

Upright versus upside-down faces: How interface attractiveness varies with orientation

KARL-HEINZ BÄUML

Universität Regensburg, Regensburg, Germany

A choice experiment is reported in which all pairs and triples of faces from a set of eight moderately attractive faces were presented, both upright and upside down, to 103 subjects. In each orientation, the subjects had to select the face that appeared more (pairs) or most (triples) attractive to them. For each orientation, the preference probabilities that arose from the pair and triple comparisons could be described by the BTL rule (Luce, 1959). Thus, each face was represented by two scores, one reflecting its attractiveness in the upright orientation and the other reflecting its attractiveness in the inverted orientation. Orientation affected the preference probabilities. Qualitatively, score ratios between faces decreased from upright to inverted orientation, suggesting that the faces became less discriminable by inversion. Quantitatively, the effect of inversion could be described by a simple rule that assumes a face's two attractiveness scores to be affinely related across orientations. This result indicates that inversion affected all faces about equally. The present findings are discussed with respect to faces' first- and second-order relational properties, a distinction emphasized in current theories of face perception. They suggest that the processing of first- and second-order relational properties is impaired by inversion to roughly the same degree.

Recognition of faces is disrupted by inversion to a far greater extent than is recognition of other classes of visual objects, such as houses, airplanes, or landscapes. In the literature, this phenomenon has been referred to as the face-inversion effect. It was first demonstrated by Yin (1969). Subsequent studies have replicated this effect under quite different experimental conditions, and in this way have demonstrated the very general and robust nature of the phenomenon (see Valentine, 1988, for a review).

Originally, the inversion effect was interpreted as evidence for specialized face-recognition processes in visual information processing (Yin, 1969, 1970). Then other approaches were proposed in order to account for the inversion effect without presuming face-specific processes (Diamond & Carey, 1986; Goldstein & Chance, 1980; Rock, 1973). The most influential of these approaches is that of Diamond and Carey. These researchers proposed a distinction between two types of spatial information that underlie the processing of visual objects: first-order relational and second-order relational properties. While first-order relational properties refer to information about the spatial relationships among parts of an object, second-order relational properties refer to information about the spatial configuration between the

parts of an object, on the one hand, and the prototypical spatial configuration of its parts, on the other (Diamond & Carey, 1986; Tanaka & Farah, 1991). Diamond and Carey hypothesized that inversion was particularly sensitive to the processing of second-order relational properties while it did not affect the processing of first-order relational properties. They attributed the inversion effect to the use of second-order relational properties that are important for, although not unique to, face recognition.¹ Indeed, the results of several studies supported the view that inversion affected the processing of second-order relational properties (Diamond & Carey, 1986; Maruyama & Endo, 1984; Sergent, 1984; Thompson, 1980; Young, Hellawell, & Hay, 1987). On the other hand, none of these studies demonstrated that inversion did not affect first-order relational properties to the same degree (Rhodes, Brake, & Atkinson, 1993; Tanaka & Farah, 1991; Valentine, 1988).

Recently, two studies reported experimental results that were interpreted as evidence against the hypothesis that second-order relational properties were responsible for the inversion effect. In a mental rotation experiment, Valentine and Bruce (1988) found response-time patterns for rotated faces which were similar to those that Shepard and Metzler (1971) found for rotated three-dimensional block drawings. The response time of same-different judgments increased linearly as a function of rotation angle when the second of a pair of faces was rotated away from the vertical. Valentine and Bruce argued that this result did not support the idea that matching two upright faces involved a process that was qualitatively dif-

The author's thanks are extended to Hede Helfrich-Hölter, Jim Tanaka, and Tim Valentine for their comments on an earlier draft of this manuscript, and also to Armin Hartinger and Maria Schnelzer for their help in the experimental work. The author's mailing address is Institut für Psychologie, Universität Regensburg, 93053 Regensburg, Germany (e-mail: heinz@rpss3.psychologie.uni-regensburg.de).

ferent from that used to match one upright and one inverted face. Instead, not only first-order relational but also second-order relational properties should have been extracted from a face when it was presented upside down.

Tanaka and Farah (1991) examined the hypothesis that second-order relational properties were disrupted by inversion while inversion did no harm to first-order relational properties in a dot-pattern experiment. They compared the effects of inversion on the identification of dot patterns that differed in the extent to which they required the encoding of second-order relational properties. They found that both first-order relational and second-order relational properties were affected by inversion. Specifically, the identification of dot patterns that required more encoding of second-order relational properties was not more vulnerable to inversion than was identification of dot patterns that mainly required encoding of first-order relational properties.

Although the results of Valentine and Bruce (1988) and Tanaka and Farah (1991) are consistent with the view that inversion affects faces' first-order relational and second-order relational properties to the same degree, their results provide only loose support for this view. First, Valentine and Bruce's (1988) linearity finding suggests that the rotation of a face in the vertical induces a very regular change in the processing of its relational properties. This finding, however, cannot exclude the possibility that it may be mainly the processing of second-order relational properties that is affected by rotation. For instance, linearity may result from the fact that an increasing rotation in the vertical does not affect first-order relational properties but has a linear impact on the processing of second-order relational properties. Second, Tanaka and Farah's (1991) result stems from experiments in which dot patterns were used. In order to generalize their result to situations involving faces, the processing of first- and second-order relational properties should remain the same for highly artificial dot patterns and photographs of real faces. However, the empirical soundness of this presupposition, though crucial, is not guaranteed. The processing of relational properties induced by artificial stimulus sets may well be different from that induced by realistic facial stimuli (e.g., Valentine, 1988). The present study presents an alternative approach to testing the hypothesis of whether the information processed from upright and inverted faces is the same. This approach is different from the ones used by Valentine and Bruce, or Tanaka and Farah, and may be more conclusive on the issue.

The approach relies on the application of a choice paradigm where preference probabilities are separately measured on sets of faces in the upright and the inverted orientations. The point to be emphasized is that the way preferences between faces are affected by orientation may tell us something about the way in which the processing of facial properties changes from upright to upside-down orientation. Consider a set of faces with given preferences on them in terms of, say, their perceived attractiveness. Suppose that the perceived attrac-

tiveness of a face depends on both its first-order relational and its second-order relational properties (see below). Suppose now that we can account for the effect of orientation on the preferences by assuming that inversion affects all faces—that is, their perceived attractiveness—in the same way. Since faces will vary in terms of the contribution of first- and second-order relational properties to their overall attractiveness, this account suggested that inversion has the same effects on faces' first- and second-order relational properties. For instance, if inversion affected only faces' second-order relational properties, inversion would have a stronger impact on faces with a higher contribution of second-order relational properties on attractiveness than on faces with a lower contribution of second-order relational properties; inversion would affect the faces differently. Based on this argument, the degree of a face-orientation interaction may indicate the extent to which inversion induces a change in the processing of face properties. Finding a substantial face-orientation interaction therefore suggests that information processed from upright and inverted faces is different; finding no substantial interaction suggests that the information processed is the same.

To bring this approach to some direct experimental application, two issues must be specified. First, a facial attribute must be chosen that relies both on faces' first-order relational properties and on faces' second-order relational properties. Second, an adequate method must be chosen to reveal the degree of the face-orientation interaction. There is good empirical evidence indicating that the perceived attractiveness of a face is based on both its first-order relational and its second-order relational properties. Several researchers (Langlois & Roggman, 1990; Perrett, May, & Yoshikawa, 1994) hypothesized that the attractiveness of a face is a function of its closeness, or deviation, from a facial prototype and thus is a function of its second-order relational properties. In fact, these researchers reported experimental results in support of their hypotheses. These hypotheses, however, can serve as only rough approximations (Alley & Cunningham, 1991) that leave room for an additional role of first-order relational properties to affect a face's attractiveness. Indeed, results from the studies of Meerdink, Garbin, and Leger (1990) and Cunningham, Barbee, and Pike (1990) indicate that not only second-order relational properties but also first-order relational properties have an impact on the perceived attractiveness of a face.

To reveal the degree of a face-orientation interaction, I deduced a theoretical scale from the preference probabilities, separately for the two orientations. Indeed, if preference probabilities in the two orientations fulfill certain regularities (see Suppes, Krantz, Luce, & Tversky, 1989), they can be used to infer two theoretical scales of perceived attractiveness. In this case, each face can be represented by two scores, one score reflecting the perceived attractiveness of a face in the upright orientation and the other reflecting its perceived attractiveness in the inverted orientation. On the basis of these two scales, I

compared a face's two scores across orientations. The present study searches for a face-independent rule that relates the faces' two scores for the two orientations. It is argued that the extent to which such a rule can account for the change in the faces' scores reveals the extent to which inversion affected faces' first-order relational and second-order relational properties to the same degree in the present experiment.

METHOD

Subjects

The subjects were 103 psychology students at the University of Regensburg. They were tested individually and were given credit for fulfilling degree requirements.

Materials

Eight frontal-view photographs of stimulus faces of moderate attractiveness were used. The pictures were all of males between the ages of 21 and 26 years, and all were very similar in terms of hair length and shadows. They wore no eyeglasses, beards, or mustaches, and the expression on their faces was neutral. All faces were unknown to the subjects. The pictures were copied individually onto monochrome slides, and copies were made at two different orientations, upright and upside-down.

Apparatus

The slides were presented using a Kodak Carousel S-RA 2500 projector that was controlled by a computer. The slides were presented on a blank wall, in front of which the subject sat at a distance of about 2 m. Each face subtended about 8° of visual angle. The subject indicated his or her response by using push buttons, and the response was recorded by the computer.

Design and Procedure

All $8!(2!6!) = 28$ different pairs and all $8!(3!5!) = 56$ different triples of faces were presented to each subject, in both upright and inverted orientations. The pairs and triples were either presented first in the upright orientation and then, about 1 week later, in the inverted orientation, or vice versa. Fifty-two subjects started with the upright orientation, and the remaining 51 subjects started with the inverted orientation. For each orientation, the presentation of the pairs preceded the presentation of the triples.

An experimental session consisted of two parts, the presentation of the pairs and the presentation of the triples. At the beginning of each experimental session, all pairs and all triples were mixed randomly, as were a pair's two faces and a triple's three faces. The two (three) slides of a pair (triple) were presented successively with a presentation time of 2 sec for each slide. The presentation of the single slides was interrupted by a 1-sec blank field. After the presentation of a pair or triple of faces, the subject immediately selected which of the two or three faces he or she preferred in terms of attractiveness by pressing one of two (pair comparison) or one of three (triple comparison) buttons. Following a 2-sec blank field, the next pair or triple of faces was presented.

Data Analysis

For each of the two orientations, a subject's pair comparisons gave rise to a 28×2 matrix, a subject's triple comparisons gave rise to a 56×3 matrix. Each cell of the pair (triple) comparison matrix was coded as 1 or 0, depending on whether a face was preferred (1) or not (0) in a pair (triple) comparison. These two matrices represented a subject's preferences with respect to the presented faces. The pair (triple) comparison matrices were summed over subjects to result in one pair (triple) comparison matrix for each orientation.

These matrices represented the data of the experiment. To formulate a statistical model of these preferences, the sequence of choices of a pair's first or second face was viewed as a sequence of Bernoulli trials with underlying parameter p_{ab} , representing the probability that, for a pair (a,b) , face a is preferred over face b . The relative frequency with which face a is chosen from the pair (a,b) is taken as an estimate of p_{ab} . Similarly, the sequence of choices of a triple's first, second, and third face was viewed as a sequence of trials with underlying parameters $p_{a,bc}$ and $p_{b,ac}$, representing the probabilities that, for a triple of faces (a,b,c) , face a is preferred over faces b and c , and face b is preferred over faces a and c . Again, relative frequencies serve as estimates for the probabilities (see Suppes et al., 1989). Based on this statistical model, an orientation's preferences were described by 140 ($28 \times 1 + 56 \times 2$) free parameters.

In order to infer a simple theoretical scale of attractiveness which will give a parsimonious account for an orientation's whole set of preference probabilities, some restrictions on the preference probabilities must be fulfilled. The Bradley-Terry-Luce rule (Bradley & Terry, 1952; Luce, 1959; Suppes et al., 1989; in the following referred to as the BTL rule) was fitted to an orientation's two matrices. This rule sets strong restrictions on the relationship between pair and triple preferences. It demands the following property, called the constant-ratio property, to be true for all triples of faces (a,b,c) :

$$\frac{p_{ab}}{p_{ba}} = \frac{p_{a,bc}}{p_{b,ac}},$$

where p_{ab} stands for the probability that given the pair (a,b) face a is preferred over face b , and $p_{a,bc}$ stands for the probability that given the triple (a,b,c) face a is preferred over faces b and c . This property, in effect, asserts that the strength of preference of the triple's face a over the triple's face b is unaffected by the other available alternative, the triple's face c .

If the constant-ratio property is satisfied, each face a can be associated with a numerical value $v(a)$, so that the preference probabilities for a pair of faces (a,b) are determined by the rule $p_{ab} = v(a)/[v(a)+v(b)]$ and the preference probabilities for a triple of faces (a,b,c) are determined by the rule $p_{a,bc} = v(a)/[v(a)+v(b)+v(c)]$. These numerical values are unique up to scalar transformations. Thus, only the numerical value of one face can be chosen freely. For the eight face stimuli employed in the present experiment, the BTL rule therefore results in only seven free parameters to describe an orientation's 140 independent data observations.

A likelihood-ratio test (cf. Lindgren, 1976) was used to determine whether the rule fitted an orientation's data. Given the data, the likelihood of the BTL rule (L_{BTL}) was compared with the likelihood of the statistical model (L_S) by using the property that the term $-2 \ln L_{BTL}/L_S$ is approximately chi-square distributed. The parameters of the BTL rule were estimated by using the iterative search procedure PRAXIS (Gegenfurtner, 1992).

RESULTS

Analysis of Preference Probabilities

Table 1 shows the estimated preference probabilities (relative frequencies) for the pair comparisons, both for the upright and the inverted orientations. The Appendix shows the estimated preference probabilities for the triple comparisons. The preference probabilities varied over a large range of values, suggesting considerable differences between faces with respect to their perceived attractiveness. This pattern of results held for both ori-

Table 1
Pair Comparison Matrices for the Two Orientations, Upright and Inverted

	F_1	F_2	F_3	F_4	F_5	F_6	F_7	F_8
Upright Orientation								
F_1	—	0.515	0.709	0.728	0.583	0.689	0.252	0.961
F_2	0.485	—	0.757	0.660	0.612	0.670	0.340	0.922
F_3	0.291	0.243	—	0.437	0.408	0.408	0.146	0.854
F_4	0.272	0.340	0.563	—	0.437	0.447	0.194	0.913
F_5	0.417	0.388	0.592	0.563	—	0.524	0.243	0.961
F_6	0.311	0.330	0.592	0.553	0.476	—	0.214	0.913
F_7	0.748	0.660	0.854	0.806	0.757	0.786	—	0.913
F_8	0.039	0.078	0.146	0.087	0.039	0.087	0.087	—
Inverted Orientation								
F_1	—	0.476	0.689	0.602	0.621	0.641	0.262	0.816
F_2	0.524	—	0.670	0.573	0.592	0.680	0.291	0.845
F_3	0.311	0.330	—	0.544	0.398	0.456	0.243	0.748
F_4	0.398	0.427	0.456	—	0.456	0.476	0.155	0.757
F_5	0.379	0.408	0.602	0.544	—	0.447	0.184	0.835
F_6	0.359	0.320	0.544	0.524	0.553	—	0.282	0.728
F_7	0.738	0.709	0.757	0.845	0.816	0.718	—	0.913
F_8	0.184	0.155	0.252	0.243	0.165	0.272	0.087	—

Note—Each cell represents the estimated probability that one of a pair of faces is preferred over the other (F_n = face number n).

entations. To test whether an orientation's preferences differed reliably from indifferent choices, I compared, for each orientation, the (perfect) fit of the statistical model that describes the 140 independent preference probabilities (see Method section) with the fit of a statistical model that restricts all preference probabilities to 1/2 (pairs) or 1/3 (triples). The likelihood-ratio tests were conducted with 140 *df*. The $\chi^2(140)$ values of 3,575.605 ($p < .0001$) for the upright orientation and 2,474.733 ($p < .0001$) for the inverted orientation demonstrate that preferences differed reliably from indifferent choices.

A first visual comparison of the preference probabilities across orientations (see Table 1 and the Appendix) suggests that the probabilities were fairly stable across orientations. However, some tendency showed up for the preferences to be closer to indifferent choices in the inverted orientation than in the upright orientation. For instance, 19 of the 28 independent preference probabilities for pairs were less close to indifference ($p = 1/2$) when the faces were presented upright than when they were presented upside-down. To test whether the preference probabilities could be assumed to be constant across orientations, I compared the (perfect) fit of a joint statistical model that described the two orientations' 280 (2×140) independent preference probabilities (see Method section) with the fit of a statistical model that restricted the probabilities to being constant across orientations. The $\chi^2(140)$ value of 184.763 ($p = .007$) demonstrates that preferences varied reliably across orientations. Thus, both faces and orientation had a reliable effect on the preferences.

Quality of Fit of the BTL Rule

To get a more detailed insight into how faces and orientation affect preferences, a more parsimonious ac-


count for an orientation's whole data set is useful. For each orientation, I examined whether the preferences could be fitted by the BTL rule. If this rule fitted the data well, each face could be represented by one positive real-valued score that reflected its perceived attractiveness for the particular orientation. The BTL rule was simultaneously fitted to an orientation's pair and triple comparison matrices. For each orientation, the BTL rule has seven free parameters to describe the 140 independent preference probabilities. The $\chi^2(133)$ values of 86.429 ($p = .999$) for the upright orientation and 103.208 ($p = .974$) for the inverted orientation demonstrate an excellent fit of the rule to the data. Thus, for each orientation, the effect of the faces on the preference probabilities can be described by the BTL rule.

In the top panel of Figure 1, the preference probabilities measured in the two orientations with the predictions of these probabilities when using the BTL rule ("two individual BTL rules") are compared. The data are merged over the two orientations. If the BTL rule held perfectly, all 280 data points would fall on the diagonal line. As suggested by the above likelihood-ratio tests, the BTL rule fits the data well. As a result, each face is represented by its two BTL scores as described below.


Analysis of BTL Scores

I compared the scores of the faces across orientations. Figure 2 provides bar charts of the faces' scores, separately for each orientation. Because the BTL rule fitted the data well, an effect of orientation on the BTL scores that was qualitatively similar to the one found for the preference probabilities would be expected. Indeed, the scores of the faces show a considerable stability across orientations. The order of the faces' scores is hardly affected by inversion, and furthermore the score ratios between faces do not change in a major way across orien-

TWO INDIVIDUAL BTL-RULES



ONE COMMON BTL-RULE



BTL-RULE + AFFINE RULE




Figure 1. Scatterplot of measured preference probabilities (P-meas) versus predicted preference probabilities (P-pred) when using differently restricted rules to fit the data. The top panel shows the quality of fit of a joint BTL rule where a face's two parameters are free to vary across orientations (7+7 = 14 parameters). The middle panel shows the quality of fit of a joint BTL rule where a face's two parameters are restricted to not vary with orientation (7 parameters). The bottom panel, finally, shows the quality of the fit of a joint BTL rule where a face's two parameters are restricted to being affinely related across orientations (7+1 = 8 parameters). The data are merged over the two orientations (upright orientation •, inverted orientation ◦). If the rules accounted perfectly for the variation in the preference probabilities, each panel's 280 data points would fall on the diagonal line.

tations. Despite this stability, there is some tendency for the score ratios between faces to be reduced by inversion. I examined whether orientation had a reliable effect on the score ratios, that is, whether the assumption that $v_{up}(a)/v_{up}(b) = v_{inv}(a)/v_{inv}(b)$ holds for all pairs (a,b) had to be rejected statistically. Because the two scales v_{up} and v_{inv} are unique only up to scalar transformations (see Method section), this question is equivalent to testing whether the two scores of a face can be assumed to be constant across orientations. I fitted a joint BTL rule to the preference matrices of the two orientations where the two scores of a face are restricted to not varying with orientation. The fit of this rule to data was compared with the fit of a joint BTL rule where the two scores of a face were free to vary with orientation. The $\chi^2(7)$ value of 112.284 ($p < .001$) rejects the hypothesis that score ratios are constant across orientations. Thus, the effect of orientation found in the preference probabilities is also reflected in the BTL scores.

The middle panel of Figure 1 depicts the comparison of the preference probabilities measured in the two orientations with the predictions of these probabilities when using a joint BTL rule where the two scores of a face do not vary with orientation ("one common BTL rule"). Again the data are merged over orientations. As suggested by the above likelihood-ratio test, the fit of this joint BTL rule is somewhat worse than the fit of a joint BTL rule where the two scores of a face are free to vary with orientation ("two individual BTL rules," top panel). Although the difference in fit is not large, the deterioration is significant. As a result, a face cannot be represented by the same BTL score for the two orientations. A rule is needed to describe how a face's two scores are related across orientations.

Accordingly, I tested a simple rule to describe the effect of orientation on the faces' scores. This rule assumes that the two scores of a face are affinely related across orientations, that is, $v_{inv}(a) = v_{up}(a) + k$ for each face a , where k is a real-valued parameter that does not depend on the faces. Equivalently, across orientations the score ratios between faces are related by the equation $v_{inv}(a)/v_{inv}(b) = [v_{up}(a) + k]/[v_{up}(b) + k]$. This rule has two interesting properties. First, the affine relationship includes the assumption that inversion affects all faces equally. Thus, no interaction is supposed to occur between faces and orientation. The whole effect of orientation is reduced to parameter k . Second, the sign of k determines how preferences are affected by inversion. If k is positive, preferences are less close to indifferent choices in the upright orientation than they are in the inverted orientation, a result suggested above by visual analyses of the preference probabilities and the score ratios. On the other hand, if k is negative, preferences are more close to indifferent choices in the upright orientation than they are in the inverted orientation. $k=0$ is equivalent to an invariance of score ratios and preferences across orientations, a hypothesis already rejected above.

To examine whether this rule describes the effect of inversion for the present data, I tested whether the pref-

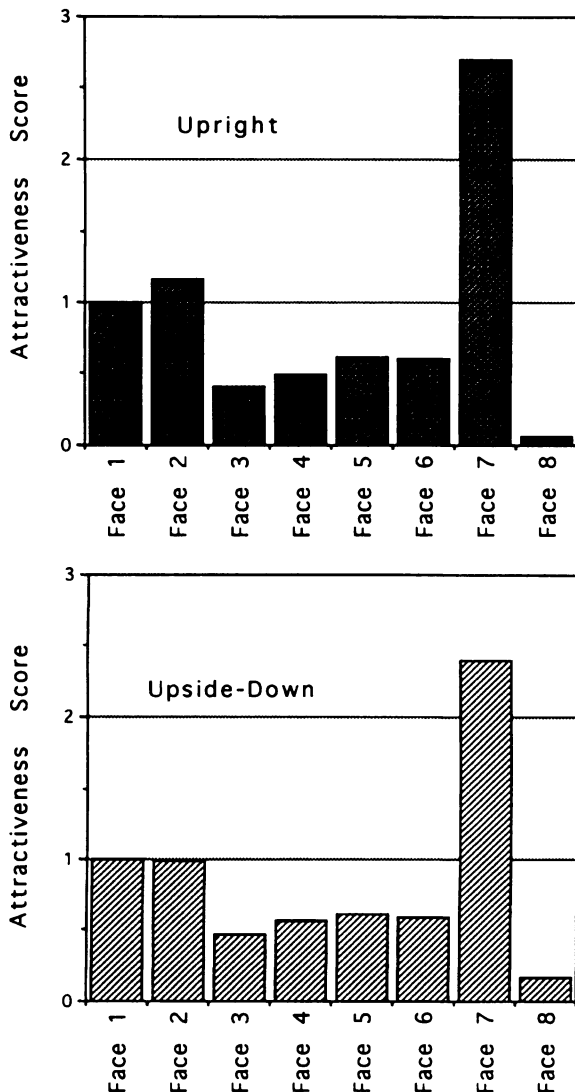


Figure 2. Bar charts showing the attractiveness scores of faces for the two orientations. The scores were estimated by fitting the BTL rule individually to an orientation's preference matrices. For each orientation, the face scores were fixed by assigning Face 1 a scale value of 1.

erence matrices of the two orientations could be fitted by a joint BTL rule where the two scores of a face were restricted to being affinely related through parameter k . Similarly, I compared the fit of this joint BTL rule with the fit of a joint BTL rule where the two scores of a face were free to vary with orientation. The $\chi^2(6)$ value of 7.241 ($p = .299$) supports the hypothesis that a face's two scores are affinely related across orientations. Thus, the introduction of the rotation parameter k reduces the chi-square value by more than 100 points, demonstrating that k plays an essential role in fitting the data. This holds true even though I found the value of k to be fairly small ($k = .135$). The fact that k is positive indicates that preferences are indeed less close to indifferent choices

in the upright orientation than they are in the inverted orientation.

The bottom panel of Figure 1 depicts the comparison of the preference probabilities measured in the two orientations with the predictions of these probabilities when using a joint BTL rule where the two scores of a face are affinely related across orientations ("BTL rule + affine rule"). Again, the data are merged over orientations. As suggested by the above likelihood-ratio test, the data are well fitted. In fact, a visual comparison with the fit of a joint BTL rule where the two scores of a face are free to vary with orientation ("two individual BTL rules," top panel) demonstrates that the more restrictive, 8-parameter model provides a fit to the data that is equal to that of the less restrictive, 14-parameter model. In this sense, the data are consistent with the hypothesis that the faces' two parameters are affinely related across orientations. In addition, the figure visualizes that the 8-parameter model provides a somewhat better fit to the data than does the 7-parameter model where the two scores of a face are restricted to not vary with orientation ("one common BTL rule," middle panel).

These analyses show that the reliable effects of faces and inversion on the preference probabilities can be described in a simple way by assuming, first, that each orientation's preference probabilities follow a BTL rule, and second, on the basis of this idea, that a face's two BTL scores are affinely related across orientations. By using these two rules, the 280 measured preference probabilities can be described by only eight parameters. Seven of these eight parameters reflect the effect of faces; the eighth parameter reflects the effect of orientation in the present data sets.

DISCUSSION

Preferences between faces were measured with regard to the perceived attractiveness of the faces. This was done with the faces presented upright and with the same faces presented upside down. Inversion affected the preferences. Qualitatively, the preferences between faces became more close to indifferent choices when the faces were presented upside down than when they were presented in the upright orientation. This indicates that the faces became less discriminable when inverted, suggesting an impairment in the processing of facial properties for this orientation. Quantitatively, the effect of inversion could be described by a simple rule that assumes the two attractiveness scores of a face to be affinely related across orientations. At the core of this rule is the assumption that inversion affects all faces equally with no substantial interaction between the faces' perceived attractiveness and their orientation. It was argued above that the degree of a face-orientation interaction reflects the extent to which information processed from upright and upside-down faces differs. The result that the effect of inversion in this study can be described equally for the single faces demonstrates a negligible amount of face-orientation in-

teraction, suggesting that about the same information is processed from upright and upside-down faces.

It is a widely held view in the literature that inversion is particularly disruptive to processing faces' second-order relational properties but that it hardly affects the processing of faces' first-order relational properties (Diamond & Carey, 1986). Although the results from a number of studies were interpreted in favor of this view, none of those studies provided unequivocal evidence for it (e.g., Rhodes et al., 1993; Valentine, 1988). Moreover, the data from two more recent studies were interpreted as evidence for an equal impairment in the processing of first- and second-order relational properties (Tanaka & Farah, 1991; Valentine & Bruce, 1988). While the relationship in these two studies between results and conclusions may still have been tentative in nature (see introduction), the results from the present study provide more direct evidence for this alternative view. Since the same information seems to have been processed from upright and upside-down faces, it is suggested that first- and second-order relational properties are affected by inversion to roughly the same degree. This finding establishes a major challenge to Diamond and Carey's (1986) proposition. It indicates that the distinction between faces' first- and second-order relational properties cannot explain why recognition of faces is disrupted by inversion to a far greater extent than is recognition of other classes of visual objects.

Most recently, Rhodes et al. (1993) reported an experiment in which they compared the effects of face inversion for detecting changes that span the continuum from first- to second-order relational properties. Using this kind of face manipulation, they found evidence that second-order relational properties are more sensitive to inversion than first-order relational properties. While their method has the desirable feature that it addresses the question of interest very directly, their method depends crucially on the assumption that faces' first- and second-order relational properties can be manipulated independently. As also outlined by Rhodes et al. (p. 50), some of their results suggest that this assumption does not hold in general: first- and second-order relational properties appear to be inherently confounded in faces. The degree to which their findings can challenge the view supported by the present study, that first- and second-order relational properties are equally affected by inversion, therefore remains unclear.

Attractiveness and Fit of BTL Rule

An orientation's preference matrices could be well described by the BTL rule, inducing a one-dimensional representation of faces with regard to their perceived attractiveness. At first, this finding might appear to conflict with some current theories of attractiveness which suggest several quite different factors as affecting a face's perceived attractiveness, including both first- and second-order relational properties (Cunningham, 1986; Cunningham et al., 1990; Meerdink et al., 1990). How-

ever, the two findings do *not* conflict with each other. Instead, the BTL representation implies only that when two faces are compared with regard to their attractiveness, all factors affecting the attractiveness of a face combine into one global score that represents the face's overall attractiveness. Indeed, this score can be interpreted as the sum of the attractiveness values of the single factors that affect the attractiveness of a face (Suppes et al., 1989; Tversky, 1972).

The present results reveal a surprisingly good fit of the BTL rule to orientation data. There are not many data sets in the literature where choice behavior could be well described by the BTL rule (Luce, 1977). Face attractiveness as investigated in the present experiment seems to be one of those. In fact, I replicated the experiment for upright faces for two other sets of faces, 10 male and 10 female. Again, I used perceived attractiveness as the facial attribute. For both the males and the females, the BTL rule led to fits to the data that compared well with those found in the present study.

Comparing Faces' Attractiveness Scores Across Orientations

This study focuses on interface relationships and the question of how attractiveness ratios between faces vary with orientation. On the basis of pair and triple comparisons of faces of equal orientation, this question could be addressed in a well-founded way. In order to include meaningful comparisons of the two attractiveness scores of a single face across orientations, however, measurements beyond those reported in the present study would be needed. Specifically, preferences between faces of *different* orientations would have to be measured. On the basis of these additional measurements, a new scale that simultaneously quantified the perceived attractiveness of the faces in their upright and their inverted orientations would have to be developed. Corresponding measurements were not conducted in the present study.

More formally, the fact that the two scores estimated for a face in this study cannot be compared in a meaningful way across orientations is a simple consequence of the fact that the two scales for upright and inverted orientation (v_{up} , v_{inv}) are unique only up to scalar transformations (see Method section), and that the units of the two scales can be fixed independently from each other. Thus, whenever the score of a face is higher in the upright orientation than in the inverted orientation—after two units for the two scales have been chosen—an appropriate change in the unit of one of the two scales can reverse the order. The development of one common scale for both orientations, if successful, would eliminate this freedom.

From Inversion to a General Rotation in the Vertical

Inversion of a face provides only a special case (180°) of a more general rotation of a face in the vertical. Indeed, rotation angles other than 180° are also known to affect face recognition (Rock, 1973), response times

(Valentine & Bruce, 1988), and/or encoding (Bäumel, 1992). As a result, the attractiveness relationships between faces can be expected to be affected not only by inversion but also by other rotation angles.

The choice paradigm employed in the present study can be used in a straightforward way to study how other rotation angles affect the relationships between faces. Specifically, if the affine rule found in the present study to describe the effect of inversion held for any rotation angle, the effect of each rotation angle on interface relationships could be represented by just one parameter, k . In this case, the way that parameter k depended on rotation angle would reveal useful information on how the rotation angle affects the relationships between faces. On the basis of Valentine and Bruce's (1988), or Rock's (1973), or Bäumel's (1992) results, a monotonic relationship between rotation angle and parameter k may be expected. More specific expectations about the functional form of the relationship between rotation angle and parameter k , however, are hard to derive from previous studies. In this sense, the question of how rotation angles other than inversion affect the attractiveness relationships between faces is open to future studies.

Three Final Remarks

First: By using attractiveness as a facial attribute, the present study provides evidence for the feasibility of using a simple rule to describe the effect of inversion on the relationships between faces. To the extent that this rule captures the whole effect of face inversion, its adequacy should not depend on the facial attribute employed. Indeed, when using other facial attributes, such as, for instance, distinctiveness or age, the pattern of results should be similar to that described above with regard to attractiveness. This expectation constitutes a strong prediction to be tested in further experiments.

Second: Recall that this study used photographs of eight moderately attractive males of about equal age. Although these experimental conditions might have favored the results obtained in this study, the faces employed provided a reasonable test of the hypothesis that the same information is processed from upright and inverted faces. Indeed, a considerable range of perceived attractiveness was spanned by the single faces. The attractiveness ratios for pairs of faces varied from less than 1.1:1 up to more than 40:1 (cf. Figure 2). Further experiments must show whether the rules found in the present study will also apply to quite different sets of facial stimuli, including those of males of very low or very high perceived attractiveness, older males, and females.

Third: The simple rule proposed in this study to account for the inversion effect led to a reasonable description of the data sets. Due to its strong restrictions, however, this rule is a simplification. Presumably, some faces do "lose" some of their properties when they are inverted, a phenomenon inconsistent with the idea that inversion affects all faces in the same way. Future studies must show whether frequency and size of these

"losses" are low enough to accept the view suggested in this study, at least as a first-order model.

REFERENCES

- ALLEY, T. R., & CUNNINGHAM, M. R. (1991). Averaged faces are attractive, but very attractive faces are not average. *Psychological Science*, *2*, 123-125.
- BÄUML, K.-H. (1992). Diskriminationslernen bei rotierten Gesichtern: Eine Markov-Analyse von Kodierungs- und Assoziationsprozessen. *Zeitschrift für Experimentelle und Angewandte Psychologie*, *39*, 1-17.
- BRADLEY, R. A., & TERRY, M. E. (1952). Rank analysis of incomplete block designs. I. The method of pair comparisons. *Biometrika*, *39*, 324-345.
- CUNNINGHAM, M. R. (1986). Measuring the physical in physical attractiveness: Quasi-experiments on the socio-biology of female facial beauty. *Journal of Personality & Social Psychology*, *50*, 925-935.
- CUNNINGHAM, M. R., BARBEE, A. P., & PIKE, C. L. (1990). What do women want? Facialmetric assessment of multiple motives in the perception of male facial physical attractiveness. *Journal of Personality & Social Psychology*, *59*, 61-72.
- DIAMOND, R., & CAREY, S. (1986). Why faces are and are not special: An effect of expertise. *Journal of Experimental Psychology: General*, *115*, 107-117.
- GENEFURTNER, K. (1992). PRAXIS: Brent's algorithm for function minimization. *Behavior Research Methods, Instruments, & Computers*, *24*, 560-564.
- GOLDSTEIN, A. G., & CHANCE, J. E. (1980). Memory for faces and schema theory. *Journal of Psychology*, *105*, 47-59.
- LANGLOIS, J. H., & ROGGMAN, L. A. (1990). Attractive faces are only average. *Psychological Science*, *1*, 115-121.
- LINDGREN, B. W. (1976). *Statistical theory* (3rd ed.). New York: Macmillan.
- LUCE, R. D. (1959). *Individual choice behavior: A theoretical analysis*. New York: Wiley.
- LUCE, R. D. (1977). The choice axiom after twenty years. *Journal of Mathematical Psychology*, *15*, 215-233.
- MARUYAMA, K., & ENDO, M. (1984). Illusory face dislocation effect and configural integration in the inverted face. *Tohoku Psychologica Folia*, *43*, 150-160.
- MEERDINK, J. E., GARBIN, C. P., & LEGER, D. W. (1990). Cross-gender perceptions of facial attributes and their relation to attractiveness: Do we see them differently than they see us? *Perception & Psychophysics*, *48*, 227-233.
- PERRETT, D. I., MAY, K. A., & YOSHIKAWA, S. (1994). Facial shape and judgements of female attractiveness. *Nature*, *368*, 239-242.
- RHODES, G., BRAKE, S., & ATKINSON, A. P. (1993). What's lost in inverted faces? *Cognition*, *47*, 25-57.
- ROCK, I. (1973). *Orientation and form*. New York: Academic Press.
- SERGENT, J. (1984). An investigation into component and configural processes underlying face perception. *British Journal of Psychology*, *75*, 221-242.
- SHEPARD, R. N., & METZLER, J. (1971). Mental rotation of three-dimensional objects. *Science*, *171*, 701-703.
- SUPPES, P., KRANTZ, D. H., LUCE, R. D., & TVERSKY, A. (1989). *Foundations of measurement* (Vol. 2, pp. 383-458). San Diego: Academic Press.
- TANAKA, J. W., & FARAH, M. J. (1991). Second-order relational properties and the inversion effect: Testing a theory of face perception. *Perception & Psychophysics*, *50*, 367-372.
- THOMPSON, P. (1980). Margaret Thatcher: A new illusion. *Perception*, *9*, 483-484.
- TVERSKY, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, *79*, 281-299.
- VALENTINE, T. (1988). Upside-down faces: A review of the effect of inversion upon face recognition. *British Journal of Psychology*, *79*, 471-491.
- VALENTINE, T., & BRUCE, V. (1988). Mental rotation of faces. *Memory & Cognition*, *16*, 556-566.
- YIN, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, *81*, 141-145.

YIN, R. K. (1970). Face recognition by brain-injured patients: A dissociable ability? *Neuropsychologia*, **8**, 395-402.
 YOUNG, A. W., HELLAWELL, D., & HAY, D. C. (1987). Configurational information in face perception. *Perception*, **16**, 747-759.

NOTE

1. Actually, Diamond and Carey (1986) not only distinguished between objects' first-order relational and second-order relational prop-

erties, but made an additional distinction between these two kinds of relational properties on the one hand and more isolated features on the other. The crucial point in Diamond and Carey's hypothesis, however, dealt with the special role of the second-order relational properties in object inversion; their hypothesis did not incorporate any distinction between first-order relational properties and the more isolated features. Thus, for sake of brevity, I use *first-order relational properties* in this paper as a summary term to refer both to first-order relational properties and to more isolated features.

APPENDIX

The triple comparison matrices for the two orientations, with each cell representing the estimated probability that one of a triple of faces is preferred over the other two.

Face Triple (a,b,c)	Upright			Inverted		
	a;bc†	b;ac	c;ab	a;bc	b;ac	c;ab
(F ₁ , F ₂ , F ₃)*	0.417	0.456	0.126	0.427	0.388	0.184
(F ₁ , F ₂ , F ₄)	0.369	0.466	0.165	0.350	0.398	0.252
(F ₁ , F ₂ , F ₅)	0.369	0.437	0.194	0.398	0.350	0.252
(F ₁ , F ₂ , F ₆)	0.388	0.447	0.165	0.495	0.340	0.165
(F ₁ , F ₂ , F ₇)	0.165	0.214	0.621	0.184	0.194	0.621
(F ₁ , F ₂ , F ₈)	0.485	0.485	0.029	0.447	0.437	0.117
(F ₁ , F ₃ , F ₄)	0.515	0.223	0.262	0.485	0.223	0.291
(F ₁ , F ₃ , F ₅)	0.495	0.243	0.262	0.515	0.194	0.291
(F ₁ , F ₃ , F ₆)	0.495	0.214	0.291	0.544	0.223	0.233
(F ₁ , F ₃ , F ₇)	0.204	0.087	0.709	0.223	0.087	0.689
(F ₁ , F ₃ , F ₈)	0.680	0.291	0.029	0.592	0.291	0.117
(F ₁ , F ₄ , F ₅)	0.485	0.204	0.311	0.495	0.252	0.252
(F ₁ , F ₄ , F ₆)	0.447	0.223	0.330	0.466	0.252	0.282
(F ₁ , F ₄ , F ₇)	0.233	0.107	0.660	0.214	0.155	0.631
(F ₁ , F ₄ , F ₈)	0.689	0.282	0.029	0.592	0.311	0.097
(F ₁ , F ₅ , F ₆)	0.398	0.320	0.282	0.456	0.272	0.272
(F ₁ , F ₅ , F ₇)	0.214	0.117	0.670	0.233	0.126	0.641
(F ₁ , F ₅ , F ₈)	0.563	0.408	0.029	0.592	0.330	0.078
(F ₁ , F ₆ , F ₇)	0.204	0.165	0.631	0.243	0.136	0.621
(F ₁ , F ₆ , F ₈)	0.592	0.359	0.049	0.602	0.301	0.097
(F ₁ , F ₇ , F ₈)	0.262	0.709	0.029	0.272	0.699	0.029
(F ₂ , F ₃ , F ₄)	0.583	0.155	0.262	0.505	0.165	0.330
(F ₂ , F ₃ , F ₅)	0.563	0.146	0.291	0.466	0.214	0.320
(F ₂ , F ₃ , F ₆)	0.524	0.204	0.272	0.563	0.214	0.223
(F ₂ , F ₃ , F ₇)	0.282	0.039	0.680	0.282	0.097	0.621
(F ₂ , F ₃ , F ₈)	0.670	0.262	0.068	0.544	0.301	0.155
(F ₂ , F ₄ , F ₅)	0.553	0.223	0.223	0.485	0.262	0.252
(F ₂ , F ₄ , F ₆)	0.524	0.233	0.243	0.427	0.311	0.262
(F ₂ , F ₄ , F ₇)	0.291	0.097	0.612	0.291	0.107	0.602
(F ₂ , F ₄ , F ₈)	0.689	0.301	0.010	0.583	0.340	0.078
(F ₂ , F ₅ , F ₆)	0.544	0.282	0.175	0.456	0.320	0.223
(F ₂ , F ₅ , F ₇)	0.243	0.126	0.631	0.320	0.097	0.583
(F ₂ , F ₅ , F ₈)	0.641	0.320	0.039	0.544	0.359	0.097
(F ₂ , F ₆ , F ₇)	0.223	0.107	0.670	0.243	0.155	0.602
(F ₂ , F ₆ , F ₈)	0.650	0.311	0.039	0.553	0.320	0.126
(F ₂ , F ₇ , F ₈)	0.262	0.718	0.019	0.311	0.680	0.010
(F ₃ , F ₄ , F ₅)	0.272	0.369	0.359	0.320	0.282	0.398
(F ₃ , F ₄ , F ₆)	0.282	0.330	0.388	0.282	0.359	0.359
(F ₃ , F ₄ , F ₇)	0.117	0.155	0.728	0.136	0.175	0.689
(F ₃ , F ₄ , F ₈)	0.456	0.495	0.049	0.369	0.505	0.126
(F ₃ , F ₅ , F ₆)	0.214	0.369	0.417	0.214	0.369	0.417
(F ₃ , F ₅ , F ₇)	0.107	0.155	0.738	0.184	0.146	0.670
(F ₃ , F ₅ , F ₈)	0.417	0.534	0.049	0.369	0.524	0.107
(F ₃ , F ₆ , F ₇)	0.087	0.184	0.728	0.126	0.233	0.641
(F ₃ , F ₆ , F ₈)	0.369	0.563	0.068	0.408	0.456	0.136
(F ₃ , F ₇ , F ₈)	0.175	0.806	0.019	0.146	0.825	0.029
(F ₄ , F ₅ , F ₆)	0.272	0.350	0.379	0.272	0.330	0.398
(F ₄ , F ₅ , F ₇)	0.146	0.136	0.718	0.136	0.155	0.709
(F ₄ , F ₅ , F ₈)	0.476	0.456	0.068	0.417	0.466	0.117
(F ₄ , F ₆ , F ₇)	0.107	0.155	0.738	0.136	0.214	0.650
(F ₄ , F ₆ , F ₈)	0.437	0.505	0.058	0.456	0.437	0.107

Appendix (Continued)

Face Triple (<i>a, b, c</i>)	Upright			Inverted		
	<i>a;bc</i> †	<i>b;ac</i>	<i>c;ab</i>	<i>a;bc</i>	<i>b;ac</i>	<i>c;ab</i>
(<i>F</i> ₄ , <i>F</i> ₇ , <i>F</i> ₈)	0.146	0.825	0.029	0.243	0.728	0.029
(<i>F</i> ₅ , <i>F</i> ₆ , <i>F</i> ₇)	0.155	0.184	0.660	0.146	0.175	0.680
(<i>F</i> ₅ , <i>F</i> ₆ , <i>F</i> ₈)	0.466	0.505	0.029	0.417	0.456	0.126
(<i>F</i> ₅ , <i>F</i> ₇ , <i>F</i> ₈)	0.184	0.777	0.039	0.223	0.748	0.029
(<i>F</i> ₆ , <i>F</i> ₇ , <i>F</i> ₈)	0.204	0.786	0.010	0.272	0.699	0.029

**F*_{*n*} = face number *n*. †*a;bc* = face *a* is preferred over faces *b* and *c*, and so forth.

(Manuscript received July 1, 1993;
revision accepted for publication January 18, 1994.)