
Qualitative cues in the discrimination of affine-transformed minimal patterns

Heljä T Kukkonen¶, David H Foster#, Jonathan R Wood§

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

Johan Wagemans

Laboratory of Experimental Psychology, Katholieke Universiteit Leuven, Tiensestraat 102, B-3000 Leuven, Belgium

Luc Van Gool

ESAT-M12, Katholieke Universiteit Leuven, Kardinaal Mercierlaan 94, B-3001 Leuven, Belgium

Received 12 July 1995, in revised form 15 December 1995

Abstract. An important factor in judging whether two retinal images arise from the same object viewed from different positions may be the presence of certain properties or cues that are 'qualitative invariants' with respect to the natural transformations, particularly affine transformations, associated with changes in viewpoint. To test whether observers use certain affine qualitative cues such as concavity, convexity, collinearity, and parallelism of the image elements, a 'same-different' discrimination experiment was carried out with planar patterns that were defined by four points either connected by straight line segments (line patterns) or marked by dots (dot patterns). The first three points of each pattern were generated randomly; the fourth point fell on their diagonal bisector. According to the position of that point, the patterns were concave, triangular (three points being collinear), convex, or parallel sided. In a 'same' trial, an affine transformation was applied to one of two identical patterns; in a 'different' trial, the affine transformation was applied after the point lying on the diagonal bisector was perturbed a short, fixed distance along the bisector, inwards for one pattern and outwards for the other. Observers' ability to discriminate 'same' from 'different' pairs of patterns depended strongly on the position of the fourth, displaced, point: performance varied rapidly when the position of the displaced point was such that the patterns were nearly triangular or nearly parallel sided, consistent with observers using the hypothesised qualitative cues. The experimental data were fitted with a simple probabilistic model of discrimination performance that used a combination of these qualitative cues and a single quantitative cue.

1 Introduction

Objects in three-dimensional space are presented to the visual system through their two-dimensional retinal projections. The retinal images change according to the point of observation, even when there are no changes within the scene. Yet these changes are usually correctly interpreted as arising from changes in viewpoint, and the shapes of the objects are perceived to be constant.

In one class of models explaining shape constancy it is assumed that the visual system bases its judgments on properties of the projected image that remain invariant under changes in viewpoint (eg Gibson 1950; Foster 1977; Cutting 1986; Biederman and Gerhardstein 1993; Pizlo 1994; Wagemans et al 1996). For example, in a pattern-discrimination task where an observer has to report whether two images presented to the eye arise from the same object viewed from different positions or from two different objects, the decision would be based on detecting whether certain invariant properties are common to the two images.

The set of properties that remain invariant under changes in viewpoint depends on the geometric transformations used to describe the changes in the projected image.

Present addresses: ¶ Department of Optometry and Vision Sciences, University of Wales College of Cardiff, Cardiff CF1 3XF, UK (author to whom correspondence should be addressed);

Department of Vision Sciences, Aston University, Birmingham B4 7ET, UK; § Department of Human Sciences, Loughborough University, Loughborough LE11 3TU, Leics, UK.

The transformations of interest in vision include perspective and affine transformations of the plane. Perspective transformations are important for describing changes in the projected image of an object when the viewing distance is small relative to the range of the object in depth, but when the relative viewing distance is large the effect of perspective is negligible: parallel edges in the object project into parallel lines in the image. In the latter condition, affine transformations may be used to describe the changes in the projected image.

Although visual invariants may include continuously varying quantities such as the ratios of distances along a line, it is possible that these invariants also include discrete or 'qualitative' cues (Foster 1980a; Wagemans 1993); for example, collinearity and parallelism (Foster 1980b; Biederman 1987; Wagemans 1993), and related properties such as concavity and convexity (Koenderink 1986; Ullman 1989). In some pattern-discrimination tasks (eg Wagemans 1992, 1993) it has been found that when two patterns have the same qualitative properties they are more likely to be judged as 'same' than when this property is present in only one of the patterns. This approach has also been exploited in a theory of 'recognition by components' (Biederman 1987; Biederman and Gerhardstein 1993), in which certain regularities in the image that are relatively stable under changes in viewpoint (eg collinearity and parallelism) have a significant effect on an observer's interpretation of the shape of the object. For example, parallel lines in a projected image of an object define an object with a constant cross section, and a straight line in the image (collinearity) is interpreted as a straight edge in the three-dimensional world (Biederman 1987).

The aim in this study was to quantify the effect of four qualitative cues, collinearity, convexity, concavity, and parallelism, on recognition of simple dot and line patterns under affine transformations. The patterns used were defined by four distinct (and, in general, noncollinear) points, the minimum number required to define uniquely an affine equivalence (Van Gool et al 1994). The shapes of these patterns were manipulated with the aid of a perturbation method (Foster 1980b) that progressively distorted pairs of 'same' patterns and degraded each of the qualitative cues. In this perturbation method, pairs of identical patterns are first identically distorted by a variable amount and then an additional, small, constant distortion (perturbation) is applied in opposite directions to the two patterns. The discriminability of the perturbed patterns from the corresponding nonperturbed patterns is determined as a function of the variable amount of distortion. In this study, the distortion altered the strength of each of the four qualitative cues; the form of the discrimination performance thus obtained was related to the characteristics of the cues. It is suggested that a combination of these qualitative cues and a continuous, quantitative cue accounts better for observed discrimination performance than either type of cue in isolation. A model employing such a combination of continuous and discrete cues is presented in the Appendix.

2 Methods

2.1 Apparatus and stimuli

Stimuli were generated under computer control on the screen of a colour monitor (Sony Corp, Tokyo, Japan; Triniton model GDM-2067S) with a refresh rate of 70 Hz. The 640 pixel \times 480 pixel display area occupied a 360 mm \times 243 mm rectangle in the centre of a 400 mm \times 300 mm screen. Pixel size was 0.56 mm \times 0.51 mm. The screen was viewed at a distance of 167 cm. The patterns comprised either four black dots ('dot' patterns), each dot of diameter 6 pixels, corresponding to a visual angle of 6.6 min arc, or four black lines forming a quadrilateral ('line' patterns), each line of width 1 pixel, corresponding to a visual angle of 1.1 min arc. In each trial, two such dot or line patterns were presented side by side on a uniform, light-grey field of luminance approximately 35 cd m⁻².

The x, y -coordinates of the four points defining each pattern were selected to fall within an imaginary circle of diameter 5.0 deg with the restriction that only one point was allowed in each quadrant of the circle. The orientation of the system of orthogonal axes defining the quadrants was selected at random. The horizontal separation between the centres of the two imaginary circles was 7.0 deg. In addition, no two points in any pattern were allowed closer than 7.7 min arc.

For tests of the parallelism cue, the first three points were generated randomly within the constraints just described and the fourth point was positioned randomly at one of fifteen possible locations on the diagonal bisector defined by the first three points (figure 1), with the centre location defining the parallel arrangement (in fact a double parallelism) of the points. The visual angle separating adjacent locations on the bisector was constant at 15 min arc (a value that was obtained from exploratory measurements).

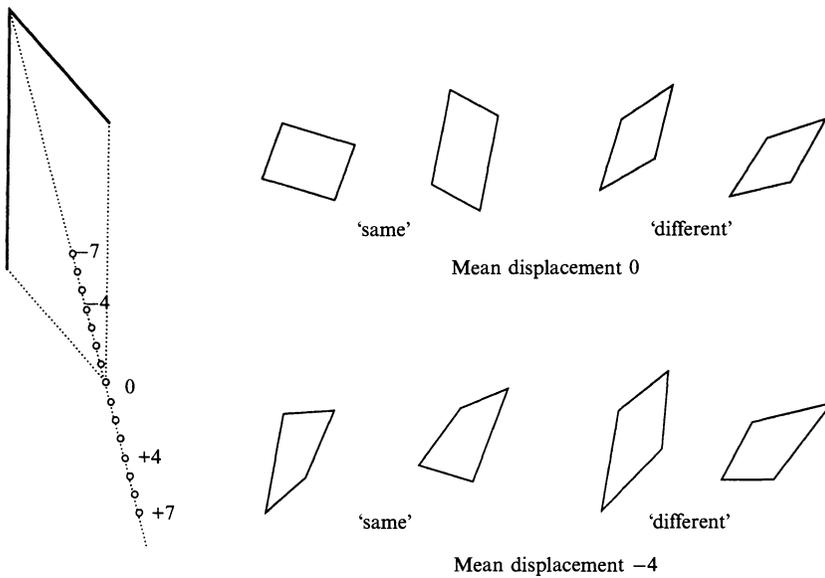


Figure 1. Distortion of parallel line patterns and examples of 'same' and 'different' pattern pairs corresponding to positions of 0 and -4 of the displaced point.

For tests of the collinearity cue, the points were generated in the same way except that the fifteen possible locations of the fourth point along the diagonal bisector were centred on the collinear arrangement of the points (figure 2). According to the position of the point on the bisector, the patterns were concave, triangular (when the fourth point was collinear with its neighbours), or convex (figure 2).

Pairs of patterns were either 'same' or 'different'. For a 'same' pair, the two patterns were generated identically (the fourth point therefore occupying the same position on the diagonal bisector), and an affine transformation was then applied to one of the patterns. For a 'different' pair of patterns, the two patterns were generated identically but the fourth point was perturbed a fixed distance (1 step, 15 min arc) inwards along the bisector in one pattern and outwards in the other; the affine transformation was then applied to one of the patterns. The transformed and untransformed patterns were randomly assigned to the left and right halves of the screen. If the fourth point of either pattern (after displacement, but before applying the affine transformations) fell outside the imaginary circle containing the pattern a completely new set of points was generated (repeatedly if necessary) until all points fell within the circle. No constraint

was imposed on the closeness of the fourth point to the opposite point, but, in the subsequent analysis, extreme cases, where these two points crossed over or were less than 1 step (15 min arc) apart, were rejected. When parallelism was the only cue present, the rare cases in which one of the patterns was not convex (enabling observers to use concavity or collinearity as a cue) were also rejected.

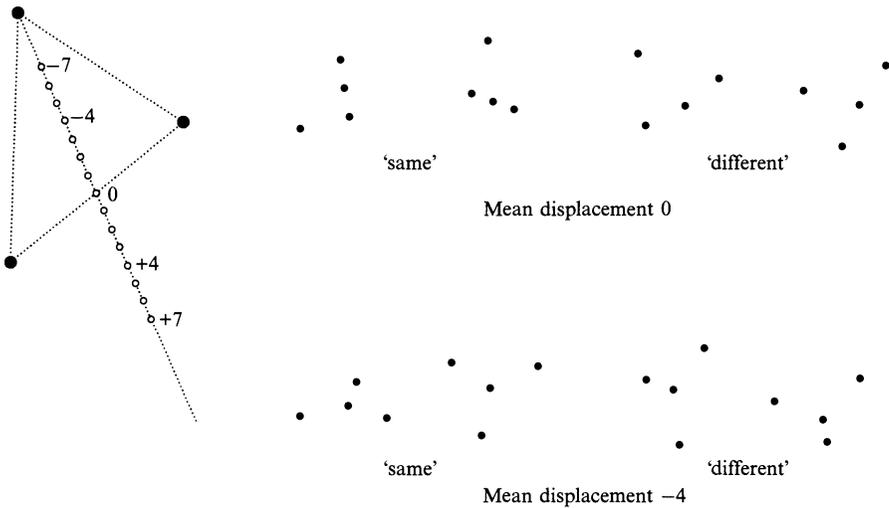


Figure 2. Distortion of collinear dot patterns and examples of 'same' and 'different' pattern pairs corresponding to positions of 0 and -4 of the displaced point.

2.2 Transformations

The affine transformations applied to the planar patterns were resolved into three components each with a clear perceptual interpretation. First a Euclidean rotation ρ in the image plane was applied. It was followed by a slant σ , that is, a vertical compression by $\cos \sigma$, equivalent to a rotation of the pattern about a horizontal axis in the image plane through angle σ , which left the horizontal dimensions unchanged. Finally, a tilt τ , a second Euclidean rotation in the image plane, was applied. Changes in scale (dilatations) and position (translations) were not of interest here (see eg Kahn and Foster 1981; Nazir and O'Regan 1990; Foster 1991), and were therefore not included in the transformations applied.

The transformation matrix was thus the matrix product

$$\begin{bmatrix} \cos \tau & -\sin \tau \\ \sin \tau & \cos \tau \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & \cos \sigma \end{bmatrix} \begin{bmatrix} \cos \rho & -\sin \rho \\ \sin \rho & \cos \rho \end{bmatrix} \\ = \begin{bmatrix} (\cos \rho \cos \tau - \sin \rho \cos \sigma \sin \tau) & (-\sin \rho \cos \tau - \cos \rho \cos \sigma \sin \tau) \\ (\cos \rho \sin \tau + \sin \rho \cos \sigma \cos \tau) & (-\sin \rho \sin \tau + \cos \rho \cos \sigma \cos \tau) \end{bmatrix}. \quad (1)$$

The planar rotation ρ took twelve values ranging over 360° in 30° steps; the slant σ (vertical compression) took five values ranging from 0° to 60° in 15° steps; and the tilt τ took four values ranging from 0° to 135° in 45° steps.

Each experimental session included exactly one instance of every combination of the three transformation components used to define a 'same' pair of patterns and exactly one instance of a corresponding 'different' pair of patterns. There were thus 12 rotation angles $\times 5$ slant angles $\times 4$ tilt angles yielding 240 pairs of 'same' patterns and an equal number of 'different' patterns. The 480 conditions occurred in random order in each experimental session.

2.3 Procedure

The screen was viewed binocularly at a distance of 167 cm in a darkened room, with ambient illumination approximately 1 cd m^{-2} . Subjects were instructed to respond 'same' if the two patterns could be regarded as the same planar object observed from two different viewpoints and at a distance which was large relative to the object's size. (As one of the objects could be assumed to be in the frontoparallel plane, its size determined an upper limit on the range of the object in depth when the object was oriented at an angle to the viewing direction.) The experimenter illustrated the effect of viewpoint by using a card on which a typical stimulus pattern, similar in size to those displayed on the screen, was held near the plane of the screen and tilted.

Trials were initiated by the observer using a push-button switch box held in the nondominant hand. Each trial consisted of a single presentation of a pair of stimulus patterns, to which the subject responded either 'same' or 'different' by pressing one of two buttons on a second push-button switch box held in the dominant hand. The patterns remained on the screen until a response button was pressed.

An experimental session lasted between 40 min and 1 h and comprised three blocks of 160 trials preceded by a shorter, practice block of 50 trials, the responses from which were subsequently discarded. Subjects were instructed to take breaks of 1–3 min between blocks and at no other time. Trial-to-trial feedback was given during the practice block; during the experiment proper, the only feedback given was the total proportion correct at the end of each block, which changed little from block to block.

2.4 Subjects

Ten observers (three female, seven male) participated in the experiment. They were aged between 20 and 28 years; all had normal or corrected-to-normal vision (Snellen acuities for each eye were 6/5 or better); and all were unaware of the aims of the experiment.

3 Results

Performance in detecting 'same' and 'different' pairs of affine transformed patterns was evaluated as a function of the distance of the displaced point along the diagonal bisectors of the patterns. Figure 3 shows the results. For ease of graphical presentation, displacement distances were measured from either the parallel or the collinear arrangement of the points. Data are presented separately for the collinearity and parallelism cues and for dot and line patterns. Proportion-correct scores were pooled over the ten observers. Binomial estimates of the standard errors of each score varied from 0.02 to 0.06. The dotted and continuous lines in the figures show the corresponding performance of a simple probabilistic model, which is described in the Appendix.

Proportion-correct scores for 'same' and 'different' patterns were not presented as a combined single measure of discrimination performance, such as the discrimination index d' from signal-detection theory, because, as becomes clear later, the effect of the qualitative cues on 'same' and 'different' responses needed separate analyses for 'same' and 'different' patterns. Nevertheless, it is worth noting that observers' performance in discriminating 'same' from 'different' patterns was above chance: average values of d' ranged from 0.37 to 1.23 in the four conditions tested (and in a fifth control condition described later).

Data for the parallelism cue for dot and line patterns are shown in figures 3a and 3b, respectively. The effect of the parallelism cue was similar for dot and line patterns, but was slightly stronger for line patterns (figure 3b). Performance in detecting 'same' pairs of patterns showed a small local maximum when the displaced point was such that the two figures were perfect parallelograms (as in figure 1). As the displaced point was moved away from this position, performance first worsened slightly, but then improved again.

For 'different' pairs of patterns, the position of the displaced point refers to its mean position over the two opposite perturbations of one step of 15 min arc (figure 1). In contrast to the data for 'same' patterns, detection performance with 'different' patterns showed a strong local minimum when the displaced point was such that the two figures were perfect parallelograms. As the displaced point was moved away from this position, performance first improved, but then worsened slightly. There was an effect of the direction in which the displaced point was moved: performance was better when it was closer to the opposite point of the figure than when it was further away. The effect was absent with 'same' patterns.

Data for the collinearity cue for dot and line patterns are shown in figures 3c and 3d, respectively. The effect of the cue was again similar for dot and line patterns, but was much stronger for line patterns (figure 3d), possibly because the greater precision with which the position of the lines could be judged provided a stronger cue for collinearity, convexity, and concavity. As with the parallelism cue, detection performance with 'same' pairs of patterns showed a small local maximum when the

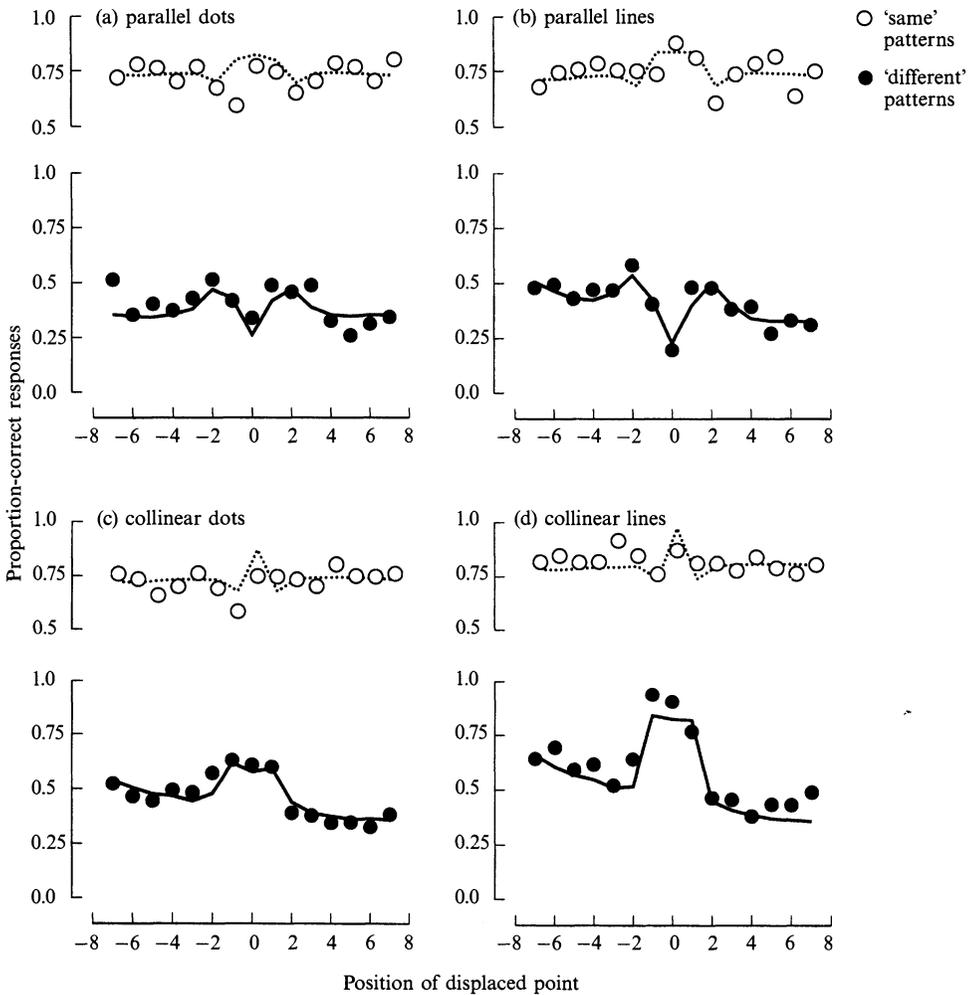


Figure 3. Proportion-correct scores plotted as a function of position of the displaced point, expressed in steps of 15 min arc, in 'same' and 'different' dot and line patterns testing the parallelism cue (a) and (b) and the collinearity cue (c) and (d). The continuous and dotted lines show performance of a simple probabilistic model of detection performance, described in the Appendix.

displaced point was such that the figures were perfect triangles (as in figure 2). As the displaced point was moved away from this position, performance first worsened slightly, but then improved again.

For 'different' patterns there was a strong local maximum when the displaced point was at or within ± 1 step of the position where the figures were perfect triangles. As the displaced point was moved further away from this position, performance worsened rapidly to the asymptotic level. As with the parallelism cue there was an effect of the direction in which the displaced point was moved: performance was better when it was closer to the opposite point of the figure than when it was further away; the effect was again absent for 'same' patterns.

As was made clear in section 2, the patterns contained the minimum amount of information required for the task; they were defined by four black dots or by four black lines, so that only the spatial arrangement of the dots or lines defined the correspondence between each dot or line in one pattern and each dot or line in the other. To test whether any additional information concerning correspondence would improve observers' performance, the experiment with the parallel dot patterns was repeated, but with one dot in one pattern and the corresponding dot in the other coloured blue, thereby resolving the potential ambiguity. The resulting data, presented in figure 4, were of the same form as the data in figure 3a for patterns with identical dots. Thus, in the conditions of this experiment, additional identifying information did not appear to be important in determining observers' performance (but see Wagemans et al 1996).

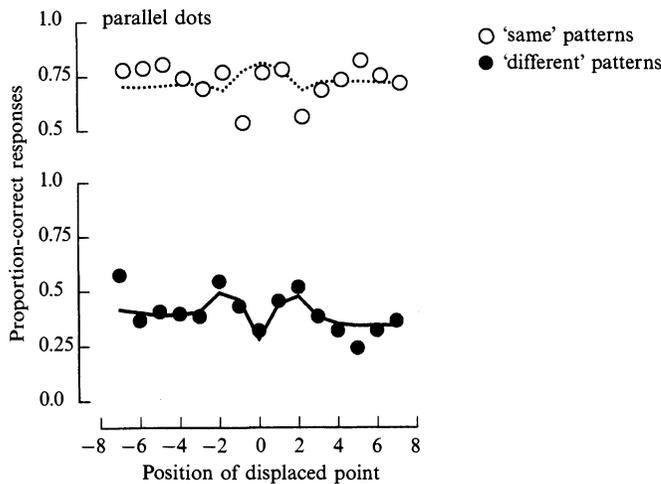


Figure 4. Proportion-correct scores plotted as a function of position of the displaced point, expressed in steps of 15 min arc, in 'same' and 'different' dot patterns testing the parallelism cue but with one dot in one pattern and the corresponding dot in the other coloured blue. The continuous and dotted lines show performance of a simple probabilistic model of detection performance, described in the Appendix.

4 Discussion

The experiment confirmed the assumption that human observers are able reliably to discriminate affine-related ('same') patterns from non-affine-related ('different') patterns defined by just four points, the minimum number required to define uniquely an affine equivalence. Irrespective of the particular effects of pattern distortion and any bias by observers towards one kind of response or another ('same' or 'different'), average values of the discrimination index d' from signal-detection theory were well above zero.

The variations in performance in detecting 'same' and 'different' patterns as a function of their distortion over regions defining collinearity, convexity, concavity, and parallelism revealed the effects of these qualitative cues. To summarise for the parallelism cue: observers were successful in detecting 'same' pairs of patterns when they were parallel; performance worsened slightly as the patterns were made less parallel, and then improved again as they were made less parallel still. This variation in performance can be interpreted in terms of the simple probabilistic model described quantitatively in the Appendix. The model made use of both qualitative (discrete) and quantitative (continuous) cues. The following comments are concerned mainly with qualitative cues. Assume that the visual system extracts from the two patterns estimates x_1, x_2 of the positions of the displaced points defined by their distances along the diagonal from some reference position, such as the opposite point; both x_1, x_2 are uncertain and can be treated as random variables (and for 'same' patterns therefore drawn from identical distributions). Consider the variation in how the patterns would be classified as they were progressively distorted. When the patterns were both parallel, there would be a high probability of both being classified as parallel and therefore 'same' (x_1, x_2 falling within a certain 'parallel' range). When the patterns were close to being parallel but not exactly so, there would be an increased probability that one of the 'same' patterns would be classified as parallel, and the other not (as the expected values of x_1, x_2 approach the boundary of the central 'parallel' range, the probability of the one falling outside and the other inside approaches a maximum); detection performance for 'same' would approach a minimum. When the patterns were far from being parallel, there would be a high probability of both being classified as nonparallel (the expected values of x_1, x_2 being far from the 'parallel' range); detection performance for 'same' would approach a maximum again.

For 'different' trials, observers performed poorly in detecting 'different' patterns when they were close to being parallel; performance improved as the patterns were made less parallel, and then worsened slightly as they were made less parallel still. In terms of the model, there would be a high probability that the 'different' pairs of patterns close to being parallel were both classified as parallel (x_1, x_2 falling within the 'parallel' range), and detection performance for 'different' would thus be at a minimum. When one of the patterns was parallel or nearly parallel and the other not (displacement positions ± 1 or ± 2), there would be a high probability that one of the 'different' patterns would be classified as parallel, and the other not (as with 'same' patterns, as the expected values of x_1, x_2 approach the boundary of the 'parallel' range, the probability of the one falling within and the other without approaches a maximum); detection performance for 'different' would approach a maximum. When the patterns were far from being parallel, there would be an increased probability of both being classified as nonparallel and detection performance for 'different' would decrease.

To summarise for the collinearity cue: as with the parallelism cue, observers were successful in detecting 'same' pairs of patterns when they were collinear; performance worsened slightly as the patterns were made slightly concave or convex, and then improved again as they were made more strongly concave or convex. The interpretation provided by the model is similar to that for the parallel cue with 'same' patterns. Thus, when the patterns were both strictly collinear, there would be a high probability of both being classified as collinear and therefore 'same'. When the patterns were close to being collinear but not exactly so, there would be an increased probability that one of the 'same' patterns would be classified as collinear, and the other not, and 'same' detection performance would approach a minimum. When the patterns were far from being collinear, there would be a high probability of both being classified as noncollinear, and therefore detection performance for 'same' would approach a maximum.

For 'different' trials, observers performed well in detecting 'different' patterns when they were close to being collinear and less well as the patterns were made more non-collinear, in contrast to their performance with the parallelism cue. In terms of the model, there would be a high probability that the 'different' pairs of patterns close to being collinear would be classified as 'different', the one concave and the other convex. The 'collinear' range of position values was thus clearly narrower than the 'parallel' range (the expected values of x_1, x_2 would then be closer to the boundaries of the 'collinearity' range, so that a parity cue, 'concave' for negative x_1 and 'convex' for positive x_2 , would be generated; see Appendix). Similarly, when one of the patterns was collinear and the other not (displacement position ± 1), there would be a slightly higher probability that one of the 'different' patterns would be classified as convex or concave, and the other as collinear; detection performance for 'different' would be at a maximum. When the patterns were far from being collinear, there would be an increased probability of both being classified as noncollinear, and detection performance for 'different' would then decrease.

As noted earlier, for both parallel and collinearity cues, there was an effect of the direction of pattern distortion in 'different' trials: performance was better when the displaced points were closer to the opposite point of the figure than when they were further away. There was no effect in 'same' trials. In the model these characteristics were accommodated by a Weber-like weighting of the continuous quantitative cue: the effect of the difference in positions $x_2 - x_1$ was weighted relative to the mean distance of the displaced points from the opposite point. The effect would therefore have been absent, on average, with 'same' patterns since the expected value of the difference cue was zero. This continuous cue was assumed to account for the residual performance in the absence of discrete qualitative cues (see Appendix for details).

The properties of collinearity and parallelism, and convexity and concavity can generally provide strong cues to identifying and discriminating random-dot and random-line patterns. The ranges over which collinearity and parallelism are visually defined, however, are different, and, depending on the amount by which patterns deviate from perfect parallelism or collinearity, the presence of qualitative cues can lead to erroneous judgments about the sameness or otherwise of patterns. Since qualitative cues are important in determining performance with these minimal patterns, such cues in conjunction with possible quantitative cues, may provide a useful basis for discriminating more complex, realistic patterns.

Acknowledgements. This research was supported by Esprit Basic Research Action No 6448 (VIVA). Johan Wagemans was supported by the National Fund for Scientific Research (NFWO, Belgium) and the Research Council of the University of Leuven (PDM-94/32). We thank Julie Savage and Laura Doherty for critical reading of the manuscript and Andreas De Troy for technical support. The work was also supported by the British Council and the NFWO, Belgium. Partial reports of the research have been presented at the Applied Vision Association Meeting on Invariance and Constancy in Vision (Reading, UK, April 1995) and the 18th European Conference on Visual Perception (Tübingen, Germany, August 1995); Heljä Kukkonen's attendance at the latter was supported by a Wellcome Trust Travel Grant.

References

- Biederman I, 1987 "Recognition by components: a theory of human image understanding" *Psychological Review* **94** 115–147
- Biederman I, Gerhardstein P C, 1993 "Recognizing depth-rotated objects: evidence and conditions for three-dimensional viewpoint invariance" *Journal of Experimental Psychology: Human Perception and Performance* **19** 1162–1182
- Cutting J E, 1986 *Perception with an Eye for Motion* (Cambridge, MA: MIT Press)
- Foster D H, 1977 "Visual pattern recognition by assignment of invariant features and feature relations" *Optica Acta* **24** 147–157
- Foster D H, 1980a "A description of discrete internal representation schemes for visual pattern discrimination" *Biological Cybernetics* **38** 151–157

-
- Foster D H, 1980b "A spatial perturbation technique for the investigation of discrete internal representations of visual patterns" *Biological Cybernetics* **38** 159–169
- Foster D H, 1991 "Operating on spatial relations", in *Pattern Recognition by Man and Machine* Ed. R J Watt (Basingstoke, Hants: Macmillan Education) pp 50–68
- Gibson J J, 1950 *The Perception of the Visual World* (Cambridge, MA: Riverside Press)
- Kahn J I, 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" *Quarterly Journal of Experimental Psychology* **33A** 155–166
- Koenderink J J, 1986 "Optic flow" *Vision Research* **25** 161–180
- Nazir T A, O'Regan J K, 1990 "Some results on translation invariance in the human visual system" *Spatial Vision* **5** 81–100
- Pizlo Z, 1994 "A theory of shape constancy based on perspective invariants" *Vision Research* **34** 1637–1658
- Ullman S, 1989 "Aligning pictorial descriptions: An approach to object recognition" *Cognition* **32** 193–254
- Van Gool L, Moons T, Pauwels E, Wagemans J, 1994 "Invariance from the Euclidean geometer's perspective" *Perception* **23** 547–561
- Wagemans J, 1992 "Perceptual use of non-accidental properties" *Canadian Journal of Psychology* **46** 236–279
- Wagemans J, 1993 "Skewed symmetry: A non-accidental property used to perceive visual forms" *Journal of Experimental Psychology: Human Perception and Performance* **19** 364–380
- Wagemans J, Van Gool L, Lamote C, Foster D H, 1996 "Minimal information to determine affine shape equivalence" *Journal of Experimental Psychology: Human Perception and Performance* (under revision)

APPENDIX

Model

The data were fitted by a simple probabilistic model of discrimination performance that involved a combination of qualitative (discrete) and quantitative (continuous) cues (Foster 1980a, 1991). It was assumed that the visual system was able to extract from the two patterns estimates x_1, x_2 of the positions of the displaced points, defined by their distances along a diagonal from some reference position, such as the opposite point. These distances were normalised with respect to the size of one of the patterns or some ensemble average. The extraction of such signals in turn assumes that a point-to-point correspondence can be set up between the patterns. There are several ways in which this correspondence could, in theory, be achieved. One way is to evaluate for each pattern a suitable affine cue, such as the four-point cross ratio, and define the 'best' correspondence as that which minimises the differences between the values of this cue for the two patterns. Another way is to find an affine transformation that maps three points of the first pattern onto the corresponding three points of the second, and define the 'best' correspondence as that which minimises the distance between the pair of points, one from each pattern, that are not brought into coincidence by such a transformation. Because of the uncertainty associated with estimates of the positions of the pattern points, there will always be at least one pair of points from the two patterns (whether they are 'same' or 'different' patterns) that can be labelled as possibly having undergone a perturbation. The difference $x_2 - x_1$ provides a continuous cue for the task.

In addition to this continuous cue, it was assumed that a discrete cue was also generated, specifying parallelism, collinearity, concavity, or convexity. Again, there are several ways in which this cue could, in theory, be generated, but, for convenience and for compatibility with the analysis of the continuous cue, it was assumed that a decision as to whether a discrete cue was generated could be expressed in terms of the values of the position estimates x_1, x_2 . It was not assumed that the discrete cue was *necessarily* generated in this way.

Since x_1, x_2 are uncertain, they were treated as random variables. For two 'same' patterns, x_1, x_2 were both drawn from a normal distribution $N(x, \sigma^2)$, having standard deviation σ and mean corresponding to the true position x of the displaced points (identical for the two patterns). For two 'different' patterns, the signals x_1, x_2 were drawn from two different distributions $N(x - 1, \sigma^2), N(x + 1, \sigma^2)$, having the same standard deviations σ but different means corresponding to the different true positions $x - 1, x + 1$ of the displaced points (± 1 step of 15 min arc away from the mean x). Having been extracted, x_1, x_2 were assumed to be processed in precisely the same way, independently of whether they originated from a 'same' or 'different' pair.

The processing steps were, then, as follows. Assume, without loss in generality, that $x_1 < x_2$.

(1) Define a continuous 'difference' cue c as the difference signal $x_2 - x_1$ weighted by the mean distance $\bar{x} = (x_1 + x_2)/2$ from the opposite point according to a (small) constant Weber-like coefficient $w \geq 0$, thus

$$c = \frac{(x_2 - x_1)}{(1 + w\bar{x})},$$

where w, \bar{x} are always such that $w\bar{x} > -1$. This formulation, modelled after normal Weber-like interval-discrimination behaviour, allows for the possibility that (if $w > 0$) the effect of the difference cue would be less when the displaced points were far from the opposite point (large \bar{x}) than when they were near to the opposite point (small \bar{x}).

(2) Define a discrete 'difference' cue d as a logical quantity whose value depends on whether x_1, x_2 fall within or without a range of values $[k_1, k_2]$, demarcating (symmetrically about the ideal position) a property such as collinearity or parallelism, thus

$$d = \begin{cases} d_0, & \text{if } x_1 < k_1 \text{ and } k_1 \leq x_2 \leq k_2; \\ d_0, & \text{if } k_1 \leq x_1 \leq k_2 \text{ and } k_2 < x_2; \\ -d_0, & \text{if } k_1 \leq x_1 \leq k_2 \text{ and } k_1 \leq x_2 \leq k_2; \\ 0, & \text{otherwise;} \end{cases}$$

where d_0 is a positive constant. The first two conditions correspond to one pattern being classified as parallel (or collinear) and the other not, thus yielding a positive cue; and the third condition corresponds to both patterns being classified as parallel (or collinear), thus yielding a negative cue. For the patterns used to test the collinearity cue, it was assumed that a positive 'parity' cue would also be generated in the following condition

$$d = d_0, \quad \text{if } x_1 < k_1 \text{ and } k_2 < x_2;$$

corresponding to one pattern being classified as convex and the other as concave.

(3) The continuous and discrete cues were added together (the latter with adjustable strength), along with an internal noise signal s (independent of the cues, and depending, for example, on cognitive factors); the noise signal s was drawn from a normal distribution $N(0, s_0^2)$ with standard deviation s_0 .

(4) In each trial, a 'different' or 'same' response was generated depending on whether the combined signal $c + d + s$ did or did not reach some threshold.

In all, there were six free parameters whose values were systematically varied in a coarse grid search so that the model provided an approximate maximum-likelihood fit simultaneously to the 30 points comprising the data for 'same' and 'different' patterns in each experimental condition. The response of the model was evaluated by numerical simulation based on 10 000 iterations for each of the thirty points in the 4 experimental conditions determined by the dot and line patterns and parallelism and collinearity cues. The approximate maximum-likelihood fits are shown in figures 3a–3d and 4 by the continuous and dotted lines, which appear to capture the trends in the data. In the light of the number of degrees of freedom of the model and its simplicity, the fits are reasonable. Values of the parameter range $k_2 - k_1$ were between 1.5 and 1.7 steps of 15 min arc for collinearity and between 3.2 and 3.8 steps of 15 min arc for parallelism. The weighting d_0 of discrete to continuous cues was between 1.9 and 2.3 for collinearity and between 1.6 and 1.9 for parallelism.