Vision Res. Vol. 33, No. 8, pp. 1067–1088, 1993
Printed in Great Britain.

0042-6989/93 $6.00 + 0.00

Higher-Order Structure in Regularity Detection

JOHAN WAGEMANS,* LUC VAN GOOL,† VALÈRE SWINNEN,† JOHAN VAN HOREBEEK‡

Received 24 February 1992; in revised form 22 October 1992

In three experiments a simple Euclidean transformation (reflection, translation, rotation) was applied to collections of twelve dots in such a way that they contained equal lower-order structure, defined on the pairwise grouping of elements with their partner following transformation (e.g. parallel virtual lines), but differed in the presence vs absence of higher-order structure, defined on pairs of pairwise groupings (e.g. virtual quadrangles with correlated angles). Based on the much better performance levels (*d'*) in the case of additional higher-order structure, we conclude that global regularities are easier to detect when the local correspondences are supported by higher-order ones formed between them. These enable the lower-order groupings to spread out across the whole pattern very rapidly (called bootstrapping). As a preliminary attempt to specify these principles, we proposed a working model with two basic components: first, a function expressing the cost of a perceptual grouping or the lack of regularity, and, secondly, an algorithm based on simulated annealing to minimize the cost function. The simulation results obtained with our current implementation of these principles showed satisfactory qualitative agreement with human regularity detection performance. Finally, the theory was shown to capture the essence of a large number of grouping phenomena taken from diverse domains such as detection of symmetry in dot patterns, global structure in Glass and vector patterns, correspondence in stereoscopic transparency and apparent motion. Therefore, we are convinced that, in principle, the mechanism used by the human visual system to detect regularity incorporates something like bootstrapping based on higher-order structure. We regard this as a promising step towards unraveling the intriguing mechanisms of classic Gestalt phenomena.

Symmetry Dot patterns Grouping Local-versus-global Euclidean transformations

The detection of regularity in visual patterns is one of those achievements human perceivers are so good at, that we do not realize how difficult the task is. For example, the bilateral symmetry created by reflecting twelve dots around a vertical axis, is immediately evident from a single glance inspection of Fig. 1(A, top panel). Psychophysical studies confirm this efficiency of the symmetry detection process (e.g. Barlow & Reeves, 1979; Jenkins, 1983a; Wagemans, Van Gool & d'Ydewalle, 1991). In a similar vein, when a random-dot pattern is transformed according to some simple rule (e.g. a translation, a sinusoidal function) and copied onto itself, the resulting Glass pattern is effortlessly perceived as being regular (Glass, 1969; Glass & Perez, 1973; Prazdny, 1984, 1986; Stevens & Brookes, 1987).

The problem with this kind of regularity is that it is inherently global yet the process has to start off somewhere locally. Mathematically, the global regularity is defined as the sum of all local correspondences, but psychologically, the mechanism by which this regularity is detected seems too fast and efficient to follow a similar route. One argument for global processing of regularity in Glass patterns is that occlusion of all but a small area leaves only an impression of random dots. Interactions between different local areas seem essential to obtain the perceptual impression of global regularity. On the other hand, small perturbations of perfect regularities such as mirror symmetry can be detected as well (Bruce & Morgan, 1975; Locher & Wagemans, 1993). As such, this problem of regularity detection is a typical instance of the classic local-vs-global dilemma which has puzzled vision research for centuries (e.g. Mach, 1959).

Previous attempts to explain regularity detection are based on what we will call *lower-order structure*, because it is defined on the pairwise grouping of elements with their partner following transformation. For example, Glass (1969) and Glass and Perez (1973) referred to the cooperativity of orientation-tuned cells in the visual cortex (see Hubel & Wiesel, 1977, for an overview) as if the local correspondences were perceptually equivalent to oriented line segments. It is interesting to note that this is the area where the notion of a *virtual line* was introduced first (by Stevens, 1978), although equivalent concepts such as *chords* (Moore & Parker, 1974) and *dipoles* (Julesz, 1971) had been hanging around for a while. Stevens' algorithm for the detection of regularity in Glass patterns operated on the basis of the relative frequencies of differently oriented virtual lines. With respect to our second example, Jenkins (1983a) noted

*To whom all correspondence should be addressed at: Laboratory of Experimental Psychology, University of Leuven, Tiensestraat 102, B-3000 Leuven, Belgium.
†Department of Electronics, ESAT-MI2, University of Leuven, Belgium.
‡Department of Mathematics, University of Leuven, Belgium.

that bilateral or mirror symmetry is defined by two characteristics: virtual lines connecting symmetrically positioned elements are oriented uniformly and their midpoints are collinear. In addition, he showed that the human visual system is very sensitive to these lower-order structures, which can, therefore, be supposed to play a role in the detection of bilateral symmetry.

Nevertheless, we are not convinced that this lower-order structure is sufficient to explain the available data. We suppose that *higher-order structure*, defined on pairs of pairwise groupings, plays an additional but essential role. Our suggestion is that global regularities are easier to detect when the local correspondences are supported by strong higher-order structures. As a complement to the concept of a virtual line, representing the grouping between two individual elements (Stevens, 1978), we, therefore, introduce the concept of a *virtual quadrangle*, representing the grouping between two virtual lines. When the angles in these virtual quadrangles are correlated pairwise, as in symmetric trapezoids or parallelograms, we call them *correlation quadrangles.* To put it somewhat more formally, assume two virtual lines, say

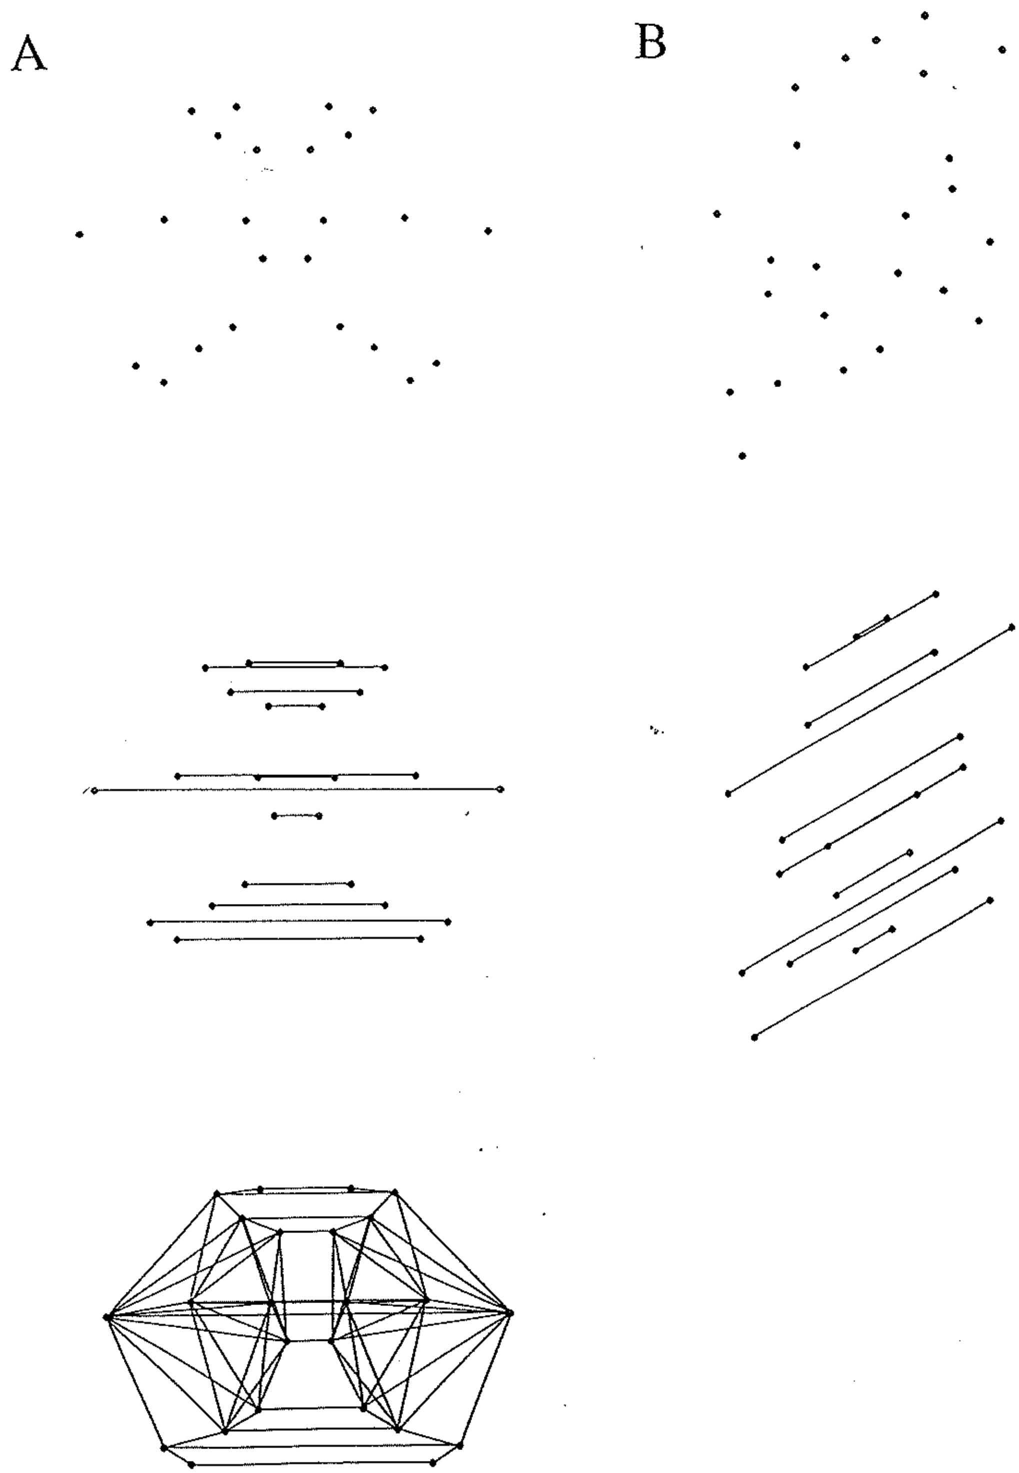

FIGURE 1. A dot pattern with (A) bilateral symmetry and (B) skewed symmetry, together with their lower-order structure (i.e. virtual line parallelism), and, for the bilateral case only, higher-order structure (i.e. correlation quadrangles).

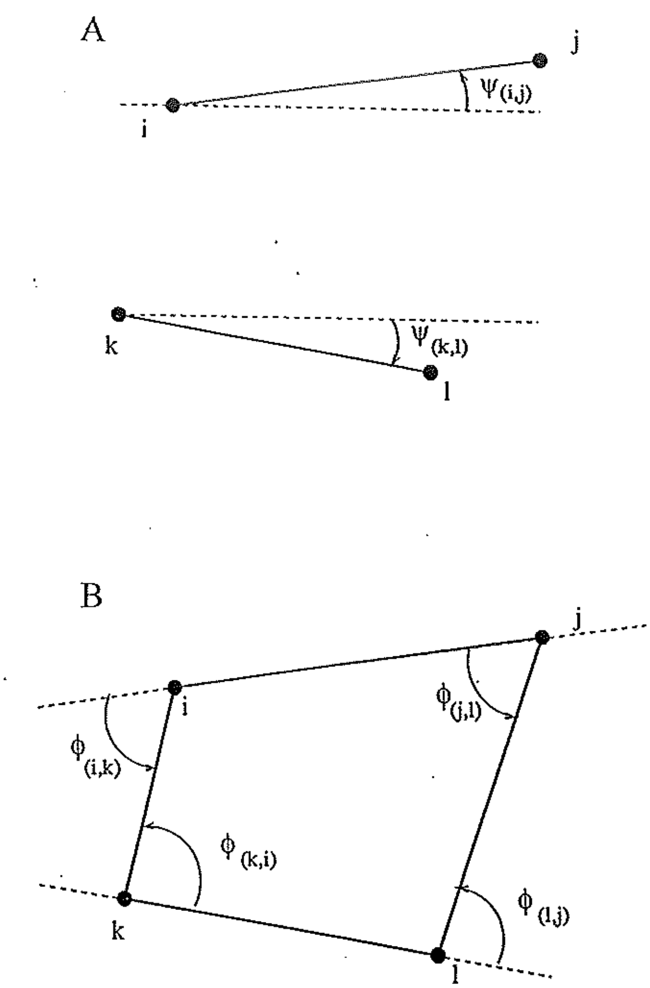

FIGURE 2. Our use of *lower-order* (A) and *higher-order* (B) structure and the way it is measured in the cost function.

(i, j) and (k, l), lying in each other's neighbourhood (see Fig. 2). The lower-order information [Fig. 2(A)] corresponds to relations between orientations of individual virtual lines (e.g. orientation differences $|\psi_{(i,j)} - \psi_{(k,l)}|$), whereas the higher-order information [Fig. 2(B)] corresponds to relations between pairs of virtual lines. The latter cannot be defined unless a virtual line *pair* has been selected (hence, higher order). The angles ϕ are defined relative to the angles ψ and always correspond to the smallest angle enclosed between the virtual lines. The higher-order structure is strong when, for example, the expression $\Sigma e^{|\phi_{(i,k)} - \phi_{(j,l)}| + |\phi_{(k,i)} - \phi_{(l,j)}|}$ tends to zero (i.e. in the case of virtual symmetric trapezoids or parallelograms). Another formalization of this criterion is contained in the Appendix.

Our suspicion about the role of higher-order structure is supported by the following observations. First, most recent research with rotational Glass patterns (e.g. Stevens, 1978; Stevens & Brookes, 1987; Prazdny, 1984, 1986) used homogeneous displacement instead of the more natural differential displacement. Instead of making a dot in the periphery travel along a larger distance than a central dot, as would be the case in a pure rotational Glass pattern, one lets them be separated by a constant distance throughout the whole pattern. Although it is not specified why this is so, we think the regularity is more salient as a result of it. By using homogeneous displacement, one adds higher-order structure to the lower-order one, because two pairwise groupings together constitute a regular virtual quadrangle (i.e. parallelogram). This seems to suggest that local pairings give more rapidly rise to a salient global regularity if they are incorporated in stronger higher-order structures.

Secondly, in a series of experiments with skewed symmetry, which results from viewing bilateral symmetry from oblique angles, substantial processing differences from bilateral symmetry were found. In fact, preattentive detection of symmetry was largely disrupted (see Wagemans, Van Gool & d'Ydewalle, 1990, 1992). Nevertheless, Jenkins' (1983a) lower-order properties are still present in skewed symmetry: the virtual lines connecting the corresponding elements are not orthogonal to the axis but still have orientational uniformity and midpoint collinearity [see Fig. 1(B)]. So, it appears that higher-order structure, defined in terms of pairs of virtual lines, which is present in bilateral symmetry (i.e. symmetric trapezoids) but destroyed by skewing [compare (A) with (B) in Fig. 1], is needed to account for this finding.

In a subsequent experiment, using a paradigm that was better suited to detect preattentive processes (d' in a random/regular discrimination task with tachistoscopic presentation), we explicitly manipulated the number and regularity of the higher-order structural elements by skewing single, double, and quadruple symmetry (i.e. symmetry about one, two, or four axes; see Wagemans *et al.*, 1991). As expected, the effect of skewing was very disruptive for single symmetry, but gradually less so for double and quadruple symmetry (see Fig. 3 for a summary across axes orientations and skewing direction). The reason appears to be that the higher-order groupings are not fully destroyed in the latter two cases (e.g. virtual parallelograms are present).

We think this higher-order structure lies at the basis of the way our visual system bridges the gap between local correspondence and global regularity. Consider the following scenario for the detection of symmetry in a dot pattern. Suppose that initially all potential groupings within a particular neighbourhood of a dot are made. From the moment that some parallelism is detected, a

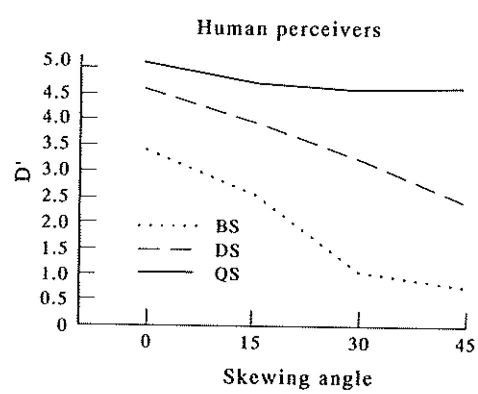

FIGURE 3. Summary of results obtained in a previous study (see Wagemans *et al.*, 1991), showing large effects of skewing on detection of single bilateral symmetry, less for double symmetry, and almost no effect for quadruple symmetry.

local reference frame can be established to it (one axis of the frame parallel to it, the other orthogonal to it), with respect to which all other angles are expressed. The propagation of this local reference frame, which we call *bootstrapping*, is much easier when the lower-order structure is supported by higher-order structure, because one of the two axes of the frame indicates the direction in which to proceed. A more detailed discussion of these notions is provided below, when we discuss a possible implementation of these ideas, and in the Appendix where it is treated more formally.

Before doing so, however, we will present three experiments which were specifically designed to test the role of higher-order structure in regularity detection. In each of them, a particular transformation (an instance of the Euclidean group of reflections, translations, and rotations), will be applied to two types of patterns, one of which will yield the higher-order structure supposed to play a role in our account, whereas the other will be devoid of them. If the regularity were easier to detect in the former case, this would constitute good evidence against the sufficiency of lower-order structure, because this will always be present in both kinds of patterns. In the General Discussion we will consider alternative explanations for some particular results, but we will argue that it is better to have one general explanatory framework for different phenomena of perceptual grouping and regularity detection.

GENERAL METHOD

In all experiments reported here, the same general method was used. The only difference between the three experiments was the regularity that had to be detected by the subjects.

Subjects

Four subjects participated in all three experiments: the first author and three volunteering undergraduate students involved in other research projects at the Laboratory of Experimental Psychology. The latter were naive with respect to the purposes of the experiments. All observers had normal or corrected-to-normal vision. Owing to the very large number of trials and conditions (see Procedure), only a few, well-motivated subjects (willing to return several times) could be run. We expected that the processes under investigation are so basic as to allow no cognitive biases or large interindividual differences. The validity of this assumption can be tested by considering whether eventual differences between observers are really qualitative instead of merely quantitative (e.g. same rank ordering of detectabilities). Therefore, results for individual subjects will be presented as well. To control for sequential effects, each subject was assigned one of all six orders in which the three experiments could be performed.

Stimuli

All stimuli used in these experiments were patterns of 24 dots located in a more or less circular area defined around the centre of the screen with a diameter of 10 cm. Half of the dot patterns were completely random (within certain limits, see below), whereas only a restricted number of dots (e.g. six or twelve) was random in the other half of the patterns. The locations of the other dots constituting the latter type of patterns were defined by applying one of the three simple Euclidean transformations to the random part: reflection, translation, or rotation. The resulting symmetries are called *mirror symmetry* (MS), *translational symmetry* (TS), and *rotational symmetry* (RS), respectively. The notion of MS is used to include skewed symmetry as well as perfect bilateral symmetry. In each type, a particular set of parameter values specifying the transformation (e.g. reflection axis, translation distance, and rotation angle, respectively) was used to obtain sufficient variation. In addition, some specific pattern aspects (e.g. collinearity) were introduced to test the role of higher-order structure. More details about the stimuli resulting from these procedures are provided below for each experiment separately.

Apparatus

The dot patterns were generated by a flexible and portable C-program on a SUN-3 Workstation with a Motorola MC 68881 floating point board. Because of the limits of the UNIX-operating system, the experiments were automated on a different system, namely, an IBM-AT-compatible with an Intel 80386 processor and a VISTA card. The stimuli were presented as black dots against a homogeneous grey background on a raster display with high temporal and spatial resolution (BARCO, type CDCT-6351B) used in PAL-mode with 50 Hz temporal resolution and 750×578 spatial resolution, noninterlaced. The experimental room was completely darkened and screen borders were covered by black cardboard with a circular aperture to reduce orientational cues. The subject was seated at a distance of 114 cm with eyeheight at the centre of the screen. At that distance, the size of the individual dots and the whole patterns subtended 5.7 arc min and 5°, respectively.

Procedure

The experiments were designed as signal detection tasks. The "signal" to be detected was symmetry or, as we explained it to the subjects, regularity. Before starting the experiments, subjects received a rather extensive introduction about what was meant by regularity. Examples of each of the different symmetries were shown and time was given to explore the dot patterns sufficiently to detect the regularity in each case. They were then instructed that they would see a lot of those patterns, mixed with random-dot patterns and presented only very shortly. The task was described as a forced-choice on the question "is it a regular pattern or not?" which had to be answered to the best of their abilities. Two of the four subjects answered *Yes* (regular) with their right index finger and *No* (random) with their left index finger, whereas the others did the reverse.

A trial was constituted from the following sequence of events. For 500 msec a black fixation circle was presented in the centre of the screen followed by a collection of 24 black dots, presented for only 100 msec, constituting a random or regular pattern on a homogeneous grey background. Immediately following this dot pattern, a masking pattern was presented for 1500 msec. This was a similar black-on-grey collection of 36 random dots sufficiently large to cover the stimulus pattern completely. A set of ten different masking patterns was used, one of which was randomly selected on each trial. This was done to avoid the use of the same mask repeatedly on all trials with the risk that it would gradually loose its masking power. From the moment the mask was on, subjects could respond by pressing one of two buttons of a response panel connected to the PC configuration used for stimulus presentation. Each answer was instantaneously evaluated by the computer so that immediate feedback could be given. A correct answer was followed by a 300 msec high-frequency tone (750 Hz), a false response by a 500 msec low-frequency tone (100 Hz). The only reason why this feedback was provided was to keep motivation and arousal of the subjects at an optimal level. Although it may have facilitated some learning during the experiments, the order of the blocks was always randomized so that no systematic effects could have resulted.

Following the introductory session in which the concept of regularity was explained and demonstrated and in which the task was described, in each experiment, subjects received a practice series containing exemplars of random and regular dot patterns of each type to be presented in that experiment. In Expts 1 and 2, the number of practice trials was 168 (i.e. three of each type). In Expt 3, we administered 160 practice trials (i.e. ten of each type). The data for these practice stimuli were not analysed. Different patterns were used in the experimental sessions.

Trials were presented in blocks of 220. In the beginning of each block of 100 random and 100 regular experimental trials, randomly intermixed, ten random and ten regular trials were given as warm-up, also randomly intermixed. Again, the patterns presented on these trials were different from the experimental ones and the data obtained with them were not analysed. The type of regularity (MS, TS and RS) was constant within and different between experiments (1, 2, and 3, respectively). In each experiment, blocks differed in the parameter values defining the transformation (e.g. reflection axis, translation distance, and rotation angle, respectively). The order of trials within a block and the order of blocks were randomized for each subject separately. After each block, which took about 10 min, subjects could choose to continue or to quit, but the large number of trials to be run encouraged them to do more than a few blocks each time they were in the lab. Subjects were advised to take short breaks between the blocks. On the average, subjects did 15–20 sessions of 6–8 blocks, spread over a period of 1 or 2 weeks. The total duration of the three experiments (i.e. introductory session, three practice sessions, and 56 + 56 + 16 experimental blocks) was about 24 hr for each subject.

EXPERIMENT 1: MIRROR REFLECTIONS

Stimuli

In Expt 1, the regular-dot patterns were created by reflecting twelve dots about a straight axis oriented vertically (V), horizontally (H), or obliquely (left, L or right diagonal, R). The input dots were localized at random within the original circle except for a minimal interdot distance to create a pattern of more or less homogeneous density. To obtain a sufficiently representative set of patterns, skewing was introduced. In this way, orthographic projection onto the image plane of planar MS oriented arbitrarily in depth was simulated (see Friedberg, 1986; Kanade & Kender, 1983), instead of restricting the results to one particular viewing position (as would be the case without skewing). The skewing angle was manipulated at seven levels: 0° (i.e. the case of perfect bilateral MS), and three 15° steps, both clockwise and counterclockwise. After skewing, interdot distances were controlled again to avoid spurious local clusters. In addition, random patterns were also skewed to avoid that a cue would be created by the form of the zone in which the dots were located (i.e. elliptic in the case of skewed patterns). Combining skewing angle with axis orientation yielded 28 different types of MS.

In the case when no further constraints were placed on the locations of the random dots used as input patterns to reflect, the resulting patterns are referred to as *ordinary MS* [examples are given in Fig. 4(A)]. The data for this type of patterns were obtained in another experiment described previously (Wagemans *et al.*, 1991). In the experiment reported here, two kinds of additional regularities were introduced by putting more constraints on the locations of the random dots used to reflect.

In the case called *equidistant MS*, the following procedure was used (see Fig. 5). One dot was randomly selected (e.g. the one indicated by 1a). This dot was then reflected about one of the four axes (e.g. V in Fig. 5; the resulting dot is called 1a′). For each random dot, another dot was chosen pseudo-randomly on the virtual line connecting the first random dot (1a) with its symmetrically positioned partner (1a′). This dot (denoted by 1b in Fig. 5) was also reflected (1b′). After reflection then, four dots were always situated on the same virtual line in two symmetric pairs. In addition, for each set of four dots generated as just described, a second set of four dots was randomly selected on a line parallel to the first while keeping the distances between the two dots at the same side of the axis constant. For example, when 1c would be selected pseudo-randomly, 1c′, 1d, and 1d′ would be fixed because of the constraint on the interdot distance and because of the reflection operation. Together all four dots generated on the basis of one random dot (1a–1b–1c–1d) constitute a higher-order

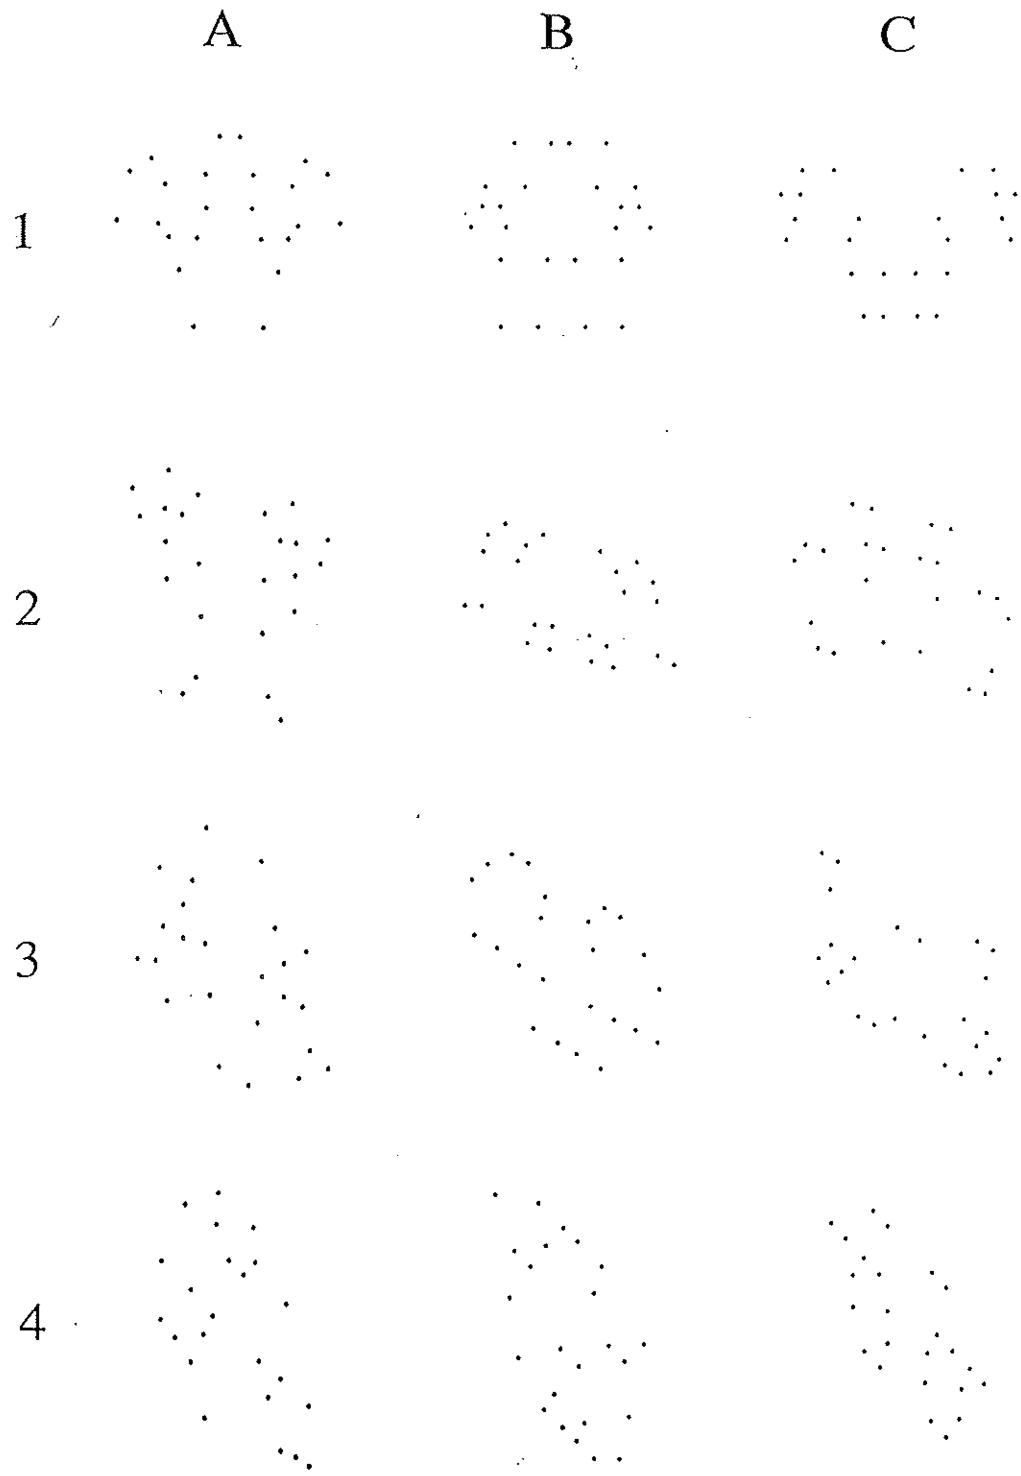

FIGURE 4. Examples of dot patterns with MS, as used in Expt 1. (A) Ordinary MS, (B) equidistant MS, and (C) collinear MS. In each case, a bilateral example is shown in (1), followed by three 15° steps of clockwise skewing in (2), (3), and (4), respectively.

structure of the parallelogram type. By definition the higher-order group of the four dots generated by reflection (1a′–1b′–1c′–1d′) form a parallelogram as well. The skewing operation destroys the symmetric trapezoids formed between two pairs of reflected dots (e.g. 1b–1b′–1d–1d′) but does not destroy these parallelograms. The procedure of generating eight dots from one randomly selected dot is repeated two times (one of which is shown in Fig. 5, viz. based on dot 2a). Examples of patterns with this type of equidistant MS are given in Fig. 4(B).

However, by introducing the higher-order structures as described above, one necessarily increases the lower-order structure as well, because the virtual lines connecting symmetric dots always contain four dots instead of two. To control for this increased first-order regularity, a case which is called *collinear MS* was tested. Here, the collinearity of four dots (in pairwise symmetric

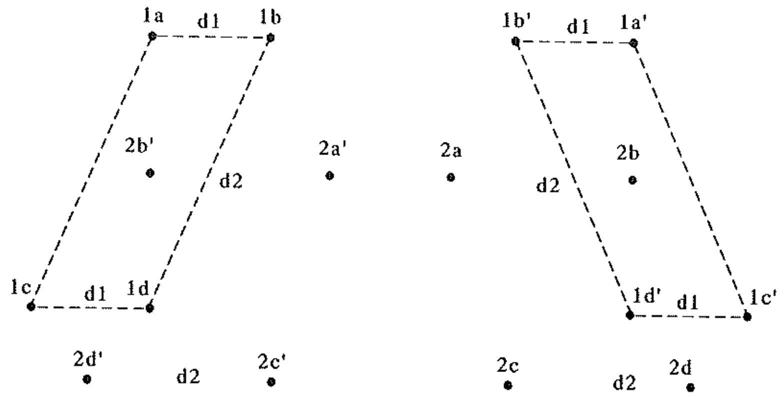

FIGURE 5. Explanation of the procedure used to create patterns with equidistant MS (see text).

positions) is present without creating higher-order regularities formed between two virtual lines at the same side of the axis. In terms of Fig. 5, the distance between 1a and 1b (denoted by dl) would not be repeated at the same side of the axis of symmetry. In order words, from each random dot only four instead of eight dots would be generated. Examples of patterns with this type of collinear MS are given in Fig. 4(C). For both equidistant and collinear MS, we used orientations and skewing angles as before, giving rise to 56 blocks of 200 trials each.

Predictions

In this way, by varying some specific aspects of the transformed patterns, the higher-order structure supposed to play an important role in regularity detection was explicitly manipulated. In all cases of MS, the lower-order structure is the same, namely, parallelisms of virtual lines with arbitrary lengths. In perfect bilateral MS (i.e. skewing angle $= 0°$), additional structure is present as a consequence of special relations between the orientations and lengths of the virtual lines connecting two parallel virtual lines (as mentioned before, these can be made explicit in virtual quadrangles, i.e. symmetric trapezoids). After skewing, this additional higher-order structure disappears. In previous experiments, it was shown that this makes the symmetry much harder to detect (Wagemans *et al.*, 1990, 1991, 1992).

In collinear MS tested here [see Fig. 4(C)], the lower-order structure is stronger because the virtual lines connecting symmetrically positioned dots can be represented more reliably (they are defined by four dots instead of two). Additional higher-order structure is introduced because all virtual lines are pairwise identical, no matter if the symmetry is skewed or not (but this additional regularity cannot be represented by quadrangles because it is created on the same virtual line). However, the virtual lines connecting two parallel virtual lines show higher-order structures (i.e. symmetric trapezoids) only in the case of zero skewing angles [compare 1 with 2–4 in Fig. 4(C)].

In equidistant MS, even more higher-order structure is present. By keeping the distances constant for two sets of four dots, not only the number of virtual lines with identical lengths is doubled, but additional parallelisms are introduced of virtual lines connecting two parallel virtual lines (giving rise to a second type of correlation quadrangles, i.e. parallelograms). Furthermore, these parallelograms are preserved even after skewing [compare 1 with 2–4 in Fig. 4(B)].

In short, we believe this manipulation of higher-order structure to be sufficiently specific to put its hypothesized role to a strong test. If higher-order structure contributes to the detection of regularity in dot patterns, the performance levels of Expt 1 with the MS patterns with special structures should be higher than with the ordinary ones tested by Wagemans *et al.* (1991). More specifically, the effect of skewing would be less drastic in the case of collinear and equidistant MS than in the case of ordinary MS. Furthermore, if a difference between the first two types would show up (i.e. if not masked by a ceiling effect), it should be in the direction of higher detectability for the equidistant ones.

Results

The data for the ordinary MS were taken from the experiment reported in Wagemans *et al.* (1991; single symmetry only) and pooled with the 56 conditions tested here. Together, they were subjected to an ANOVA with a $3 \times 4 \times 7$ design with four d's (one per subject) in the smallest cell: pattern type with three levels (ordinary, collinearity, and equidistance), orientation with four levels (V, H, L, R), and skewing with seven levels (0° and three $\pm 15°$ steps). The main effects of pattern type and skewing were highly reliable, $F(2,6) = 121.76$, $P < 0.0005$, and $F(6,18) = 85.97$, $P < 0.00001$, respectively. The main effect of orientation did not reach statistical significance, $F(3,9) = 1.48$, $P > 0.25$. Three of the four interactions were reliable at $P < 0.00001$: $F(12,36) = 7.67$, for pattern type × skewing; $F(18,54) = 8.58$, for orientation × skewing; and $F(36,108) = 3.64$, for pattern type × orientation × skewing. The remaining second-order interaction (i.e. pattern type × orientation) was not statistically significant, $F(6,18) = 1.61$, $P > 0.20$.

Because there was no systematic difference between clockwise and counter-clockwise skewing angles, averages obtained by pooling across them will be presented. In addition, the specifications of the reliable effects will be restricted to those that are relevant to the main purpose of the experiment (i.e. a test of the role of higher-order structure). Orientation was varied only to allow some sufficiently general conclusions to be drawn (orientational effects were the focus of the set of experiments in Wagemans *et al.*, 1992).

Average detectability scores (d's) were 1.70, 3.16, and 3.76 for ordinary, collinear, and equidistant MS, respectively. As indicated by Tukey's HSD-tests, the latter two differed from the former one at $P < 0.001$, whereas the difference between both special types was reliable at $P < 0.05$. As can be seen in Fig. 6, the pattern type × skewing interaction is mainly due to the fact that skewing had a much more drastic effect on the detectability of MS when there was no special grouping enhancing the lower- or higher-order structure. Moreover, it is interesting to note that the effect of pattern type was not reliable at zero skewing.

Discussion

In general, the results (see Fig. 6) clearly support the predictions. Human detection of regularity in dot patterns with MS is better in conditions which contain

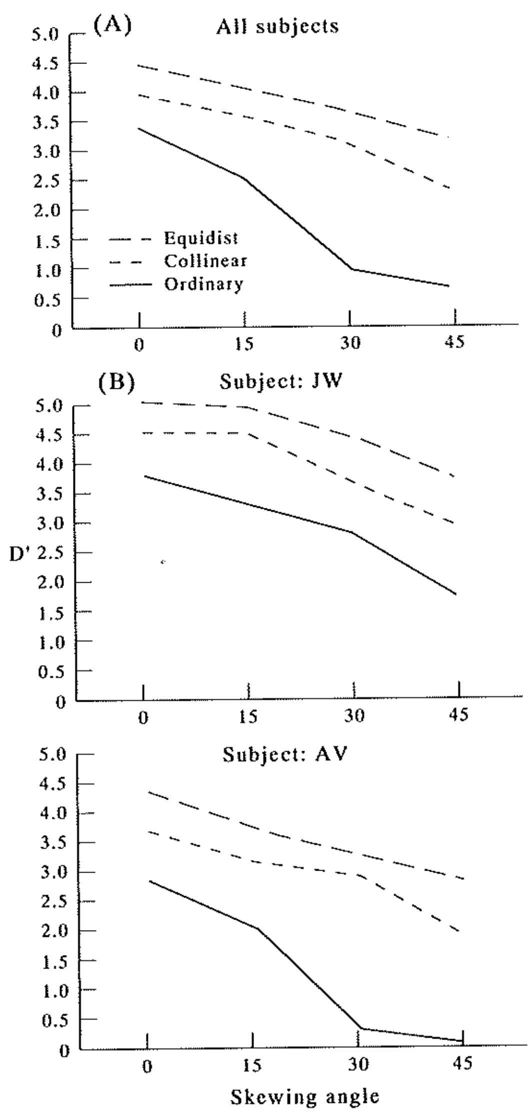

FIGURE 6. Results of Expt 1 with MS, averaged across subjects (A) and for two individual subjects (B).

higher-order structure. More specifically, skewed MS, which is normally hard to detect preattentively (see Fig. 6, solid line), is easy to detect when the dot patterns contain collinear or equidistant groupings (Fig. 6, dashed and dotted line, respectively). The higher detectability of equidistant MS as compared with collinear MS suggests that the improvement is not mainly due to the increased saliency of the lower-order parallelism. So, the higher-order structure formed between two parallel virtual lines of equal length (represented by virtual parallelograms) seems to play an additional role. As a consequence of the fact that skewing destroys the virtual symmetric trapezoids but preserves the virtual parallelograms, its effect is much smaller for equidistant MS.

It must be noted that the larger difference between collinear and ordinary MS, as compared with the one between equidistant and collinear MS, poses no serious problems for our account. First, we do *not* claim that higher-order structure (which is only present in conditions with zero skewing or with equidistant groupings) is the *only* factor contributing to detectability of a global regularity. In contrast, we propose to incorporate higher-order structure as an additional factor on top of lower-order one. Secondly, as suggested by the notion of bootstrapping, the propagation of a local frame of reference, which is possible on the basis of symmetric trapezoids but not to the same degree with parallelograms, is equally important as the mere presence of the correlation quadrangles themselves. This explanation is corroborated by the absence of reliable differences between grouping conditions when MS is perfect. In all these cases, the lower-order correspondences can spread out very easily because of the cooperative support by the higher-order structure (i.e. the parallel virtual lines are all incorporated in symmetric trapezoids) suggesting the global axis of MS.

EXPERIMENT 2: TRANSLATIONS

Stimuli

In Expt 2, the regular-dot patterns were created by translating twelve dots in a single direction along a particular distance which is only a small proportion of the pattern size, as it is done to create Glass patterns (Glass, 1969; Glass & Perez, 1973). This distance was either fixed or subject to some variability. The directions used were poles of the four orientations used in previous experiments (i.e. down, left, down-left, and down-right). The random patterns were located within the same zones on the screen to avoid position cues. The translation distance was varied within a sufficiently wide range (four 10-pixels steps starting from 15). The basic manipulation with respect to our predictions (see below) was to keep the distance fixed or not. In case when it was varied, it was done by a Gaussian distribution with, as the mean, the same value as the unvaried distance, and, as the standard deviation, all multiples of 10 smaller than the average distance. As a result, some of the obtained interdot distances were quite small. In contrast with the stimuli of Expt 1, these were not avoided here. The combination of these three variables (i.e. direction, distance, and variation) yielded 56 conditions of regular dot patterns with translation as the underlying regularity. Examples of the stimuli are provided in Fig. 7.

Predictions

Previous research with TS was aimed at finding a spatial limit to the detection of structure in those patterns (Jenkins, 1983b). He distinguished between three possible percepts which are clearly determined by the translation distance. With small translation distances, patterns with TS give rise to a strong impression of *pairedness* which is present throughout the display. Increasing the separation makes it much harder to find the point-pairs, but leaves the global impression of *striation* intact, until a particular critical distance is reached beyond which only *randomness* is seen. A series of experiments was devoted to finding the threshold between the latter two percepts and testing how it is effected by variables such as size of the field, density, etc.

Jenkins (1983b) interpreted his results as support for the explanation offered by Glass (1969) and Glass and Perez (1973) of how the visual system detects the

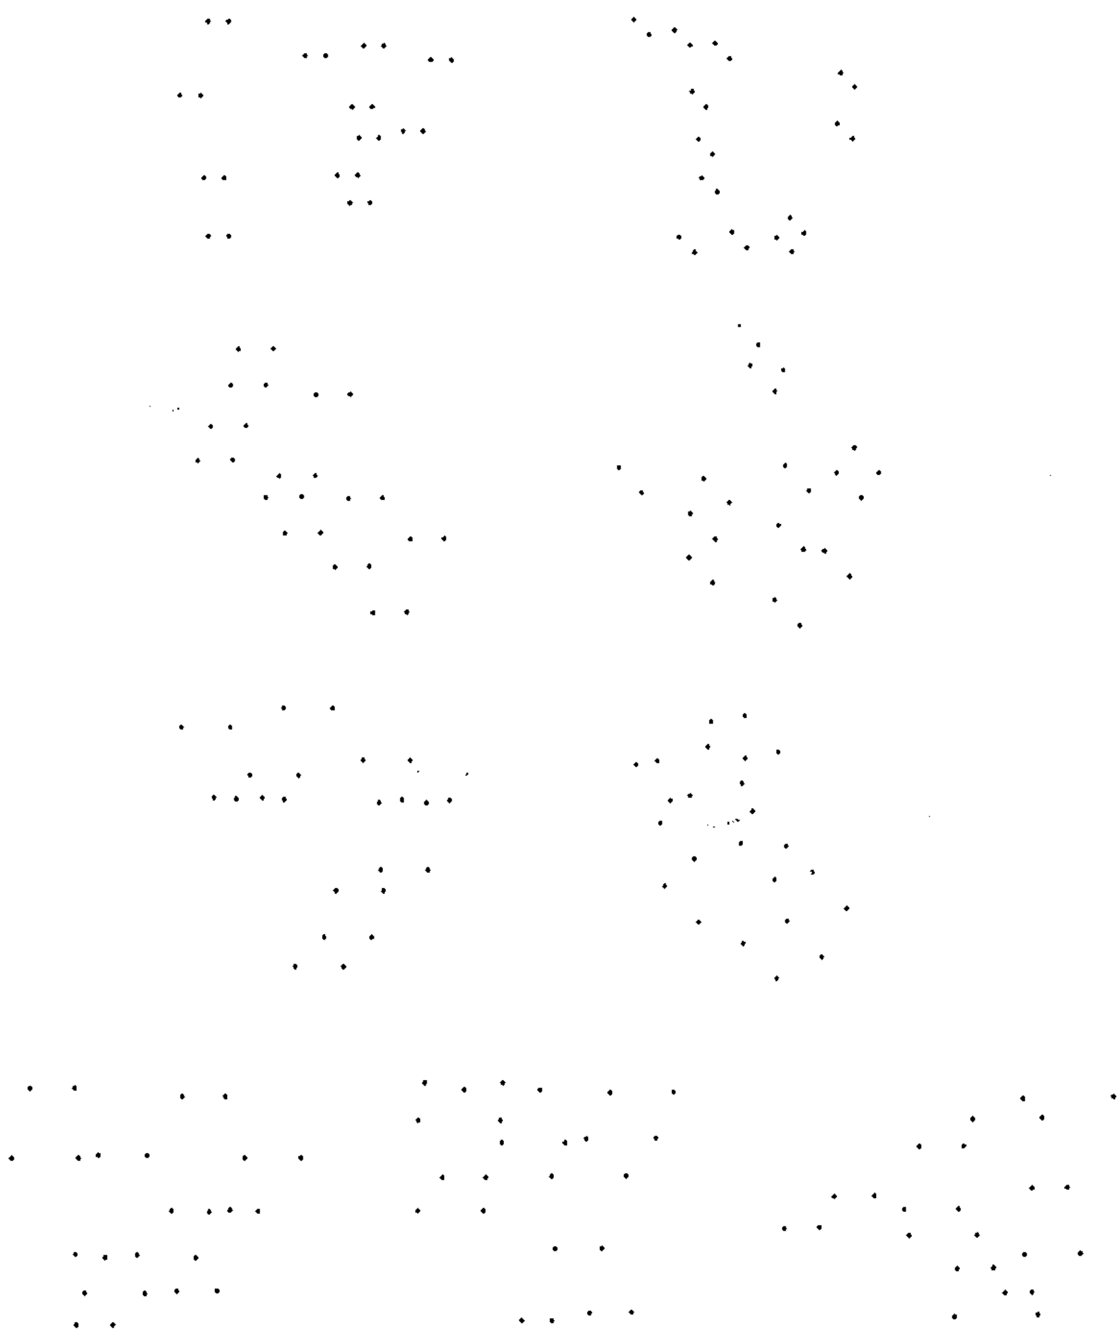

FIGURE 7. Examples of dot patterns with TS, as used in Expt 2. The six examples in the upper half are TS patterns with fixed translation distances of 15, 25, and 35, from top to bottom, respectively. The direction is H-left and down-R in the left and right column, respectively. The three examples in the lower half are TS patterns with variable translation distances. The mean distance is always 35 and the direction is H-left, but the variability increases in 10-pixels steps from 10 to 30, from left to right, respectively.

correlations in TS patterns. Coactivation of orientation-selective cells in the visual cortex (Hubel & Wiesel, 1977) is less when several different orientations are present in the neighbourhood of the correct point-pairs. Naturally, this is to be expected when the translation distance increases, because neighbour dots which are not partner dots become quite likely then. Experiments indicated that the perception of pairedness, which is locally defined, becomes more easily disturbed by closer non-partner dots [Stevens (1978) regards 2 or 3 as a limit] than the perception of striation, which is defined more globally (e.g. Glass & Switkes, 1976; Jenkins, 1983b). Similar data on the disturbing effects of different orientations in a close neighbourhood were obtained by adding local rotations on the point-pairs (Glass & Switkes, 1976).

Variability of the distances was introduced because it created interesting material to test the proposed role of higher-order structure. When the translation direction and distance are constant, the dot pattern is full of virtual parallelograms (see Fig. 7). However, these become gradually more disturbed by adding noise to the translation distance. If lower-order structure (as supposed by all previous models) would be the only determinant of regularity detection, one would not expect a difference between the conditions with fixed and variable translation distance. In fact, several arguments could be given to predict better performance with larger variation. For example, having a large variation gives quite often rise to very small interdot distances, which are known to enhance the grouping considerably (e.g. Jenkins, 1982; Smits & Vos, 1986; Stevens, 1978). The fact that there are also large interdot distances does not matter much for these models. They need some point-pairs to start off with and it is these initiating pairs that are more salient when interdot distances are small. Also, Lowe's incorporation of density as a factor to weigh groupings (see Lowe & Binford, 1982), implies that the perceptual groupings would be less salient when the densities are most homogeneous (i.e. when the interdot distances are fixed across the whole dot pattern). The role of higher-order structure as an additional factor determining the ease of regularity detection clearly leads to the contrary prediction, namely that increasing variability of the translation distances causes decreasing detectability of the regularity.

Results

Because more levels of variation were introduced on the larger distances, both variables relevant to the model (which were, therefore, not manipulated orthogonally) were combined to give 14 distance-variation levels. The 56 conditions resulting from using four directions were subjected to a 14 × 4 ANOVA, again with four d's (one per subject) in the smallest cell. The main effects of distance-variation and direction were both highly reliable (i.e. $P < 0.00001$), $F(13,39) = 82.87$, and $F(3,9) = 47.16$, respectively. The interaction between both factors also reached statistical significance, $F(39,117) = 1.98$, $P < 0.005$.

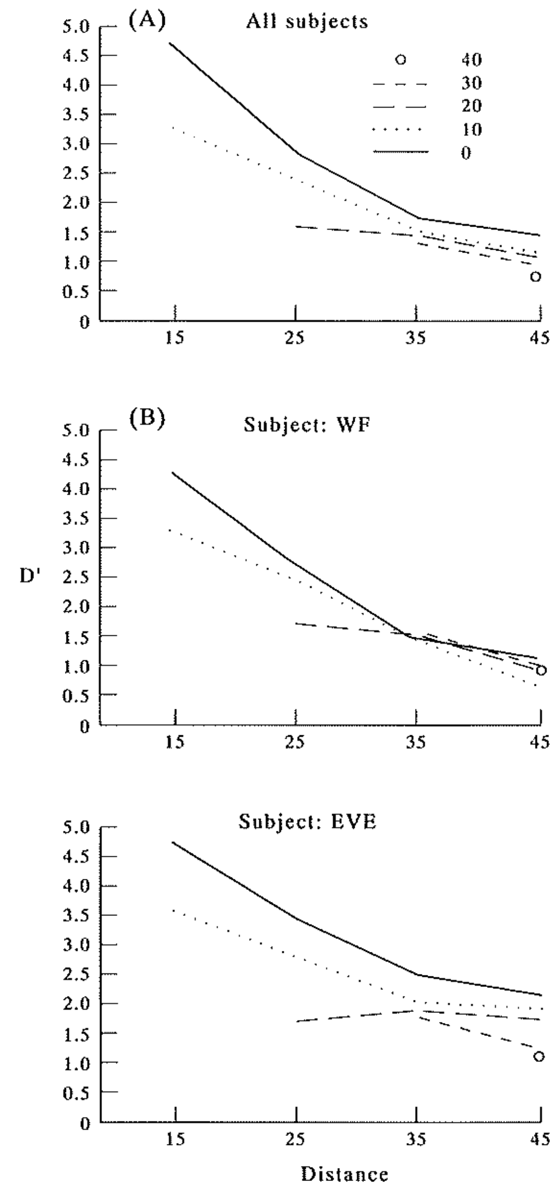

FIGURE 8. Results of Expt 2 with TS (averaged across directions), averaged across subjects (A) and for two individual subjects (B).

The main effect of distance-variation is quite straightforward. Averaged across variation, larger distances produced lower TS saliency (d' are 4.02, 2.31, 1.08, and 1.06, respectively); averaged across distance, larger variation produced a similar decrease (2.71, 2.08, 1.39, 1.12, and 0.78, respectively). Also, for each translation distance, the detectability of TS was higher if the variation was smaller or even absent (see Fig. 8). Tukey's HSD-tests indicated that these variation effects were more pronounced for the smaller translation distances (e.g. no reliable differences for distance = 35 or 45).

The main effect of direction is completely due to the fact that H translation (d' is 2.39) yielded higher detectability performances than all other directions (d's are 1.78, 1.68, and 1.66), which did not differ from each other. The two-way interaction is represented graphically by plotting the distance-variation effects for the four directions separately in diagrams (A–D) of Fig. 9. Two peculiar results were obtained that specify this

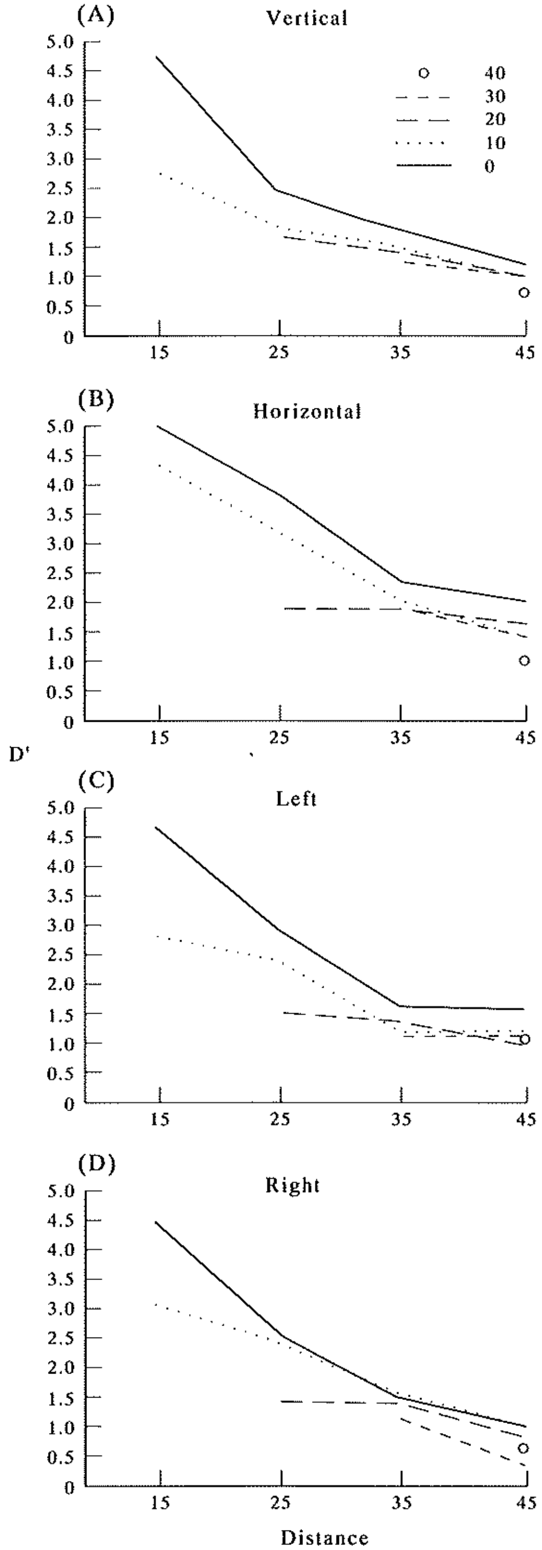

FIGURE 9. Results of Expt 2 with TS (averaged across subjects), for the four directions (A = V, B = H, C = L, and D = R).

interaction. First, at a small distance (i.e. 15), H translations were less effected by a small variation (i.e. 10) than translations in other directions. Secondly, the difference between variations of 10 and 20 pixels as standard deviations of the Gaussian distribution with a mean of 25, which were significant for the three other directions, disappeared for V translations.

Discussion

The results on the effects of introducing variability of the translation distances are clearly in agreement with the predictions derived from the additional role of higher-order structure, not with those derived from models restricted to lower-order structure. More specifically, the detectability of the regularity decreases as the translation distance and the variation of it increase (see Fig. 7). Although some lower-order groupings would definitely be stronger in the latter case, the fact that the higher-order structures are disrupted by it, seems to influence the mechanism of regularity detection more.

With respect to the effects of translation directions, it is interesting to note that classic anisotropy effects (e.g. Appelle, 1972; Essock, 1980) were only partially replicated. Results were better for H translations, not for V ones. As indicated by Wagemans *et al.* (1992), similar findings (H advantage) have been reported previously in the literature (e.g. Jenkins, 1983a, Expt 4; Pashler, 1990, Expt 4), but they were not taken seriously in light of the very often quoted V advantage (e.g. in the context of MS: Barlow & Reeves, 1979; Fisher & Fracasso, 1987; Palmer & Hemenway, 1978; Rock & Leaman, 1963; Royer 1981; in the context of TS: Jenkins, 1985). Our previous experiments with skewed symmetry (see Wagemans *et al.*, 1992) clarified that this ambiguity is due to the confounding between axis- and virtual-lines orientation in bilateral symmetry. In addition, it was shown that the latter factor yields much more pronounced effects on the detection of MS (bilateral or skewed).

This would be consonant with our suggestion that the detection of global regularity starts off with local virtual line parallelism supported by additional higher-order structure. From that perspective, it is not surprising that H translations are generally more salient than all others, and that detection of H TS is, therefore, more robust to small variations on the translation distances. Such an anisotropy of a correspondence-sensitive mechanism is to be expected for a biological vision system which has evolved in a natural environment in which, for example, stereo disparity occurs in almost H planes and motions often have a dominant H component (see below).

EXPERIMENT 3: ROTATIONS

Stimuli

In Expt 3, the regular-dot patterns were created by rotating twelve dots in a single direction along a particular angle. Eight different rotation angles were used: 10, 15, 20, 25, 30, 45, 90 and 180°. The last has been tested by Foster *et al.* who referred to it as *point inversion* (e.g. Bischof, Foster & Kahn, 1985; Foster, 1978; Foster & Kahn, 1985; Kahn & Foster, 1986). The small rotation angles would give rise to Glass patterns (e.g. Glass, 1969; Glass & Perez, 1973), if a sufficient number of dots would be used. To remain comparable with the other two experiments, the relatively small number of dots was retained.

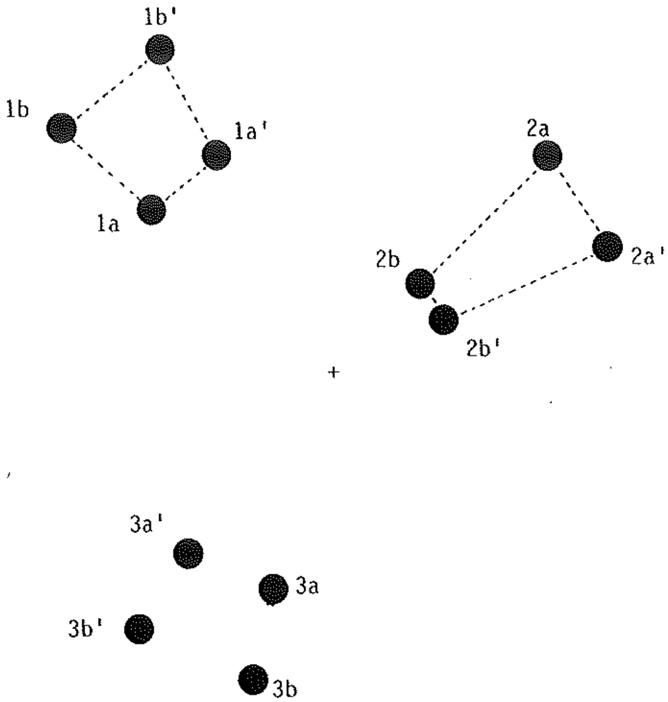

FIGURE 10. Explanation of the procedure used to create patterns with collinear RS (see text).

The critical manipulation with respect to higher-order structure was the randomness of the input patterns (i.e. the twelve dots to be rotated). In half of the cases, called *ordinary RS*, a completely random collection of twelve dots (i.e. even without constraints on the interdot distances) was used to start from. In the other half of the cases, only six dots were located randomly, and for each random dot, a second dot was selected randomly on the line going through the random point and the centre of the rotation. More details are given in Fig. 10, where the rotation angle is 20°. Starting from a randomly selected dot (e.g. 1a), one chooses another one (e.g. 1b) pseudo-randomly on the virtual line going through 1a and the centre of rotation (indicated by a plus sign) and rotates both through 20° (giving 1a′ and 1b′, respectively). Together these four dots constitute a higher-order structure (i.e. a symmetric trapezoid). This procedure is then repeated five times (two of which are shown in Fig. 10) to obtain a pattern of 24 dots with so-called *collinear RS*. As a side effect of this manipulation, this second type of RS patterns contained virtual lines that are radial to the rotation midpoint. To control for this, collinear random dot patterns were used as alternative stimuli in these blocks of trials. The total number of conditions in Expt

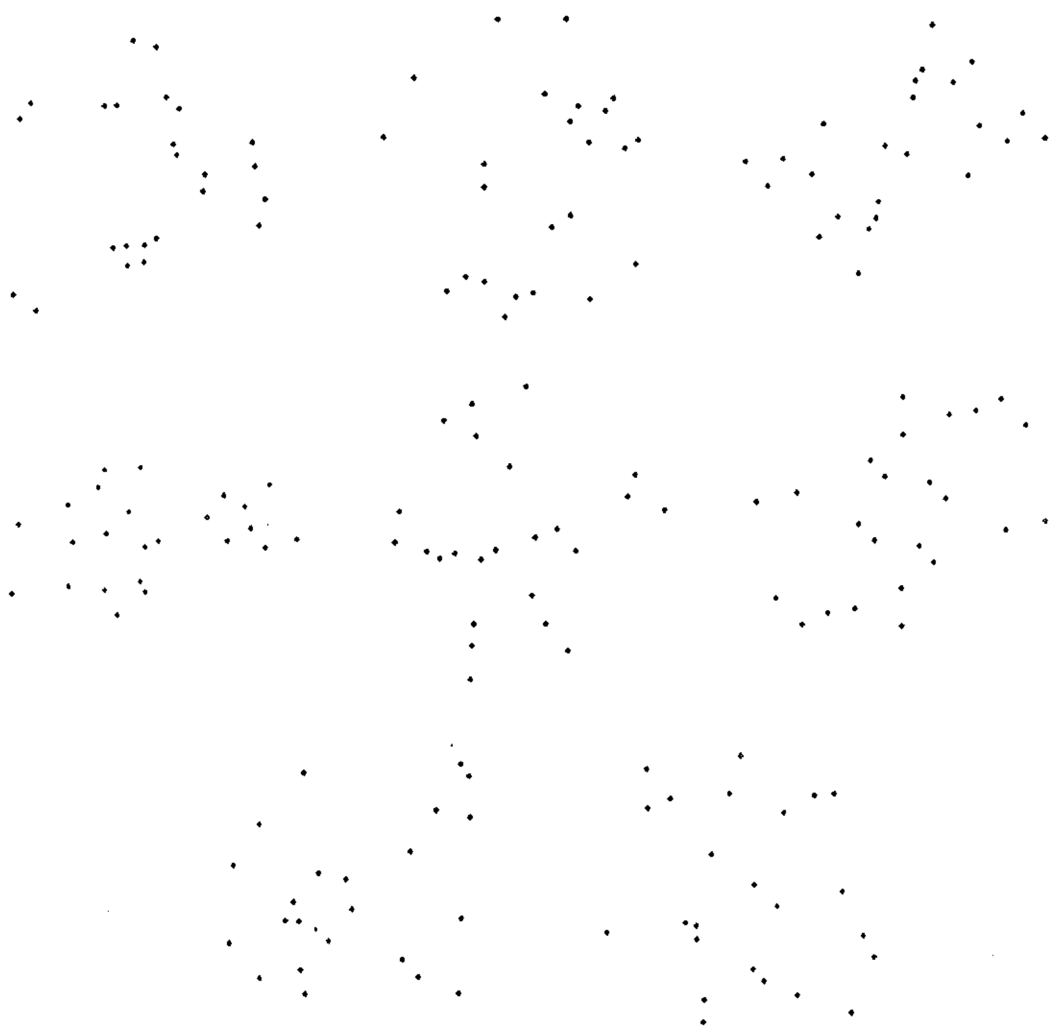

FIGURE 11. Examples of dot patterns with RS, as used in Expt 3. The six examples in the upper half are RS patterns. The first row are ordinary ones, the second are collinear. Rotation angle increases from 15 to 45 to 180°, from left to right, respectively. The two examples in the lower half are random-dot patterns. The right one is collinear, the left is an ordinary one.

3 was 16 (i.e. 2 × 8). Examples of the stimuli are shown in Fig. 11.

Predictions

In general, it is to be expected that this type of regularity is hard to detect unless the rotation angle is small, in which case virtual parallelograms are present, or additional regularity such as collinearity is included, in which case virtual symmetric trapezoids are present (see Fig. 11). A purely transformational approach would not predict that two types of patterns resulting from rotation could differ largely in the saliency of the regularity caused by it. As indicated above, however, the fact that experimenters started to use patterns with homogeneous displacement (e.g. Prazdny, 1984, 1986; Stevens, 1978), instead of applying the same rotation angle globally, probably has some consequence on the saliency of the Moiré structure afforded by it. In addition to arriving at homogeneous patterns (which is these authors' purpose), one disturbs the pure character of the transformation (it becomes a mixture of translation and rotation) and unexpected side-effects are obtained. In terms of our reasoning, one creates virtual parallelograms by it which, if the account is valid, enhances the regularity.

Because quantitative results of these implicit manipulations have not been reported, higher-order structure was manipulated more explicitly to test its effect. As hinted at above, varying the rotation angle from small to large gradually destroys the resemblance of the virtual quadrangles to parallelograms [more formally, it increases the pairwise angular differences, cf. Fig. 2(B)]. If no further regularities such as collinearity would be introduced, the prediction would be a corresponding gradual decrease in detectability, except for the 180° rotation, which again allows for virtual parallelograms (always spanning the centre of rotation). Manipulation of collinearity should have a generally positive effect on the regularity's saliency, because of the additional correlation quadrangles (i.e. symmetric trapezoids) afforded by it.

Results

All 16 d's obtained for each of the four subjects were subjected to an ANOVA with two factors: pattern type with two levels (ordinary and collinearity) and rotation angle with eight levels (four 5° increments from 10 to 30°, and three very large angles, i.e. 45, 90, and 180°). All effects were reliable: $F(1,3) = 27.44$, $P < 0.05$, for pattern type; $F(7,21) = 30.60$, $P < 0.00001$, for rotation angle; and $F(7,21) = 7.15$, $P < 0.0005$ for their interaction.

Interpretation of these effects is facilitated by Fig. 12. The difference between RS with and without collinearity was in the predicted direction (average d's are 2.22 and 1.63, respectively). Tukey's HSD-tests showed that the main effect of rotation angle was caused by a gradual decrease in detectability of RS with increasing rotation angle, except for point inversion (180°), where RS became more salient again [this increased performance for point inversion had been obtained by Foster (1978) as well]. The means are 4.17, 2.76, 2.23, 1.55, 1.23, 0.84, 0.41 and 2.21, respectively. As can be seen in Fig. 12, the pattern type × rotation angle interaction was due to the disappearance of a difference between ordinary and collinear RS at the smallest (i.e. 10°) and the largest (i.e. 180°) rotation angle.

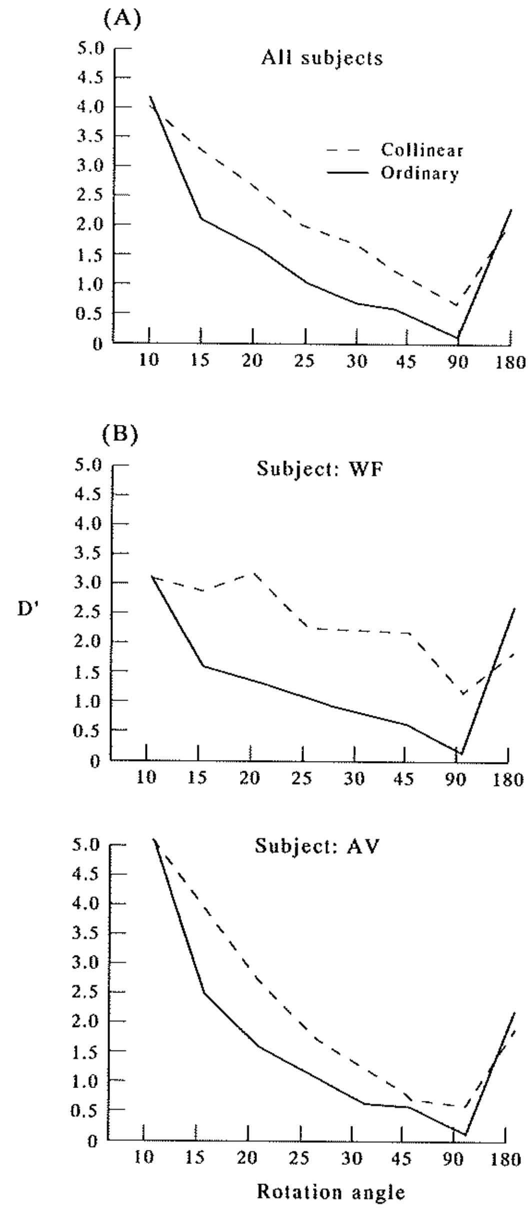

FIGURE 12. Results of Expt 3 with RS, averaged across subjects (A) and for two individual subjects (B).

Discussion

Both predictions on the detectability of RS in dot patterns, as derived from our suggetions about the role of higher-order structure, received clear experimental support. First, increasing the rotation angle up to 90° yields a pronounced disruptive effect because of the resulting increase of the orientation difference between pairs of virtual lines. At 180°, where the higher-order structures become correlation quadrangles again (i.e. parallelograms), RS is quite salient again. Secondly, introducing collinearity is generally very helpful. The suggestion is that this is so because it allows for a second type of regular higher-order structures to be formed (i.e. symmetric trapezoids). The fact that collinearity gives rise to virtual lines which are radial to the centre of rotation cannot be the explanation for this effect, because the random-dot patterns used as discriminative stimuli in these blocks of trials contained this feature also.

Two additional observations deserve some discussion. First, ordinary and collinear RS were equally detectable at 180°. As with the general effects, this specific finding can be understood on the basis of our view on regularity detection. Introducing collinearity in patterns with point-inversions simply enhances the lower-order virtual lines (as with collinear MS, four dots define the line instead of two), without giving rise to additional higher-order virtual quadrangles, which were not already present in the ordinary type. In accordance with the fact that radiality cannot be the explanation for the results obtained in Expt 3, it is quite natural that enhancing this lower-order structure does not add anything to the detection of the global regularity.

A second interesting finding is a similar absence of a difference between ordinary and collinear RS at the smallest rotation angle. The reason might be that with small rotation angles, the difference between the two types of correlation quadrangles is only negligibly small because both approach a rectangle. Perhaps, the higher-order structures present in ordinary RS (with 10° rotation) are strong enough already so that additional symmetric trapezoids afforded by collinear RS are superfluous.

GENERAL DISCUSSION

Summary of experimental data

Three experiments were designed in each of which one of the three simple Euclidean transformations was applied to collections of 12 dots which differed only in the number and type of higher-order structures afforded by them. Whereas models restricted to lower-order structure (e.g. Jenkins, 1983a for MS; Lowe & Binford, 1982 for TS; Stevens, 1978 for RS) would not expect different detectability performances resulting from this, specific predictions derived from the supposed role of additional higher-order structure were corroborated empirically. First, skewing MS was shown to be less detrimental if particular equidistant groupings yielding virtual parallelograms were included in the patterns. This result could not be explained by lower-order structure only because mere collinearity proved less efficient. Secondly, introducing variability of translation distances in dot patterns with TS caused a corresponding decrease in detectability. Whereas all higher-order structures disappear as a result of this manipulation, some very salient pairwise groupings would make some theories restricted to lower-order parallelism of virtual lines expect the reverse finding. Finally, RS patterns were more easy to detect if transformed dots were collinear so that symmetric trapezoids could be formed, except for very small rotation angles (i.e. 10°) where noncollinear patterns contain virtual parallelograms already, and for very large rotation angles (i.e. 180°) where collinearity does not add anything to the existent higher-order structure. Again, these specific findings seem difficult to explain by a pure lower-order based theory of regularity detection or grouping model because possible cues (such as radiality) were carefully avoided.

Alternative explanations

For some specific results obtained in this set of three experiments alternative explanations could be offered which do not rely on the notion of higher-order structures. For example, as one of the reviewers suggested, the finding that regularity in patterns with TS is harder to see when interdot distances vary is not so surprising if the visual system would represent the pairwise distances in a histogram and would respond to peaks in the latter. Similarly, the effect of skewing on the detection of regularity in patterns with MS could simply result from a special tuning or affinity of the visual system to orthogonal angles between the virtual lines and the axis of symmetry. As another reviewer suggested, the increased performance in the detection of regularity in patterns with RS of 180° (i.e. point inversion) is consonant with the idea of a fixed visual association between local features and certain spatial relations between these local features (Foster, 1978).

However, we strongly believe that these alternative schemes fail to explain other specific results. For example, apart from the fact that a "special affinity for orthogonality" seems to cry for a deeper explanation, it does not say anything about the difference between the different types of skewing. It might be regarded as "an explanation" of the difference between perfect bilateral symmetry on the one hand and skewed symmetry on the other, but it remains silent on the different effects observed when skewing different kinds of patterns with MS (e.g. ordinary, collinear, and equidistant). Likewise, the difference between 90 and 180° rotation might be easy to explain, but why is the regularity in patterns with RS much easier to detect when there is collinearity? In the Result sections of the three experiments described above, other findings could be identified which are difficult or impossible to explain by the suggested alternative mechanisms.

Furthermore, even if the alternative explanations could be adapted to incorporate all results of a particular experiment, they would remain only valid for one particular kind of regularity. For example, we do not see how the orientation histogram which might play a role in Expt 2 with TS could underly the results in Expts 1 and 3 with MS and RS, respectively. We think the strongest argument in favour of a mechanism based on higher-order structure is that it fits with *all* of the results obtained in these experiments. To have one single explanation for a diverse set of results is usually considered much better than to have to refer to different mechanisms for each specific result. In the following paragraphs we will offer a more specific suggestion for a general mechanism of regularity detection based on higher-order structure. Finally, we will give additional indications of the model's general validity by showing how the same account could explain some previously published results from a wide variety of studies.

A model of regularity detection based on higher-order structure

Given that the results show that lower-order structure is insufficient to explain regularity detection, the question arises as to why higher-order structure would help. In defining the notions of lower- and higher-order structure above, we already hinted at a possible mechanism of regularity detection in which higher-order structure was seen as enhancing the spread-out of local correspondences across the global pattern by a process called bootstrapping. In the Appendix, we treat these principles somewhat more formally and we outline the steps of a current implementation of them. Although we are not yet fully satisfied with the implementation, we are convinced of the psychological plausibility of the principles themselves. We describe the model and its implementation for illustrative purposes only, that is, to give to reader a feeling of how regularity detection might be based on higher-order structure. In a subsequent section, we will also show some simulation results with a remarkable equivalence to human regularity detection and perceptual grouping.

The model consists of two basic components: first, a function f that computes the goodness or salience of the groupings. The task of perceptual grouping or regularity detection can be formulated as an optimization problem over the set of all possible interpretations [for a general introduction to the use of minimization of energy functionals in vision, see Poggio, Torre and Koch (1985)]. Since this set is quite large, an iterative approximation algorithm is resorted to. This constitutes the second basic component of our model. Several parameters in the cost function take care that it will converge to a solution more rapidly, if more higher-order structure is present and more bootstrapping takes place.

The two essential terms of the functional, orientation differences between two virtual lines and angular differences between the smallest included angles between two pairs of virtual lines, correspond with lower- and higher-order structure, respectively. In order to avoid being trapped in local minima, a stochastic method, called *simulated annealing*, is used [for a good technical introduction, see Aarts and Korst (1989); a more accessible account is given by Hinton and Sejnowski (1986)] as a technique to minimize the cost function.

To understand why this feature of the mechanism is interesting from a psychological point of view, compare it with the evolution in time of *a spin-glass system*. This is a model used in statistical physics to study the global magnetization characteristics of some materials in terms of the local interactions between the outer electrons of each atom, represented by a vector called *spin* (e.g. Kirkpatrick, Gelatt & Vecchi, 1983). In the case here, with each couple of dots that can be connected, a spin is associated with two possible states (signs): active (+1) means that the connection is made, passive (−1) denotes that it is not made.

Furthermore, each spin has not only a state but also a weight. The spins can have a range of values between +1 and −1 so that the probability that a spin changes its sign can be proportional to its value. For the active, positive spins representing the connections made, this value is proportional to the resistance of the neighbourhood on the connection (represented in the first term in the energy function), whereas for the passive, negative spins, this value denotes the potential support of the neighbourhood for this connection if it would be made (represented in the second term in the energy function). In other words, if a spin becomes active (i.e. a pairwise grouping is made), the chance that another promising spin (in terms of correlation quadrangles) will become activated, as well as the chance that another less interesting connection will be removed, increases proportionally to the weight of the first one.

Suppose one has found the good quadrangle $(i,j)(k,l)$ by adding the line (i,j). If the quadrangle is of a nonaccidental nature (i.e. the genuine result of the regularity in the pattern), then (k,l) can serve as a starting point for building additional quadrangles (see Fig. 13). In this way, the grouping and regularity detection spread out very easily. Moreover, it should be noted that the chance of finding accidental higher-order structure (i.e. quadrangles) is smaller than the chance of finding accidental lower-order regularity (i.e. accidentally parallel virtual lines). In that way, eventually after some misleading directions, the process can propagate quickly in the right direction.

Finally, all this is regulated by a parameter so that in the beginning (at high temperature) the spins change their states rather arbitrarily, whereas gradually (with decreasing temperature) each change depends more on the orientation of the neighbours. A more formal treatment of the model and its current implementation is given in the Appendix.

Simulation results

So far, we tested some predictions derived from our theory about the role of higher-order structure by experiments on human regularity detection. Having an algorithm implementing the model's principles, allows a second type of test by letting the algorithm try to detect the regularity in the same patterns. The algorithm's

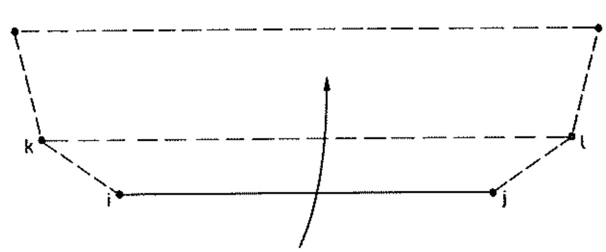

FIGURE 13. Bootstrapping. The correlation quadrangle formed between $(i,j)(k,l)$ suggests a direction in which to proceed.

output (i.e. its indication of regularity) was a number of connections found in the dot patterns, represented by virtual lines. Because each of the different types of regularities possesses at least 12 of these pairwise connections, the algorithm's efficiency was expressed as a ratio of the discovered versus the available ones. Because a kind of categorical decision was needed (to compare it with the observers' *Yes/No*-answers), 8/12 was arbitrarily selected as sufficient to answer positively. The relevant data point for each condition was then the number of regular input patterns (i.e. 100, as with the human subjects), correctly categorized as such on the basis of the prefixed criterion.

We could have chosen other criteria, such as the perceptual cost or lack of regularity of a grouping as computed on the basis of the cost function above, or the time it takes to minimize the functional. However, these other measures are still more sensitive to the specific model parameters which bear no relevance whatsoever to the basic principles behind it (e.g. sequential instead of parallel). In fact, the absolute efficiency of the algorithm is, in this context, not as important as the relative rank orderings of the different conditions. More specifically, we hope that these differences correspond with those obtained for human perceivers. To enable this comparison, the algorithm's results will be plotted in the same way. Only two sets of patterns have been tested until now.

Multiple and skewed symmetries. First, an attempt was made to simulate the results obtained for the multiple and skewed symmetries (see Wagemans *et al.*, 1991). A brief summary of the psychophysical results was given above (Fig. 3). The simulation results are shown in Fig. 14. As for human perceivers, the effect of skewing was very disruptive for single symmetry, but gradually less so for double and quadruple symmetry. Again, this is interpreted as evidence for the role of higher-order structure in the detection of this type of regularity. Although skewing preserves lower-order parallelism between virtual lines, it destroys higher-order correlation quadrangles (i.e. symmetric trapezoids) in patterns with single MS. In skewed double symmetry, a second type of correlation quadrangles is available (i.e. parallelograms),

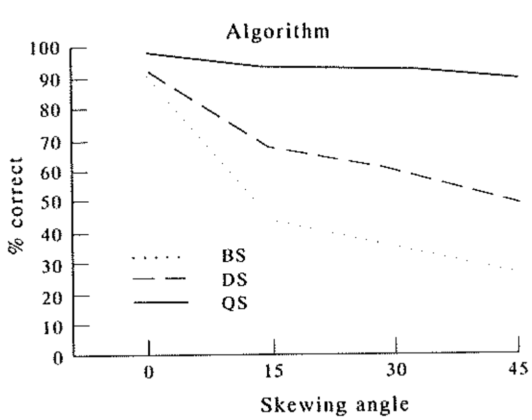

FIGURE 14. Results of Simulation 1 with multiple skewed symmetry, to be compared with Fig. 3 for human perceivers.

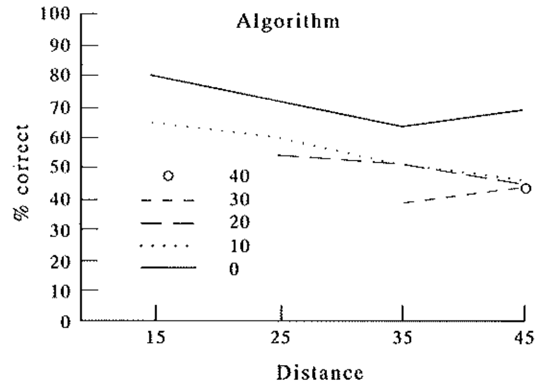

FIGURE 15. Results of Simulation 2 with TS, to be compared with Fig. 8 for human perceivers.

whereas in skewed quadruple symmetry the bootstrapping-facilitating higher-order structures are still preserved.

Although the general agreement with the psychophysical data is quite good (i.e. the same rank ordering of average performance levels), some minor differences deserve some comment. First, for the zero skewing angles, additional symmetry, which is very helpful for human perceivers (e.g. Humphrey & Humphrey, 1989; Palmer & Hemenway, 1978; Royer, 1981; Wagemans *et al.*, 1991), apparently does not support the detection of regularity by the algorithm (90, 91 and 97%, for one, two and four axes of symmetry, respectively, to be compared with $d' = 3.41$, 4.60 and 5.09, respectively). An obvious explanation might be a kind of ceiling effect, which could have been avoided by using a more stringent criterion for the categorical decision (e.g. 10/12). Casual testing with different criteria showed, however, that all conditions yielded correspondingly lower or higher performance levels, so that the relative differences were preserved. Nevertheless, other performance measures, such as the time it takes to optimize the cost function, did show a difference between single, double, and quadruple (unskewed) symmetry.

A second salient feature in the pattern of results obtained with the algorithm is the rapid decrease in detectability of single MS, even for the smallest skewing angle. Human regularity detection is more robust and seems to tolerate larger orientation differences (see Fig. 4 and Wagemans *et al.*, 1991). Most probably, this larger robustness (stressed previously by Barlow & Reeves, 1979) is enabled by the way in which orientations (and orientation differences) are measured in our visual system, that is, by populations of cells with more-or-less narrowly tuned functions. In contrast, the algorithm works with single exact values.

Translations. The 56 different sets of patterns with TS, as defined by different variabilities on the translation distances and four directions (see Expt 2), were used to test the grouping algorithm. In general, the results of these simulations (see Fig. 15) are in close agreement with the psychophysical ones (see Figs 8 and 9). Although the algorithm does not seem to be subject to the same dramatic distance-variation effects as humans, one has to be careful to compare the different performance measures (% CRT vs d'). As for human perceivers, detection of TS becomes increasingly more difficult with increasing translation distance (i.e. 72, 61, 50 and 47%) and increasing variation (i.e. 70, 54, 49, 40 and 41%).

A small but interesting difference between both data sets is that the algorithm is less affected by larger translation distances (e.g. 45) than human regularity detection, at least when it is fixed throughout the pattern (i.e. zero variation). This result indicates that the formation of virtual lines by the algorithm, even relatively long ones, are probably less influenced by neighbouring dots than we are. As suggested previously, the human visual system is constrained to show serious disturbance effects by spurious point-pairs in the local surround (e.g. Glass & Switkes, 1976; Jenkins, 1983b; Stevens, 1978). Different choices of relevant neighbourhoods in the algorithm could reduce this discrepancy. In fact, psychophysical data on neighbourhood issues are available with respect to other grouping phenomena such as collinearity or curvilinearity (e.g. Foster, 1979; Zucker & Davis, 1988) and are incorporated in algorithms designed to detect them (e.g. Lowe & Binford, 1982; Smits & Vos, 1986; Smits, Vos & van Oeffelen, 1985).

Despite these minor differences, we are quite happy with the qualitative equivalence between some psychophysical results and the simulation data. This shows that the way we think higher-order structure helps is at least a candidate mechanism to be considered for further research.

An indication of the model's generality

A further strong argument in favour of our approach is that it is able to integrate findings about perceptual processes that differ considerably at first sight. A closer look, however, reveals that they are all concerned with grouping, and, more specifically, the correspondence problem that has to be solved in it. Considering evolutionary pressures, it seems quite natural that some common mechanism exists to fulfill tasks that can be traced back to the same kind of preattentive grouping mechanism. Indeed, it would be very uneconomical to have devoted modules for the detection of, for example, symmetry in dot patterns, global structure in Glass or vector patterns, and correspondence between two images as in stereo or motion. That is the reason why we now want to give a feeling of the model's generality by indicating some phenomena that can be captured by it.

Symmetry detection. Because the model was inspired by our results on symmetry detection, only one additional phenomenon taken from this area will be given as an example of its integrative power. It was discovered and further investigated by Palmer and Hemenway (1978). As one of the kinds of symmetry, the detection of which was tested in a reaction time paradigm, they presented their subjects polygons with double or quadruple symmetry, constructed by reflecting a basis pattern not once, but twice or even four times. The fact that this multiple symmetry was easier to detect is not surprising [it has often been replicated since then (e.g. Humphrey & Humphrey, 1989; Royer, 1981)]. An additional finding, however, was that the presence of an additional symmetry about a H axis helped in the detection of symmetry about a V axis. Palmer and Hemenway's suggestion that this enhancement is caused by the better Gestalt of the two pattern halves to be compared, is unsatisfactory because it introduces another factor to be explained (why is it a better Gestalt?). Our model explicitly permits these kinds of interactions between symmetries to take place.

In addition to the fact that the number of symmetric trapezoids is much (two or four times) larger, some virtual quadrangles are rectangles in double and quadruple symmetry. Assuming that our visual system has a special affinity for higher-order structural elements such as symmetric trapezoids and parallelograms, it is quite natural to suppose that a quadrangle which is both at the same time has a special status. Moreover, these special correlation quadrangles allow the local reference frames used in bootstrapping to be swapped, so that the spread-out from local correspondences throughout the whole pattern can be even more efficient.

Global structure in vector patterns. As with symmetry detection, there are abundant data on the perception of global structure in these types of patterns. Vector patterns have been introduced to demonstrate the validity of Hoffman's (1966) Lie Transformation Group model of the Neuropsychology of Perception, LTG/NP for short (Caelli & Dodwell, 1982, 1984; Dodwell & Caelli, 1985). These patterns consist of oriented line segments positioned to suggest the tangent vectors on the orbits created by the simple Lie operators [for a clear introduction to this terminology see Dodwell (1983)]. The basic rationale behind these studies was to show that local perturbations of the orientation and/or position of the elements were harder to detect when the global structure created by them was more salient (i.e. in agreement with the basic Lie germs). More recently, Moraglia (1989) showed that a horizontal line segment could be detected preattentively (as suggested by the flat search functions) when it was incorporated in a global vector pattern, whereas it required serial, attentive search when it was on the same circular position of an incoherent structure.

Although the set of results fits the general predictions of LTG/NP quite well, there is nothing in the data that is incongruent with our account presented here. For example, a starwise or circular grouping of parallel or tangential line segments is full of correlation quadrangles (pairwise connecting the line segments at their endpoints yields paralellograms, whereas connecting four line segments at their midpoints gives rise to symmetric trapezoids, see Fig. 16), which become gradually destroyed by perturbing the local properties of the line segments.

Correspondence in stereo. Working with random-dot stereograms, Prazdny (1985) showed that human perceivers can group dots belonging to the same depth plane, even when two "transparent" random-dot planes are simultaneously present. This finding cannot be explained by lower-order structure only, because all

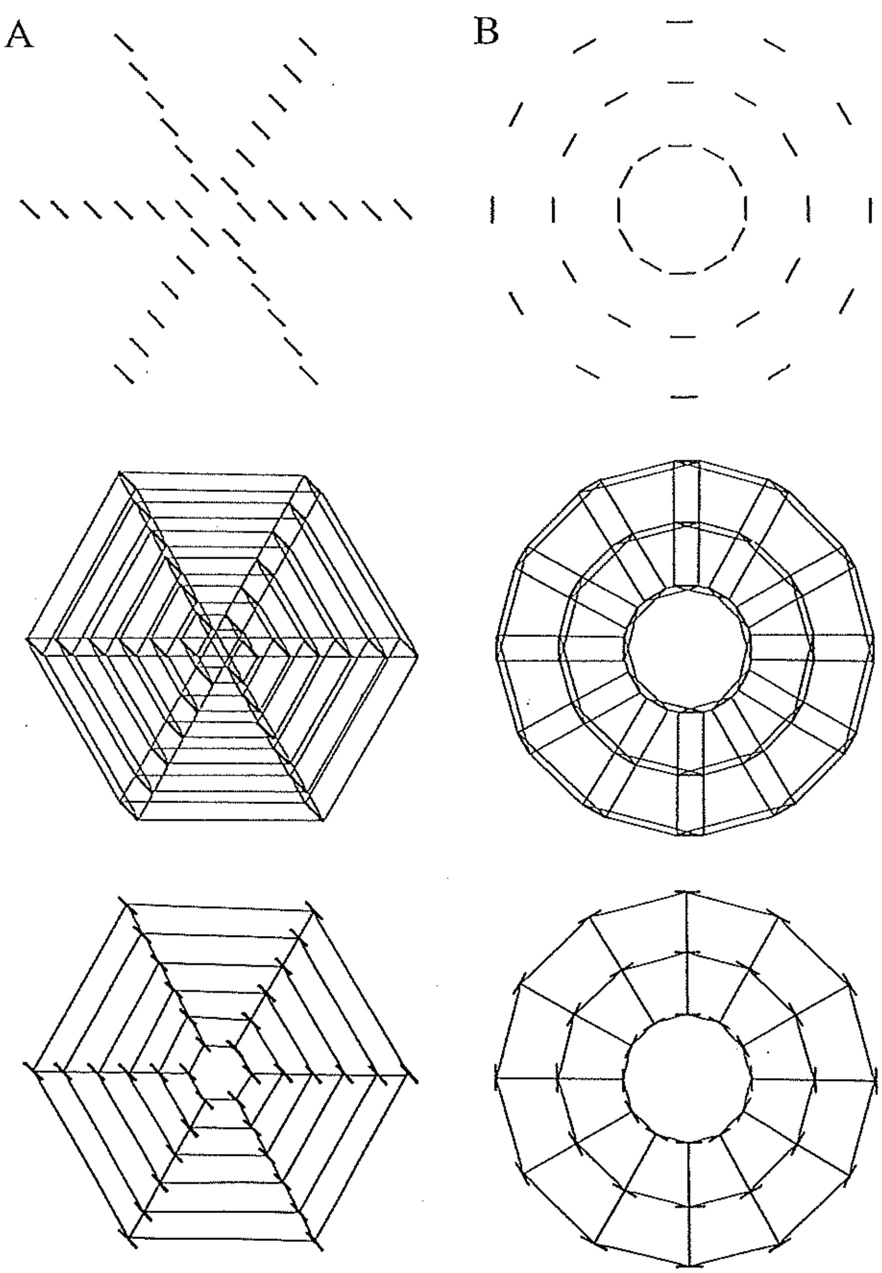

FIGURE 16. Two vector patterns together with their correlation quadrangles. For the starlike pattern in (A) as well as for the circular pattern in (B), correlation quadrangles can be formed by connecting endpoints of two vectors or, at a coarser level, by connecting mere positions of four line segments.

point pairs in the two depth planes are connected by more or less horizontal (and thus parallel) virtual disparity lines. Separating the planes on the basis of lower-order information alone is, therefore, impossible. We suggest that higher-order groupings formed between pairs of virtual lines belonging to the same depth plane account for this perception of stereoscopic transparency.

Some indirect support for this explanation is provided by more recent experimental work by Akerstrom and Todd (1988). They attempted to facilitate the perception of stereoscopic transparency by manipulating the similarity of the corresponding elements with respect to several attributes. Although this manipulation was successful for some attributes (such as color), it did not enhance stereoscopic transparency when corresponding elements (i.e. line segments) were given the same orientation. However, only two orientations were used and they were mirror symmetric with respect to the symmetry axis (i.e. $\pm 45°$). Accidentally, these can give rise to spurious correlation quadrangles between the wrong elements (i.e. different depth planes). We suspect that introducing similar line orientations for elements belonging to the same depth plane would be helpful if they were not mirror symmetric, so that no quadrangles of the correlational type can be formed (see Fig. 17).

Correspondence in motion. A similar explanation could also be offered for the nice results which were recently obtained by Werkhoven, Snippe and Koenderink (1990) on the effects of element orientation on apparent motion perception. Essentially, their paradigm allowed to test the degree of correspondence (as measured by the relative preference for one of two directions in an ambiguous

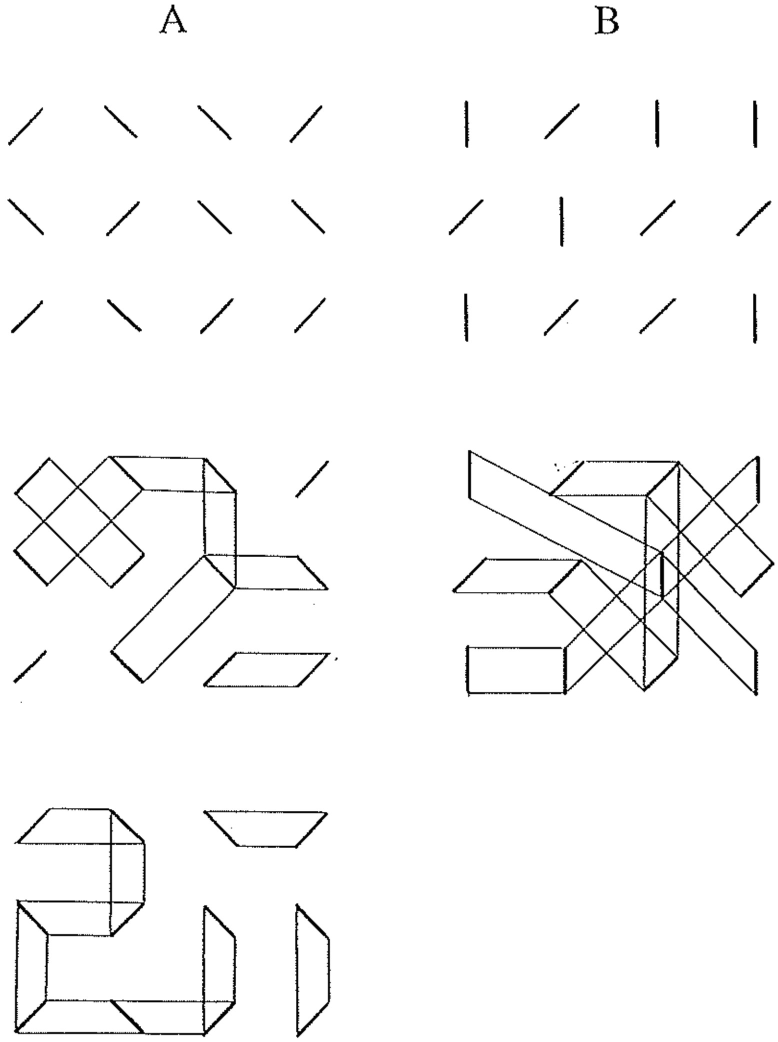

FIGURE 17. A very simplified sketch of the explanation offered for stereoscopic transparency phenomena (Akerstrom & Todd, 1988). In (A), where two orthogonal line orientations are used to allow segregation of the planes, correct matches (in the second row), as well as incorrect ones (in the third row) can be formed on the basis of correlation quadrangles (parallelograms vs symmetric trapezoids, respectively). In (B), where two different but nonorthogonal line orientations are used, only the correct correspondences can be established by higher-order grouping. In realistic displays, these properties hold only in statistical terms of course.

apparent motion situation) between elements that were oriented the same or differently. In two experiments, this correspondence measure was a function of the angular difference between the line orientation ($\pm 180°$) and the motion direction (which is a kind of higher-order structure, see below): independent of whether the absolute line orientations were the same or not (which constitutes lower-order correspondence), the motion direction was preferred that was as orthogonal as possible to the line orientation. When the absolute line orientations were different in the horizontal motion direction, this correspondence was nevertheless chosen if the higher-order grouping was giving rise to strong symmetric trapezoids (i.e. with a small difference between the four angles). Likewise, the vertical correspondence between similar line orientations was preferred if the higher-order grouping was giving rise to strong parallelograms (see Fig. 18).

CONCLUSION

In this paper, we presented experimental data showing the insufficiency of lower-order structure to explain regularity detection. We suggested that a global regularity is detected more easily if local pairwise groupings (represented by virtual lines) are supported by higher-order ones formed between them (represented by correlation quadrangles). These enable the lower-order groupings to spread out across the whole pattern very rapidly (called bootstrapping). As a preliminary attempt to specify these principles, we proposed a working model with two basic components: first, a function expressing the cost of a perceptual grouping or the lack of regularity, and, secondly, an algorithm based on simulated annealing to minimize the cost function. The simulation results obtained with our current implementation of these principles showed satisfying qualitative agreement with human regularity detection performance. Finally, the theory was shown to capture the essence of a large number of grouping phenomena taken from diverse domains such as detection of symmetry in dot patterns, global structure in Glass and vector patterns, correspondence in stereoscopic transparency and apparent motion. Therefore, we are convinced that, in principle, the mechanism used by the human visual system to

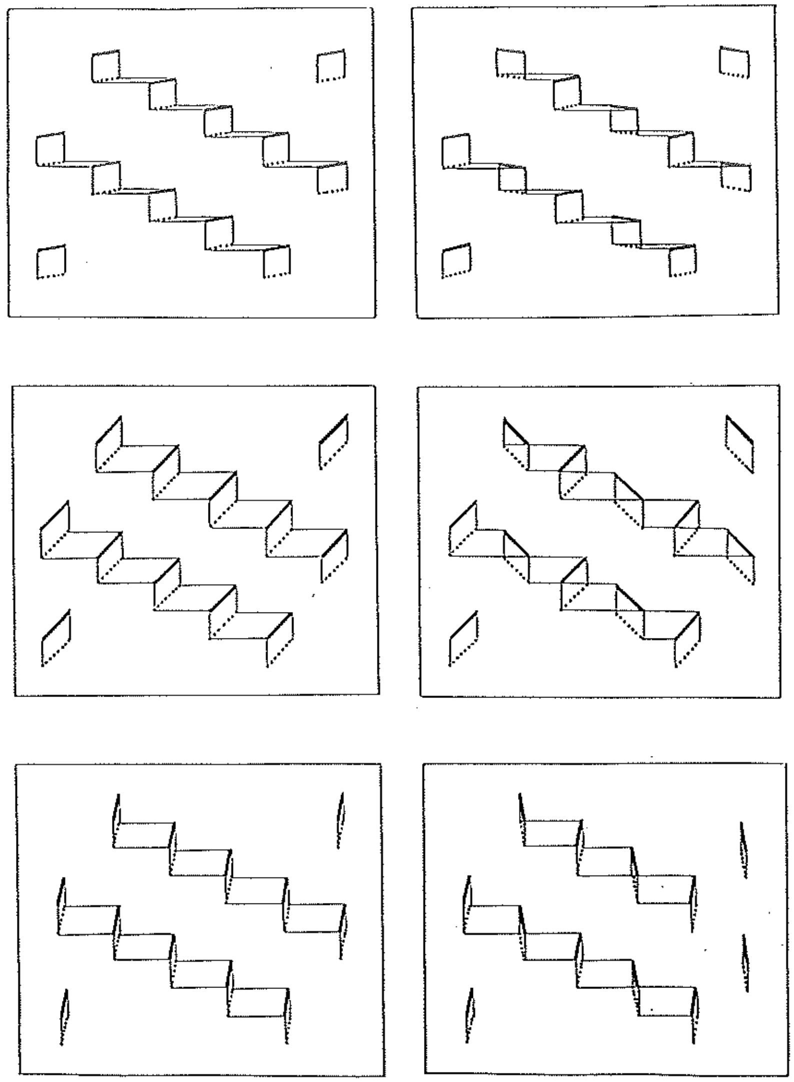

FIGURE 18. A sketch of the explanation offered for apparent motion phenomena (Werkhoven *et al.*, 1990). In (A) with same orientation segments as well as in (B) with different orientation ones, a small angle (1) favours a vertical motion and a large angle (3) a horizontal one in displays where motion direction is ambiguous (thick solid lines indicate elements at frame t, dotted ones at $t + 1$; thin solid lines represent the correlation quadrangles that can be formed). At intermediate angles (2), ambiguity is maximal. The suggestion is that it all depends on the strength of the correlation quadrangles that can be made.

detect regularity incorporates something like bootstrapping based on higher-order structure. We regard this as a promising step towards unraveling the intriguing mechanisms of classic Gestalt phenomena.

REFERENCES

Aarts, E. & Korst, J. (1989). *Simulated annealing and Boltzmann machines: A stochastic approach to combinatorial optimization and neural computing.* New York: Wiley.

Akerstrom, R. A. & Todd, J. T. (1988). The perception of stereoscopic transparency. *Perception and Psychophysics, 44,* 421–432.

Appelle, S. (1972). Perception and discrimination as a function of stimulus orientation: The "oblique effect" in man and animals. *Psychological Bulletin, 78,* 266–278.

Barlow, H. B. & Reeves, B. C. (1979). The versatility and absolute efficiency of detecting mirror symmetry in random dot displays. *Vision Research, 19,* 783–793.

Bischof, W. F., Foster, D. H. & Kahn, J. I. (1985). Selective internal operations in the recognition of locally and globally point-inverted patterns. *Spatial Vision, 1,* 179–196.

Bruce, V. G. & Morgan, M. J. (1975). Violations of symmetry and repetition in visual patterns. *Perception, 4,* 239–249.

Caelli, T. & Dodwell, P. (1982). The discrimination of structure in vectorgraphs: Local and global effects. *Perception and Psychophysics, 32,* 314–326.

Caelli, T. & Dodwell, P. (1984). Orientation-position coding and invariance characteristics of pattern discrimination. *Perception and Psychophysics, 36,* 159–168.

Dodwell, P. C. (1983). The Lie transformation group model of visual perception. *Perception and Psychophysics, 34,* 1–16.

Dodwell, P. C. & Caelli, T. (1985). Recognition of vector patterns under transformations: Local and global determinants. *Quarterly Journal of Experimental Psychology, 37A,* 1–23.

Essock, E. A. (1980). The oblique effect of stimulus identification considered with respect to two classes of oblique effects. *Perception, 9,* 37–46.

Fischer, C. B. & Fracasso, M. P. (1987). The Goldmeier effect in adults and children: Environmental, retinal, and phenomenal influences on judgments of visual symmetry. *Perception, 16,* 26–39.

Foster, D. H. (1978). Visual comparison of random-dot patterns: Evidence concerning a fixed visual association between features and feature-relations. *Quarterly Journal of Experimental Psychology, 30,* 637–654.

Foster, D. H. (1979). Discrete internal pattern representations and visual detection of small changes in pattern shape. *Perception and Psychophysics*, *26*, 459–468.

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

Friedberg, S. A. (1986). Finding the axes of skewed symmetry. *Computer Vision, Graphics, and Image Processing*, *34*, 138–155.

Glass, L. (1969). Moiré effect from random dots. *Nature*, *223*, 578–580.

Glass, L. & Perez, R. (1973). Perception of random dot interference patterns. *Nature*, *246*, 360–362.

Glass, L. & Switkes, E. (1976). Pattern recognition in humans: Correlations which cannot be perceived. *Perception*, *5*, 67–72.

Hinton, G. E. & Sejnowski, T. J. (1986). Learning and relearning in Boltzmann machines. In Rumelhart, D. E., McCulland, J. L. & the PDP Research Group (Eds), *Parallel distributed processing: Explanation in the microstructure of cognition* (Vol. 1, pp. 282–317). Cambridge, Mass.: MIT Press/Bradford Books.

Hoffman, W. C. (1966). The Lie algebra of visual perception. *Journal of Mathematical Psychology*, *3*, 65–98.

Hubel, D. H. & Wiesel, T. N. (1977). Functional architecture of macaque monkey visual cortex. *Proceedings of the Royal Society of London B*, *198*, 1–59.

Humphrey, G. K. & Humphrey, D. E. (1989). The role of structure in infant visual pattern perception. *Canadian Journal of Psychology*, *43*, 165–182.

Jenkins, B. (1982). Redundancy in the perception of bilateral symmetry in dot textures. *Perception and Psychophysics*, *32*, 171–177.

Jenkins, B. (1983a). Component processes in the perception of bilaterally symmetric dot textures. *Perception and Psychophysics*, *34*, 433–440.

Jenkins, B. (1983b). Spatial limits to the detection of transpositional symmetry in dynamic dot textures. *Journal of Experimental Psychology: Human Perception and Performance*, *9*, 258–269.

Jenkins, B. (1985). Orientational anisotropy in the human visual system. *Perception and Psychophysics*, *37*, 125–134.

Julesz, B. (1971). *Foundations of cyclopean perception*. Chicago, Ill.: University of Chicago Press.

Kahn, J. I. & Foster, D. H. (1986). Horizontal-vertical structure in the visual comparison of rigidly transformed patterns. *Journal of Experimental Psychology: Human Perception and Performance*, *12*, 422–433.

Kanade, T. & Kender, J. R. (1983). Mapping image properties into shape constraints: Skewed symmetry, affine-transformable patterns, and the shape-from-texture paradigm. In Beck, J., Hope, B. & Rosenfeld, A. (Eds), *Human and machine vision* (Vol. 1, pp. 237–257). New York: Academic Press.

Kirkpatrick, S., Gelatt, C. D. & Vecchi, M. P. (1983). Optimization by simulated annealing. *Science*, *220*, 671–680.

Locher, P. & Wagemans, J. (1993). Effects of element type and spatial grouping on symmetry detection. *Perception*. In press.

Lowe, D. G. & Binford, T. O. (1982). Segmentation and aggregation: An approach to figure-ground phenomena. In *Proceedings of image understanding workshop* (pp. 168–178). Palo Alto, Calif.: DARPA.

Mach, E. (1959). *The analysis of sensations*. New York: Dover. (Originally published in German, 1886.)

Moore, D. J. H. & Parker, D. J. (1974). Analysis of global pattern features. *Pattern Recognition*, *6*, 149–164.

Moraglia, G. (1989). Display organization and the detection of horizontal line segments. *Perception and Psychophysics*, *45*, 265–272.

Palmer, S. E. & Hemenway, K. (1978). Orientation and symmetry: Effects of multiple, rotational, and near symmetries. *Journal of Experimental Psychology: Human Perception and Performance*, *4*, 691–702.

Pashler, H. (1990). Coordinate frame for symmetry detection and object recognition. *Journal of Experimental Psychology: Human Perception and Performance*, *16*, 150–163.

Poggio, T., Torre, V. & Koch, C. (1985). Computational vision and regularization theory. *Nature*, *317*, 314–319.

Prazdny, K. (1984). On the perception of Glass patterns. *Perception*, *13*, 469–478.

Prazdny, K. (1985). Detection of binocular disparities. *Biological Cybernetics*, *52*, 93–99.

Prazdny, K. (1986). Some new phenomena in the perception of Glass patterns. *Biological Cybernetics*, *53*, 153–158.

Rock, I. & Leaman, R. (1963). An experimental analysis of visual symmetry. *Acta Psychologica*, *21*, 171–183.

Royer, F. L. (1981). Detection of symmetry. *Journal of Experimental Psychology: Human Perception and Performance*, *7*, 1186–1210.

Smits, J. T. S. & Vos, P. G. (1986). A model for the perception of curves in dot figures: The role of local salience of "virtual lines". *Biological Cybernetics*, *54*, 407–416.

Smits, J. T. S., Vos, P. G. & van Oeffelen, M. P. (1985). The perception of a dotted line in noise: A model of good continuation and some experimental results. *Spatial Vision*, *1*, 163–177.

Stevens, K. A. (1978). Computation of locally parallel structure. *Biological Cybernetics*, *29*, 19–28.

Stevens, K. A. & Brookes, A. (1987). Detecting structure by symbolic constructions on tokens. *Computer Vision, Graphics, and Image Processing*, *37*, 238–260.

Wagemans, J., Van Gool, L. & d'Ydewalle, G. (1990). Visual search in dot-patterns with bilateral and skewed symmetry. In Brogan, D. (Ed.), *Visual search* (pp. 99–114). London: Taylor & Francis.

Wagemans, J., Van Gool, L. & d'Ydewalle, G. (1991). Detection of symmetry in tachistoscopically presented dot patterns: Effects of multiple axes and skewing. *Perception and Psychophysics*, *50*, 413–427.

Wagemans, J., Van Gool, L. & d'Ydewalle, G. (1992). Orientational effects and component processes in symmetry detection. *Quarterly Journal of Experimental Psychology*, *44A*, 475–508.

Werkhoven, P., Snippe, H. P. & Koenderink, J. J. (1990). Effects of element orientation on apparent motion perception. *Perception and Psychophysics*, *47*, 509–525.

Zucker, S. W. & Davis, S. (1988). Points and endpoints: A size/spacing constraint for dot grouping. *Perception*, *17*, 229–247.

Acknowledgements—The authors are supported by grants from the Belgian National Fund for Scientific Research (NFWO) to JW and JVH, from the Belgian Institute for the Encouragement of Scientific Research in Industry and Agriculture (IWONL) to LVG, and from the University of Leuven (OT/89/11) to VS. The interdisciplinary collaboration is enabled by OT/89/11 and by the Belgian Fund for Joint Basic Research (FKFO). The equipment used is sponsored by the European Economic Community (BRA and PCA), the Belgian Ministry of Science (AI-Impulse Program) and the University of Leuven (GOA). The writing of the paper was partly sponsored by BRA 6448. We gratefully acknowledge the technical support by D. Christiaens, M. Lenaerts, and J. Van Rensbergen. The comments and suggestions made by D. Foster and an anonymous reviewer greatly improved the manuscript.

APPENDIX

Whereas an intuitive feeling of our model of regularity detection was given above, we now provide more formal details with respect to its two components, the cost function and the optimization method. As indicated previously, we see this version of our model as only one of the possibilities, which still have to be evaluated on the basis of further research.

Cost function

Given a random-dot pattern with dots $\{i, j, k \ldots\}$, each grouping x is represented as a set of virtual lines V. Each virtual line is determined by its endpoints $\{(i_1, j_1), (i_2, j_2), \ldots\}$. With the convention that a high value of $f(\,)$ corresponds with a bad grouping (hence, cost function), it is defined as:

$$f(x) = \sum_{r} (\exp(|\phi_{(i,j)} - \phi_{(k,l)}|) - 1) - \sum_{Q} (\exp(q_{(i,j)(k,l)}) - 1)$$

$$+ \sum_{\text{dots}} \exp(g(i)).$$

The three terms have the following meaning. The first term expresses the degree of parallelism of the virtual lines in x (i.e. lower-order structure). $|\phi_{(i,j)} - \phi_{(k,l)}|$ represents the orientation difference as the absolute value of the smallest included angle between the virtual line (i, j) and (k, l). The second term reflects the higher-order structure dependent on the lower-order one between the virtual lines (i, j) and (k, l). Together, these lines form a quadrangle.

$$(q_{(i,j)(k,l)}) \text{ is defined as } (1 - |\phi_{(i,j)} - \phi_{(k,l)}|)(1 + \theta_{(i,j)(k,l)}).$$

The weighting factor θ is proportional to the quality of the corresponding quadrangle: it is maximal when the absolute value of the difference of the smallest included angles determined by the lines $(i, j)(i, k)$ and $(i, j)(j, l)$ (and vice versa) is minimal. One sums over all quadrangles (i, j) with k respectively l being a neighbour of i respectively j. By definition, the set of neighbours of a dot are the α nearest dots (α is a constant). For each virtual line, one always assures that at least one quadrangle from each corresponding half plane is considered. This set of quadrangles is denoted by Q.

For both lower- and higher-order terms only restricted neighbourhoods of virtual lines are taken into account. In addition to increasing the algorithm's efficiency, this has a clearcut psychological underpinning (e.g. Barlow & Reeves, 1979; Jenkins, 1982; Julesz, 1971). Note, however, that the summation sets for lower- and higher-order structure differ. This reflects a different scope of influence: the lower-order interaction of two virtual lines is weakened as the distance between them increases, whereas the higher-order influence is strong when not too many points disturb the considered quadrangle. Both distance and density have been used to define neighbourhoods in previous models (for examples of the former, see Smits & Vos, 1986; Smits *et al.*, 1985; for examples of the latter, see Lowe & Binford, 1982; Stevens, 1978; Zucker & Davis, 1988). However, common to all these models is that they are restricted to lower-order regularity.

The last term imposes continuity restrictions on all virtual lines v (if any) that have a given endpoint i in common. At the same time, it penalizes for the presence of points not contributing to any virtual line (isolated elements reduce the saliency of the grouping or regularity).

Optimization method

It generates iteratively a sequence of approximations x_t, which are realizations of a Markov chain converging as t goes to infinity to a probability distribution having non-zero values at optima of $f(\ )$ only. The basic mechanism is the following. (1) Given x_t, generate a candidate y for x_{t+1} by deleting or adding one virtual line. (2) If $f(y) < f(x)$, take $x_{t+1} = y$; otherwise, accept the change with probability

$$\sim \exp\left(\frac{f(x_t) - f(y)}{T_t}\right).$$

T_t is a decreasing sequence to 0, usually referred to as the *temperature*
The general outline of the algorithm is as follows:

```
determine T_0;
determine x_0;
WHILE not-good-enough DO
  WHILE equilibrium-not-reached DO
    given x_t, generate a candidate y for x_{t+1} by deleting
    or adding a line;
    if f(y) < f(x_t), take x_{t+1} = y;
    else, accept the change with probability ~exp((f(x_t) - f(y))/T_t)
  END;
calculate T_{n+1};
END.
```

At the beginning (when T is high), the spins (virtual lines, cortical cells) change their states rather arbitrarily. As T decreases, each change of a spin depends gradually more on the orientation of the neighbours. Because of the form of the cost function, the computation $\Delta f() = f(x_t) - f(y)$ requires only information about the neighbouring lines. In other words, $\exp(-f(x)/T)$, that is, the equilibrium distribution of the Markov chain generated by the inner loop, is a Markov Random Field (e.g. Poggio *et al.*, 1985). This would allow a large degree of parallelism in the computations to be carried out. Although this seems natural to enhance both the computational efficiency and the psychological plausibility, the current implementation is purely sequential (an implementation on transputers is in progress).

Parameter choices

Generation step. The usual mechanism, that is, generate x_{t+1} from the uniform distribution over the neighbours of x_t, is much less efficient than it can be by considering the enormous possibility of bootstrapping. Therefore, a quantity $a_{(i,j)}$ is associated with each possible virtual line:

$$\begin{cases} \sum_{\text{all lines}} (\exp(|\phi_{(i,j)} - \phi_{(k,l)}|) - 1) & \text{if } (i,j) \in x \\ -\sum_{\text{quadr.}} (\exp(q_{(i,j)(k,l)}) - 1) & \text{otherwise.} \end{cases}$$

In summary, the generation step incorporating these characteristics becomes:

(1) With probability P: propose a virtual line, chosen from the set of interesting connections with a probability proportional to $a_{(i,j)}$.

(2) With probability Q: delete a virtual line, chosen from the set on uninteresting connections in x with a probability proportional to $-a_{(i,j)}$.

(3) Otherwise (i.e. with probability $1 - P - Q$): select a point randomly and try to find a closest neighbour in a region centred around it as defined by two random corners (i.e. one between 0 and 2π, and the other between $\pi/8$ and $\pi/4$).

Equilibrium-not-reached criterion. The inner loop is terminated after a fixed number of steps. The present implementation does not seem to require an accurate determination whether or not the equilibrium is reached.

Decreasing the temperature. The current implementation used $T_{n+1} = 0.9\ T_n$. However, it often turns out that in dot patterns with a salient global structure such as BS, the global minimum of $f(\ )$ is much smaller than the local minima. In these cases, it would be useful to decrease the temperature dramatically after the detection of a sudden jump of the cost function (e.g. $T_{n+1} = 0.1\ T_n$). It remains to be seen how the algorithm would behave if this kind of variability in the decrease of the temperature would be incorporated.

Not-good-enough criterion and determination of T_0. The former is true as long as the percentage of accepted changes during a cycle of the inner loop is lower than the square root of the number of dots. The latter is determined in such a way as to get at least one chain at fixed temperature where the ratio of the accepted vs rejected moves is > 0.5. More work would be needed to enhance the psychological plausibility of these criteria.