Spatial Vision, Vol. 9, No. 1, pp. 33-55 (1995)

A model for global symmetry detection in dense images*

FRANÇOYS LABONTÉ,[1] YERUCHAM SHAPIRA,[1] PAUL COHEN[1] and JOCELYN FAUBERT[1,2]

[1]*Perception and Robotics Research Group, Department of Electrical and Computer Engineering, Pavillon André-Aisenstadt, Ecole Polytechnique de Montréal, PO Box 6079, Station 'Centre-Ville', Montréal, Québec, Canada, H3C 3A7*

[2]*Ecole d'Optométrie, Université de Montréal, PO Box 6128, Station 'Centre-Ville', Montréal, Québec, Canada, H3C 357*

Received 28 June 1993; revised 14 June 1994; accepted 10 August 1994

Abstract—In this paper, a model is proposed for bilateral symmetry detection in images consisting of dense arrangements of local features. The model is elaborated on the basis of a psychophysical experiment showing that grouping precedes and facilitates symmetry detection. The proposed computational model consists of three stages: a grouping stage, a symmetry-detection stage, and a symmetry-subsumption stage. Reliance upon a preliminary grouping stage enables a significant reduction of the computational load for detecting symmetry. An implementation of the model is described, and results are presented, showing a good agreement of the model performance with human symmetry perception.

1. INTRODUCTION

In this paper, we propose a model for global bilateral symmetry detection in dense images. Dense arrangements of local features, such as dots or oriented segments, are encountered in various situations such as images of textured scenes, brightness gradient fields, stereo disparity fields, motion flow fields, and Glass patterns.

Different classes of symmetry may be present in an image: bilateral (also known as mirror symmetry), rotational, or repetitive. Multiple symmetries, i.e. identities under more than one transformation, may be also present. Unlike its mathematical equivalent, visual symmetry may be characterized by a variable degree of exactness. It may be approximate, incomplete (missing elements), and apply only to restricted regions of the image.

Symmetry relations can be local or global. Local symmetry refers to small neighborhoods of the image, and can usually be detected through a limited number of comparisons between nearby features. Global symmetry concerns regions of larger

*Part of the work in this paper was presented at the IEEE 4th International Conference on Computer Vision, Berlin 1993.

extents (although usually not affecting the entire image) and requires an examination of far-away image elements, followed by a comparison stage that involves the establishment of symmetry relations between compatible elements.

Since introduced by Mach (1906/1959), symmetry detection in images has drawn considerable attention in terms of both experimental studies and computational models. This interest probably comes from the fact that symmetry could be helpful in the accomplishment of a large number of visual tasks.

For instance, symmetry might be used for shape representation, either by providing a means for efficient encoding of pictorial information (Attneave, 1954), or by facilitating the establishment of object-centered representations involving perceptual reference frames (Palmer, 1983). Several representations based on local symmetry have been proposed in computer vision, such as the Symmetrical Axis Transform (Blum and Nagel, 1978), the Smoothed Local Symmetry Transform (Brady, 1983), and the Local Rotational Symmetry Transform (Fleck, 1986). Recently, Subirana-Vilanova (1990) proposed a method to extract curved axes of symmetry from an image, for the efficient description of object shapes.

Symmetry might also facilitate the inference of the 3D structure of symmetrical objects from their image projections. Kanade and Kender (1983) suggest a method that uses the amount of symmetry skewness of projected object contours. Nalwa (1989) uses local bilateral symmetry in line drawings, and its invariance under changes of viewpoint under orthographic projection, to infer the presence of surfaces of revolution. Ulupinar and Nevatia (1988) analyze the constraints of two kinds of symmetry, parallel and bilateral, under orthographic projections, to characterize zero-Gaussian curvature surfaces.

Symmetry has been also used for several other tasks, such as detection of interest points in an image (Reisfeld *et al.*, 1990), and efficient image encoding (Kumar *et al.*, 1983). It was also found to be useful in other cognitive tasks, such as evaluation of pattern interest, complexity and pleasingness (Day, 1968), numerosity judgments in patterns (Howe and Jung, 1987), and memory encoding (Attneave, 1955). Recently, Leyton (1992) proposed a general theory of human perception according to which symmetry plays a central role in the cognition of shape. Through a series of principles relating such notions as symmetry, curvature and morphogenesis, he showed that symmetry constitutes a fundamental element in shape-history recovery, and thus in shape memorization.

Symmetry has been also considered by many to be a fundamental grouping property of perceptual organization, as suggested by the Gestalt psychologists. Elements sharing symmetry relations tend to aggregate and be perceived as figure rather than ground. However, the domination of symmetry by other grouping properties, such as convexity and good continuation, demonstrates the weakness of the organizational power of this property. Recent studies (Jenkins, 1983; Pomerantz and Kubovy, 1986; Pashler, 1990) call into question the role of symmetry as a strong grouping property. They suggest rather that symmetry detection could operate among features obtained from other grouping properties, thus implying a precedence of grouping over symmetry perception.

In this paper, we propose to elucidate this question by studying the strategies used by the visual system to detect global bilateral symmetry in dense images. As Section 2.1 reveals, two separate mechanisms may exist for symmetry detection, one for simple shapes and another one for dense patterns (Julesz, 1971); here we concentrate on the latter. The problem of detecting global symmetry in dense images has not received much attention, and some interesting issues have not been studied; in particular the strategy followed by the visual system to integrate non-central, non-contiguous symmetry information in dense images. In Section 2.2, we present a psychophysical experiment that we conducted to investigate this issue. Results of our experiment, in agreement with recent studies (Jenkins, 1983; Pomerantz and Kubovy, 1986; Pashler, 1990), show that grouping facilitates and actually precedes symmetry detection.

We also propose a computational model for global symmetry detection in dense images. This model, elaborated on the basis of the psychophysical experiment reported in Section 2.2, is presented in detail in Section 3. The proposed model consists of three stages: a grouping stage, in which neighboring compatible elements are aggregated, a comparison stage, in which clusters formed in the grouping stage are systematically compared to detect local axes of symmetry, and a subsumption stage, in which local axes are merged into global ones. The implementation details of the model and results are reported in Section 4, showing that the model performance agrees in principle with human symmetry perception. The paper ends with a discussion of further issues that might be investigated.

2. HUMAN SYMMETRY PERCEPTION IN DENSE IMAGES

2.1. Previous work

Some aspects of human symmetry perception in the case of complex high-frequency patterns have been studied by Julesz (1966, 1971; see also Julesz and Chang, 1979), and in the case of moderately complex random-dot patterns (100 dots) by Barlow and Reeves (1979). Barlow and Reeves' results indicate that approximate symmetry can still convey the visual impression of symmetry, provided the deviation from exact symmetry is not too large. They suggest that the mechanism of symmetry detection is not a highly accurate one, and considers dots positioned at ± 6 arcmin of visual angle from the exact symmetrical position as contributing to the symmetry effect. The authors also studied the effects of introducing violations (i.e. suppression) of symmetry at different locations in the pattern, and of changing the location of the axis of symmetry. They found that violations are more easily detected near the axis of symmetry, and that symmetry perception becomes less sensitive when the location of the axis is not central. Jenkins (1982) also discusses the importance of the central region near the axis of symmetry. According to his conclusions, for completely symmetrical images, only the symmetry information located within a strip of 1.1 deg of visual angle centered on the axis of symmetry is exploited by the human visual system. Outside of this strip, symmetry information does not contribute to symmetry detection when the symmetrical portions are contiguous. All the preceding results

were obtained in experimental conditions involving brief exposure times (of the order of 100 ms). Other findings reported by Julesz (1971) indicate that bilateral symmetry is detected more quickly than rotational symmetry, and that vertical symmetry is detected faster than horizontal symmetry. Also, multiple symmetry (with respect to more than one axis at the same time) is detected faster than simple symmetry (one axis only). Zucker (1986) suggests the presence of two separate mechanisms in low-level vision, one for contour detection, the other for texture analysis, thus implying that symmetry detection in textures would be accomplished by a different process than the one used in contour images. Julesz (1971) also suggests the existence of two different mechanisms for symmetry perception.

Given that global symmetry relations often involve non-adjacent symmetrical regions, it would be of particular interest to determine how non-central, non-contiguous symmetry information is detected by the visual system. In such cases, the preceding results about the importance of the region near the axis of symmetry probably do not apply, since special grouping configurations along the axis of symmetry are not created with non-contiguous symmetry regions. Another important issue is the processing time required for symmetry detection. It might be that, for certain conditions, symmetry detection requires longer periods of time (of the order of 1 s). Even if the aforementioned experiments show that, for particular conditions, symmetry can be detected with brief exposure times, detection performance might be different for more typical vision situations, involving longer processing times.

2.2. Experiment with human symmetry detection

Bilateral symmetry can be defined as the reflection of a region about a given axis. In symmetry-detection studies with dense patterns (e.g. random dots), symmetrical images generally are created by reflecting one half of the image about a central vertical axis. Patterns generated according to this method usually contain salient grouping configurations along the central axis. An example of a figure containing such configurations can be seen in Fig. 1. Results of an experiment on symmetry detection using this kind of image indicate that subjects probably were using the presence of central coherent features as a clue to discriminate between symmetrical and asymmetrical patterns (Jenkins, 1983). This observation suggests that a preliminary cluster formation (grouping) stage might precede, or at least facilitate, symmetry detection.

In natural environments, symmetrical regions are not always contiguous, and therefore do not always contain central grouping configurations to facilitate symmetry detection. On the other hand, non-contiguous regions of symmetry often can be segregated from the background, and thus contain some form of grouping. It is reasonable to suppose that texture-related grouping could be used by the visual system to facilitate the symmetry-detection process for non-contiguous regions. In this section, we report the results of a psychophysical experiment that we conducted in order to investigate the effect of grouping on symmetry detection. The goal of the study was to establish the strategy followed by the visual system in order to elicit non-central, non-contiguous symmetry information from dense images, and to discover how the system deals with the computational load imposed by global symmetry detection.

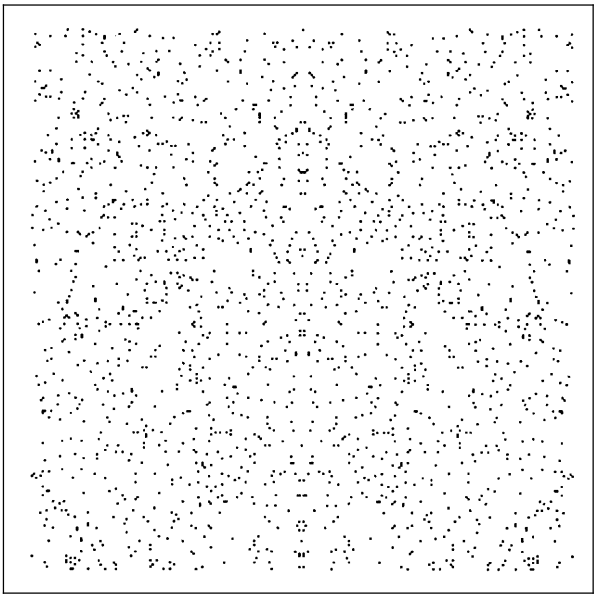

Figure 1. Dense random-dot image with grouping configurations along the central vertical axis of symmetry.

The experiment compared symmetry-detection performance when a significant amount of grouping was possible in the regions of symmetry by means of texture segregation (grouping condition), and when it was not (no-grouping condition), thus reducing the grouping possibilities. A reasonable amount of time (of the order of 1 s) was given to the subject in order to obtain results applicable to typical vision situations, as opposed to more constrained situations with brief exposure times (of the order of 100 ms). The stimuli consisted of dense images of oriented segments for which grouping was performed on the basis of the similarity of orientation of neighboring segments. Texture segregation, and thus grouping, was produced when segments inside a symmetry region had a different orientation from the background segments. Therefore, for the no-grouping condition, background segments and segments inside regions of symmetry had the same orientation, while for the grouping condition they had different orientations. The segment orientations in corresponding regions were always symmetrical and the subjects had to verify whether or not the position of the segments inside the regions were symmetrical.

The detection task was not simply a case of orientation discrimination of the texture segregation components; segment positions inside the regions of symmetry had to be compared. Detection performance was measured at three different distances from the central axis of symmetry. In the first case, the regions of symmetry were adjacent. In this situation, in addition to the controlled grouping/no-grouping factor, grouping configurations were created along the central axis of symmetry for the two conditions. Therefore, a significant difference in detection performance between the two conditions was not expected. In the second case, the regions of symmetry were positioned in the middle of each half image. In this situation, a significant difference in

detection performance was expected since only the controlled grouping/no-grouping factor was involved. For the third case, the regions of symmetry were located at the outer extremities of the image. In this situation, additional grouping effects might have been caused by segregation at the texture/non-texture borders of the image.

2.2.1. Method.

Stimuli. Computer-generated images were presented, on the center of a VGA monitor (640 by 480 pixels) measuring 42 cm by 31 cm, to subjects at a viewing distance of 57 cm from a fixation point placed in the middle of the screen. The mean luminance of the screen was 50 cd m^{-2}. An 80386 IBM computer running at 33 MHz with an 80387 math coprocessor was used to generate the images in real time. A chin-rest facilitated stabilization of head position and an I-SCAN system with infrared camera was used to monitor eye movements. Subjects' responses were recorded with the mouse of the computer: the left button corresponded to a *yes* response and the right button to a *no* response.

Images were made of 1000 white oriented segments on a black background. Image size was 26.4 by 26.4 deg of visual angle and the two regions inside each image were 6.6 by 13.2 deg. The three values for the distance between the regions were 0, 6.6, and 13.2 deg. Each stimulus could represent one of the two following types of images with one of the two types of grouping. Figure 2 shows examples of the stimuli that were used.

- *Symmetrical image:* the position and orientation of segments inside the regions were symmetrical about the central vertical axis of symmetry, and background segments were randomly located outside the regions.
- *Asymmetrical image:* the orientation but not position of segments inside the regions were symmetrical about the axis, and background segments were randomly located outside the regions.
- *Grouping condition:* the background segments were vertical and segments inside the regions were oriented at 45 or 135 deg, depending on which region of symmetry they belonged to.
- *Non-grouping condition:* the background segments and segments inside a region were both vertical.

Control. To control for possible orientation effects, another grouping condition, where background segments were either oriented at 45 or 135 deg, and segments inside the symmetrical region were vertical, was also used for the distance of 6.6 deg. Figure 3 shows an example of a stimulus that was used to control for orientation effects. It is not possible to perform the same type of control for the non-grouping condition since inclined segments (45 or 135 deg) cannot be used at the same time in the background and in the regions of symmetry without creating texture segregation.

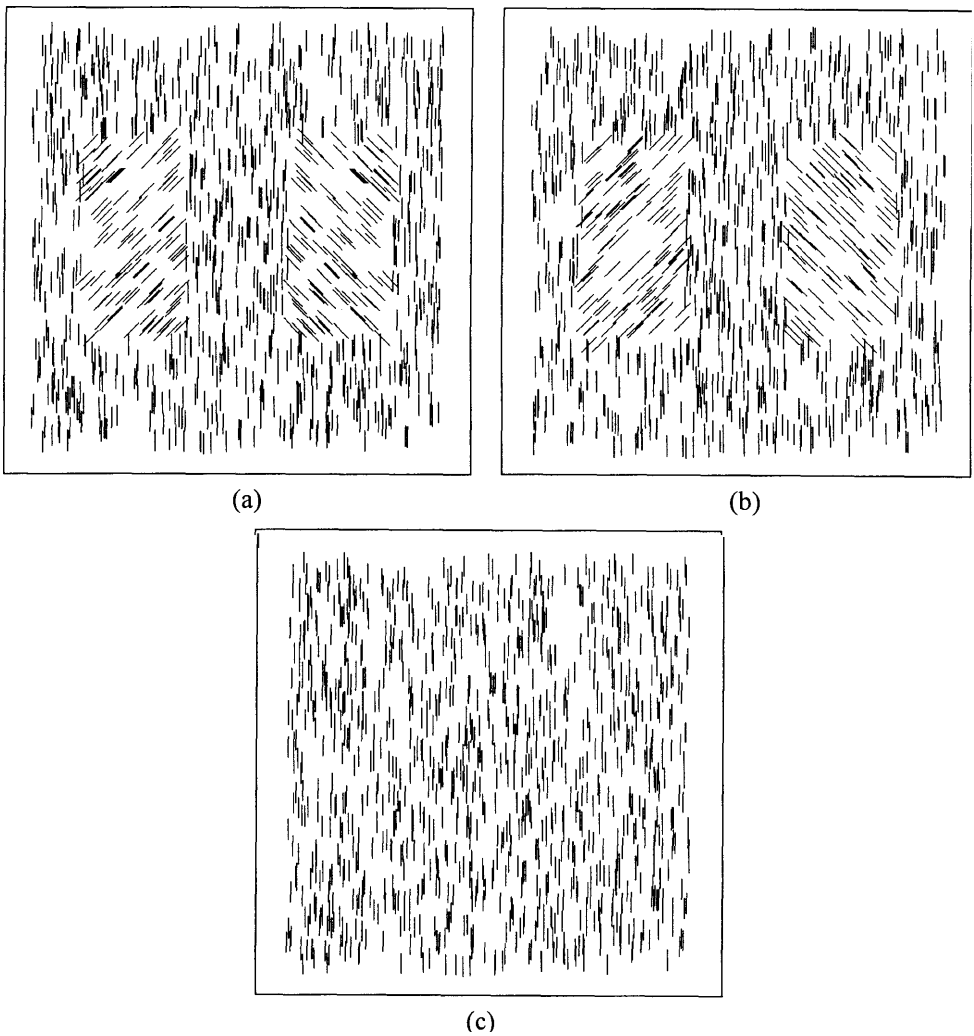

Figure 2. Examples of stimuli used: (a) symmetrical with grouping, (b) asymmetrical with grouping, and (c) symmetrical without grouping. The stimuli are shown with a distance of 6.6 deg of visual angle between the regions of symmetry.

Procedure. Five right-handed subjects with normal or corrected-to-normal vision were tested. They all agreed to perform the experiment without any remuneration. Subjects were students or staff members with the Perception and Robotics Laboratory at Ecole Polytechnique de Montréal or with Ecole d'Optométrie at Université de Montréal.

The experiment consisted in testing each subject with 7 independent blocks of 50 images: 3 for the grouping condition, with distances between the regions of 0, 6.6, and 13.2 deg respectively, 3 for the no-grouping condition, with the same distances as in the grouping condition, and 1 to control for orientation effects, with a distance of 6.6 deg (Fig. 3). In each block, the probabilities of presentation for symmetrical

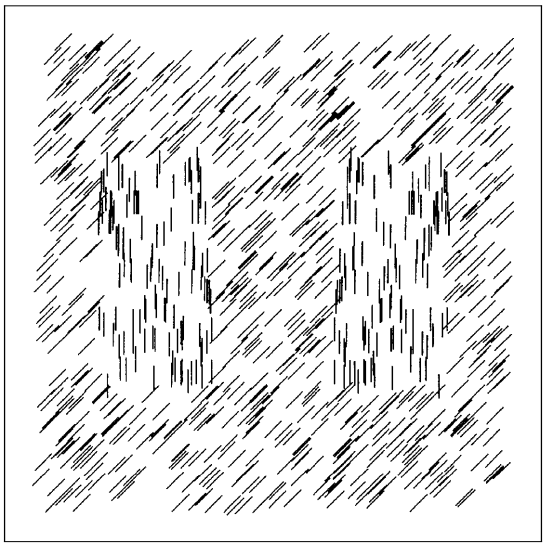

Figure 3. Example of stimuli used to control for orientation effects: symmetrical with grouping at a distance of 6.6 deg of visual angle.

and asymmetrical images were equal. The images were randomly generated but, for each block, all subjects were shown the same set of images in the same order. However, the order of presentation of blocks varied between subjects. This technique of presentation, for which the distance between the regions of symmetry is constant in each block, permits the elimination of position uncertainty since the subjects knew the exact position of the target regions at all times.

Subjects were asked to detect the presence or absence of symmetry in the images that were presented. They were told that the accuracy of their responses was important and that reaction times were not measured. They had to look at the fixation point in the middle of the screen and were told not to move their eyes. The presentation time of each image was 1.5 s. A new image was not shown until a response was recorded. The experimenter monitored the eye movements of the subjects and told them whenever they were moving their eyes. Before each block, example images were shown. During this presentation, the experimenter told the subjects whether or not the images were symmetrical. Subjects then performed practice trials during which they had to assess if the 10 presented images were symmetrical or not. Finally, for each block, subjects were tested with 50 images. It took, on average, 1 h and 10 min to perform the experiment (10 min per block).

Data analysis. The individual data at each distance, in terms of symmetry-detection rate as a function of false-positive response, are shown in Fig. 4. Figure 5 presents the averaged symmetry-detection rate of the 5 subjects as a function of grouping condition and distance. For the grouping condition, performance as a function of the distance between the regions of symmetry slightly decreases with increasing distance, while for the no-grouping condition, performance is V-shaped: high at a distance of 0 deg,

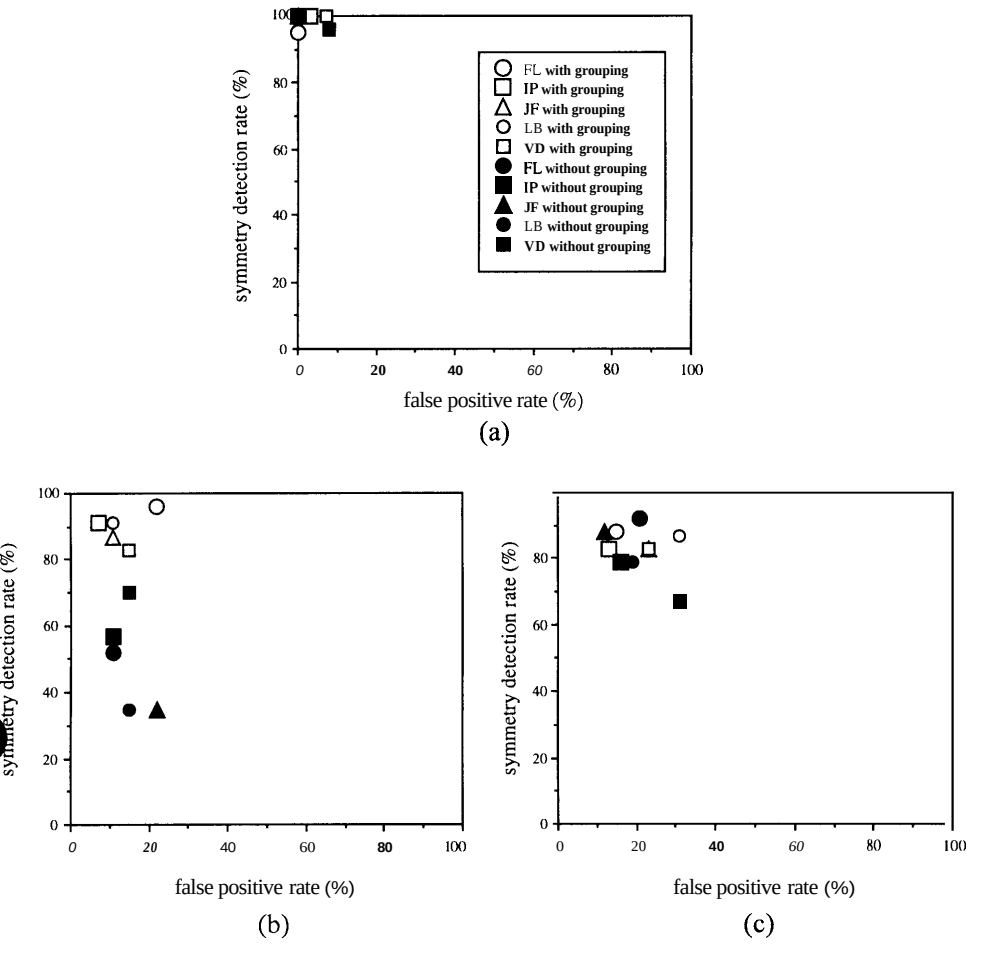

Figure 4. Comparison of symmetry detection rate between the grouping and the non-grouping conditions as a function of false positive detection rate at distances of 0, 6.6, and 13.2 deg of visual angle (a, b and c respectively).

significantly lower at 6.6 deg, and high again at 13.2 deg, consistent with the shape but not the range of Barlow and Reeves, (1979) data. Results of a two-factor ANOVA indicate that the effect of grouping is significant ($F(1, 24) = 26.35$, $p < 0.001$), the effect of distance is significant ($F(2, 24) = 36.40$, $p < 0.001$) and the interaction grouping-distance is significant ($F(2, 24) = 20.37$, $p < 0.001$). Results of a post-hoc Tukey test also show that, even if there is a monotonic decrease of symmetry detectability as a function of the distance in the grouping condition, this effect is not significant.

At a distance of 0 deg, performance is almost perfect for all subjects, regardless of the condition. At a distance of 6.6 deg, performance for the grouping condition is much better than for the non-grouping condition. This difference in detection cannot be attributed to a criterion shift because the false-positive rate remained fairly constant.

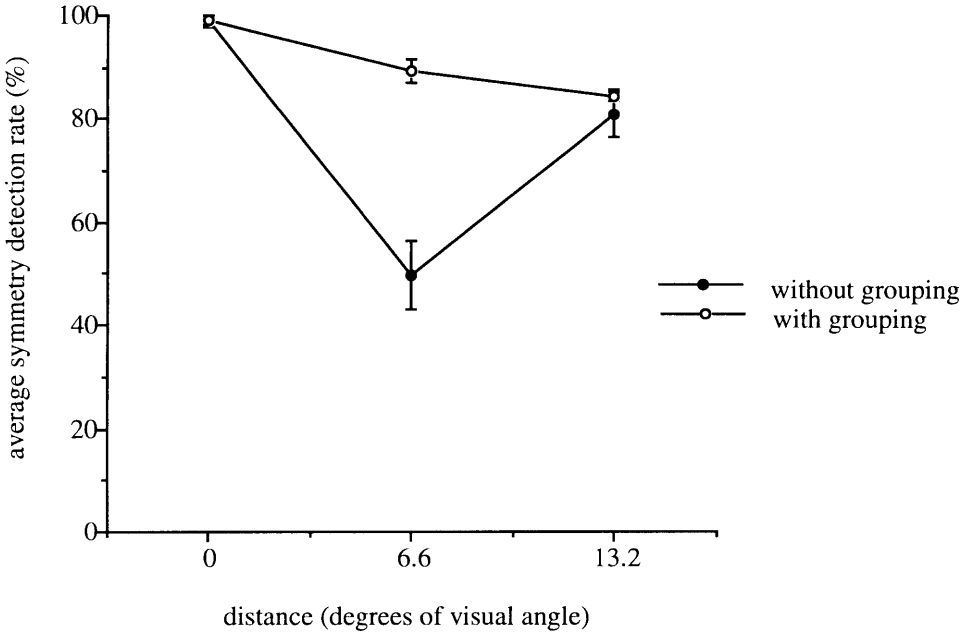

Figure 5. Averaged symmetry detection rate for the 5 subjects as a function of grouping conditions and of distance between regions of symmetry. Error bars represent ±1 SE.

At a distance of 13.2 deg, no major difference between the two conditions can be observed. Post-hoc comparisons between the conditions by a Tukey test indicate significant differences between the no-grouping condition at a distance of 6.6 deg and the five remaining conditions. A paired t-test comparison between the two grouping conditions at a distance of 6.6 deg (the control for orientation effects) shows no significant effect ($t(4) = 0.22$); thus orientation is not a factor.

2.2.2. Discussion. Symmetry-detection performance in dense images was measured as a function of the distance and of grouping of oriented segments. As expected, a significant advantage in detection performance for the grouping condition was observed at a distance of 6.6 deg, since only the controlled grouping factor was involved in the detection task at this distance. At a distance of 0 deg, the results show no significant difference in detection performance between the grouping and the non-grouping condition. As discussed at the beginning of Section 2.2, it has been suggested that central grouping configurations help symmetry detection. We suggest that, at a distance of 0 deg, regardless of the controlled grouping conditions, additional grouping configurations were already present along the central axis of symmetry. Thus, the facilitation effect of central grouping configurations was equally present for both conditions. At a distance of 13.2 deg, no significant difference in detection performance was observed. Barlow and Reeves (1979) mention that, in addition to the vivid impression created by the paired dots along the central axis of symmetry, the outline or the shape of the pattern is also important for symmetry detection. However, this explanation does not apply with the kind of stimuli that we used since the global shape of the image did

not vary. We propose rather that grouping effects were present at the outer borders of the image. These groupings, due to segregation at the texture/non-texture border, helped the detection of symmetry at the edges.

In summary, we propose that at distances of 0 and 13.2 deg, groupings were already present, either along the central axis of symmetry or along the image borders, regardless of the grouping condition. Therefore the introduction of another grouping factor was superfluous. On the other hand, at a distance of 6.6 deg, no texture-related grouping effect based on orientation was available. Therefore, the introduction of a grouping factor improved symmetry-detection performance. We propose, based on our results, that grouping plays a significant role for symmetry detection not only near the central axis of symmetry, but also for non-adjacent regions of symmetry. Alternative explanations cannot account for the obtained performance for the following reasons. First, grouping was not an important factor for the localization of the regions of symmetry since, for each block of stimuli images, regions of symmetry were at the same position. Second, orientation effects induced by groupings did not play a significant role in the detection task since subjects had also to verify the segments' positions in the regions of symmetry. Third, low spatial-frequency components, which are known to significantly affect symmetry perception, were not involved since the spatial-frequency content of the stimuli was similar for both conditions. Fourth, small groupings created by the proximity of neighboring segments did not significantly affect detection performance since they were present equally in both conditions. On the other hand, texture segregation, which was facilitated in the grouping condition, improved symmetry-detection performance.

Our findings have important theoretical implications for the elaboration of a model for symmetry detection. They imply that, in addition to the fact that symmetry is a fundamental grouping property, as suggested by Gestalt psychologists, grouping could be used to detect symmetry. If Julesz (1971) is correct that there are two different processes for symmetry detection it could be that, in the case of dense images, general-purpose grouping mechanisms already used for other visual tasks perform the first stage of the symmetry detection process. Symmetry detection would then be a comparison of higher-order features obtained from grouping mechanisms. This hypothesis agrees with recent studies. (1) Owing to the weak organizational power of symmetry, other grouping factors could precede it. These factors would dominate symmetry, and symmetry would be used for grouping only when it is not in competition with them. Symmetry could be useful not at the level of grouping formation, but at a higher level where it would be detected among groupings created on the basis of other factors (Pomerantz and Kubovy, 1986). (2) Grouping mechanisms already used for other visual functions could perform the most important part of the symmetry detection task (Pashler, 1990). (3) Symmetry detection in random-dot images could be accomplished in three stages, including a grouping stage (Jenkins, 1983). (4) Symmetry detection could be accomplished by general grouping mechanisms, and therefore, no special symmetry-detection mechanism needs to be postulated (Wagemans *et al.*, 1991).

In the next section, we propose a computational model for symmetry detection in dense images. Our model agrees in principle with the results presented in this section

since symmetry detection is preceded by a grouping stage which permits a significant reduction in the number of comparisons to be performed. Before presenting our model, the possible strategies for symmetry detection in images are reviewed.

3. COMPUTATIONAL MODELS

3.1. The possible strategies for symmetry detection

A number of computational models for symmetry detection have been proposed, either to model the human visual system or to provide efficient means of detecting symmetry in images. Most of the proposed models for the human visual system are not detailed enough to elaborate algorithms that could be simulated and tested on a computer. For instance, Palmer and Hemenway (1978) suggest that symmetry is detected in two stages. The first stage consists of a fast and global analysis to select a potential axis of symmetry by considering all the orientations of the axis at the same time. The second stage consists of an explicit comparison of the two halves of the stimuli to detect symmetry. From a computational point of view, simulation of the first stage of this model is not a trivial task.

Two principal strategies are used by most of the proposed methods for symmetry detection in images. The first one necessitates a systematic comparison of all the potentially symmetrical elements in the image (usually the individual pixels). Bigün (1988) illustrates this approach by presenting a convolution operator which works on the brightness values, and detects various situations of local symmetry in neighborhoods of circular, linear, hyperbolic and parabolic shapes. Reisfeld *et al.* (1990) propose, in a similar strategy, a method for the detection of various interest points in an image. In an attempt to model the human visual system, Royer (1981) proposes that symmetry relations present in a stimulus are represented integrally in a code consisting of classes and subclasses of symmetry. Two channels sensitive to vertical and horizontal orientations are used. By computing many transformations in parallel on the output of these channels, the type of symmetry present in an image can be identified. Palmer (1983) also proposes a model for the human visual system. The visual field is covered by a multitude of local spatial analyzers working in parallel. Invariance under Euclidean similarity transformations applied to the output of these analyzers indicates symmetry relations.

The second strategy is accomplished in two stages. First, more elaborate (and less dense) features than brightness values are extracted. Second, these features are systematically compared, in the search for symmetry relations. Wilson (1991), for example, presents a system to detect local symmetry among contours extracted from an image by comparing key-points on each contour (local extrema of curvature, middle points of straight sections and points of change in curvature sign). Jenkins (1983) proposes, as a plausible model for the human visual system, a three-stage model for symmetry detection in random-dot images. First, orientation uniformity of pairs of points is detected. Second, salient pairs are fused to form a new representation that corresponds to virtual lines joining the two points of each pair. Third, from this

new representation, symmetry is detected. The orientation of the axis of symmetry is determined by joining the middle points of the virtual lines. Wagemans *et al.* (1991) suggest that, in addition to the first-order regularities consisting of orientation uniformity and middle point collinearity, second-order relations between pairs of symmetrical elements are used to detect symmetry. These second-order regularities correspond to geometric regularities (symmetric trapezoid and parallelogram connectivities) between virtual lines joining pairs of points. To justify their proposal, they note that the same first-order relations are present for bilateral symmetry and skewed symmetry. Since skewed symmetry is much more difficult to detect than bilateral symmetry, other factors (second-order regularities) must affect symmetry detection.

In the case of global symmetry detection in dense images, the first strategy is inadequate because of the excessive amount of computation required. Since global symmetry involves aggregates of local elements sharing compatible symmetry relations, it seems more appropriate to use, in accordance with the second strategy, a procedure that does not look immediately for instances of symmetry among local elements, but tries first to elicit aggregates of elements that constitute plausible candidates for global symmetry relations. Symmetry is then tested among these aggregates only. The advantages of such a procedure are clear: by making grouping a prerequisite to symmetry detection, a substantial reduction in computational cost is achieved and more global levels of symmetry relations may be achieved.

3.2. The proposed model

In view of the experimental results reported in Section 2.2 illustrating the precedence and the facilitation of grouping over symmetry detection, we propose a strategy for bilateral symmetry detection in dense images that consists of three successive stages:

1. *a grouping stage* in which clusters are formed among local elements presenting a sufficient level of mutual affinity;
2. *a symmetry-detection stage* in which pairs of symmetrical clusters are discovered and their axes of symmetry determined;
3. *a symmetry-subsumption stage* in which an attempt is made to detect even more global symmetries by comparing the various axes of symmetry previously found.

The strategy is computationally efficient and suitable for parallel implementation because only local computations are performed in the first stage, and the subsequent stages necessitate only a small number of comparisons (with respect to a systematic comparison of every element with all other elements) of higher-level features. The first stage involves local computations to evaluate the affinity between each local element and its various neighbors (relying on such properties as proximity, collinearity, and similarity of orientation), and the formation of clusters of elements by relaxation labeling. The second stage involves systematic comparisons between pairs of clusters in order to discover symmetrical pairs and compute the positions and orientations of the corresponding axes of symmetry. Clusters are compared on the basis of their compatibilities in orientation and dimensions. At this stage, the number of comparisons should be limited due to the small number of clusters. The third stage involves

a fusion of the axes of symmetry previously determined. Each axis being characterized by its location and orientation, a Hough-transform type of procedure (Duda and Hart, 1973) is used to subsume the pairs of symmetrical clusters into more global symmetrical structures. It is important to observe that the plausibility of our model for the human visual system relies on experimental results for the first stage only. We do not claim that the human visual system necessarily performs the same computations as the ones we propose in our model. However, our model is plausible at a computational level: a first stage implying massive parallel computations performed by a general-purpose mechanism is followed by more specific subsequent stages that perform computations on a smaller number of higher-level features.

Before describing our model in more detail, it is worth mentioning that our model is in agreement with a recent research trend questioning the role of symmetry as a strong property of perceptual organization and suggesting rather that symmetry is detected among more abstract features obtained from general purpose grouping mechanisms (see Sections 2.2 and 3.1). Our model fits into this scheme, since the cluster-detection process can be seen as a general-purpose grouping mechanism useful for other visual tasks, and the cluster comparison process as a tool more specific to symmetry detection, which takes advantage of the small number of higher-level features that were created. It is also important to note that the validity of the model does not depend on the density of the features in the image. In fact, the term *dense images* is used to refer to texture-like images consisting of a large number of similar local features.

4. IMPLEMENTATIONAND RESULTS

To illustrate the relevance and the performance of the above model we present an implementation for the particular case of dense images consisting of oriented line segments. Such images are common as intermediate representations in several fundamental vision mechanisms such as shape from shading, shape from texture, shape from stereo or shape from motion. It is furthermore believed that, in such cases, global symmetry may prove to be a powerful tool of scene interpretation.

4.1. The grouping stage

We use here a relaxation-labeling procedure in order to coalesce clusters of mutually compatible oriented segments. First proposed by Rosenfeld *et al.* (1976), relaxation labeling is a parallel, local, cooperative and iterative process that assigns labels to objects, together with weights indicating the levels of confidence of each association between an object and a label.

Let $\mathbf{A} = \{a_1, a_2, \ldots, a,\}$ denote the set of oriented segments to be labeled and $\mathbf{A;} = \{\lambda^0, \lambda^1\}$ the set of possible labels for segment a_i, where $\lambda_i = \lambda^0$ indicates that a_i belongs to a cluster and $\lambda_i = \mathrm{A}'$ that it does not. Let Λ_{ij} designate the set of compatible pairs of labels for each pair of segments (ai,a_j), $i \neq j$. The four possible pairs represent the three following situations: a_i and a_j both belong to a cluster

($\lambda_i = \lambda_j = \lambda^0$), one of a_i or a_j belongs to a cluster but not the other ($\lambda_i \neq \lambda_j$), and neither a_i nor a_j belong to a cluster ($\lambda_i = \lambda_j = \lambda^1$).

Let $r_{ij}(\lambda_i, \lambda_j)$ designate the compatibility function between λ_i and λ_j assigned to the neighboring segments a_i and a_j. The neighbors of segment a_i are all the Voronoi neighbors a_j that fall within an elliptic region centered on a_i and having the same orientation as a_i. The use of an elliptic restricting region is meant to favor Voronoi neighbors that are positioned longitudinally with respect to the current segment. A possible expression for the compatibility function might be:

$$r_{ij}(\lambda_i, \lambda_j) = \begin{cases} \dfrac{1}{1 + e^{K(|\Delta\theta| - W)}} & \lambda_i = \lambda_j = \lambda^0, \\ 1 - \dfrac{1}{1 + e^{K(|\Delta\theta| - W)}} & \text{otherwise,} \end{cases} \tag{1}$$

where $\Delta\theta$ designates the difference in orientation between segments a_i and a_j, K is a positive constant, and W represents the maximal orientation difference producing a mutual affinity.

According to the above expression, segments with a small orientation difference mutually encourage each other to belong to the same cluster, while those with a large orientation difference mutually discourage each other. Consequently, in a situation such as the one of Fig. 6, the two adjacent clusters would be merged into one single cluster. Such an undesirable effect may be avoided by introducing border elements into the labeling process, in order to separate adjacent clusters.

Let $B = \{b_{1,2}, b_{1,3}, \ldots, b_{i,j}, \ldots, b_{n,n-1}\}$ be the border elements defining cluster boundaries. A border element b_{ij} is defined between each pair of segments (a_i, a_j), $i \neq j$. The role of border elements is not to delineate exactly each cluster; it is

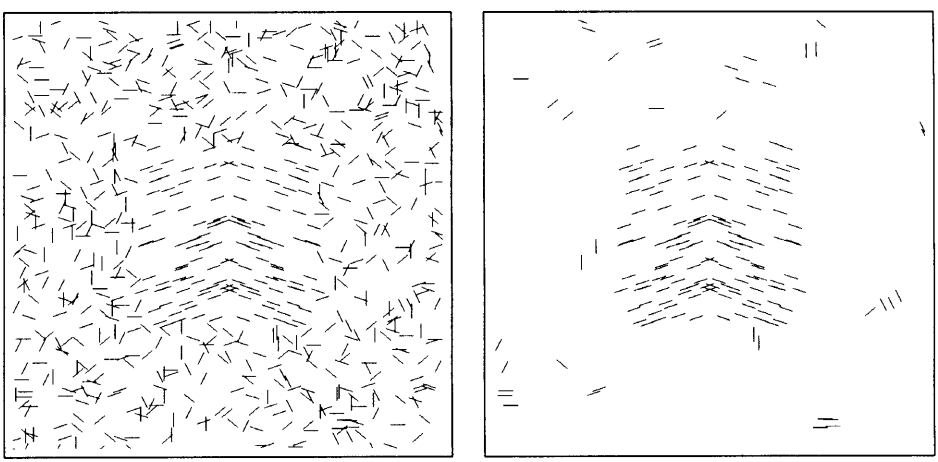

(a) (b)

Figure 6. (a) Original image and (b) result of the grouping process (without border elements).

rather to create a separation between segments belonging to different clusters. Let $\Omega_{ij} = \{\omega^0, \omega^1\}$ be the set of possible labels for each border element b_{ij}, ($\omega_{ij} = \omega^0$ and $\omega_{ij} = \omega^1$ indicate, respectively, the absence and the presence of a border element between segments a_i and a_j). The product $\Lambda_{ij} \times \Omega_{ij}$ represents the ensemble of compatible label triplets for each pair of segments (a_i, a_j), $i \neq j$. Five of the eight possible triplets represent the four following situations (the three other triplets represent incoherent situations): (1) a_i and a_j belong to the same cluster and there is no border element between them ($\omega_{ij} = \omega^0, \lambda_i = hj = \lambda^0$); *(2)* a_i and a_j belong to different clusters and there is a border element between them ($\omega_{ij} = \omega^1, hi = hj = \lambda^0$); *(3)* a_i or a_j belongs to a cluster while the other one does not and there is a border element between them ($\omega_{ij} = \omega^1, \lambda_i \neq \lambda_j$); (4) neither a_i nor a_j belong to a cluster and there is no border element between them ($\omega_{ij} = \omega^0, \lambda_i = hj = \lambda^1$).

Let $s_{ij}(\omega_{ij}, hi, \lambda_j)$ represent the new compatibility function characterizing the associations of ω_{ij} to b_{ij}, λ_i to a_i and λ_j to a_j. We propose the following expression for $s_{ij}(\omega_{ij}, hi, \lambda_j)$:

$$s_{ij} = \begin{cases} \dfrac{1}{1 + e^{K(|\Delta\theta| - W)}} & \omega_{ij} = \omega^0,\ \lambda_i = \lambda_j = \lambda^0 \\ 1 - \dfrac{1}{1 + e^{K(|\Delta\theta| - W)}} & \omega_{ij} = \omega^1,\ \lambda_i = \lambda_j = \lambda^0 \\ 0 & \omega_{ij} = \omega^0 \text{ with } \lambda_i \neq \lambda_j \\ & \omega_{ij} = \omega^1,\ \lambda_i = \lambda_j = \lambda^1 \\ 1 & \omega_{ij} = \omega^1 \text{ with } \lambda_i \neq \lambda_j \\ & \omega_{ij} = \omega^0,\ \lambda_i = \lambda_j = \lambda^1. \end{cases} \quad (2)$$

Neighboring elements with a small orientation difference mutually encourage each other to belong to the same cluster and discourage the appearance of a border element between them ($\omega_{ij} = \omega^0$, $\lambda_i = hj = \lambda^0$). Neighboring elements with a large orientation difference, which are each encouraged by their other neighbors to belong to a cluster, favor the appearance of a border element between them ($\omega_{ij} = \omega^1$ and $\lambda_i = \lambda_j = \lambda^0$). Whenever a_i is encouraged by its neighbors to belong to a cluster, and a_j is discouraged by its neighbors from belonging to a cluster (or the opposite), they favor the appearance of a border element between them ($\lambda_i \neq hj$). When neighboring elements are both discouraged by their neighbors from belonging to a cluster, they discourage the appearance of a border element between them ($hi = hj = \lambda^1$).

Let $p_i^k(\lambda)$ be the probability of associating λ with a_i, and $u_{ij}^k(\omega)$ be the probability of associating ω with b_{ij} at the k^{th} iteration. These probabilities are changed in a parallel and iterative fashion according to the following expressions:

$$p_i^{k+1}(\lambda) = \frac{p_i^k(\lambda) q_i^k(\lambda)}{\sum_\lambda p_i^k(\lambda) q_i^k(\lambda)}, \quad \text{where } q_i^k(\lambda) = \sum_j \sum_{\lambda'} r_{ij}(\lambda, \lambda') p_j^k(\lambda') u_{ij}^k(\omega^0); \quad (3)$$

$$u_{ij}^{k+1}(\omega) = \frac{u_{ij}^k(\omega) t_{ij}^k(\omega)}{\sum_\omega u_{ij}^k(\omega) t_{ij}^k(\omega)}, \quad \text{where } t_{ij}^k(\omega) = \sum_\lambda \sum_{\lambda'} p_i(\lambda) p_j(\lambda') s_{ij}(\omega_{ij}, \lambda, \lambda'). \quad (4)$$

The expression for $q_i^k(\lambda)$ includes the term $u_{ij}^k(\omega^0)$, which represents the probability of absence of a border element. If this term was not present, two neighboring segments a_i and a_j with a small orientation difference would mutually encourage each other to belong to the same cluster. However, two neighboring elements a_i and a_j with a large orientation difference, which are each encouraged by their other neighbors to belong to different clusters, would mutually discourage each other from belonging to a cluster. This discouraging effect is undesirable and it would be preferable that each segment be encouraged to belong to its own cluster. Introducing the term $u_{ij}^k(\omega^0)$ has precisely the desired effect: when a_i and a_j have a small orientation difference, $u_{ij}^k(\omega^0)$ is high and its presence does not modify the probability values. However, when a_i and a_j have a large orientation difference, $u_{ij}^k(\omega^0)$ is low and attenuates the undesirable discouraging effect.

Figure 7 illustrates the results of the grouping process by the relaxation-labeling method, in the three following cases: two non-touching isolated clusters, two adjacent clusters, and two clusters with a smooth variation of orientation. The results are shown after 10 iterations, starting from a uniform label probability distribution ($p^0(\lambda^0) = p^0(\lambda^1) = 1/2$). In this experiment, K and W were set to $0.05\ \text{deg}^{-1}$ and 20 deg respectively, and the label probability threshold was set to 0.9. As one can observe, the clusters are properly separated for the three cases. It is important to note that the role of border elements is not to delineate exactly the shape of the clusters but rather to create a separation between them. Even if some isolated lines also show up, they will be eliminated when clusters are identified.

The result of relaxation labeling is an ensemble of segments and an ensemble of border elements, from which clusters must be identified. Elements a_i and a_j belong to the same cluster when they are neighbors, when $p_i(\lambda_i = \text{A})$ and $p_j(\lambda_j = \text{A})$ are above the label probability threshold, and $u_{ij}(\omega^1)$ is under this threshold. All pairs of segments verifying these three conditions belong to the same cluster. Clusters that are too small (less than 5 segments) are eliminated.

4.2. The symmetry-detection stage

To detect instances of symmetry, the clusters resulting from the grouping stage have to be compared. Computational efficiency imposes the requirement to base the comparisons on a representation of each cluster in terms of a limited number of global parameters. In the case of clusters consisting of similarly oriented line segments, these parameters can be the size of each cluster, as well as the mean orientation and the mean position of its segments. However, clusters with smooth variations of orientations, such as the ones in Fig. 7f, cannot be represented by the mean orientation of their segments. Consequently, one must subdivide such clusters into smaller components with reduced orientation variations, while retaining their common identity for further restoration. Clusters are then systematically compared on the basis of location, orientation and size, in order to detect eventual symmetrical pairs.

4.2.1. Cluster subdivision. The eventual subdivision of clusters into subcomponents of small orientation variations is performed by means of a relaxation-labeling

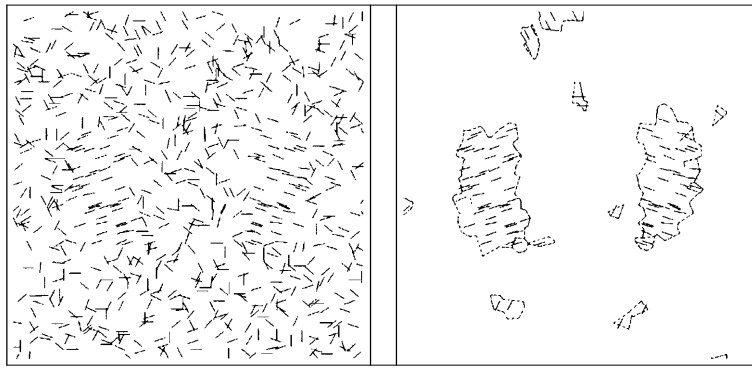

(a) (b)

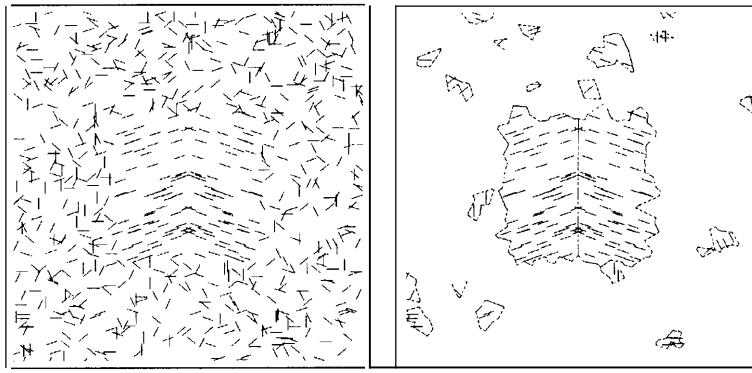

(c) (d)

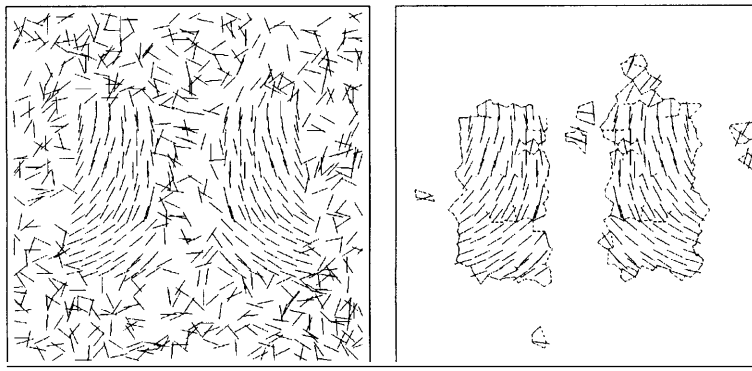

(e) (f)

Figure 7. Original images (a, c, e), and the corresponding results of the grouping process (with border elements) (b, d, f).

method similar to the one presented in Section 4.1. Let $\mathbf{A} = \{a_1, a2, \ldots, a,\}$ now designate only the line segments belonging to the cluster under consideration, and $\Lambda_i = \{\lambda^0, \ldots, \lambda^k\}$ be a set of labels representing $\boldsymbol{k}$ successive orientation intervals of width $\boldsymbol{a} = \boldsymbol{n}/\boldsymbol{k}$ between 0 and π (we choose here $\boldsymbol{k} = 12$). At the end of the relaxation procedure, each label will eventually characterize a component whose mean orientation falls within the range of orientations it represents. The compatibility function between neighboring segments a_i and a_j is defined as:

$$r_{ij}(\lambda_i, \lambda_j) = \begin{cases} \dfrac{1}{1 + e^{K(|\Delta\theta| - W)}} & \lambda_i = \lambda_j, \\ 1 - \dfrac{1}{1 + e^{K(|\Delta\theta| - W)}} & \text{otherwise.} \end{cases} \tag{5}$$

The parameters are the same as in Section 4.1. For each line segment, the initial probability distribution is a normal distribution centered on the label corresponding to the segment orientation, with a small standard deviation. Figure 8 illustrates the results of the cluster-subdivision process applied to Fig. 7f, after 20 iterations. As desired, the original clusters with smooth variations of orientations have been subdivided into components with smaller orientation variations. Figures 7b and 7d are not modified by the cluster-subdivision process.

4.2.2. Cluster comparison. Each cluster is represented by its size as well as the mean position and the mean orientation of its segments. The size information is made of two measurements, namely the largest distances between two line segments in the mean orientation direction and in the direction normal to the mean orientation.

Clusters are systematically compared by searching for pairs of clusters of comparable sizes, whose orientations are identical under a reflection transformation around the medial axis of the line joining the cluster centers. Since the strategy is meant to detect approximate symmetry, exactness is not imposed for the preceding relations. Orientations and sizes need to correspond within certain tolerances. A reasonable tolerance on the orientations is the parameter W used in the compatibility function, since it represents the maximal amount of orientation difference characterizing the affinity between line segments. The tolerance used for size comparisons is twice the length of the line segments.

4.3. The symmetry-subsumption stage

Having compared all clusters or subcomponents, the last stage consists of trying to elicit global symmetry relations involving several pairs of clusters or components. In order to do so, one must simply compare the axes of symmetry established during the symmetry-detection stage, grouping collinear axes and the corresponding cluster pairs. A Hough-transform technique is used for that purpose (Duda and Hart, 1973), each axis being parameterized according to its distance r to a fixed point and its orientation Θ. The resolution of the Hough transform accumulator is chosen in order to allow sufficient tolerances on r and Θ (in our experiments 1/30 of the image size and 0.1 rad, respectively).

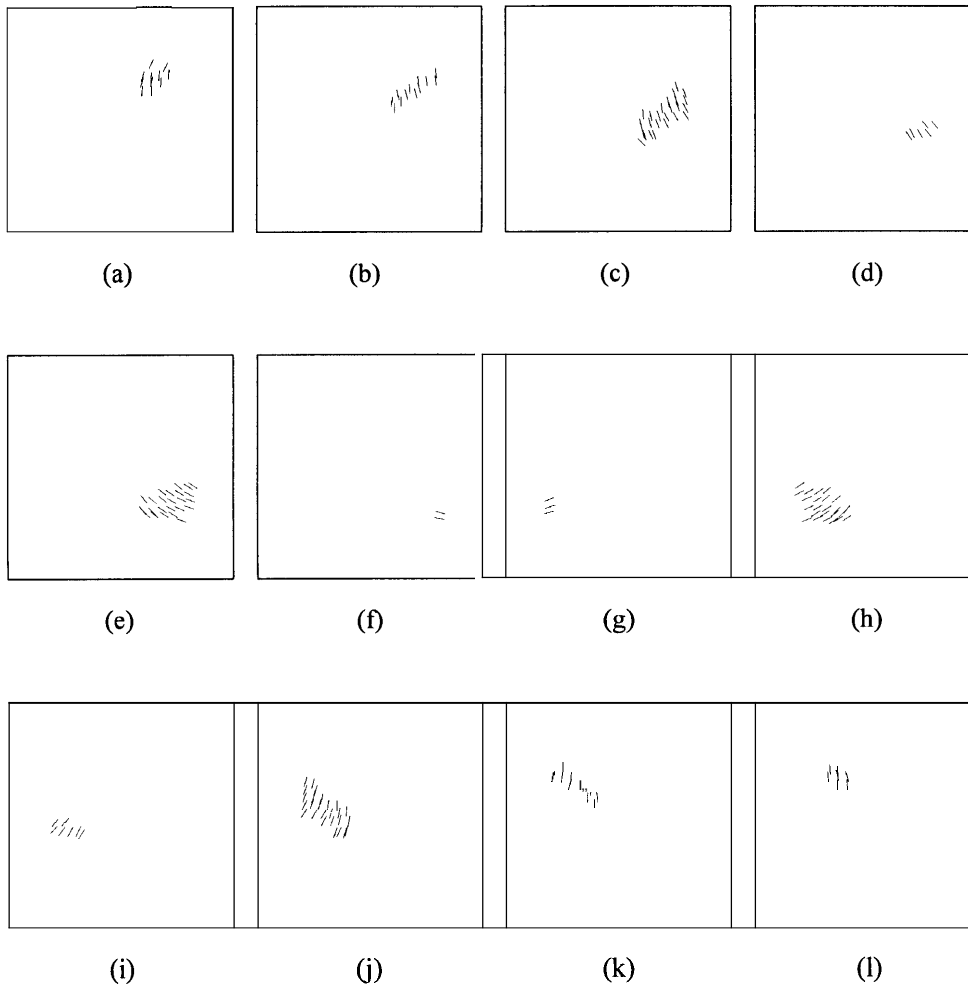

Figure 8. Results of the cluster subdivision process applied to the right cluster (a–f) and to the left cluster (g–l) of Fig. 7f.

Figure 9 illustrates the final results of the method. As one can observe, symmetry was detected in the case of Figs 7a, 7c, and 7e. Figure 9d illustrates an image containing two symmetrical regions, in which grouping cannot be achieved. Symmetry was not detected in this image, a fact which is in agreement with human symmetry perception.

5. CONCLUSION

Detecting symmetry in images is interesting because of the many possible roles that this property could play. One controversial issue concerns the role of symmetry as a grouping property. Gestalt psychologists suggest that symmetry is a fundamental

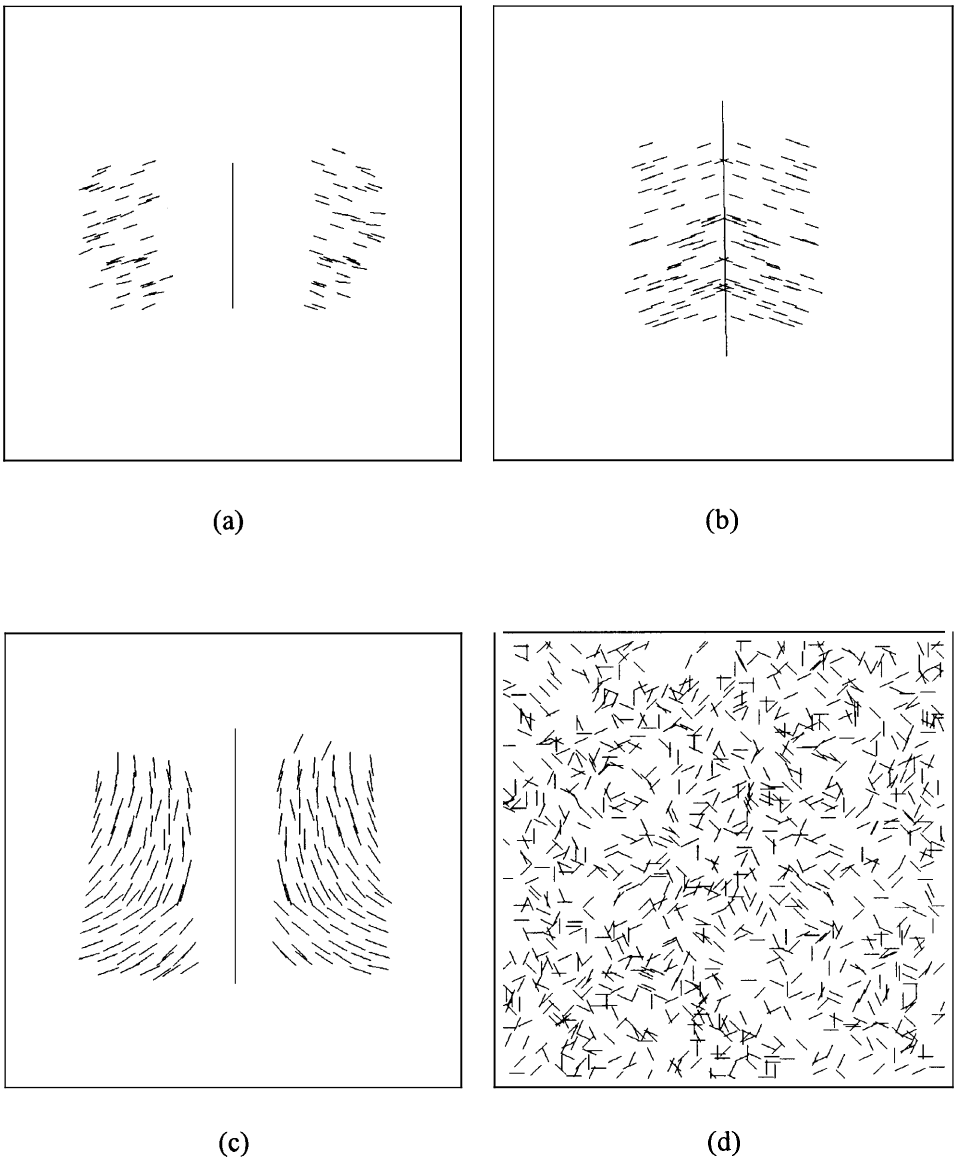

Figure 9. Final results of the method applied to images 7a, 7c and 7e respectively (a, b, c); (d) an image containing symmetrical regions in which grouping cannot be achieved and for which symmetry was not detected.

grouping property of perceptual organization, while recent studies suggest rather that symmetry might be detected among higher level features, created on the basis of other grouping properties.

We have reported here a psychophysical experiment conducted in order to investigate this question. Results of the experiment show that the task of detecting symmetry is significantly easier when the elements in the stimuli can be grouped together according

to properties such as similarity of orientation and proximity, and thus support the hypothesis that grouping precedes symmetry detection.

As in the case of other visual tasks (e.g. motion perception, shape perception, object recognition), more than a single mechanism may underly the ability of the human visual system to detect symmetry. Julesz (1971) has suggested two mechanisms for symmetry detection, one underlying detection in dense stimuli and the other suited for larger-scale stimuli.

We have proposed a three-stage computational model for the detection of symmetry in field-type dense images. In the first stage, relaxation labeling is used to group primitive elements of similar properties into a small number of clusters. This grouping stage may also be useful for other visual tasks. Pairs of clusters are then compared in the second stage of the model in order to establish local symmetry axes. These symmetry axes are then aggregated in the third stage (using Hough transform) to form global symmetry relations.

The model is in good agreement with human symmetry perception in dense images, in the sense that it deals not only with exact symmetry, but also with approximate symmetry. The reliance upon a preliminary grouping stage, consisting of local, cooperative computations, enables a significant reduction of the computational load entailed for the process as a whole (compared to a more direct approach to the problem).

The model has been presented and implemented in the context of texture images, providing satisfactory results. Other possible domains of relevance are gradient fields, optical flow fields or $2\frac{1}{2}$D representations. Indeed, interesting questions relate to the use of symmetry properties in the segmentation and interpretation of motion fields, or in the extraction of lighting conditions and shape information from needle diagrams produced by shape from shading algorithms. In the context of human symmetry perception, further studies are also required to investigate the relationship between grouping and symmetry detection, and the effect grouping may have on the complexity of the processes involved.

Acknowledgement

This work was supported by the Canadian Program of Networks of Centers of Excellence (IRIS Network, project A-4). JF is supported by NSERC operating grant OGP0121333 and FCAR grant 93-NC-0903.

REFERENCES

Attneave, F. (1954). Some informational aspects of visual perception. *Psychol. Rev.* **61,** 183–193.

Attneave, F. (1955). Symmetry, information and memory for patterns. *Am. J. Psychol.* **68,** 209–222.

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

Bigün, J. (1988). Recognition of local symmetries in gray value images by harmonic functions. In: *Proc. 9th Int. Conf. on Pattern Recognit.* pp. 345–347.

Blum, H. and Nagel, R. N. (1978). Shape description using weighted symmetric axis features. *Pattern Recognit.* **10,** 167–180.

Brady, M. (1983). Criteria for representation of shape. In: *Human and Machine Vision.* J. Beck, B. Hope and A. Rosenfeld (Eds). Academic Press, New York.

Day, H. (1968). The importance of symmetry and complexity in the evaluation of complexity, interest, and pleasingness. *Psychonomic Sci.* **10,** 339–340.

Duda, R. O. and Hart, P. E. (1973). *Pattern Classification and Scene Analysis.* Wiley, New York.

Fleck, M. M. (1986). Local rotational symmetries. In: *Proc. Con. Vision and Pattern Recognit.* pp. 332–337.

Howe, E. and Jung, K. (1987). Judgement of numerosity: Effects of symmetry and goodness in dot pattern arrays. *Acta Psychologica 64,* 3–11.

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

Jenkins, B. (1983). Component processes in the perception of bilaterally symmetric dot textures. *Percept. Psychophys. 34,* 433–440.

Julesz, B. (1966). Binocular disappearance of monocular symmetry. *Science* **153,** 657–658.

Julesz, B. (1971). *Foundations of Cyclopean Perception.* University of Chicago Press, Chicago.

Julesz, B. and Chang, J. J. (1979). Symmetry perception and spatial-frequency channels. *Perception* **8,** 711–718.

Kanade, T. and Kender, J. R. (1983). Mapping image properties into shape constraints: Skewed symmetry, affine-transformable patterns, and the shape-from-texture paradigm. In: *Human and Machine Vision.* J. Beck, B. Hope and A. Rosenfeld (Eds). Academic Press, New York.

Kumar, N., Mukherjee, A. K. and Chakraborti, N. B. (1983). On using symmetry properties for selecting transform components for image coding. *IEEE Trans. Acoustics, Speech, and Signal Processing* **31,** 749–752.

Leyton, M. (1992). *Symmetry, Causality, Mind.* MIT Press, Cambridge, Massachusetts.

Mach, E. (1906/1959). *The Analysis of Sensations and the Relation of the Physical to the Psychical.* Dover, New York. (Original publication in 1906).

Nalwa, V. S. (1989). Line-drawing interpretation: Bilateral symmetry. *IEEE Trans. Pattern Anal. and Machine Intell.* **11,** 104–108.

Palmer, S. E. (1983). The psychology of perceptual organization: A transformational approach. In: *Human and Machine Vision.* J. Beck, B. Hope and A. Rosenfeld (Eds). Academic Press, New York.

Palmer, S. E. and Hemenway, K. (1978). Orientation and symmetry: Effects of multiple, rotational, and near symmetries. *J. Exp. Psychol: Human Percept. Perform.* **4,** 691–702.

Pashler, H. (1990). Coordinate frame for symmetry detection and object recognition. *J. Exp. Psychol: Human Percept. Perform.* **16,** 150–163.

Pomerantz, J. R. and Kubovy, M. (1986). Theoretical approaches to perceptual organization. In: *Handbook of Perception and Human Performance.* K. R. Boff, L. Kaufman and J.P. Thomas (Eds). Wiley, New York.

Reisfeld, D., Wolfson, H. and Yeshurun, Y. (1990). Detection of interest points using symmetry. In: *Proc. 3rd Int. Conf. Computer Vision.* pp. 62–65.

Rosenfeld, A., Hummel, R. and Zucker, S. (1976). Scene labeling by relaxation operations. *IEEE Trans. Systems, Man, and Cybernet.* **6,** 420–433.

Royer, F. L. (1981). Detection of symmetry. *J. Exp. Psychol: Human Percept. Perform.* **7,** 1186–1210.

Subirana-Vilanova, J. B. (1990). Curved inertia frames and the skeleton sketch: Finding salient frames of reference. In: *Proc. 3rd Int. Conf Computer Vision.* pp. 702–708.

Ulupinar, F. and Nevatia, R. (1988). Using symmetry for analysis of shape from contour. In: *Proc. 2nd Int. Conf Computer Vision.* pp. 414–426.

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

Wilson, S. (1991). *Perceptual Organization and Symmetry in Visual Object Recognition.* Masters Thesis, University of British Columbia, Canada.

Zucker, S. W. (1986). Early processes for orientation selection and grouping. In: *From Pixels to Predicates.* A. P. Pentland (Ed.). Ablex Publishing Corporation, Norwood, New Jersey.