inversion effect plays a crucial role in discussing the representation and processing of faces, as it has often been argued that the explanation of the effect will reveal which kind of information is critical in face processing.

A number of studies has shown that the perception of local face features is hardly disrupted by inversion (e.g. Searcy & Bartlett, 1996; Leder & Bruce, 1998). Searcy and Bartlett made faces grotesque by either changing local features (such as blackening teeth) or distorting the spatial relationships (such as moving eyes apart). The local feature manipulation made faces more grotesque in both orientations, whereas the spatial manipulation affected the perception of upright faces but hardly made any difference to inverted faces. In a similar approach, Leder and Bruce found that increased face distinctiveness resulting from local manipulations (such as darker eyebrows) was rather insensitive to inversion, both in terms of rated distinctiveness and recognition memory performance, but effects based on changing spatial relationships almost vanished in inverted faces.

Diamond and Carey (1986) also produced evidence that the FIE did not result from the processing of local information from faces. They found that verifying the presence of local features was faster for upright than for inverted items, but that there was no difference in the size of this inversion effect for faces, houses, and landscape pictures. However, the position is rather less clear following results reported by Rhodes, Brake, and Atkinson (1993) who found that manipulations of individual face features gave inversion effects somewhat greater than those obtained when spatial relationships were manipulated. Their somewhat tautological conclusion was that such feature manipulations gave rise to configural processing. This highlights a major problem that plagues this field, which is that almost any manipulation made to a face affects both local and configural (however defined) properties. A different shaped nose affects the holistic pattern of the face and may also affect the spatial relationship existing between, say, the tip of the nose and the top of the lips. One aim of the experiments presented here was to specify local information in a way that minimized any confound between local and configural properties, in order to re-examine the effects of inversion on local face features.

Although effects of inversion on local features are still not clear, all researchers agree that inversion affects configural processing. Diamond and Carey (1986) stress the importance of second-order information for the occurrence of large inversion effects. They predicted that inversion effects occur whenever (a) all members of the class share a common configuration, (b) configural features allow the individuation of individual members, and (c) participants have the expertise to exploit these features. In support of their theory they provided evidence that inversion effects are also found in another field of expertise (using dog experts) when the three conditions are met. However, their experiments do not distinguish the relational from the holistic position. Moreover, Tanaka and Farah (1991) provided evidence that it is probably not the degree of reliance on second-order information in general that causes face-inversion effect as they did not find inversion effects for patterns (made of dots) that differed according to a second-order prototype.

Experiments such as those of Searcy and Bartlett (1996) and Leder and Bruce (1998) described above, showing that configural processing is impaired by inversion, do not readily distinguish the relational from the holistic position, as either could have caused the reported effects. Moreover, these studies do not tell us whether it is relational information manipulated in these experiments that is crucial to determine whether inversion effects are found.

The role of holistic processing in the inversion effect was recently investigated by Tanaka and Farah (Farah, Tanaka, & Drain, 1995; Tanaka & Farah, 1993). They defined holistic as "reduced part decomposition" and showed that whole faces were better recognized upright than faces learned as parts only, and that part-based learning of a face made its recognition insensitive to inversion. However, these experiments do not rule out the suggestion that relational information, rather than "holistic" information, is crucial in face recognition. A further problem is that the presentation of isolated facial features does not produce a face-like stimulus and therefore cannot be used to explore local processing within the context of real faces. Presumably one component of the strong holistic notion must be that the holistic processing should take effect whenever a whole face is encountered; thus the critical items to test these ideas are faces, not their constituent parts.

To summarize, although there is evidence that disrupted processing of configural information causes the face-inversion effect, it is still not clear whether it is the processing of spatial relationships between local features or holistic (non-decomposed) processing that is responsible. Support for the holistic version so far stems from experiments in which a lack of holistic processing has always been confounded with a lack of relational information. Similarly, when relational information has been investigated then in many empirical studies relational information has been confounded with both local and holistic information.

To tease apart the contributions played by different sorts of information, the present studies use sets of faces that differ in terms of local, relational or both sorts of information. In this approach local or relational information can be manipulated while the other dimension is kept constant. Moreover, specific hypotheses about the retrieval of critical information that distinguish relational from holistic processing were also tested.

In all five experiments, the participants learned the identities of a small set of faces, which they later recognized in complete or part versions, in upright or inverted orientation. In Experiment 1, two sets of faces were compared: one in which the faces have exactly the same local features but differed in terms of local relational information, and the other in which the faces to be identified shared exactly the same configuration but differed in terms of individual combinations of local information. Both sets of faces, therefore, incorporate differences in the "holistic" face pattern. If inversion effects are obtained only for the first set of faces, this would be strong evidence for the relational processing hypothesis and would rule out at least the strongest form of the hypothesis of holistic processing.

EXPERIMENT 1

In Experiment 1, the identities of faces were learned in an upright orientation and tested in upright or inverted orientation. Inversion effects were examined for two separate sets of faces. Within one set the faces differed from each other in respect to local feature information only (local), and within the other they differed from each other in respect to relational information only (relational). To avoid the confound of local with configural aspects, all faces in the first version (local) had the same shaped mouths, eyes, and noses, and the same style of hair, but differed in the colour of these features. Colour has been shown to affect recognition memory, though not identification, of faces by Kemp, Pike, White, and Musselman (1996). Recently, J.K. Lee and Perrett (1997) showed that colour also plays a role in the identification of familiar faces and concluded that "color information is helpful when differentiating between stimuli" (page 748). However, so that attention could not be focused on these colours in an isolated way, each face had a unique combination of different colour values, but each individual value was shared by at least one other face. The hues used were those that are naturally occurring in faces (e.g. yellow, pink, brown, red), but we selected values of these hues that could easily be distinguished from each other.[1]

Thus each face was identifiable as a specific ensemble of local information. In the second set of faces, each face again had identically shaped local features (eyes, mouth, etc.) but variation arose from the spatial relationships between these different features.

If the FIE is caused by the specific disruption of the processing of second-order spatial relationships then only the relational versions should reveal an inversion effect. On the other hand, if we take an extreme version of the "holistic" processing hypothesis (and/or if "local" information is also made more difficult to process by inversion) then we might expect inversion to disrupt the local versions as well, because even faces that differ only in terms of their colours form different holistic patterns.

Method

Participants

Sixteen graduate students and undergraduates from the University of Stirling participated; they either received course credit or were paid for their participation. All participants were tested individually.

[1] A coloured version of the stimuli of Experiments 1 and 3 is available on the web site http: //userpage.fu-berlin.de/-lederh

Materials

Two stimulus sets were created. Each set contained six faces and names. The faces were made using Mac-a-Mug features. The first set (local) contained faces that all had the same shaped features in the same spatial relationship to each other, but differed in respect of local properties that do not affect their configuration. The features used to create these faces differed in colour only. Figure 1L shows the stimuli used in this condition. Each feature (hair, mouth, eyes) varied in two colours. Importantly, each stimulus had a specific combination of local information, but each single feature value was shared by another face. Thus a single feature was not sufficient to identify a face—but each unique combination of different local features did identify a face. For example, there were two blond faces (Figure 1L, right column), which differed in terms of mouth colour and eye colour.

All faces were presented in an eight-bit colour format to reveal good colour quality and included facial texture to encourage configural processing (see Leder, 1996).

To create stimuli that differed in relational information all stimuli were made using the same set of eyes, nose, and mouth within the same facial outline. To exclude variation in terms of local information all faces in the relational set had exactly the same colours (all presented in greyscale).

To induce critical relational information in each face, features were displaced by moving them up or down and varying the distances between them. Figure 1R shows all the stimuli used in this set. Each stimulus of set R comprised a specific configural arrangement. For example, Face 1 has a raised nose, and in Face 2 the eyes have been moved closer together.

A total of 12 three-letter names were selected, which were randomly assigned to one of two sets. The names we used were: Bob, Don, Ian, Max, Rex, Sam, Ted, Joe, Guy, Ken, Tim, Les.

Procedure and Design

Participants were tested individually on each of the two sets in two sessions, which were separated by at least 30 min, during which the participants performed an unrelated filler task or other experiments.

The order in which the local or the relational faces were learned was counterbalanced across participants; thus participants started with either of the two versions. The allocation of names to faces was also counterbalanced across the two sets, so that each face was learned under two different names (by different participants). During all experiments the order of trials within each block was randomized by the experimental programme.

Study Phase. Pre-tests revealed that participants had difficulties noticing that the faces in the relational versions (which differed only in terms of critical configural information) differed from each other when they were presented one after the other. Moreover, learning the relational set to criterion proved more difficult than learning the local set. To get around this problem, all faces of each set were shown at the beginning of the experiment on the screen simultaneously. The instruction in this pre-exposure stressed that a set of "difficult" faces had to be learned in this experiment. Because of the greater difficulties learning the relational faces we exposed all stimuli in this condition for 2 min and increased the number of learning blocks (without feedback) from three to four. In the local condition all stimuli were also shown at the beginning of the experiment, but for about 30 sec only.

At the beginning of each experimental session the participants were told that they would be exposed to six different persons' faces, which they should try to learn and later recognize. The session started with three blocks (four for the relational set) in which each of the six stimuli was presented on the screen for 5 sec together with a short sentence saying "This is. . . " plus the name. Within each of these blocks the stimuli were presented in a randomized order. The training block was followed by a block in which each face was exposed, and the participants' task was to tell the

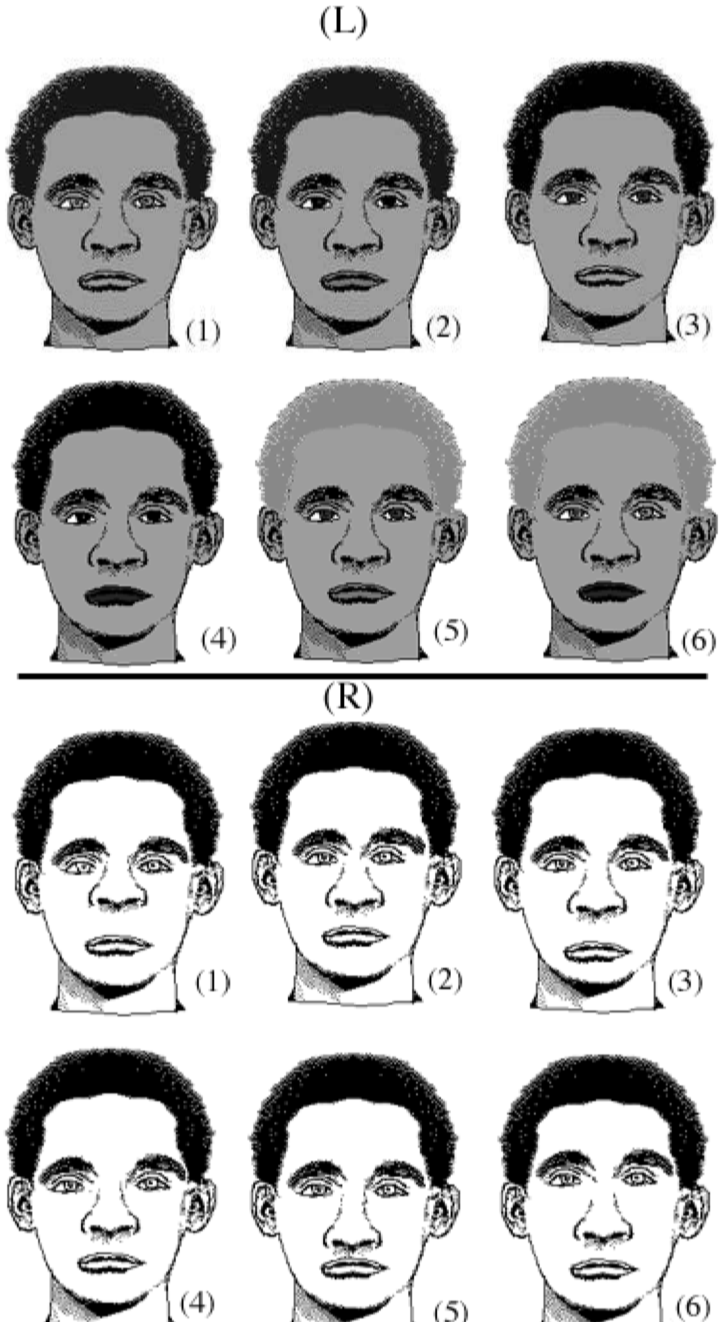

FIG. 1. Stimuli used in Experiment 1. The upper six faces represent the Local (L) versions, the lower six the Relational (R) versions.

experimenter the name of each face. During this phase each face was shown for 3 sec together with the question "Who is this?". After 3 sec the correct name was shown beneath the stimuli to provide feedback. If participants got fewer than four faces right in this block, the block was repeated until each participant met this criterion.

Test Phase. After a short break the test block began. During the test block all six names were shown together with a number of 1 to 6 beneath. The numbers were added to indicate which number to press on the keyboard for each name. In each trial one test face was presented beneath the list of names when the participants pressed the space bar. Participants were instructed to press the number attached to the name that they thought was the stimulus person's name. Each face was shown twice in each orientation (upright and inverted) yielding a total of 24 trials at test for each session (local and relational). The order of presentation of the stimuli at test was randomized for each participant.

Results and Discussion

Figure 2 shows the mean proportion of correct identifications at test in Experiment 1 in both conditions and both orientations. The by-subjects analysis of variance (ANOVA), F_1, used version (local versus relational) and orientation (upright versus inverted) as within-subjects factors, and the by-items analysis, F_2, had version as between- and orientation as within-items factors. The analysis revealed a main effect of orientation, $F_1(1, 15) = 12.164$, $p < .01$, $F_2(1, 10) = 16.72$, $p < .01$, and a significant interaction between both factors, $F_1(1, 15) = 9.638$, $p < .01$; $F_2(1, 10) = 10.6$, $p < .001$.

T tests revealed that there was no effect of inversion in the local versions, $t(1, 15) = 0.512$, $p < .31$ by subjects, $t(1, 5) = 0.698$, $p < .26$ by items, but a significant effect in the relational versions, $t(1, 15) = 4.855$, $p < .001$ by subjects, $t(1, 5) = 4.579$, $p < .01$ by items.

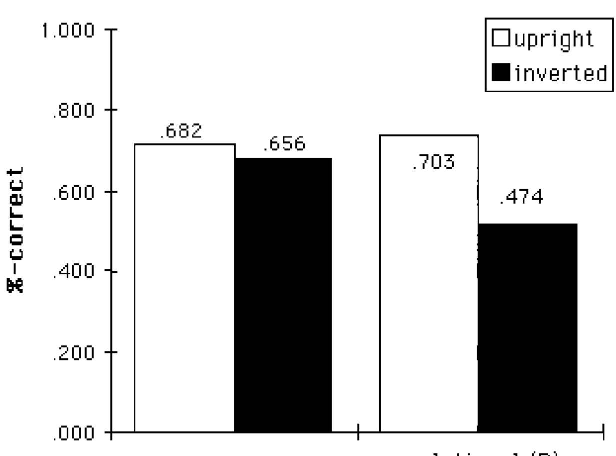

FIG. 2. Results of Experiment 1. The mean proportions (sampled over participants) of correct recognition are shown for upright and inverted presentation of the relational and local versions.

Thus the analyses reveal that inversion effects do not occur if faces differ only in terms of local information, but that large inversion effects occur if faces are individuated in terms of relational information. The lack of an inversion effect in the local faces suggests that not every face is processed holistically. From Experiment 1 it seems that relational information alone causes the FIE.

However, the experiment revealed that the discrimination of faces on the basis of relational information alone is difficult. This might be interpreted as weakening the argument that face processing is specialized for this kind of information. However, a set of more than two faces that differ only in terms of configural information in reality rarely occurs. No theory of face processing so far proposes that the distinctive local information that will normally also vary between items is not used to distinguish faces. However, the experiment does suggest that the processing of such local information does not contribute to the effect of inversion, even though it may normally contribute to the recognition of upright faces.

The local versions in Experiment 1 did not show any inversion effect. Whether this is really due to the lack of any variation in relational information was tested in Experiment 2. In Experiment 2 we included one condition in which relational variation was added (redundantly) to variation in local information. Comparing the outcomes of Experiments 1 and 2 will also reveal whether the extent to which face recognition relies on relational information determines the size of the FIE. Experiment 2 will reveal whether the availability of local information can compensate for the inversion deficits that accompany configural information: If that is the case then we expect inversion effects in the relational+local condition, which might be smaller than in the pure relational versions of Experiment 1.

EXPERIMENT 2

In Experiment 2, inversion effects were examined for faces that differ in respect to either isolated feature information only (local) or isolated feature information plus relational information (relational+local).

Method

Participants

Sixteen graduate students and undergraduates from the University of Stirling participated; they either received course credit or were paid for their participation. All participants were tested individually, and none had taken part in Experiment 1.

Materials

The local versions were the same as those in Experiment 1. The relational+local faces were produced by combining the two sets of Experiment 1. The features in the set of local faces were displaced to the same amount as in the relational versions of Experiment 1. All faces were presented in an eight-bit colour format to reveal good colour quality. The same names as those in Experiment 1 were used.

To reduce any confusion due to the similarity between the two sets, we used two different basic textures (a lighter and a darker one, see Footnote 1) to make the two sets appear more different. The assignment of the two basic skin colours was counterbalanced between the conditions.

All other aspects of the procedure were the same as those in Experiment 1, except that the pre-exposure and number of training blocks were now the same for both conditions and the same as in the local versions of Experiment 1.

Results and Discussion

Figure 3 shows the mean proportion of correct identification at test in both orientations sampled over all 16 participants. Again the local versions did not produce an inversion effect. Faces that differed in relational properties as well as in local features (relational+-local) do show an inversion effect.

The mean values of each participant were submitted to an ANOVA using version (local vs. relational+local) and orientation (upright vs. inverted) as within-subjects factors. The analysis revealed a main effect of orientation, $F(1, 15) = 7.535$, $p < .05$, and a significant interaction between both factors, $F(1, 15) = 6.713$, $p < .05$. Planned comparisons revealed that there was no effect of inversion in the local versions, $t(1, 15) = 0.59$, $p < .282$, but a significant inversion effect in the relational+local versions, $t(1, 15) = 3.381$, $p < .01$. The means of local-upright and relational+local-upright versions did not differ significantly, $t(1, 15) = -1.0$, $p = .33$. For each of the two sets the means were also sampled over items. An initial ANOVA on these means did not yield the significant main effects found in the analysis by subjects: orientation, $F(1, 10) = 2.87$, $p = .12$; Orientation × Version, $F(1, 10) = 1.65$, $p = .227$. However, t-tests on the item means supported the subjects' analysis revealing that inverting the local items did not produce an effect, $t(1, 5) = .225$, $p < .415$, but that there was an inversion effect for the relational+local items, $t(1, 5) = 3.471$, $p < .05$.

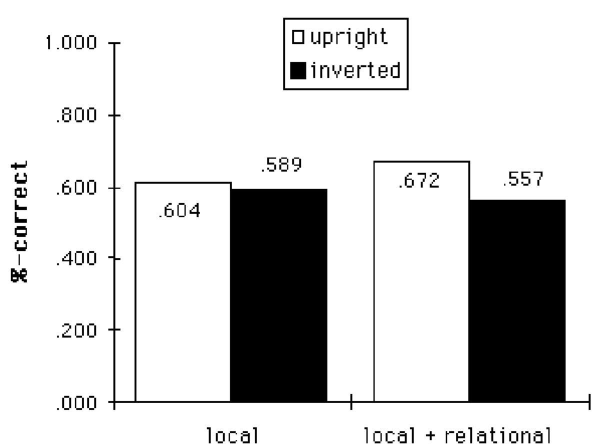

FIG. 3. Results of Experiment 2. The mean proportions (sampled over participants) of correct recognition are shown for upright and inverted presentation of the local and the relational-plus-local versions.

Thus Experiment 2 revealed again that without the presence of critical relational information there are no inversion effects, even with stimuli that all share the same arrangement of features. This clearly suggests that to get an inversion effect it must be possible to differentiate the exemplars with respect to relational information. Consistent with this, adding relational information to local information in Experiment 2 had the effect of inducing an inversion effect (a significant effect by subjects, and on a t test by items). The availability of distinctive spatial relationships did not significantly increase the recognition of upright faces, although the results show a trend for a slightly superior recognition in upright presentation.

Colour is a variable that has been shown to affect the recognition of familiar faces (J.K. Lee & Perrett, 1997) and faces differ in the texture and colour of local features such as eyes and facial hair. Furthermore, V.H. Lee, Kim, and Park (1996) have presented results that show that colour might be useful to segment regions of the face by thresholding the hue space of colour distributions. However, the faces differing in colour in Experiments 1 and 2 were easier to learn than the configural faces; and the colour variations, although natural hues were used, were more extreme than those found in most faces. It is possible, therefore, that the effects demonstrated using local variations of colour are unusual in some way. We therefore decided to replicate the findings using another sort of "local" information in Experiment 3 to increase the scope of our conclusions. In Experiment 3 differences in brightness values were used instead of colour variations. Kemp, Pike, White, and Musselman (1996) have demonstrated that it is luminance rather than the values that play the more important role in face recognition – and that it is such differences in luminance that may account for the difficulty of recognizing faces from photo-negatives. Other evidence on the importance of luminance values in face processing stems from an investigation reported by Bruce, Burton, Carson, Hanna, and Mason (1994) who found that the amount of repetition priming between pictures of familiar faces was affected by the extent to which test faces duplicated the precise pattern of grey-scale values shown in the prime pictures.

Thus, in Experiment 3 luminance level of "local" information was compared against relational information in upright and inverted presentation.

EXPERIMENT 3

In Experiment 3, inversion effects were examined for faces, which differed either in respect to local information (brightness values) or in terms of relational information only (as in Experiment 1). The faces used in Experiment 3 were constructed from a different facial outline and different features from those used in Experiments 1 and 2.

Method

Participants

Sixteen graduate students and undergraduates from the Free University of Berlin participated for course credit.

Materials

The local versions were constructed similarly to those in Experiment 1. The information discriminating between the faces now consisted of brightness values of the features, with each face again consisting of a specific combination of local features. The following grey values (ranging from 0, which is black, to 255, which is white on the RGB grey value scale) for each feature were used (hair: 153, 77, 0; eyebrows: 177, 106, 29; mouth: 228, 152, 75). None of the features was diagnostic for one individual face—it was combinations of brightness values that were critical (see Footnote 1). The stimuli used in the local version of Experiment 3 are shown in Figure 4.

The relational variations were made by displacing features by an amount very similar to that in the relational versions of Experiment 1. The same names as those in Experiment 1 were used.

All other aspects of the procedure were the same as those in Experiment 1 except that the pre-exposure and number of training blocks were now the same for both conditions, as they turned out to be of a similar level of difficulty for most participants. For the pre-exposure all stimuli were presented for about 15 sec. In the learning phase all faces were presented four times with their names, and participants received two additional learning blocks if they did not reach the learning criterion of five correctly named faces.

Results and Discussion

Faces that were different in local features were correctly recognized in 74% (SD = 0.20) in upright and 72% (SD = 0.18) in inverted presentation. The relational versions were recognized in 74% (SD = 0.21) in upright and 57% (SD = 0.22) in inverted orientation. Again the local versions did not produce an inversion effect. As in the previous experiments, faces that differed in terms of relational information did show an inversion effect.

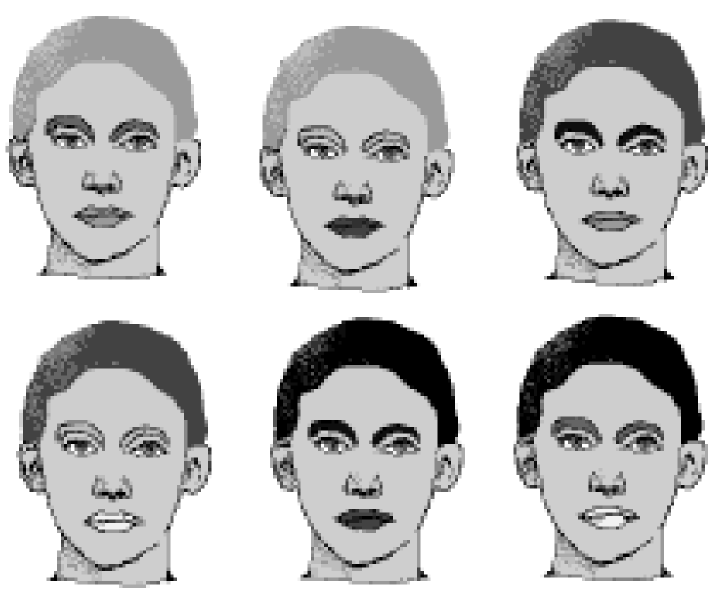

FIG. 4. Stimuli used in the local versions of Experiment 3.

ANOVAs showed a main effect of orientation, $F_1(1, 15) = 11.693, p < .01$; $F_2(1, 10) = 5.207$, $p < .05$ by items, and a (near)-significant interaction between both factors, $F_1(1, 11) = 10.804, p < .01$; though $F_2(1, 10) = 3.591$, $p = .087$. *T*-tests revealed that there was no effect of inversion in the local versions, $t(1, 15) = 0.476$, $p = .32$ by subjects, $t(1, 5) = 0.542$, $p = .306$ by items, but a significant inversion effect in the relational versions, $t(1, 15) = 4.398$, $p = .001$ by subjects, $t(1, 5) = 2.236$, $p < .05$ by items.

Thus Experiment 3 again revealed that without the presence of critical relational information there are no inversion effects, even with stimuli that all share the same arrangement of features. This clearly suggests that to get an inversion effect it must be possible to differentiate the exemplars with respect to relational information. Consistent with this, relational information in Experiment 3 had the effect of inducing an inversion effect.

To summarize, Experiments 1, 2, and 3 suggest that the inversion effect in faces arises from the processing of critical relational cues from upright faces. The use of these cues is disrupted when the faces are presented upside-down.

However, the results of Experiments 1, 2, and 3 do not tell us much about how relational information is processed. One possibility is that individual spatial relationships are processed and made explicit—a kind of local processing of relational information, in which, for example, "distance between the eyes" becomes one of the features listed in the description of a face. Alternatively, spatial relationships may be more embedded within the overall pattern of the face—a kind of "holistic" representation of relationships. Experiments 1 to 3 have allowed us to exclude one extreme view of the holistic hypothesis. Those stimuli in Experiments 1 to 3 that did not show inversion effects (the local versions) obviously were face stimuli and were presented as facial wholes. Moreover, as they differed in terms of unique combinations of colour or brightness features there should be no bias towards processing them as lists of unrelated features. As these items did not give rise to inversion effects, we can discount certain interpretations of such effects in terms of holistic processing. However, it is possible that specific combinations of colour or brightness are not among the most diagnostic features normally important for face processing, and so it may be premature to dismiss all possible versions of a holistic processing account at this stage.

In Experiment 4 the nature of relational information was further investigated, contrasting two interpretations. First, relational information could consist of isolated spatial relationships between local features. Evidence for the use of this kind of information stems from Haig (1986). He reported that face recognition is sensitive to simple spatial relations, such as nose–mouth distances. Configural information extracted from the face on this view could be encoded as a list of values of discrete spatial relationships. This kind of information would be available within the stimuli used by Tanaka and Farah (1991) and Farah et al. (1995) in cases where more than one feature was present. In these conditions, inversion effects were found. Moreover, it is this kind of information that is omitted in those stimuli that produced no inversion effects in previous studies (Farah et al., 1995). If such isolated spatial relationships are made explicit in descriptions of faces then it would be expected that people should be able to recognize isolated spatial relationships between local features without the presence of a facial context. Moreover, if such spatial relationships were difficult to process in upside-down faces, then even isolated spatial cues should be even harder to process when inverted.

According to the alternative view, however, the relational information used in face processing depends upon the facial context, which might produce a more complex pattern of relations between all elements that make up one individual face. In its strongest version, such a theory would suggest that "isolated" or "decomposed" recognition of spatial relationships would not be possible. In a weaker version the availability of the facial context, such as the facial outline and the hair, should significantly help to retrieve relational information. Moreover, the facial context could be used either to retrieve local relations, such as the distance between the eyes, or even to retrieve the exact location of each feature. According to this view, spatial relationships will be easier to retrieve at test when an appropriate facial context is provided, and only under these circumstances will an inversion effect be found.

EXPERIMENT 4

Experiment 4 was designed to investigate whether the relational information in faces is explicitly represented and useful for recognition. It also tested whether the retrieval of relational information requires a facial context by testing whether critical relational information is better recognized when a redundant context is added that favours holistic processing. All conditions were tested in both orientations to reveal which kind of information is (most) disrupted by inversion.

In a recognition paradigm similar to that in the previous experiments, participants learned stimuli that differed from each other in one relational feature only. At test, participants saw one of the following: (1) the isolated parts that contained the individual relational information (IsoRel); (2) the isolated parts plus a highly redundant context (CtxRel); (3) the same context, which included one of the facial features only that was involved in the critical relational information (CtxPart), but which excluded direct information about the relational information itself. The recognition rate of the isolated relational features reveals whether relational features are explicit in face recognition. A comparison between the IsoRel and CtxRel reveals the role of a context, and the CtxPart reveals to what extent the context is useful, if the relational information is not completely available. Tests of recognition of the full faces were also included for comparison.

Method

Participants

Sixteen undergraduates from the University of Fribourg participated for course credit. All participants were tested individually, and none had participated in any of the previous studies.

Materials

In Experiment 4 a new set of stimuli was used. The eight faces were all constructed from the same set of Mac-a-Mug features. Each face was manipulated to be distinctive only through one relational feature. Figure 5 shows the faces and the different part-versions that were used at test. Faces displayed the following relational features: The eyes were moved closer together in Face 1 and

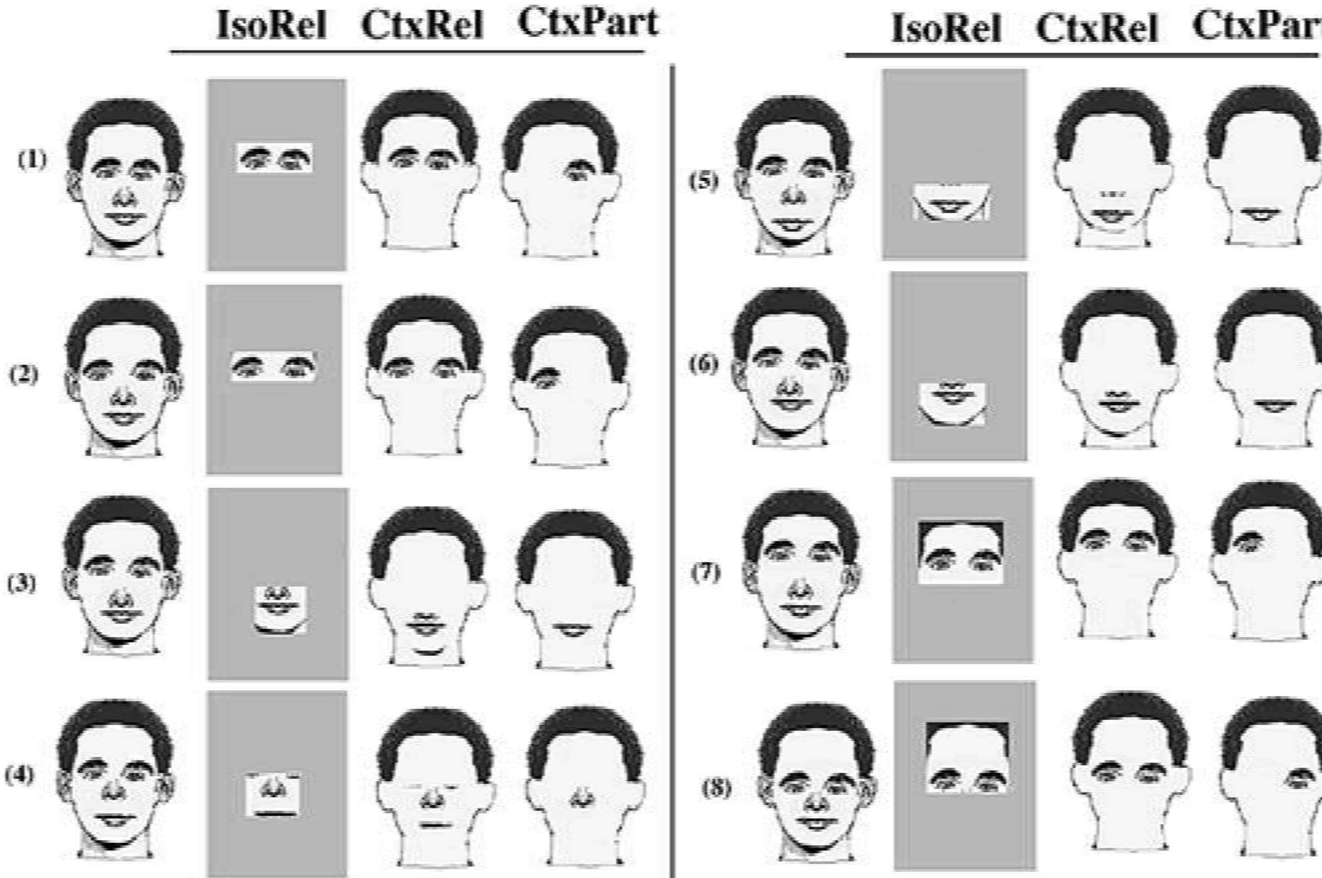

FIG. 5. Stimuli used in Experiment 4. In each set of four columns, the left-hand column shows the face, and the remaining columns show the different stimuli used for each face: IsoRel contains the critical relational feature in isolation; CtxRel has a face outline added; and in CtxPart only one of the features that are part of the critical relational feature is shown at its location within the face context.

further apart in Face 2; The nose was moved down in Face 3 and up in Face 4; The mouth was moved down in Face 5 and up in Face 6; and the eyes were moved up in Face 7 and down in Face 8.

In addition to these whole faces, three different test items were created for each face. The first version contained the critical distinctive feature only (IsoRel), whereas in the second version the same feature plus a redundant context was presented (CtxRel). If it is the facial context that is important for the processing of a stimulus as a face, and hence the spatial relationships within the face, then recognition performance should be much better with the CtxRel versions. However, if is the individual spatial relationships that determine performance, then performance in the IsoRel should be equally good. The CtxPart version (also shown in Figure 4) included the same context as the CtxRel versions, but one of the elements of the features that together determine the relational information was omitted. These versions thus contain much more holistic information that the IsoRel condition, but lack the primary spatial relationships themselves. Nonetheless there is sufficient information in the CtxtPart faces to identify each face uniquely. The full faces were also included in the test.

Because Experiment 4 was run in Fribourg, a new set of eight German names was used. The names were Klaus, Simon, Bernd, Erich, Georg, Willi, Frank, Heiko.

Procedure and Design

In Experiment 4, participants learned a single set of eight faces. The procedure was similar to that in the previous experiments. As in Experiment 2, all faces were shown at the beginning for 20 sec, but as a printout on paper.

Learning Phase. The pre-exposure was followed by eight learning blocks (without feedback), in which each face was presented for 5 sec. After the faces had been learned, a first test revealed whether the learning had been successful. When all faces were named correctly in a free naming task the experimenter started the test block. When not all the faces were correctly named, the importance of "knowing them all for the purpose of the experiment" was stressed, and all faces were shown in one or more blocks of three repetitions (until the criterion was met) before the tests started.

Test Phase. All four versions were used in the recognition test. Each stimulus was shown twice in each version in both orientations. All 128 trials were presented in a randomized order. Participants always saw the list of names and indicated their choice by pressing a number attached to each name of the face to which they thought the stimulus would belong.

Results and Discussion

The results of Experiment 4 are presented in Figure 6. The mean recognition rates (proportion correct) sampled over all participants are shown. In the CtxPart condition it turned out that one stimulus by mistake did not include the moved features (right-hand panel for Face 3 in Figure 5). This stimulus was therefore excluded from the analyses. Thus in the analyses by subjects this stimulus was excluded from the CtxPart condition. For the item analysis the stimulus was completely excluded in all conditions. Figure 5 shows that performance in all conditions was better than chance, and that in all four versions performance was better when faces were presented in upright orientation. The mean values of each participant were submitted to a 4×2 ANOVA, using the four versions and the two orientations as within-subject factors. For each condition the means

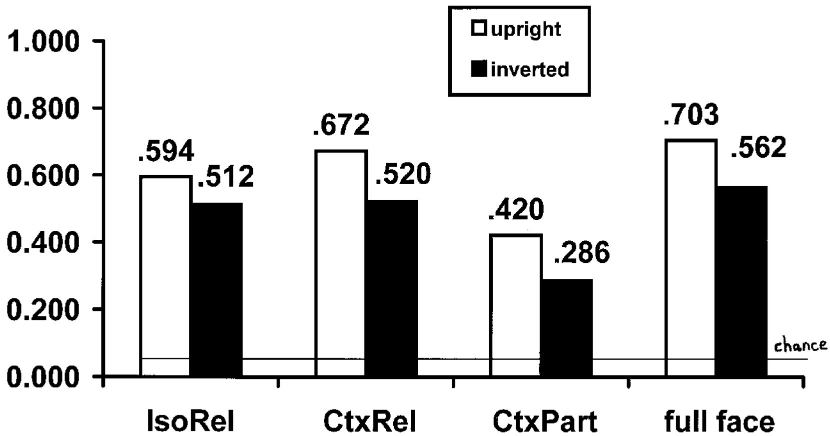

FIG. 6. Mean results of Experiment 4. The mean proportions (sampled over participants) of correct recognition are shown for the critical relations presented in isolation (IsoRel); with the facial context (CtxRel); and when single elements (which were part of the critical relations) were presented in the facial context (CtxPart). Recognition of the complete stimuli in upright and inverted orientation is also shown.

were also sampled over items. Both analyses used orientation as within-subjects factors. The analysis revealed a main effect of version, $F_1(3, 45) = 26.364$, $p < .001$; $F_2(3, 18) = 8.044$, $p < .001$, and a main effect of orientation, $F_1(1, 15) = 22.970$, $p < .001$; $F_2(1, 6) = 32.112$, but no interaction between both factors.

Simple comparisons (t tests) revealed that there were inversion effects in all four conditions. T tests also revealed that performance in the CtxPart condition was lower than that in all other conditions, which did not differ from each other, although there was a trend for the full faces to be better recognized than the IsoRel faces, $t(1, 15) = 2.101$, $p < .053$, which was not found in the CtxRel condition. However, IsoRel and CtxRel did not differ from each other, $t(1, 15) = 1.195$, $p < .250$.

The experiment thus clearly showed that information about critical spatial relationships is represented in memory. Performance was well above chance both when the distinctive relational features were presented in isolation and when they were embedded in the facial outline. The results of Experiment 4 therefore show that relational information is an important part of the memory representation of a face.

Another aim of Experiment 4 was to examine further the notion of holistic processing. Farah et al. (1995) accept that the crucial issue is whether it is assumed that the face pattern is decomposed into parts whose spatial relationships are described with respect to each other, or whether the pattern is analysed as a non-decomposed whole. In the former case, people should be able to recognize isolated spatial relationships between depicted local features as confirmed in Experiment 4. According to the latter view, face recognition should improve as more of the whole pattern is shown, even if the additional information is redundant, and Experiment 4 suggests this too (in the trend for whole faces to be better recognized than those in the IsoRel condition). However, the comparison between IsoRel and CtxRel reveals that the condition that was expected to favour holistic processing did not increase the performance significantly. Moreover, when the CtxPart version is considered then it is obvious that the exclusion of the elements that constitute the critical relation information cannot be compensated for by the context. Although performance is still better than chance, this condition reveals rather poor performance. In the distinction between relational and holistic versions of the configural processing hypothesis, this result again favours relational information. Thus, Experiments 1 to 4 taken together appear to allow us to reject the notion of holistic processing, at least in its strong versions.

Moreover, the results of Experiments 1 to 4 all suggest that the critical information in face processing is by its nature relational rather than higher order. The FIE according to these results is caused by a disruption of relational information.

However, there are still some loose ends in past research, where some inversion effects involving local features have been reported. Specifically, Rhodes et al. (1993) observed strong inversion effects when faces whose features had been exchanged had to be recognized. Rhodes et al. interpret this result to mean that local features are processed configurally, but this is a rather circular account. Experiments 1 to 3 strongly suggested that local features alone do not give rise to inversion effects, but they may be criticized because the use of colour and brightness as discriminating dimensions may limit the strength of the conclusions drawn. Experiment 5 therefore returns to the issue of pitting local feature against spatial relationships, to test further the proposal that only the spatial relationships become difficult to process following inversion. Experiment 5 also follows up Experiment 4 by investigating further the retrieval of isolated relational information and comparing this with isolated local cues.

EXPERIMENT 5

In Experiment 5, two questions were addressed. First, we investigated whether the retrieval of isolated relational information can be demonstrated when the faces also differ in their distinctive local information. The experiment thus adds to the information provided by Experiments 2 and 4. In Experiment 5 each face was distinguishable by a specific relational feature, shown in Figure 7. For example, the uppermost face in the left row of Figure 7 has a lowered eye region, but the local elements constituting this relational information are shared by another face. Therefore, the retrieval of relational information is impossible from the recognition of the local features used. However, in addition, each face was unique by at least one local feature (eyes, nose, mouth) that was not shared by another face. We can thus assess how relational (but not local) and local (with minimal relational) cues are affected by inversion when both these cues are available but separable within the same faces.

Second, Experiment 5 extends the findings of Experiments 1 to 3 (comparing relational with local) by using local information that varies in terms of feature shape rather than colour or brightness. A lack of an inversion effect in this condition thus would strengthen the conclusions drawn from the results of Experiments 1 to 3.

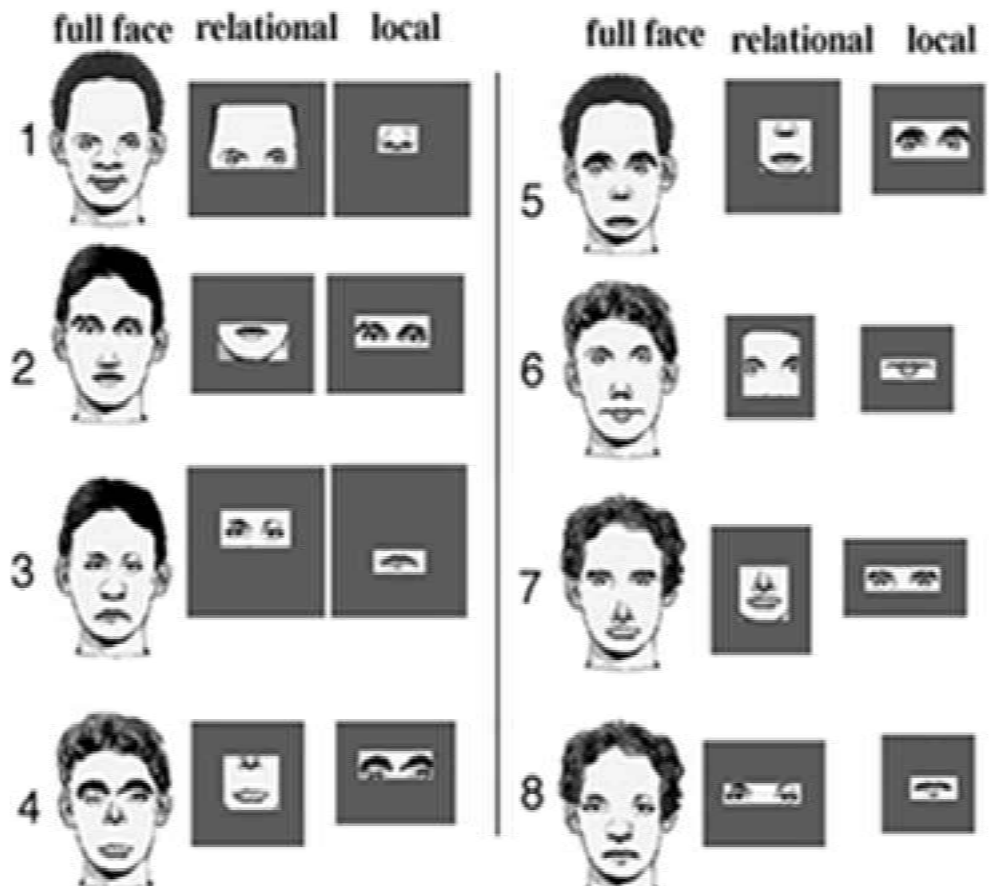

FIG. 7. Stimuli used in Experiment 5. Each face contained one critical local feature (eyes, nose, or mouth) and a critical relational feature.

Method

In Experiment 5, retrieval of relational and local information is tested in faces that differ from each other in terms of both these types of information.

Participants

Eighteen undergraduates from the University of Fribourg participated for course credit. All participants were tested individually, and none had participated in any of the previous studies.

Materials

The stimuli used in Experiment 5 are shown in Figure 7. Each face was composed of eyes, nose, mouth, and hair, and some of the features were shared between faces. Most important, each face had one unique local and one unique relational feature. For example, Face 1 has the eyes moved down (relational feature) and a unique nose. To exclude the possibility that the unique relational feature might be recognized by the local features involved, there was always at least one other face that shared the local features that made up a unique relational feature. Thus Face 1 has the unique high forehead but the same eyes as Face 6. This kind of relation holds for all other relational features. The detailed description of the critical features are listed in the Appendix. Experiment 5 was run in Fribourg, therefore the same name set as that in Experiment 4 was used.

Procedure and Design

The procedure was similar to that in the previous experiments. In Experiment 5 all participants learned the complete set of eight faces. Each face was shown with its name five times for 5 sec during the learning phase. Before participating in the test all faces had to be named once correctly when presented without the name. Otherwise three additional learning trials were run until this criterion was met.

After all faces were named correctly in free naming, the experimenter started the test block. Figure 7 shows the versions used in the test phase. The complete faces were used as well as versions in which the relational information only or the critical local feature only was shown. All versions were presented twice in upright and inverted orientation, which resulted in 96 test trials. Participants always saw the list of names and indicated their choice by pressing a number attached to each name. Trials were presented in a randomized order. As in the previous experiments, two different assignments of names to faces were balanced across participants.

Results and Discussion

The results of Experiment 5 are presented in Figure 8. The figure shows the mean proportions of correct identification sampled over the participants. ANOVAs with repeated measures on the two factors (the three versions at test, and orientation) revealed a main effect of the version, $F_1(2, 34) = 143.481$, $p < .001$; $F_2(2, 14) = 60.706$, $p < .001$, a main effect of orientation, $F_1(1, 17) = 15.94$, $p < .001$; $F_2(1, 7) = 15.244$, $p < .01$, and an interaction between the two factors, which was significant by subjects, $F_1(2, 34) = 3.670$, $p < .036$, though not by items, $F_2 = 2.623$, $p = .108$. Planned t tests revealed that there was an inversion effect with the full faces, $t(1, 17) = 3.887$, $p < .001$ by subjects;

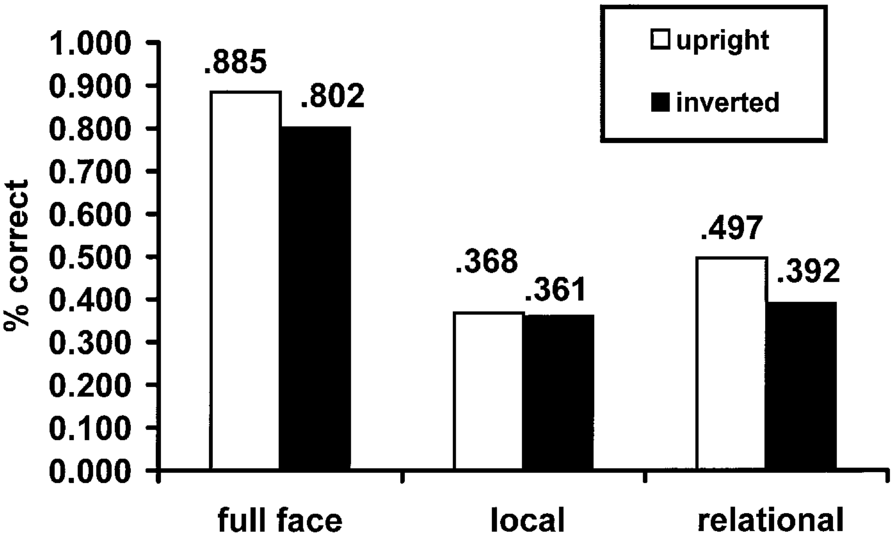

FIG. 8. Mean results of Experiment 5. The mean proportions (sampled over participants) of correct recognition are shown for the full faces, the relational features, and the local features, in both orientations.

$t(1, 7) = 3.384$, $p < .01$ by items; and with the relational versions, $t(1, 17) = 3.896$, $p < .001$ by subjects; $t(1, 7) = 3.230$, $p < .01$ by items, but not in the local versions, $t(1, 17)$ $p = .215$, by subjects, and $t(1, 7)$ $p = .194$, by items. Performance was higher with the full faces than with both isolated versions: compared with local versions, $t(1, 17) = 15.486$, $p < .001$ by subjects, and $t(1, 7) = 9.128$, $p < .001$ by items; and the relational versions, $t(1, 17) = 12.435$, $p < .001$ by subjects, and $t(1, 7) = 8.816$, $p < .001$ by items. The difference between the relational versions and the local versions was significant by subjects, $t(1, 17) = 2.808$, $p < .01$, but not by items, $t(1, 7) = 1.809$, $p < .113$.

These results reveal that the retrieval of relational information is possible even in faces that also differ distinctively in terms of local information. Indeed there is a trend (significant in the participants analysis) for performance in the relational condition to exceed that in the local condition.

The second issue investigated in Experiment 5 was concerned with the inversion effect in the retrieval of facial information, returning to the questions explored in Experiments 1, 2, and 3. Consistent with these earlier experiments, the full faces showed an inversion effect, and the same was found when relational information was presented alone. When local information was presented, however, there was no inversion effect.

One possible artefact must be considered. The construction of stimuli made it possible that some of the errors in the relational condition arose because the participants wrongly identified the "other" faces whose features were shared with those shown in the isolated relational condition. An analysis of errors showed that the proportion of errors that were of this type was 23% in the upright relational conditions and 27% when relational features were presented upside down. This shows that the confusion errors do not affect the occurrence of inversion effect in the relational conditions. Moreover, the inversion effect in this condition was still significant: with items, $t(1, 7) = 2.184$, $p < .05$,[2] even when these cases were eliminated.

Thus Experiments 1, 2, 3, and 4 taken together show that inversion effects do not occur in the recognition of local features, but when relational information has to be recognized. Experiment 5 shows that this applies even when the local features differ in form and not just colour or brightness values.

GENERAL DISCUSSION

The aim of the present study was to clarify the contribution and nature of configural information in the face-inversion effect and therefore to enhance our understanding of the processes of face recognition more generally. In Experiment 1 we found that faces that differed in terms of local feature information did not show an inversion effect, whereas faces that differed only in terms of relational information did show inversion effects.

As both sets of faces consisted of whole facial patterns, and neither sort could be processed by selected focus on isolated features (as it was combinations of local features that were important in the local condition), the experiment rules out the possibility that all whole faces are processed in a mandatory holistic fashion and that it is the disruption of this kind of holistic processing that yields inversion effects. Rather the results allow a simple explanation: The critical information that is used in face recognition and that is disrupted by inversion consists of relations between single features.

The broader conclusion that the extent to which faces show inversion effects is related to the extent that they may be distinguished on the basis of relational information was investigated in Experiment 2. Although the presence of local information somewhat reduced the size of the inversion effect, clear inversion effects did emerge when relational information was used to distinguish faces from each other. These results are well in accordance with Diamond and Carey's (1986) conditions for disproportionate inversion effects and provide a direct test of their claims.

A second point is concerned with the nature of configural information and whether relations between features are (a) directly encoded or (b) retrieved only by a holistic process (Farah et al., 1995). We chose an indirect test for this question (Experiment 4). Faces were used that differed from each other only in terms of individual relational features. We proposed that if (a) is correct then participants should be able to recognize faces when the relational features were presented in isolation. On the other hand, if (b) is correct, then embedding the relational features in a highly redundant face context should considerably help the recognition. We found a remarkable ability to recognize isolated relational information, and there was no significant increase when a redundant context (which was included to encourage holistic processing) was added. In all conditions that included the critical relational information performance was high and sensitive to orientation. When the two elements that constituted the relational feature were incomplete then performance was strongly disrupted although the context still was available.

[2] Due to an administrative error it was not possible to analyse for this effect by subjects.

The present results support the notion that relational information is processed in a local and possibly independent way. This is in accordance with recent findings of Macho and Leder (1998) who reported that this kind of relational information (the eye distance) did not interact with the availability of other local features in a face similarity decision task. In a recent approach Rakover and Teucher (1997) report findings according to which the inversion effect is relatively well predicted by the performance on local features. This apparent contradiction to our conclusions might also be removed when configural information is seen as consisting of local relationships and thus also contained within such local features as eye distances and shapes of foreheads.

A third aim was to examine the relative contribution of different classes of information. Adding local information to the relational information in Experiment 2 compensated for the disruption of relational information—at least in part. However, inversion effects remained. In Experiment 5 a different sort of local information was used: Again, inversion effects were found only in the relational condition.

In sum, if facial stimuli are used that avoid the confounding of local with configural information, this can help us to understand the role that is played by each kind of information in face information processing. The results of our experiments suggest that local and configural information may both contribute to upright face processing, but that inversion did not significantly affect local but configural information. In all our experiments we found that it is indeed the relational notion of configuration that provides the best account of the critical information used to recognize faces. Moreover, when faces vary only on local features, their recognition is unaffected by inversion.

REFERENCES

Bartlett, J.C., & Searcy, J.H. (1993). Inversion and configuration of faces. *Cognitive Psychology*, *25*, 281–316.

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94*(2), 115–147.

Bruce, V., Burton, A.M., Carson, D., Hanna, E., & Mason, O. (1994). Repetition priming of face recognition. In C. Umiltà and M. Moscovitch (Eds.), *Attention and performance XV* (pp. 179–201). Cambridge, MA: MIT Press.

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

Enns, J.T., & Shore, D.I. (1997). Separate influences of orientation and lighting in the inverted-face effect. *Perception & Psychophysics*, *59*(1), 23–31.

Farah, M.J., Tanaka, J.W., & Drain, H.M. (1995). What causes the face inversion effect? *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 628–634.

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

Haig, N. (1986). Exploring recognition with interchanged facial features. *Perception*, *15*, 235–247.

Kemp, R., Pike, G., White, P., & Musselman, A. (1996). Perception and recognition of normal and negative faces—the role of shape from shading and pigmentation cues. *Perception*, *25*, 37–52.

Leder, H. (1996). Line drawings of faces reduce configural processing. *Perception*, *25*, 355–366.

Leder, H., & Bruce, V. (1998). Local and relational aspects of face distinctiveness. *The Quarterly Journal of Experimental Psychology*, *51A*, 449–473.

Lee, J.K., & Perrett, D. (1997). Presentation-time measures of the effects of manipulations in color space on discrimination of famous faces. *Perception*, *26*, 733–752.

Lee, V.H., Kim, J.S., & Park, K.H. (1996). Automatic human face location in a complex background using motion and color information. *Pattern Recognition*, *29*(11), 1877–1889.

Macho, S., & Leder, H. (1998). Your eyes only? A test of interactive influence in the processing of facial features. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 1486–1500.
Rakover, S.S., & Teucher, B. (1997). The facial inversion effects: Parts and whole relationship. *Perception & Psychophysics*, *59*(5), 752–761.
Rhodes, G., Brake, S., & Atkinson, A.P. (1993). What's lost in inverted faces? *Cognition*, *47*, 25–57.
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.
Tanaka, J., & Farah, M. (1991). Second-order relational properties and the inversion-effect: Testing a theory of face perception. *Perception & Psychophysics*, *50*, 367–372.
Tanaka, J., & Farah, M. (1993). Parts and wholes in face recognition. *Quarterly Journal of Experimental Psychology*, *46*, 225–245.
Valentine, T. (1988). Upside-down faces: A review of the effect of inversion upon face recognition. *British Journal of Psychology*, *79*, 471–491.
Yin, R.K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, *81*, 141–145.

Original manuscript received 9 February 1998
Accepted revision received 10 May 1999

APPENDIX

The table shows which local and relational features made each face in Experiment 5 distinctive and which other face shared the local features that were part of the relational feature in each face.

Face	*Local*	*Relational*	*Shared*
1	nose	eyes–down	6
2	eyes	mouth–up	5
3	mouth	eyes–closed	8
4	eyebrows	nose–up	7
5	eyes	mouth–down	2
6	mouth	eyes–up	1
7	eyes	nose–down	4
8	mouth	eyes–wide	3