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, Tienestraat 102, B-3000 Leuven, Belgium.

†Department of Electronics, ESAT-MI2, University of Leuven, Belgium.

‡Department of Mathematics, University of Leuven, Belgium.

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 $\sum e^{|\phi_{(i,k)} - \phi_{(j,l)}| + |\phi_{(k,i)} - \phi_{(l,j)}|}$ 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 hom-

ogeneous 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.

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 lose 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 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 $d1$) would not be repeated at the same side of the axis of symmetry. In other 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 $\pm 15^\circ$ 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 $\times$ skewing; $F(18,54) = 8.58$, for orientation $\times$ skewing; and $F(36,108) = 3.64$, for pattern type $\times$ orientation $\times$ skewing. The remaining second-order interaction (i.e. pattern type $\times$ 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

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 dis-

tances, 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.

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.

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

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

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 $\times$ 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.

Discussion

Both predictions on the detectability of RS in dot patterns, as derived from our suggestions 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

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.

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 parallelograms, 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 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 ($\pm 180^\circ$) 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

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

$$f(x) = \sum_r (\exp(|\phi_{(i,j)} - \phi_{(k,l)}|) - 1) - \sum_Q (\exp(q_{(i,j)hk,h}) - 1) + \sum_{\text{dots}} \exp(g(i)).$$