Spatial Vision, Vol. 1, No. 2, pp. 179–196 (1985)

Selective internal operations in the recognition of locally and globally point-inverted patterns

WALTER F. BISCHOF, DAVID H. FOSTER* and JEREMY I. KAHN

Department of Communication and Neuroscience, University of Keele, Keele, Staffordshire ST5 5BG, UK

Received 11 January 1985; revised 3 July 1985; accepted 16 July 1985

Abstract—Performance in discriminating rotated 'same' patterns from 'different' patterns may decrease with rotation angle up to about 90° and then increase with angle up to 180°. This anomalously improved performance under 180° pattern rotation or *point-inversion* can be explained by assuming that patterns are internally represented in terms of local features and their spatial-order relations ('left of', 'above', etc.), and that, in pattern comparison, an efficient internal sense-reversal operation occurs (transforming 'left of' to 'right of', etc.). Previous experiments suggested that local features and spatial relations could not be efficiently separated in some pattern-comparison tasks. This hypothesis was tested by measuring 'same-different' discrimination performance under four transformations: point-inversion ι of the whole pattern, point-inversion ι_F of local features alone, point-inversion ι_P of local-feature positions alone, and identity transformation Id. The results suggested that internal sense-reversal operations could be applied selectively and efficiently, provided that local features were well separated. Under this condition performances for ι_F and ι were about the same whereas performance for ι_P was significantly worse, the latter performance resulting possibly from an attempt to apply internal global and local sense-reversal operations serially.

INTRODUCTION

In general, rotation of a pattern or figure in the plane leads to reduced probability of its recognition (Mach, 1897; Dearborn, 1899; Aulhorn, 1948; Arnoult, 1954; Kolers and Perkins, 1969, 1975; Rock, 1973, Chapter 3; Foster, 1978; Kahn and Foster, 1981). The variation in visual performance may not, however, be monotonic with angle of pattern rotation: performance in discriminating rotated 'same' patterns from 'different' patterns has been found to decline with rotation angle for angles up to about 90°, and then increase again with rotation angle for angles up to 180° (Dearborn, 1899; Aulhorn, 1948; Rock, 1973; Foster, 1978; Kahn and Foster, 1981). The apparently anomalous improvement in performance under 180° rotation, or *point-inversion* ι, has been demonstrated with randomly contoured shapes (Dearborn, 1899; Rock, 1973), with random-dot patterns (Foster, 1978; Kahn and Foster, 1981), and with alphabetic characters (Aulhorn, 1948).

This kind of discrimination performance, often elicited with displays of limited duration, should be distinguished from the typically monotonic dependence of reaction time on rotation angle observed by R. N. Shepard and his colleagues in 'mental rotation' experiments where highly accurate discriminations were made of rotated patterns from rotated *reflected* patterns (see e.g. Shepard and Metzler, 1971; Cooper and Shepard, 1973; Shepard, 1975; Shepard and Cooper, 1982). Reaction times were usually of the order of seconds.

* To whom all correspondence should be addressed.

Non-monotonic dependence of 'same-different' discrimination performance on rotation angle has been explained in terms of a relational-structure scheme for visual pattern recognition (Foster, 1978; Foster and Mason, 1979; Kahn and Foster, 1981; Foster and Kahn, 1985). Because of the central importance of this scheme to the experiments reported here, its main elements are reviewed below.

INTERNAL REPRESENTATIONS AND OPERATIONS

It was assumed (after Sutherland, 1968, and Barlow *et al.*, 1972) that a pattern A was represented internally by the visual system in terms of, at least, the following

(a) local features f_i, $i = 1, 2, \ldots, m$, which for random-dot patterns could be dot clusters of a particular density and shape; and

(b) spatial-order relations r_x, r_y that specified how one local feature f_j was related to another f_k, $1 \leqslant j < k \leqslant m$, in a horizontal-vertical reference system, thus

$$r_x(f_j, f_k) = \begin{cases} 1, & \text{for } f_j \text{ 'left of' } f_k, \\ -1, & \text{for } f_j \text{ 'right of' } f_k, \\ 0, & \text{otherwise;} \end{cases}$$

$$r_y(f_j, f_k) = \begin{cases} 1, & \text{for } f_j \text{ 'above' } f_k, \\ -1, & \text{for } f_j \text{ 'below' } f_k, \\ 0, & \text{otherwise.} \end{cases}$$

(The value 0 would arise for r_x for example when f_j was vertically in line with f_k.) It was further assumed that two patterns A_1, A_2 were judged to be the same if their internal representations could be brought into coincidence. In the special case of patterns related by a 180° rotation or point-inversion, $A_2 = \iota(A_1)$ (shown in Figs 1a and d), recognition was supposed to occur by a simple internal global reversal σ of the sign or sense of the spatial-order relations r_x, r_y, so that r_x was replaced by $-r_x$ and r_y by $-r_y$. Because sense-reversal in the x- and y-directions is equivalent to a 180° rotation, the effects of pattern point-inversion were compensated for precisely. More formal details are given in Appendix 1. For arbitrary planar pattern rotations, Foster and Mason (1979) were able to predict, by counting the number of altered spatial relations and allowing for

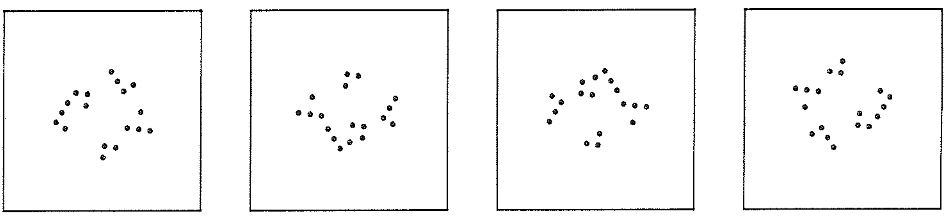

a b c d

Figure 1. Illustrations of examples of patterns used to determine the effects of various point-inversion transformations on 'same-different' judgements. Each pattern was composed of five subpatterns chosen pseudo-randomly, with replacement, from the set shown in Fig. 2, with the orientations of the subpatterns chosen pseudo-randomly. The positions of the subpatterns were chosen pseudo-randomly subject to the constraint that the centres of the subpatterns fell within an imaginary circle of diameter 0.5° visual angle and that the minimum centre-to-centre separation of the subpatterns was 0.17°. 'Same' pattern pairs were related by one of the following transformations: identity transformation Id (pattern a and its duplicate), position point-inversion ι_P (patterns a and b), local-feature point-inversion ι_F (patterns a and c), and global point-inversion ι (patterns a and d). 'Different' pattern pairs were obtained (in Experiments 1 and 2) by pairing completely unrelated patterns. The patterns were presented sequentially, centred on the point of fixation. Stimuli appeared brighter than the background.

possible internal global sense-reversals, the detailed variation with rotation angle of recognition performance obtained with a large, fixed repertoire of random-dot patterns (Foster, 1978).

In the data referred to above concerning recognition of rotated patterns, performance was determined for pairs of patterns presented symmetrically about the fixation point. Previous experiments showing the upturn in performance at 180° rotation also involved a symmetric presentation of the stimuli; patterns were viewed typically, one at a time, centrally in the visual field (Dearborn, 1899; Rock, 1973). But it was found by Kahn and Foster (1981) and Foster and Kahn (1985) that the upturn in performance for point-inverted patterns was dependent on the symmetry of the display; when symmetry of the pattern positions was disturbed, performance for point-inverted patterns was reduced. To explain this result, the above scheme for internal pattern representations and operations was modified (Kahn and Foster, 1981; Foster and Kahn, 1985) in the following way. Patterns were assumed to be represented in terms of local features f_i, the spatial relations r_x, r_y between those local features, *and* the approximate global positions d_x, d_y of the patterns in the field defined within a horizontal-vertical reference system centred on the point of fixation. The internal sense-reversal operation σ was assumed to be applicable not only to the r_x, r_y, but also to the d_x, d_y. When applied, σ was assumed to operate uniformly on *all* the relations specifying a particular order (e.g. to both r_x and d_x, or to both r_y and d_y). A second internal operation which operated on individual components of the internal representation in a progressive, but increasingly inefficient, fashion was also postulated. This 'continuous-modification' operation is not developed further here.

SELECTIVE INTERNAL OPERATIONS

Experimental data obtained by Kahn and Foster (1981) and Foster and Kahn (1985) on the structure of internal representations suggested that it was not possible to apply efficiently the internal sense-reversal operation σ to the spatial relations r_x, r_y alone. The latter seemed to be in a fixed association with the global-position components d_x, d_y. A fixed association of a different kind, involving local features and spatial relations, had also been suggested previously (Foster, 1978). Experiments had been carried out on 'same-different' discrimination of random-dot patterns according to different judgemental criteria for equality: the one based on pattern shape, the other on pattern dot-number or numerosity. From these experiments, it was argued (Foster, 1978) that in an internal representation a fixed association was established between local features and the spatial relations between these local features. Consequently, when spatial-relation information was in principle irrelevant to the pattern-comparison task, as in judgements of numerosity, this information could not be separated efficiently from relevant local-feature information, thereby leading to reduced performance.

Taken together, these studies appeared to imply a cohesiveness of internal visual representations. Spatial relations between local features and global-position relations seemed not to be separated efficiently in applying internal point-inversions, and local features and spatial relations seemed not to be separated efficiently in different internal-comparison tasks. In this study, a different approach to the testing of the association between local features and spatial relations was undertaken. This approach involved specifically the application of two kinds of non-uniform point-inversion transformations to a pattern: ι_F, which inverted all the hypothesized local features about each of their local centres, leaving their positions unaltered; and ι_P, which inverted the positions of all the local features about the centre of the pattern, leaving the orientations of the local features intact.

The rationale, formal details of which are given in Appendix 1, was as follows. First, consider local-feature point-inversion ι_F. Let A_1, A_2 be patterns related by transformation ι_F, that is, $A_2 = \iota_F(A_1)$ (compare Figs 1a and c). If there existed an efficient way of accessing and matching the local features f_i in the internal representation of A_1 to the point-inverted local features $\iota_F(f_i)$ in the internal representation of A_2, independently of the orientations of the f_i, or if there existed an efficient way of accessing local features in the internal representation and by the operation σ reversing the sense of the 'orientation relations' associated with the f_i, without reversing the sense of the spatial relations between them, then 'same'-detection performance for patterns related by transformation ι_F should be high.

Second, consider position point-inversion ι_P. Let A_1, A_2 be patterns related by transformation ι_P, that is, $A_2 = \iota_P(A_1)$ (compare Figs 1a and b). If there existed an efficient way of accessing and matching the local features f_i in the internal representations of A_1 and A_2, independently of the spatial relations between the f_i, or if there existed an efficient way of reversing the sense of the spatial relations without modifying the orientation relations of the local features themselves, then 'same'-detection performance for patterns related by transformation ι_P should be high.

If, however, none of these selective, independent internal operations was feasible, then the only way in which these patterns could be detected as 'same' would be by the use of other inefficient procedures, such as the individual-component 'continuous-modification' operation mentioned earlier. Note that the composition of local-feature point-inversion ι_F and position point-inversion ι_P results in global point-inversion ι; that is, $\iota = \iota_F \circ \iota_P = \iota_P \circ \iota_F$, which may be compensated for precisely and efficiently by global application of the internal sense-reversal operation σ to the internal representation.

In this study, as in others concerned with the effects of pattern rotation (Dearborn, 1899; Rock, 1973), randomly formed figures were used as stimuli. Random-dot patterns in particular were used to allow comparison with previous studies (Foster, 1978; Kahn and Foster, 1981; Foster and Kahn, 1985). Randomly formed stimuli have the advantage of avoiding confounding effects due to semantic content and conventional orientation and handedness. When stimuli with conventional orientation and meaning are used, striking effects may be obtained by application of global point-inversion ι and local-feature point-inversion ι_F, as Thompson (1980) has demonstrated with a photograph of a currently well-known face.

EXPERIMENT 1: LOCAL AND GLOBAL PATTERN POINT-INVERSIONS

Methods

Apparatus. The stimuli for the experiment were produced on the screen of an $X-Y$ display oscilloscope (Hewlett-Packard, Type 1300A here and 1321A in subsequent experiments) with P4 sulphide phosphor (decay time 60 μs), controlled by a minicomputer (CAI Alpha LSI-2) and graphics generator (Sigma Electronic Systems QVEC 2150). The screen was viewed binocularly at a distance of 1.7 m through a view-tunnel and optical system which produced a uniform white background field subtending at the eye 7.4° × 6.2° (38° × 44° in subsequent experiments) and of luminance approximately 100 cd m^{-2}. The stimuli were white and appeared superimposed on the background field. The intensity of the stimuli was adjusted by each subject at the beginning of each experimental session to 10-times luminance increment threshold. This setting was achieved by introducing a 1.0-log-unit neutral-density filter over the stimulus dots: the intensity of the dots was then adjusted to increment threshold on the unattenuated background.

Fixation was aided by four computer-generated white lines, each approximately 0.25° long, pointing to the centre of the display, and by five dots forming a cross. The lines were displayed throughout each presentation; the cross was extinguished at the start of each trial. The subject controlled the start of each trial and gave his responses on a hand-held push-button box linked to the computer.

Stimuli. Stimuli were composed of dot subpatterns, as shown in Fig. 2. Each dot subtended about 0.02°. These subpatterns, chosen to be representative of the local features that might be found in random-dot patterns, consisted of three-dot figures forming chevrons with (a) obtuse, (b) right, and (c) acute internal angles, and four-dot figures forming (d) T-shaped and (e) L-shaped figures. These subpatterns were chosen specifically for their asymmetry under point-inversion. The extent to which each subpattern gave rise to a well-defined local feature is discussed later. Each pattern was generated by pseudo-random selection, with replacement, of five of the subpatterns. The positions of the subpatterns were chosen pseudo-randomly under the constraint that the centres of the subpatterns should all lie within an imaginary circle of diameter 0.5° and that the centres of no two subpatterns should be closer than 0.17°. The orientation of each subpattern was chosen pseudo-randomly. Illustrations of typical patterns generated in this way are given in Fig. 1.

Pattern transformations. There were four possible transformations relating the patterns in each 'same' pair:

(a) Id (identity transformation): the two patterns were identical (Fig. 1a and its duplicate);

(b) ι_P (position point-inversion): one pattern was obtained from the other by inversion of the positions of the subpatterns about the centre of the pattern, the orientations of the subpatterns remaining the same in the two patterns (Figs 1a and b);

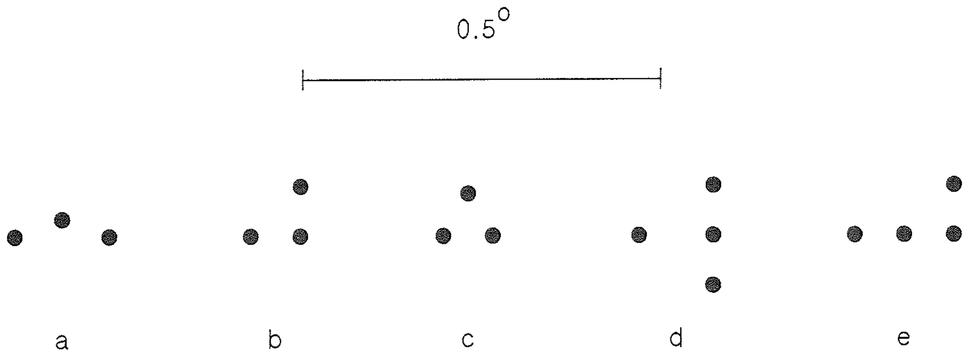

Figure 2. Illustrations of the subpatterns used to form the 'local features' of the patterns in Fig. 1. The scale of the subpatterns is indicated by the marker, which represents 0.5° visual angle.

(c) ι_F (local-feature point-inversion): one pattern was obtained from the other by inversion of each of the subpatterns about their local centres, the positions of the subpatterns remaining the same in the two patterns (Figs 1a and c); and
(d) ι (global point-inversion): one pattern was obtained from the other by inversion about the centre of the pattern (Figs 1a and d).

For 'different' pairs, the two patterns were generated independently of each other. A fresh pair of patterns was generated for every trial.

Instructions. At the beginning of the experiment, subjects were informed of the nature of the stimuli and of the types of transformation involved. Subjects were instructed to indicate after the presentation of each pair of patterns whether they were 'same' or 'different' according to the above transformations. It was emphasized that steady fixation was to be maintained throughout each presentation period and that responses should be made as quickly as possible whilst preserving accuracy. No feedback on performance was given to subjects.

Presentation sequence. Patterns in each pair were presented sequentially, centred on the fixation cross. Following initiation of the trial by the subject, the fixation cross was extinguished, and, after a 1.0-s delay, the first stimulus pattern appeared for 100 ms; after a 1.0-s delay, the second stimulus pattern appeared for 100 ms. The subject's response was recorded by the computer. As a control, the time taken to make the response was also recorded. After a further 1.0-s delay, the fixation cross was redisplayed indicating that the next trial could be started.

Experimental design. There were 40 trials in each experimental run. In each run, each 'same' pattern transformation (Id, ι_P, ι_F, ι) occurred five times; there were also 20 'different' pattern transformations, so that a run consisted of 20 'sames' and 20 'differents'. Each subject performed 10 runs in one session. The order of the pattern transformations in a run was chosen pseudo-randomly but balanced over runs to offset stimulus order and carry-over effects.

Subjects. Five male students, aged 22–27 years, participated in the experiment. Each had normal or corrected-to-normal vision. All subjects except one (co-author J.I.K.) were unaware of the purpose of the experiment.

Results

Discrimination performance. Figure 3 shows 'same-different' pattern-discrimination performance as a function of pattern transformation: (a) identity transformation Id, (b) position point-inversion ι_P, (c) local-feature point-inversion ι_F, and (d) global point-inversion ι. As discrimination performance was determined by responses to both 'same' and 'different' patterns, scores were expressed in terms of the discrimination index d' from signal detection theory (Green and Swets, 1966). The index d' is zero when performance is at chance level and increases monotonically with improvement in performance. It has a number of advantages as a performance measure (Swets, 1973), including the properties of being bias-free and additive (Durlach and Braida, 1969). The d' data in Fig. 3 were weighted by variances and averaged over subjects (Appendix 2). Chi-squared tests (Appendix 2) on individual subjects' data showed significant dif-

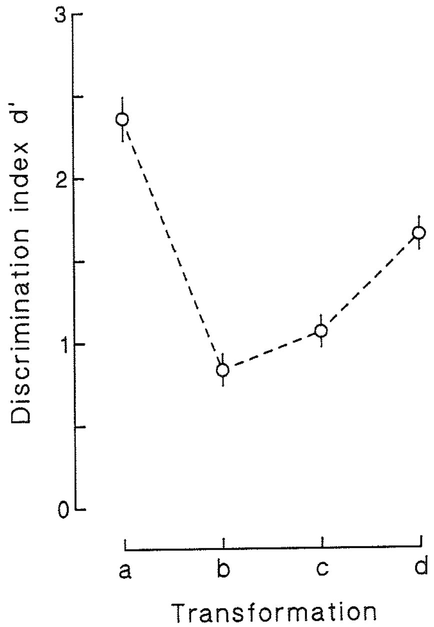

Figure 3. 'Same-different' discrimination performance as a function of pattern transformation. Mean discrimination index d' is shown for (a) identity transformation Id, (b) position point-inversion ι_P, (c) local-feature point-inversion ι_f, and (d) global point-inversion ι. The vertical bars indicate ± 1 SEM.

ferences between subjects' performances ($\chi^2_{16} = 64.0$, $P < 0.001$), but these differences disappeared after normalizing performances with respect to each subject's mean level (Appendix 2) ($\chi^2_{11} = 14.0$, $P > 0.2$).

Comparisons (Appendix 2) on the averaged d' data yielded the following results (two-tailed significance tests in all cases):

(a) 'same'-detection performance for patterns related by transformation Id was significantly higher than that for patterns related by transformation ι ($z = 4.24$, $P < 0.0001$);

(b) 'same'-detection performance for patterns related by transformation ι was significantly higher than that for patterns related by either of the transformations ι_F, ι_P($z \geqslant 4.23$, $P < 0.0001$ in both cases);

(c) there was no significant difference between 'same'-detection performance for patterns related by transformation ι_F and those related by transformation ι_P ($z = 1.70$, $P > 0.05$); and

(d) 'same'-detection performance for patterns related either by ι_F or by ι_P was significantly greater than chance level ($d' = 0$) ($z \geqslant 8.5$, $P < 0.0001$).

Reaction times. Averaged over subjects and conditions, correct responses (mean $\pm$ 1 SEM) were numerically faster than incorrect responses, 663 ± 49 ms vs 778 ± 72 ms, but not significantly ($t_{48} = 1.32$, $P > 0.1$). Correct 'same' responses were not significantly faster than correct 'different' responses, 670 ± 56 ms vs 633 ± 110 ms ($t_{23} = 0.30$, $P > 0.5$). All tests were two-tailed.

There was no trade-off between performance (per cent correct) and reaction time (RT). Over subjects and conditions, RTs for correct 'same' responses tended to be negatively correlated with performance, with gradient (mean ± 1 SEM) -2.15 ± 1.38 ms (per cent)$^{-1}$, although this did not reach significance ($z = 1.56$, $P > 0.1$).

Discussion

The relatively low 'same'-detection performance obtained here under local-feature point-inversion ι_F and position point-inversion ι_P could have arisen as a result of two distinct mechanisms: the one concerned with the internal accessibility of the hypothesized local features and spatial relations, the other concerned with the existence of the putative internal operations applied to local features and spatial relations. First, if local features were not adequately separated and accessed in the internal representations, then for patterns related by ι_F the selective internal operations of matching local features independently of their orientations or of reversing the sense on the 'orientation relations' associated with local features could not have been applied, even though the operations may have been available. Analogously, if spatial relations were not adequately separated and accessed in the internal representations, then for patterns related by ι_P the internal operations of matching local features independently of the spatial relations between them or of reversing the sense of the spatial relations without modifying the orientation relations of the local features could not have been applied, even though these operations may also have been available. Second, although local features and spatial relations may have been adequately separated in the internal representations, it is possible that the internal operations were not available, at least in the selective form required to give high 'same'-detection performance under the transformations ι_P and ι_F.

From this experiment, it was not possible to determine which of these two mechanisms may have been the cause of the reduced performance. Two further experiments were accordingly undertaken. In Experiment 2, the influence of the separation of the hypothesized local features on 'same'-detection performance was determined under the pattern transformations considered in Experiment 1. In Experiment 3, a specific test was carried out to determine how accurately 'same'-detection occurred under local-feature point-inversions.

EXPERIMENT 2: POINT-INVERSIONS OF LARGE PATTERNS

As already noted, one prerequisite for the efficient recognition of patterns related by local-feature point-inversions ι_F or position point-inversions ι_P is the ability to separate and access efficiently the local features of the internal representations. In the stimulus patterns used in Experiment 1, the ratio of centre-to-centre separation to diameter for the subpatterns assumed to give rise to local features was 2.0 in over 87% of presentations. This may have been too small for adequate internal separation of local features. In this experiment, the overall extent of stimulus patterns was doubled, and with it the separation of subpatterns. The size of the subpatterns remained the same. Since a new group of subjects was used, Experiment 1 was replicated as a control.

Methods

Stimuli and procedure were similar to those in Experiment 1. The positions of the subpatterns were again chosen pseudo-randomly under the constraint that their centres should lie within an imaginary circle of 1° diameter (below referred to as 'large' patterns) or 0.5° (below referred to as 'small' patterns). The minimum separation of subpattern centres was fixed at 0.33° for large patterns and at 0.17° for small patterns. Four subjects, three male and one female, aged 23–36 years, participated in the experiment. Each had normal or corrected-to-normal vision. All subjects except one (co-author W.F.B.) were unaware of the purpose of the experiment.

Results

Discrimination performance. Figure 4 shows 'same-different' pattern-discrimination performance as a function of pattern transformation: (a) identity transformation Id, (b) position point-inversion ι_P, (c) local-feature point-inversion ι_F, and (d) global point-inversion ι, for small patterns (open circles) and for large patterns (solid circles). The d' data were weighted by individual variances and averaged over subjects (Appendix 2). Chi-squared tests on individual subjects' data showed no significant differences between subjects' performances ($\chi^2_{12} = 13.7$, $P > 0.2$, for small patterns, and $\chi^2_{12} = 18.9$, $P > 0.05$, for large patterns) without normalizing with respect to each subject's mean level of performance (Appendix 2). Results for small patterns, used as a control, were not significantly different from results obtained in the first experiment after

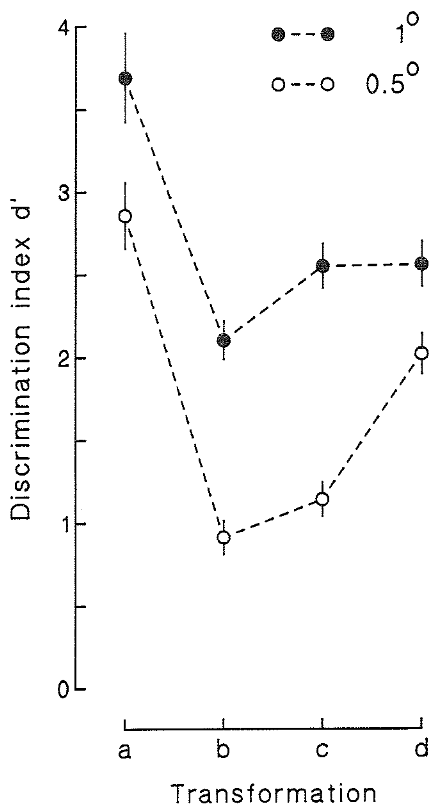

Figure 4. 'Same-different' discrimination performance as a function of pattern transformation. Mean discrimination index d' is shown for (a) identity transformation Id, (b) position point-inversion ι_P, (c) local-feature point-inversion ι_F, and (d) global point-inversion ι. Solid circles show performance for patterns with diameter of 1° visual angle and open circles for patterns with diameter of 0.5° visual angle. 'Different' patterns were completely unrelated. The vertical bars indicate $\pm$ 1 SEM.

allowance was made for differences in overall performance level ($\chi_3^2 = 4.68$, $P > 0.1$). There were, however, significant differences between results for small patterns and those for large patterns, even after allowance for overall performance level ($\chi_3^2 = 13.6$, $P < 0.01$). Comparisons (Appendix 2) on averaged d' data obtained with large patterns yielded the following results (two-tailed significance tests in all cases):

(a) 'same'-detection performance for patterns related by transformation Id was significantly higher than that for patterns related by transformation ι, ι_P, or ι_F($z \geqslant 3.71$, $P < 0.01$);

(b) there was no significant difference between 'same'-detection performance for patterns related by transformation ι and for patterns related by transformation ι_F ($z = 0.04$, $P > 0.5$);

(c) 'same'-detection performance for patterns related by transformation ι_P was significantly lower than that for patterns related by transformation ι or ι_F ($z \geqslant 2.46$, $P < 0.05$); and

(d) 'same'-detection performance for patterns related by any of the transformations was significantly higher than chance level ($d' = 0$) ($z \geqslant 13.6$, $P < 0.0001$).

Reaction times. For small and large patterns reaction times varied similarly. Averaged over subjects and conditions, correct responses (mean ± 1 SEM) were significantly faster than incorrect responses (1061 ± 19 ms vs 1328 ± 48 ms, $t_{48} = 9.57$, $P < 0.0001$, for small patterns, and 823 ± 16 ms vs 1124 ± 39 ms, $t_{48} = 14.7$, $P <$ 0.0001, for large patterns). Correct 'same' responses were significantly faster than correct 'different' responses (1012 ± 26 ms vs 1105 ± 26 ms, $t_{23} = 3.41$, $P < 0.01$, for small patterns, and 792 ± 21 ms vs 857 ± 23 ms, $t_{23} = 3.17$, $P < 0.01$, for large patterns).

There was no trade-off between performance (per cent correct) and reaction time. Over subjects and conditions, RTs for correct 'same' responses tended to be negatively correlated with performance, with gradient (mean ± 1 SEM) -7.56 ± 3.22 ms (per cent)$^{-1}$ ($z = 2.35$, $P < 0.05$) for small patterns and -8.54 ± 5.20 ms (per cent)$^{-1}$ ($z = 1.64$, $P > 0.1$) for large patterns.

Discussion

There were two major differences between 'same-different' discrimination performances with small and large patterns. First, performance averaged over all transformations was much higher for large patterns, consistent with subjects' reports that they found the task much easier. Second, and more important, was the result that differences in 'same'-detection performance for patterns related by global point-inversion ι and local-feature point-inversion ι_F disappeared for large patterns. This appears to suggest that local features were adequately accessed in internal representations and that the selective matching of local features independently of their orientations or the selective reversal of the sense of orientation relations associated with local features was indeed applied successfully for patterns related by ι_F.

There is, however, another possible interpretation of these results. Pattern pairs defined as 'different' were completely unrelated, that is, they differed in their local features f_i as well as in their spatial-order relations r_x and r_y. Therefore, if information about the structure of local features were suppressed and the internal matching operations were performed on the low-spatial-frequency content of the patterns alone, then the experimental task would be reduced effectively to determining whether the 'low-pass-filtered' patterns were the same (for transformations Id and ι_F) or point-inverted (for transformations ι and ι_P). Under this hypothesis, 'same'-detection performance should be the same for patterns related by transformations Id and ι_F and should be the same for patterns related by transformations ι and ι_P. The fact that there were significant differences in 'same'-detection performance for these two pairs of transformations suggests that complete suppression of local-feature information was not possible. This suppression interpretation did, however, match some subjects' reports that their strategy for solving the experimental task entailed looking only at the overall shape of the patterns without attending to the details of the subpatterns.

The main conclusion from this experiment is that some selective internal operations may be applied efficiently to internal representations, provided that the separation of subpatterns forming local features is large enough. These operations were mainly the selective matching of local features independently of their orientations or the selective reversal of the sense of local-feature orientation relations, or, most probably, the selective suppression of high-spatial-frequency information associated with local features and matching on the basis of low-spatial-frequency content alone. Experiment 3 was undertaken in order to determine whether selective local-feature matching or selective sense-reversal of orientation relations and spatial relations could be effected when the hypothesized suppression strategy was rendered inappropriate.

EXPERIMENT 3: SELECTIVE OPERATIONS ON SMALL AND LARGE PATTERNS

To determine whether local features could be matched independently of their orientations, it was necessary to eliminate the hypothesized low-spatial-frequency cues differentially present in 'same' and 'different' patterns. This was achieved by appropriately extending the class of pattern pairs defined as 'different'.

To describe the new transformations, an operator θ is introduced to denote a random transformation of the patterns. Thus two patterns A and θ(A) were completely unrelated (i.e. they were generated independently); patterns A and θ_P(A) had the same subpatterns, but unrelated spatial relations between subpatterns; and patterns A and θ_F(A) had the same positions of subpatterns in the patterns but the subpatterns themselves were unrelated. The enlarged class of transformations defining the patterns in each 'different' pair were:

(a) θ_F (local-feature randomization): subpattern positions were identical but subpatterns were unrelated;
(b) $\iota_P \circ \theta_F$ (position point-inversion with local-feature randomization): subpattern positions were point-inverted and subpatterns were unrelated;
(c) θ_P (local-feature-position randomization): subpattern positions were unrelated but subpatterns were identical;
(d) $\theta_P \circ \iota_F$ (local-feature-position randomization with local-feature point-inversion): subpattern positions were unrelated and subpatterns were point-inverted; and
(e) $\theta = \theta_P \circ \theta_F$ (global randomization): both subpattern positions and subpatterns were unrelated.

(Note that all operator compositions in (b), (d), and (e) were commutative.) In the two preceding experiments, only the last 'different' pattern relation θ was used to define 'different' pattern pairs. The rationale behind extending the class of 'different' pattern pairs was as follows. If subjects attended only to local-feature positions by suppressing information concerning the structure of local features, then they would fail to detect pattern pairs related by the transformations θ_F and by $\iota_P \circ \theta_F$ as 'different'. On the other hand, if subjects attended to local-feature structure alone and ignored local-feature position information, then they would fail to detect pattern pairs related by the transformations θ_P and by $\theta_P \circ \iota_F$ as 'different'.

To determine how well selective internal operations could be applied to the orientation relations associated with local features and the spatial relations between local features, pattern pairs to be discriminated included those with point-inverted and unrelated local features but with identical spatial relations, i.e. related by transformations ι_F and by θ_F, and those with the same and unrelated local features and both transformed by position point-inversion, i.e. related by transformations ι_P and by $\iota_P \circ \theta_F$.

Methods

Stimuli and experimental procedure were similar to those in Experiments 1 and 2. For small patterns, the maximum stimulus diameter and minimum separation of sub-pattern centres were 0.5° and 0.17°, respectively; for large patterns, these distances were 1.0° and 0.33°, respectively. Four possible transformations related patterns in each 'same' pair: Id, ι_P, ι_F and ι, and five possible transformations related patterns in each 'different' pair: θ_F, $\iota_P \circ \theta_F$, θ_P, $\theta_P \circ \iota_F$ and θ. There were 18 trials in each experimental run, and, in each run, each transformation occurred twice. The order of pattern transformations in each run was chosen pseudo-randomly but balanced over runs to offset stimulus order and carry-over effects. Each subject performed 25 runs for each of the pattern sizes, thus totalling 50 responses for each of the pattern transformations and pattern sizes. Subjects were the same as in Experiment 2.

Results

Discrimination performance. Figure 5 shows 'same-different' pattern-discrimination performance as a function of pattern transformation for small patterns (open circles) and for large patterns (solid circles): (a) identity transformation Id versus local-feature randomization θ_F, (b) position point-inversion ι_P versus position point-inversion with local-feature randomization $\iota_P \circ \theta_F$, (c) local-feature point-inversion ι_F versus local-feature randomization θ_F, and (d) global point-inversion ι versus position point-inversion with local-feature randomization $\iota_P \circ \theta_F$. The d' data were weighted by individual variances and averaged over subjects. Chi-squared tests on individual subjects' data showed no significant differences between subjects' performances ($\chi^2_8 = 3.40$, $P > 0.5$, for small patterns, and $\chi^2_8 = 7.73$, $P > 0.2$, for large patterns) after performance was normalized (Appendix 2) with respect to each subject's mean level.

Comparisons (Appendix 2) on averaged d' data for small patterns yielded the following results (two-tailed significance tests in all cases):

(a) 'same'-detection performance for patterns related by transformation Id was significantly higher than that for patterns related by any of the transformations ι, ι_P, and ι_F ($z \geqslant 2.87$, $P < 0.01$);

(b) 'same'-detection performance for patterns related by transformation ι was significantly higher than that for patterns related by either of the transformations ι_F and ι_P ($z \geqslant 5.84$, $P < 0.0001$); and

(c) 'same'-detection performance for patterns related by transformation ι_F was just different from chance level ($z = 2.23$, $P < 0.05$) whereas that for transformation ι_P did not exceed chance level ($z = 1.03$, $P > 0.1$), but the difference between the two performances was not significant ($z = 0.81$, $P > 0.1$).

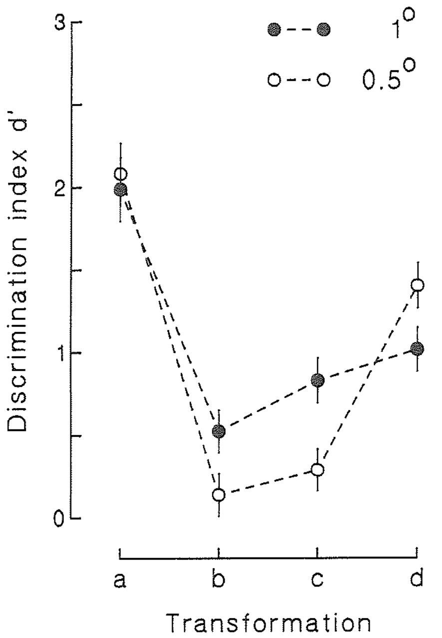

Figure 5. 'Same-different' discrimination performance as a function of pattern transformation. Mean discrimination index d' is shown for (a) identity transformation vs local-feature randomization, (b) position point-inversion vs position point-inversion with local-feature randomization, (c) local-feature point-inversion vs local-feature randomization, and (d) global point-inversion vs position point-inversion with local-feature randomization. Filled circles show performance for patterns with diameter 1° visual angle and open circles for patterns with diameter 0.5° visual angle. The vertical bars indicate ± 1 SEM.

As in Experiment 2 a significant difference in 'same'-detection performance was obtained for small and large patterns, even after allowance for overall performance level ($\chi^2_3 = 14.3$, $P < 0.01$). For large patterns the results were also similar to those of Experiment 2:

(a) 'same'-detection performance for patterns related by transformation Id was significantly higher than that for patterns related by any other transformation ($z \geqslant 4.07$, $P < 0.001$);
(b) 'same'-detection performance for patterns related by transformation $\imath$ did not differ significantly from that for patterns related by transformation $\imath_F$ ($z = 0.98$, $P > 0.1$);
(c) 'same'-detection performance for patterns related by transformation $\imath_P$ was significantly lower than that for patterns related by transformation $\imath$ ($z = 2.64$, $P < 0.01$) although not significantly lower than that for patterns related by transformation $\imath_F$ ($z = 1.62$, $P > 0.05$);
(d) 'same'-detection performance for patterns related by any of the transformations was significantly above chance level ($z \geqslant 4.0$, $P < 0.001$).

Reaction times. Averaged over subjects and conditions, correct responses (mean $\pm$ 1 SEM) were significantly faster than incorrect responses (722 ± 9 ms vs 839 ± 19 ms for small patterns, and 902 ± 14 ms vs 1117 ± 55 ms for large patterns, $t_{48} \geqslant 7.24$,

$P < 0.001$), but correct 'same' responses were not significantly faster than correct 'different' responses (725 ± 15 ms vs 721 ± 10 ms for small patterns, and 895 ± 20 ms vs 907 ± 19 ms for large patterns, $t_{23} \leqslant 0.62$, $P > 0.1$).

As in the preceding experiments, there was no tradeoff between performance (per cent correct) and reaction time (RT). Over subjects and conditions, RTs for correct 'same' responses were negatively correlated with performance, with gradient (mean $\pm$ 1 SEM) -4.15 ± 0.81 ms (per cent)$^{-1}$ for small patterns and -8.55 ± 1.28 ms (per cent)$^{-1}$ for large patterns ($z \geqslant 5.12$, $P < 0.001$).

Discussion

Comparison of 'same'-detection performance for local-feature point-inversion ι_F and position point-inversion ι_P for small and large patterns shows clearly that the separation for subpatterns in the 0.5°-diameter patterns was too small for the hypothesized local features to be accessed and operated on selectively. For the 1°-diameter patterns, however, the separation of subpatterns was large enough for selective operations to be applied successfully.

In Experiment 2 'same' patterns had to be discriminated from completely unrelated 'different' patterns, whereas in the present experiment they had to be discriminated from 'different' patterns with unrelated local features but with the same overall shape (i.e. the same local-feature positions). Comparison of the results of the two experiments reveals a remarkable stability in variation of discrimination performance across transformations. The introduction of the special 'different' patterns led to a general worsening in performance but did not affect differentially the internal matching operations. The worsening in discrimination performance is consistent with the interpretation that internal matching operations were in part performed on the basis of the low-spatial-frequency content of the patterns. This suppression strategy was, however, not applicable in the present experiment and it appears that the hypothesized internal operations were indeed applied selectively to local features. Thus local features were matched independently of their orientations or the sense of orientation relations was reversed without affecting the sense of the spatial relations between local features. Rather less efficiently, local features were also matched independently of their spatial relations or the sense of spatial relations was reversed without modifying the orientation relations of local features.

SUMMARY AND CONCLUSION

The experiments described in this study have been aimed at investigating the association between local features and spatial relations assumed to be involved in the internal representations of visual patterns. Previous studies (Foster, 1978; Kahn and Foster, 1981; Foster and Kahn, 1985) have led to the hypothesis that these two components of internal representations may be so closely linked that it may be impossible to apply efficiently internal operations that selectively manipulate either local features or the spatial relations associated with local features. The main conclusion from this study is that selective internal operations can be applied efficiently, but only if certain conditions are met.

One condition for the efficient application of selective internal operations appears to be a large separation of the subpatterns giving rise to the local features in the internal representation. For the small patterns used in this study, the centre-to-centre separation of subpatterns was twice the subpattern diameter in over 87% of cases, but performance under local-feature point-inversion ι_F and position point-inversion ι_P was relatively poor. Two possible mechanisms were considered: first, that the separation of subpatterns was such that local features and spatial relations were not adequately separated and accessed in the internal representation; second, that although local features and spatial relations were adequately separated and accessed, the hypothesized internal sense-reversal operations could not be applied effectively. In Experiment 2 it was shown that for larger patterns with proportionally greater subpattern separations performance improved overall, and, in particular, performance with patterns related by local-feature point-inversion reached the level of that with patterns related by global point-inversion. Performance with patterns related by position point-inversion remained the poorest. It was noted, however, that performance under local-feature point-inversion and position point-inversion could have resulted from a strategy of low-spatial-frequency sampling of the patterns in which information concerning local-feature structure was suppressed. In Experiment 3, it was shown that with the introduction of special 'different' pattern pairs the variations in performance under the three point-inversion operations could not be attributed to this strategy. Performances under global point-inversion ι and local-feature point-inversion ι_F, although lowered, were still equal and performance under position point-inversion ι_P, although not reduced to chance level, was still inferior to both.

One possible explanation for the relatively poor performance under position point-inversion ι_P exploits the fact that ι_P is formally equivalent to the composition of global point-inversion and local-feature point-inversion, $\iota_P = \iota \circ \iota_F$ (compare Figs 1b and c). Let us suppose that the composition of two internal operations results in some less efficient internal operation, perhaps through increased computational complexity or time required to effect the operations serially. Let us further suppose that the internal sense-reversal operation σ discussed in the Introduction can be applied to the whole internal representation and selectively to the orientation relations associated with local features, but not selectively to the spatial relations between local features. Then performances with patterns related by ι and by ι_F should be relatively good, and equal, and should be better than performance with patterns related by ι_P, for 'same'-detection under the last could be achieved only by the composition of the two separate internal operations.

Since performance with large patterns was relatively good under local-feature point-inversion ι_F, implying the efficiency of at least this selective internal operation, it seems likely that the relatively poor performance under ι_F with small patterns resulted from a failure adequately to separate and access the transformed local features in the internal representation. This failure may have, in turn, resulted from the creation of more complex local features as part of a hierarchy of structures depending on the density and proximity of subpatterns in the stimulus displays (Foster, 1980; Koenderink, 1984). If subpatterns were composed of connected lines rather than disconnected points, it is possible that fewer or weaker higher-order structures would have been introduced and the internal sense-reversal operation σ could have been applied selectively to small patterns as efficiently as it was to large patterns.

The results of this study contrast with those of Foster (1978) who had used 'same-different' judgements of shape and of numerosity in random-dot patterns to argue for the inseparability of local features and spatial relations between local features. An important difference between these two studies may, however, have been that, in the numerosity judgements, *whole* patterns had to be processed, whereas in the 'same-different' judgements involved here, correct responses could have been made on the basis of processing parts of patterns alone.

Acknowledgement

The authors are grateful to S. R. Pratt for critical reading of the manuscript and R. J. Watt for useful discussion. W.F.B. was supported by grant 81.166.0.84 of the Swiss National Science Foundation. J.I.K. was supported by an award from the Science and Engineering Research Council. R. Knapper provided technical assistance.

REFERENCES

Arnoult, M. D. (1954). Shape discrimination as a function of the angular orientation of the stimuli. *J. Exp. Psychol.* **47**, 323–328.

Aulhorn, O. (1948). Die Lesegeschwindigkeit als Funktion von Buchstaben und Zeilenlage. *Pflügers Archiv* **250**, 12–25.

Barlow, H. B., Narasimhan, R., and Rosenfeld, A. (1972). Visual pattern analysis in machines and animals. *Science* **177**, 567–575.

Cooper, L. A. and Shepard, R. N. (1973). Chronometric studies of the rotation of mental images. In: *Visual Information Processing.* W. G. Chase (Ed.). Academic Press, New York (pp. 75–176).

Dearborn, G. V. N. (1899). Recognition under objective reversal. *Psychol. Rev.* **6**, 395–406.

Durlach, N. I. and Braida, L. D. (1969). Intensity perception. I. Preliminary theory of intensity resolution. *J. Acoust. Soc. Am.* **46**, 372–383.

Foster, D. H. (1978). Visual comparison of random-dot patterns: evidence concerning a fixed visual association between features and feature-relations. *Quart. J. Exp. Psychol.* **30**, 637–654.

Foster, D. H. (1980). A description of discrete internal representation schemes for visual pattern discrimination. *Biol. Cybern.* **38**, 151–157.

Foster, D. H. and Kahn, J. I. (1985). Internal representations and operations in the visual comparison of transformed patterns: effects of pattern point-inversion, positional symmetry, and separation. *Biol. Cybern.* **51**, 305–312.

Foster, D. H. and Mason, R. J. (1979). Transformation and relational-structure schemes for visual pattern recognition. *Biol. Cybern.* **32**, 85–93.

Gourevitch, V. and Galanter, E. (1967). A significance test for one parameter isosensitivity functions. *Psychometrika* **32**, 25–33.

Green, D. M. and Swets, J. A. (1966). *Signal Detection Theory and Psychophysics.* John Wiley, New York.

Kahn, J. I. and Foster, D. H. (1981). Visual comparison of rotated and reflected random-dot patterns as a function of their positional symmetry and separation in the field. *Quart. J. Exp. Psychol.* **33A**, 155–166.

Koenderink, J. J. (1984). The structure of images. *Biol. Cybern.* **50**, 363–370.

Kolers, P. A. and Perkins, D. N. (1969). Orientation of letters and their speed of recognition. *Percept. Psychophys.* **5**, 275–280.

Kolers, P. A. and Perkins, D. N. (1975). Spatial and ordinal components of form perception and literacy. *Cognitive Psychol.* **7**, 228–267.

Mach, E. (1897). *Contributions to the Analysis of the Sensations.* Open Court, Chicago. Republished (1959) in revised and supplemented form as *The Analysis of Sensations.* Dover, New York.

Rock, I. (1973). *Orientation and Form.* Academic Press, New York.

Shepard, R. N. (1975). Form, formation, and transformation of internal representations. In: *Information Processing and Cognition: The Loyola Symposium.* R. L. Solso (Ed.). Lawrence Erlbaum, Hillsdale, NJ (pp. 87–122).

Shepard, R. N. and Cooper, L. A. (1982). *Mental Images and their Transformations.* MIT Press, Cambridge, Mass.

Shepard, R. N. and Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science* **171**, 701–703.

Sutherland, N. S. (1968). Outlines of a theory of visual pattern recognition in animals and man. *Proc. R. Soc.* **171B**, 297–317.

Swets, J. A. (1973). The relative operating characteristic in psychology. *Science* **182**, 990–1000.

Thompson, P. (1980) Margaret Thatcher: a new illusion. *Perception* **9**, 483–484.

APPENDIX 1

Notation here follows that used in the main text.

Global point-inversion ι

Let R_1, R_2 be the respective internal representations of the patterns A_1, A_2 with $A_2 = \iota(A_1)$. The relevant components of R_1, R_2 may be displayed as pairs of finite sequences, thus

$$R_1 = ((f_i{:}1 \leqslant i \leqslant m), (r_x(f_j, f_k), r_y(f_j, f_k){:}1 \leqslant j < k \leqslant m)),$$
$$R_2 = ((\iota(f_i){:}1 \leqslant i \leqslant m), (r_x(\iota(f_j, f_k)), r_y(\iota(f_j, f_k)){:}1 \leqslant j < k \leqslant m)),$$

where $\iota(f_i)$, $\iota(f_j, f_k)$ are the transformed local features and their transformed spatial relationships. (A more detailed description is given in Foster, 1980.) Application of the internal sense-reversal operation $\sigma = (\sigma_x, \sigma_y)$ to the spatial-relation component $r_x(\iota(f_j, f_k))$ in R_2 results in $r_x(f_j, f_k)$, thus

$$\sigma_x(r_x(\iota(f_j, f_k))) = \sigma_x(-r_x(f_j, f_k)) = r_x(f_j, f_k).$$

Application of σ to $r_y(\iota(f_j, f_k))$ and to the orientation relations associated with the $\iota(f_i)$ in R_2 has analogous effects. By this process, the two internal representations R_1, R_2 are brought into coincidence.

Local-feature point-inversion ι_F

Let R_1, R_2 be the respective internal representations of the patterns A_1, A_2 with $A_2 = \iota_F(A_1)$. Then

$$R_1 = ((f_i{:}1 \leqslant i \leqslant m), (r_x(f_j, f_k), r_y(f_j, f_k){:}1 \leqslant j < k \leqslant m)),$$
$$R_2 = ((\iota_F(f_i){:}1 \leqslant i \leqslant m), (r_x(f_j, f_k), r_y(f_j, f_k){:}1 \leqslant j < k \leqslant m)).$$

In the experiment, the subpatterns were chosen so that $\iota_F(f_i) \neq f_i$. If the internal sense-reversal operation σ may be applied selectively to the orientation relations associated with the f_i, then R_1 and R_2 are brought into coincidence.

Position point-inversion ι_P

Let R_1, R_2 be the respective internal representations of the patterns A_1, A_2 with $A_2 = \iota_P(A_1)$. Then

$$R_1 = ((f_i{:}1 \leqslant i \leqslant m), (r_x(f_j, f_k), r_y(f_j, f_k){:}1 \leqslant j < k \leqslant m)),$$
$$R_2 = ((f_i{:}1 \leqslant i \leqslant m), (-r_x(f_j, f_k), -r_y(f_j, f_k){:}\ 1 \leqslant j < k \leqslant m)).$$

If the internal sense-reversal operation σ may be applied selectively to the spatial relations between the f_i, then R_1 and R_2 are brought into coincidence.

APPENDIX 2

The scores of each subject were converted into the discrimination index d' using the false-alarm rate (that is, the proportion of incorrect 'same' responses). Variances were estimated using the method described by Gourevitch and Galanter (1967).

Chi-squared test for differences between subjects

The discrimination indices d'_{ij} and variances v_{ij}, where $i = 1, \ldots, n_s$ specifies the subject and $j = 1, \ldots, n_t$ specifies the 'same' transformation, were used to compute the quantity

$$\Sigma_{ij}(d'_{ij} - d'_{.j})^2/v_{ij},$$

where $d'_{.j} = (\Sigma_i d'_{ij}/v_{ij})/(\Sigma_i 1/v_{ij})$, which has variance $v_{.j} = (\Sigma_i 1/v_{ij})^{-1}$. Under the hypothesis that there are no differences between subjects' performances, the computed quantity should be distributed as chi-squared with $(n_s - 1)n_t$ degrees of freedom.

Chi-squared test for differences between subjects allowing for each subject's overall performance level

Let the notation be as above. The mean performance level $d'_{i.} = (\Sigma_j d'_{ij})/n_t$ for each subject $i = 1, \ldots, n_s$ was subtracted from his d' scores to give a normalized value $e_{ij} = d'_{ij} - d'_{i.}$. Under the hypothesis that there are no differences between subjects' performances when each of these is expressed relative to the subject's mean performance level, the quantity

$$\Sigma_{ij}(e_{ij} - e_{.j})^2/v_{ij},$$

where $e_{.j} = (\Sigma_i e_{ij}/v_{ij})/(\Sigma_i 1/v_{ij})$, should be distributed as chi-squared with $n_s n_t - n_s - n_t$ degrees of freedom.

Contrasts for differences between conditions

Let $c_j, j = 1, \ldots, n_t$ be the weights of a contrast. Under the hypothesis that there is no effect, the quantity

$$\Sigma_j c_j d'_{.j}/(\Sigma_j c_j^2 v_{.j})^{1/2}$$

should be distributed as the standard normal variable z.