International Journal of Computer Vision, 14, 119–130 (1995)

Context-Free Attentional Operators: The Generalized Symmetry Transform

DANIEL REISFELD, HAIM WOLFSON AND YEHEZKEL YESHURUN
Computer Science Department, Sackler Faculty of Exact Sciences, Tel Aviv University

Received December 10, 1992; Revised February 10, 1994.

Abstract. Active vision systems, and especially foveated vision systems, depend on efficient attentional mechanisms. We propose that machine visual attention should consist of both high-level, context-dependent components, and low-level, context free components. As a basis for the context-free component, we present an attention operator based on the intuitive notion of symmetry, which generalized many of the existing methods of detecting regions of interest. It is a low-level operator that can be applied successfully without a priori knowledge of the world. The resulting *symmetry edge map* can be applied in various low, intermediate- and high- level tasks, such as extraction of interest points, grouping, and object recognition. In particular, we have implemented an algorithm that locates interest points in real time, and can be incorporated in active and purposive vision systems. The results agree with some psychophysical findings concerning symmetry as well as evidence concerning selection of fixation points. We demonstrate the performance of the transform on natural, cluttered images.

1 Introduction

Biological vision is foveated, highly goal oriented, and task dependent. This observation, which is rather clear if we trace the behavior of practically every vertebrate, is now being taken seriously into consideration by the computer vision community. This is evident from recent work on active vision systems and heads (Brunnstrome et al. 1992; Crowley 1991; Rimey and Brown 1992) and general active vision concepts and algorithms (Aloimonos et al. 1987; Bajcsy 1988; Aboot and Ahuja 1988; Ballard 1990; Culhane and Tsotsos 1992). One of the fundamental features of active vision is the use of space-variant vision and sensors (Yeshurun and Schwartz 1989; Tistarelli and Sadini 1990; Rojer and Schwartz 1990), that allow, the case of the log-polar representation, data reduction as well as a certain degree of size and rotation invariance.

The use of such sensors requires efficient mechanisms for gaze control, which are, in turn, directed by attentional algorithms. Using psychophysical terms, these algorithms are either *overt*, analyzing in detail the central foveated area, or *covert*, analyzing various regions within the field of view that are not necessarily in the central foveated area.

Like many other issues in computational vision, the attention problem seems to be trapped in the typical top-down–bottom-up cycle, as well as in the global local cycle: global processes are necessarily based on local features and processes, whose crucial parameters, in turn, depend on global estimates. This is the case for recognition tasks, where, for example, thresholds and size tuning of local feature detectors are optimally determined by the model of the object of the system expects. In curve and edge detection, local discontinuities are classified as signals or as noise according to global matching based on these very local estimates (Zucker et al. 1989). Similarly, attention is undoubtedly a concurrent top-down *and* bottom-up process: computational resources are assigned to regions of interest. But detection of regions of interest is both *context dependent* (top down), since the system is task oriented, and *context free* (bottom up), since one of the most important aspects of such a system is detection of unexpected signals. Thus, attention must be based on highly coupled low-level and high-level processes. While we do not offer a solution to this fundamental problem, we propose here to begin this cycle with a low-level attentional mechanism.

Visual processes, in general, and attentional mechanisms, in particular, seem effort less for humans. This introspection, however, is misleading. Psychophysical experiments show that infants (age 1–2 months) tend for fixate around an arbitrary single distinctive feature of the stimulus, like the corner of a triangle. (Haith et al. 1977; Salapatek and Kessen 1973). Moreover, when

presented with line drawings, children up to age 3–4 spend most of their time dwelling only on the internal details of a figure, and in general, children make more eye movements and are less likely than adults to look directly at a matching target in their first eye movements (Cohen 1981).

In comparison, adults display a strong tendency to look directly at forms that are informative, unusual, or of particular functional value (Antes 1974; Loftus and Mackworth 1978). Thus, it seems that gaze control in adults is indeed task- and context-dependent, but it is probably based on natal (hardwired) low-level local and context-free attentional mechanisms. At first, only the low-level context-free mechanisms are available. Gradually, as more information regarding the environment is being learned, higher-level processes take their place.

Active vision definitely needs high-level, context-dependent attentional algorithms; but these should be adaptive, trainable algorithms based on acquired knowledge, that use lower-level context-free attentional modules. Considering the fact that this research area is rather new, we believe that robust and efficient low-level attentional algorithms are the basic building blocks for machine visual attention. Here, we present an attentional algorithm, that is context-free, does not depend on object segmentation, and seems to generalize many of the existing cues used in previous computer vision models for detecting regions of interest. While some of the motivation for our work is based on psychophysical observations, and our results are also in good agreement with psychophysical data relating overt attention, we do not present here our work as a model for overt and covert human attention.

2 Attentional Mechanisms in Computer Vision

An early attentional operator based on grey-level variance (Moravec 1977) is still being widely used in many systems. Other researchers suggested to measure "*busyness*"—the smoothed absolute value of the Laplacian of the data, rapid changes in the gray levels, and edge junctions (Brunnstrome et al. 1992). Following early psychophysical findings (Attneave 1954; Cohen 1981), interest points can be regarded also as points of high curvature of the edge map (Lamdan et al. 1988; Yeshurun and Schwartz 1989).

Instead of selecting one of these methods for each application, we show that they can all be derived from a more generalized concept, and suggest that local generalized symmetry is the principle underlying context-free attention. In addition, the operator we propose more closely fits psychophysical evidence, and it is more useful in detecting interesting features in complex scenes.

Natural and artificial objects often give rise to the human sensation of symmetry. Our sense of symmetry is so strong that most man-made objects are symmetric, and the Gestalt school considered symmetry as a fundamental principle of perception. Looking around us, we get the immediate impression that practically every interesting area consists of a qualitative and generalized form of symmetry.

In computer vision research, symmetry has been suggested as one of the fundamental non-accidental properties, which should guide higher-level processes. This sensation of symmetry is more general than the strict mathematical notion. For instance, a picture of a human face is considered highly symmetric by the layman, although there is no strict reflectional symmetry between both sides of the face.

Symmetry is being widely used in computer vision (Davis 1977; Nevatia and Binford 1977; Blum and Nagel 1978; Brady and Asada 1984; Atallah 1985; Bigun 1988; Marola 1989; Xia 1989; Zabrodsky et al. 1992). However, it is mainly used as a means of convenient shape representation, characterization, shape simplification, or approximation of objects that have been already segmented from the background. A schematic (and simplified) vision task consists of edge detection, followed by segmentation, followed by recognition. A symmetry transform is usually applied after the segmentation stage. We present a symmetry transform which is inspired by the intuitive notion of symmetry and assigns a *symmetry magnitude* and a *symmetry orientation* to every pixel in an image at a low-level vision stage which follows edge detection. Specifically, we compute a *symmetry map*, which is a new kind of an edge map, where the magnitude and orientation of an edge depends on the symmetry associated with the pixel. Strong symmetry edges are natural interest points, while linked lines are symmetry axes.

Since our symmetry transform can be applied immediately after the stage of edge detection, it can be used to direct higher-level processes, such as a segmentation and recognition, and can serve as a guide for locating objects. We shall now define the transform, and demonstrate these ideas on natural images.

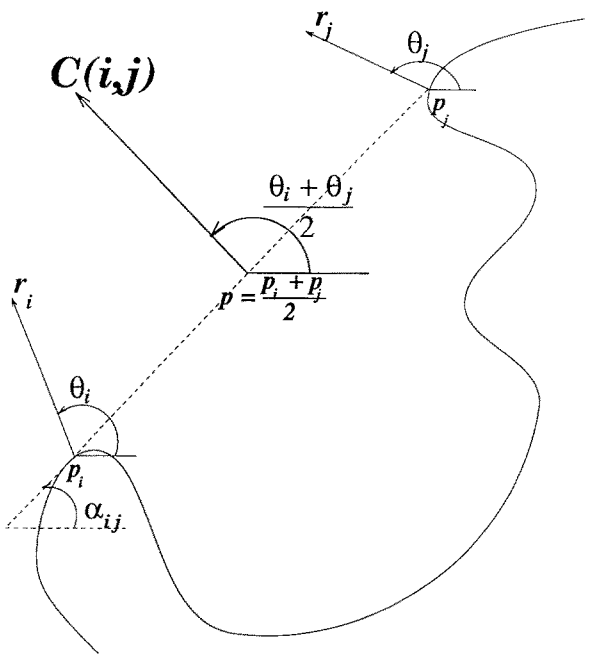

Fig. 1. The contribution to symmetry of the gradients at p_i and p_j.

3 Definition of the Transform

In the usual mathematical notion, an object is regarded symmetric if it is invariant to the application of certain transformation, called symmetry operations. A typical symmetry operation is the well-known reflectional (mirror) symmetry. In order to use these symmetry operations it is necessary to know the shape of an object before we can estimate whether it is symmetric or not. However, the process of finding interest points must precede complex processes of detecting the objects in the scene. Even if the object's shapes are known, truly symmetric objects are rare in natural scenes, and therefore any attempt to formulate an interest operator based on the strict mathematical notion of symmetry is doomed to fail.

Our symmetry transform does not require knowledge of the object's shape. It performs local operations on the edges of the image. Moreover, it assigns a continuous symmetry measure to each point in the image, rather than a binary symmetry label.

We first define a symmetry measure of each point. Let $p_k = (x_k, y_k)$ be any point $(k = 1, \ldots, K)$, and denote by

$$\nabla_{p_k} = \left(\frac{\partial}{\partial_x} p_k, \frac{\partial}{\partial_y}\right)$$

the gradient of the intensity at point p_k. We assume that a vector $v_k = (r_k, \theta_k)$ is associated with each p_k such that $r_k = \log(1 + \|\nabla p_k\|)$ and $\theta_k = \arctan(\frac{\partial}{\partial_y} p_k / \frac{\partial}{\partial_x} p_k)$. For each two points p_i and p_j, we denote by l the line passing through them, and by α_{ij} the angle counterclockwise between l and the horizon. We define the set $\Gamma(p)$, a distance weight function $D_\sigma(i, j)$, and a phase weight function $P(i, j)$ as

$$\Gamma(p) = \left\{(i, j) \middle| \frac{p_i + p_j}{2} = p\right\}$$

$$D_\sigma(i, j) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{\|p_i - p_j\|}{2\sigma}\right)$$

$$P(i, j) = [1 - \cos(\theta_i + \theta_j - 2\alpha_{ij})] \times [1 - \cos(\theta_i - \theta_j)]$$

We define the contribution of the points p_i and p_j as

$$C(i, j) = D_\sigma(i, j) P(i, j) r_i r_j$$

This measure can be easily normalized, and reflects the fact that each of its components modulates the other ones. The *symmetry magnitude* or *isotropic symmetry* $M_\sigma(p)$ of each point p is defined as

$$M_\sigma(p) = \sum_{(i,j)\in\Gamma(p)} C(i, j)$$

which averages the symmetry value over all orientations. We define the direction of the contribution of p_i and p_j as

$$\varphi(i, j) = \frac{\theta_i + \theta_j}{2}$$

A complementary definition for $\varphi(i, j)$ is given in Section 6.

The *symmetry direction* is defined as $\phi(p) = \varphi(i, j)$ such that $C(i, j)$ is maximal for $(i, j) \in \Gamma(p)$. Thus, the symmetry of the point p is defined as

$$S_\sigma(p) = [M_\sigma(p), \phi(p)]$$

The demand that the symmetry transform be local is reflected by the Gaussian distance weight function, $D_\sigma(i, j)$. Different values for σ imply different scales, thus enabling convenient implementation of multiresolution schemes. Note that the Gaussian defined above has circular isotherms, i.e., it has no preferred orientation. However, one can also define Gaussians with elliptic isotherms. This is useful when the transform is applied as a feature detector of elliptic regions such as eyes in human faces (Reisfeld and Yeshurun 1992; Edelman et al. 1992). In the experimental results

Fig. 2. Two opposite situations with the same symmetry value.

presented in this article we have used only circular Gaussians.

The phase weight function, $p(i, j)$ is composed of two terms. The first term, $1 - \cos(\theta_i + \theta_j - 2\alpha_{ij})$, allows to achieve maximum symmetry when $(\theta_i - \alpha_{ij}) + (\theta_j - \alpha_{ij}) = \pi$, i.e. when the gradients at p_i and p_j are oriented in the same direction toward each other. This is consistent with the intuitive notion of symmetry. This expression decreases continuously as the situation deviates from the ideal one. Notice that the same measure is achieved when the gradients are oriented toward each other or against each other. The first situation corresponds to symmetry within a dark object on a light background, and the second corresponds to symmetry within a light object on a dark background. It is easy to distinguish between the two cases. (See Fig. 2). One can simultaneously process both types of the transform obtaining two classes of interest points. However, for the simplicity of exposition, in all the experimental results presented in this article, we consider only the case of gradients facing each other (which correspond mainly to dark objects on a brighter background). We emphasize that in general, both a dark and a bright object can be simultaneously detected. Obviously, the amount of false alarms will increase, but the basic idea behind any interest operator is dimensionality reduction: in the case of a 1000 by 1000 image, computational resources should be directed at, say, 10^2 locations, instead of 10^6. Thus, the increase in potential regions of interest caused by including bright and dark symmetry peaks is not significant in our context.

The second term of $P(i, j)$, $1 - \cos(\theta_i - \theta_j)$, is introduced because the first term attains its maximum whenever $(\theta_i - \alpha_{ij}) + (\theta_j - \alpha_{ij}) = \pi$. This includes the case $\theta_i - \alpha_{ij} = \theta_j - \alpha_{ij} = \pi/2$, which occurs on a straight edge, and which we do not regard as interesting. The current expression compensates for this situation.

The term $r_i r_j$ is high when there is a strong correlation between two large gradients. We use gradients rather than intensities since we are mainly interested in edges that relate to object borders. For instance a uniform-intensity wall is highly symmetric but probably not very interesting. In natural scenes, we prefer to use the logarithm of magnitude instead of the magnitude itself, since it reduces the differences between high gradients, and therefore the correlation measure is less sensitive to very strong edges.

Note that the transform we have defined detects reflectional symmetry. It is invariant under 2-D rotation and translation transformation. Moreover, it is quite effective in detecting skewed symmetry (Kanade and Kender 1983) as well. Since skewed symmetry results in an affine transformation of the 2-D picture, the midpoint of a segment and parallelism are conserved. Thus the location of the symmetry edge is preserved under skewed symmetry, but the direction of the symmetry edge should not be the average of the two directions. In Section 6 we give the precise derivation of the correct direction. One should note, however, that since the weight of the direction prefers parallel edges, even the above defined direction will be still effective under skewed symmetry.

Sometimes it is necessary to detect points that are highly symmetric in multiple distinct orientations rather then in a principal one. We define such a symmetry as *radial symmetry* $-RS(p)$ and its value can be evaluated using the formula

$$RS_\sigma(p) = \sum_{(i,j)\in\Gamma(p)} C(i, j) \sin^2[\varphi(i, j) - \phi(p)]$$

This expression emphasizes contribution in directions perpendicular to the main symmetry direction, and attains its maximum in a point that is surrounded by edges. Notice that due to the continuous nature of the operator, the radial symmetry is not sensitive to gaps in the contour that surrounds the point p, and does not require this contour to be uninterrupted. A preliminary definition of this symmetry form was suggested in (Reisfeld et al. 1990).

Although the symmetry transform is presented here as a context-free mechanism, it can obviously use additional knowledge, if available. For example, if we look for a symmetry in a given direction, we define the *symmetry projection* at a point p and an orientation ψ as

$$PS_\sigma(p, \psi) = M_\sigma(p) \cos[\phi(p) - \psi]$$

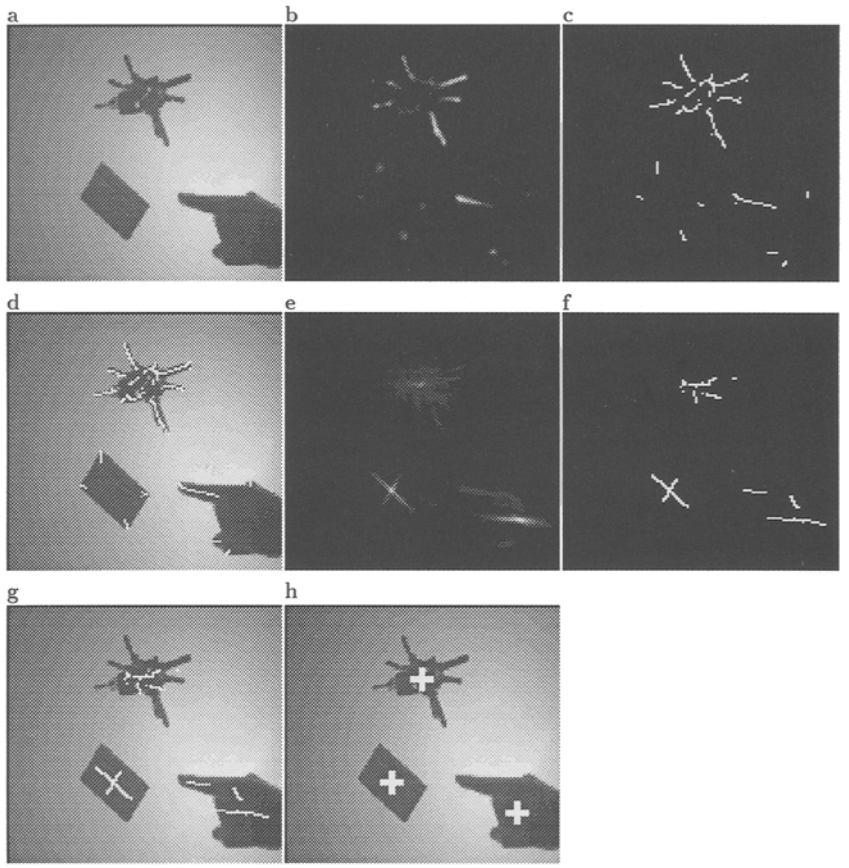

Fig. 3. A credit card, keys, and a hand. (a) Original image; (b) Isotropic symmetry map for small σ; (c) Nonmaximal suppression of (b) followed by thresholding. (d) Superposition of (a) and (c); (e) Isotropic symmetry map for large σ. Nonmaximal suppression of (e) followed by thresholding. (g) Superposition of (a) and (f); (h) Maximal values of the radial symmetry with large σ marked by crosses on the original image.

The symmetry transform can thus be turned into an efficient feature detector. We have recently demonstrated this in Reisfeld et al. (1990), where the transform was used for detection of facial features.

The complexity of our algorithm is relatively low. We first discuss serial implementation. Suppose the operator is applied to a picture composed on n pixels, and the Gaussians in the weight function almost vanish in radius $r < \sqrt{n}$. Each two pixels, whose distance is less than or equal to $2r$, contribute a value which can be computed using look-up tables. Therefore the time complexity is $O(nr^2)$.

Space complexity is small too. The results of the operator are the symmetry magnitude and symmetry orientation maps, each occupies n pixels. In addition, a temporary space for holding the maximal response is needed. For the radial symmetry, another n-pixels map is needed.

The complexity of a parallel implementation depends on the architecture. For an architecture where there is a processor allocated for each pixel and each processor is connected to its neighbors up to radius r, we can achieve maximal speed-up of the algorithm and reduce time complexity to $O(r^2)$ [with $O(nr^2)$ messages if no global memory is available]. If we have at our disposal a sufficiently large neural network, we can easily implement the look-up table and summation operations needed and then perform the operation in constant time. A different realization by "cortical-like" elements would be closely related to the co-circularity detectors suggested by Zucker et al. (1989).

4 Operation on Natural Images

The symmetry transform can be applied successfully on intricate natural scenes. It should be emphasized that all the images in Figs. 4–8 have been processed

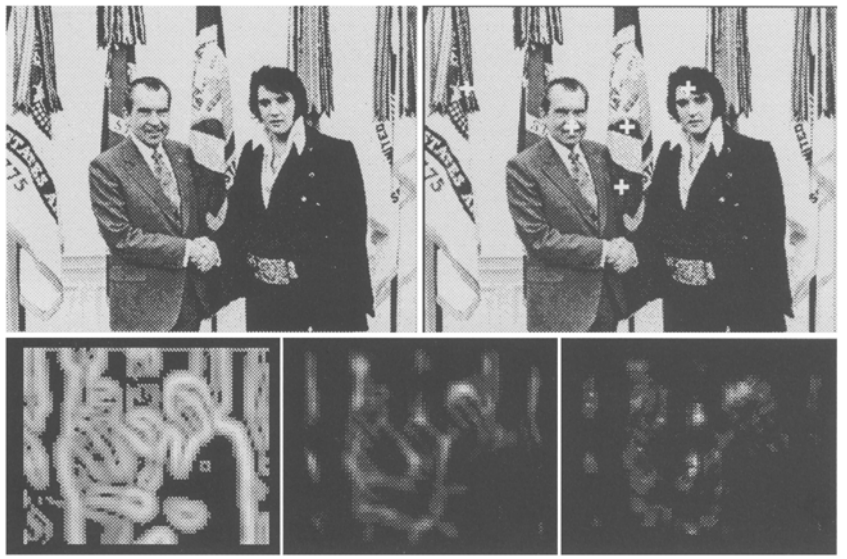

Fig. 4. Two persons with flags. Top: Original image (left) and the peaks of the radial symmetry marked by crosses. Bottom: Edge detection (left), isotropic symmetry (middle), radial symmetry (right).

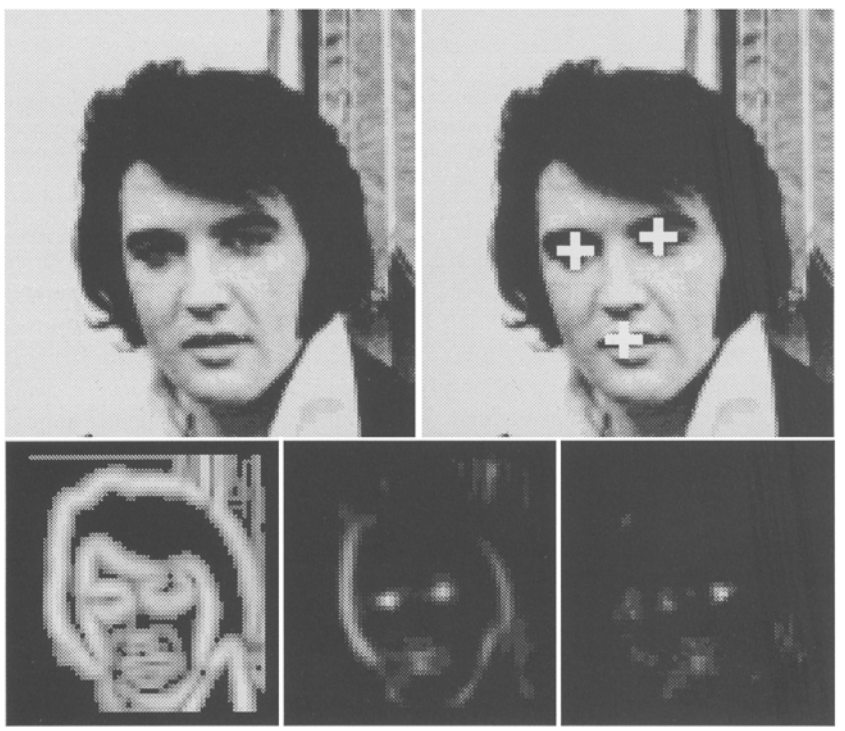

Fig. 5. Further processing the previous figure. Top: Area around the highest radial-symmetry peak in finer resolution and the peaks of the radial symmetry. Bottom: Edge detection (left), isotropic symmetry (middle), and radial symmetry (right).

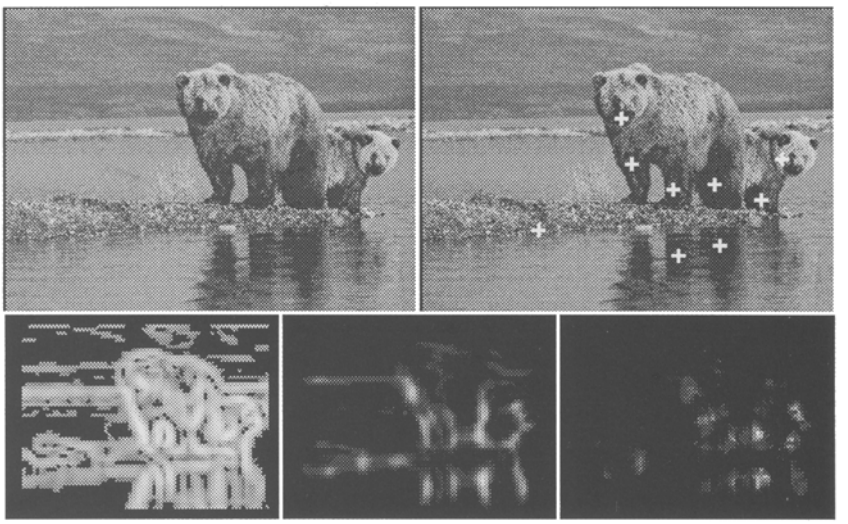

Fig. 6. Two bears with reflections. Top: Original image and the peaks of the radial symmetry. Bottom (left to right): Edge detection, isotropic symmetry, and radial symmetry.

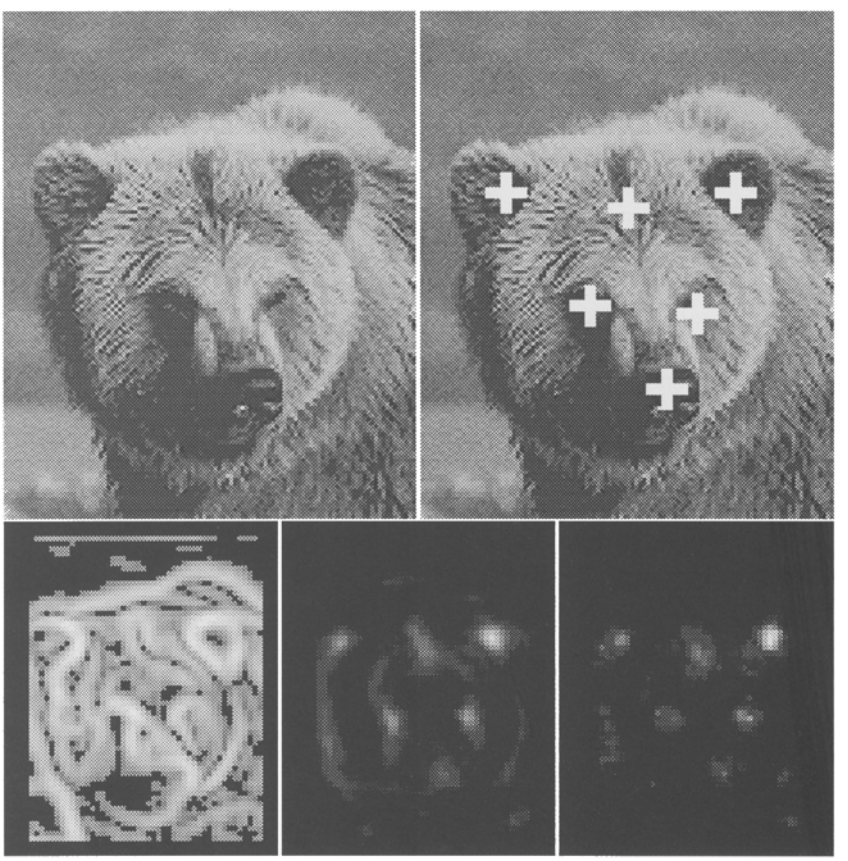

Fig. 7. Further processing the previous figure. Top: Area around the highest radial symmetry peak in finer resolution and the peaks of the radial symmetry marked by crosses. Bottom (left to right): Edge detection, isotropic symmetry, and radial symmetry.

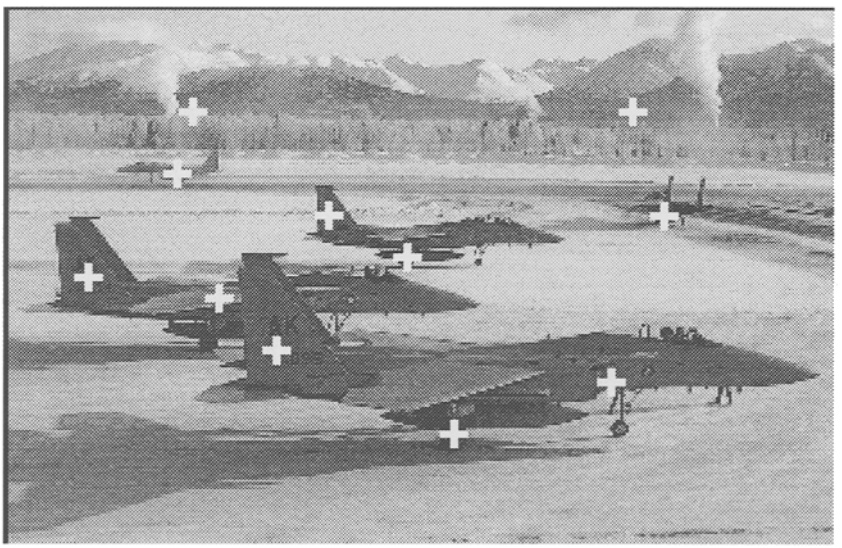

Fig. 8. The peaks of the radial symmetry marked on an image of F-15a Eagles.

using the same parameters of the symmetry transform. Figure 3 demonstrates that the transform performance is not affect by the existence of several objects in the scene. One can also see that linking the strong symmetry edges results in natural symmetry axes of the parts involved. Notice that the credit card is a skewed rectangle.

Figure 4 demonstrate the processing of an image of two persons on a noisy background. Note that the faces of the persons are marked among the highest-symmetry peaks. Since the transform is designed to respond best to symmetric edges, this is not surprising. However, experimenting with many natural images, we have observed that, in spite of the context-free, low-level nature of the transform, it detects regions that humans would consider interesting.

Zooming in on Fig. 4 at the vicinity of the highest radial-symmetry peak, we demonstrate (Fig. 5) the results of applying the transform with exact the same parameters. The location of the strongest peaks of the radial symmetry transform are marked by crosses and are located on the eyes and mouth. This figure can also demonstrate the superiority of the symmetry transform over the direct search for dark blobs: such a search would mainly detect the hair, while we are interested in more intricated areas, which are designated by the local symmetric configuration of edges.

Notice that the active-vision task typically involves exactly these steps: detecting an interest point, zooming in (which is equivalent to modifying the scale parameter), and using the same transform to detect regions of interest at the next level. This is also the stage where higher processes could use the low-level modules for feature and object detection. In this regard, it might be useful for a higher-level process to analyze in detail the scale-space patterns of the symmetry edge maps that are formed while zooming.

Figures 6 and 7 have been processed in the same manner as Figs. 4 and 5 with exactly the same parameters. Notice that this image includes many high-frequency edges (e.g., on the fur), yet the transform responds in a similar manner to Fig. 4, detecting the interesting objects on the coarse level, and the facial features on the fine level. This phenomenon should be further investigated, however, by analyzing the appropriate scale space.

Again, processing Fig. 8 with the same symmetry parameters as before, one can see the effectiveness of the symmetry transform in a different scene. Due to the continuous nature of the generalized symmetry transform, it gracefully degrades; and this is evident especially in cases of partial occlusion of symmetric objects, as can also be seen in this figure. Notice also that the same features are detected for the three foremost planes, though their size varies considerably.

Since symmetry is a powerful nonaccidental property, significant information can be extracted from the image, based on the symmetry map. In particular, one can attempt direct segmentation of an image based on a backward transform which detects the image edges

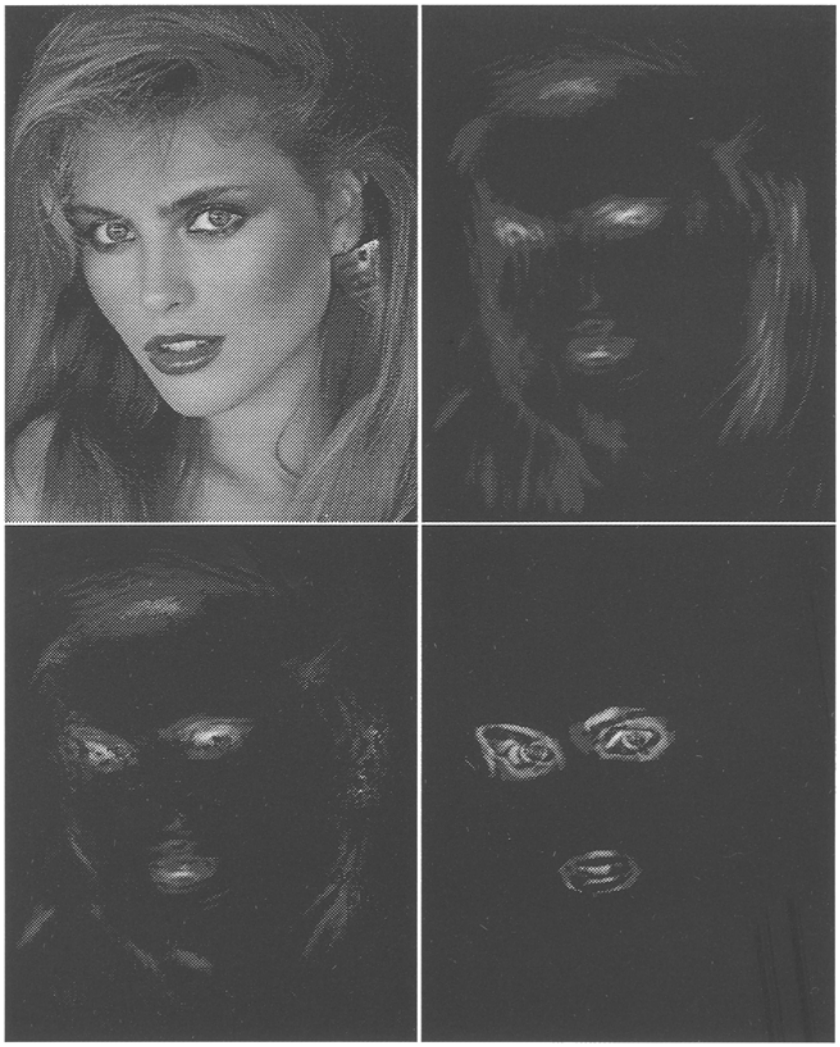

Fig. 9. Back projection of the symmetry. Top left to bottom right: Original image. Isotropic symmetry. Horizontal symmetry projection. The result of voting for the edges which contributed to high symmetry values.

contributing to high symmetry. In Fig. 9 we show the result of ranking the edges according to their contribution to the symmetry maps. [See Reisfeld of Yeshurun (1992) for more details.] The edges that had high symmetry value can be used to segment the objects that were the source of this symmetry.

5 Conclusion

Stressing the role of low-level context-free attentional mechanisms, we have introduced a interest operator that associates a symmetry magnitude and direction with each pixel in an image. As opposed to previous paradigms the symmetry map is computed directly from the edge map of the image, without prior segmentation.

The symmetry map can serve as an efficient low-level process for indexing attention to the regions that are likely to be of high interest in a picture. Other processes can then direct computational resources to these regions and interpret the data in them. The symmetry transform agrees with some psychophysical data and can be further investigated to find its relation to models for human visual perception (Bonneh et al. 1993), and for indexing attention in biological vision, such as

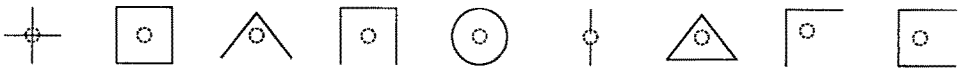

Fig. 10. Mean spontaneous fixation positions for various figures, each subtending about 2° visual angle on the retina, reported by Kaufman and Richards. The small dotted circles indicate the region in which 86% of the fixations occurs.

Fig. 11. Applying the symmetry transform on a simple figure. Left to right Original image; its edges; isotropic symmetry with a large σ; radial symmetry with the same σ; superposition of the isotropic symmetry with the edges—the skeleton detected; superposition of the radial symmetry with the edges (compare to previous figure); isotropic symmetry with a small σ—the corners detected; superposition of the fine symmetry with the edges.

the models used by Posner and Peterson (1990) and by Ullman (1984).

There are many psychophysical works investigating humans ability to detect bilateral symmetry [e.g., Locher and Nodine (1986)]. It is accepted that humans are able to detect mirror symmetry and that symmetry detection has a local nature. However, there was almost no link suggested between these capabilities and the choice of attention points. Kaufman and Richards (1969) studied spontaneous fixation tendencies of the eyes when they are confronted with very simple forms. Fig. 10 shows the results for a collection of nine small figures. These results hold for forms that subtend less than 5° on the retina. As the forms become larger, fixation becomes scattered. Thus, it seems that the natural fixation mechanisms has a local nature. These results are in intriguing agreement with the results of applying the symmetry transform to similar forms. As an example, consider the square in Fig. 11.

Another result, that might be more directly related to our operator, shows that symmetric patterns are more efficient in attracting attention (Locher and Nodine 1986). It was found that when symmetric patterns were presented to subjects, both the survey and examination fixations were concentrated along the axes of symmetry. It is interesting to note that in this research, subjects were shown abstract art works in order to eliminate context.

Some psychophysical findings (Attneave 1954; Kaufman and Richards 1969) lead to the assumption that interest points can be regarded also as points of high curvature. However, using the symmetry transform at low resolution results in detection of points with high curvature (as shown in Fig. 11). This is partially due to the fact that the symmetry axis passes through the points with maximal curvature (Yuille and Leyton 1990).

Our operator generalizes most of the existing methods for detection of interest points. The symmetry-transform value is high near points of high curvature, since the edges that construct the high curvature are highly symmetric according to our definition. This is in agreement with the theoretical analysis of Yuille and Leyton (1990), which shows that axes of symmetry terminate at maximal curvature points. Corners also give rise to high symmetry values, since they consist of edges that point toward each other. Busyness measures can also be regarded as a rough estimate of the symmetry transform, since both are influenced by edge intensities. Note that busyness is insensitive to orientation, and thus will yield poor results in noisy scenes. Since most symmetry axes of an object pass through its center of gravity, the symmetry transform, and especially its radial form, highly responds to the center of gravity. However, as opposed to methods based on center of gravity, our operator does not require any segmentation or specification of the object's vertexes.

There is, however, much more to attention than simple, context-free mechanisms. As we have argued, attention, like almost every other visual task, is trapped in the local-global, top-down– bottom-up, and context-free–context-dependent vicious circle. A possible path to follow should consist of a three-level paradigm:

- A context-free direct computation of a set of simple and early mechanisms, like color, motion, or the generalized symmetry.
- An analysis of the geometry of this early map, based on general and task-dependent knowledge.
- A conventional object detection and recognition (Lamdan et al. 1988) performed only in locations indexed by the previous stages.

This approach is not necessarily bottom up, as it seems to be on first sight, since the specific set of

context-free features used in the lower level could be selected and modified by the higher levels, once a specific task is performed. These modifications can include, for example, use of nonsymmetrical (oriented) Gaussians in the lower level, specification of the appropriate combination of dark and bright symmetries, and the use of the symmetry projection. In the higher levels, the symmetry-edge map obtained by the lower level can be operated on as an edge map applying standard computer-vision tools. In particular, one can apply a Hough transform for line detection, thus detecting significant straight-symmetry axes. By applying other standard edge-linking procedures, one obtains all the symmetry axes in the image, which may be curved. Another analysis that can be useful at this stage is a detailed analysis of the scale space formed by the symmetry maps while zooming in.

The approach could be demonstrated, for example, in the task of detecting persons and facial features in images. We first compute the context-free generalized symmetry map, and then we look for a geometric pattern where the symmetry peaks are vertically arranged, below a circular symmetry peak (persons), or three symmetry peaks that form a triangle (facial features). Indexed locations could then be specifically analyzed by edge-based or intensity-based recognition schemes.

Our main point in this work, was to propose a scheme for the first level of the intricated process of attention, that generalizes most of the existing methods. The approach is stable, is proved to work on real, cluttered images, can tolerate partial occlusion of objects, can be implemented in real time, and thus is a promising method for early vision-attentional mechanisms in purposive and active vision systems.

6 Appendix: Direction of the Skewed Symmetry Edges

Consider an ideal object that possesses a reflectional symmetry. Let p_i and p_j be two contour points which are mirror symmetric with respect to the (symmetry axis) point p. Then, $p = (p_i + p_j)/2$. Let θ_i and θ_j be the orientation of the gradients at p_i and p_j respectively, then $\theta_i + \theta_j = \pi$, and the direction of the symmetry axis at p is $(\theta_i + \theta_j)/2$. The generalized-symmetry transform presented here provides a continuous measure of symmetry and is designed also for cases where $\theta_1 + \theta_2 \neq \pi$. We have not assumed any cause for the deviation from the ideal situation, and therefore we have defined the symmetry orientation to be $(\theta_i + \theta_j)/2$ even in the nonideal case. Note that this definition is invariant under translation, rotation, and change of scale (similarity transformation).

A perspective image of an ideal symmetric object is in general nonsymmetric. Thus, one would like to define a symmetry transform that is invariant to the perspective transformation. A common approximation to the perspective transformation is the affine transformation. This kind of symmetry is known as *skewed symmetry* (Kanade and Kender 1983).

Let us consider first the position of a skewed-symmetry edge. Since a midpoint of a segment is an affine invariant, the position is the same as in the previous case. The definition of a skewed-symmetry edge orientation is, however, different. It follows from basic trigonometric considerations (Reisfeld 1994) that the natural skewed-symmetry edge orientation is

$$\varphi(i, j) = \arctan \frac{2\cos(\theta_i - \alpha_{i,j})\cos(\theta_j - \alpha_{i,j})}{\sin(\theta_i + \theta_j - 2\alpha_{i,j})}$$

where α_{ij} is defined as in Section 3. One may need to add π to be consistent with the ideal case.

Acknowledgments

We thank M.D. Levine and S.W. Zucker for helpful comments on this manuscript. Research supported in part by grant No. 89-00481 from the US-Israel Binational Science Foundation (BSF) to HW, by grant No. 89-00078 from the US-Israel Binational Science Foundation (BSF) to YY, and grant No. 4478192 from the French-Israeli MOST.

References

Aboot, A.L. and Ahuja, N., 1988. Surface reconstruction by dynamic integration of focus, camera vergence and stereo, *Proc. 2nd Intern. Conf. Comput. Vis.*, Tampa, FL.

Aloimonos, J.Y., Weiss, I., and Bandyopadhyay, A., 1987. Active vision, *Intern. J. Comput. Vis.*, 1:334–356.

Antes, J.R., 1974. The time course of picture viewing, *J. Psychol.*, 103:62–70.

Atallah, M.J., 1985. On symmetry detection, *IEEE Trans Comput.* C-34:663–666.

Attneave, F., 1954. Informational aspects of visual perception, *Psychological Review*, 61:183–193.

Bajcsy, R., 1988. Active perception. *Proc. IEEE*, 76(8):996–1006.

Ballard, D., 1990. Animated vision, Tech. Rep. TR 61, University of Rochester, Department of Computer Science, 1990.

Bigun, J., 1988. Pattern recognition by detection of local symmetries. In E.S. Gelsma and L.N. Kanal, eds., *Pattern Recognition and Artificial Intelligence*, Elsevier, North Holland, pp. 75–90.

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

Bonneh, Y., Reisfeld, D., and Yeshurun, Y., 1993. Texture discrimination by local generalized symmetry, *Proc. 4th Intern. Conf. Comput. Vis.*, Berlin.

Brady, M. and Asada, H., 1984. Smooth local symmetries and their implementation, *Intern. J. Robot. Res.* 3(3):36–61.

Brunnstrome, K., Lindeberg, T., and Eklundh, J.O., 1992. Active detection and classification of junctions by foveation with a head-eye system guided by the scale-space primal sketch, *Proc. 2nd Europ. Conf. Comput. Vis.*, Santa Margherita, Ligure, Italy, 7:701–709.

Cohen, K., 1981. The development of strategies of of visual search, in eye movements. In D. Fisher, R. Monty, and J. Senders, Ed., *Cognition and Visual Perception*, Laurence Erlbaum Assoc: Hillsdale NJ, pp. 299–314.

Crowley, J.L., 1991. Towards continuously operating integrated vision systems for robotics applications, *SCIA-91, 7th Scandinavian Conf. Image Anal.*, Aalborg.

Culhane, S.M. and Tsotsos, J.K., 1992. An attentional prototype for early vision, *Proc. 2nd Europ. Conf. Comput. Vis.*, S. Margherita, Ligure, Italy, May, pp. 551–560.

Davis, L.S., 1977. Understanding shape: Ii. symmetry, *IEEE Trans. Syst. Man, Cybern.* 7:204–211.

Edelman, S., Reisfeld, D., and Yeshurun, Y., 1992. Learning to recognize faces from examples, *Proc. 2nd Europ. Conf. Comput. Vis.*, Santa Margherita, Ligure, Italy, 7:787–791.

Haith, M.M., Bergman, T., and Moore, M.J., 1977, Eye contact and face scanning in early infancy, Science 198:853–855.

Kanade, J. and Kender, J.P., 1983. Mapping image properties into shape constraints: skewed symmetry, affine-transformable patterns, and the shape-from-texture paradigm. In Beck, Hope and Rosenfeld, eds., *Human and Machine Vision*, Academic Press: New York.

Kaufman, L. and Richards, W. 1969. Spontaneous fixation tendencies for visual forms, *Perception and Psychophysics*, 5(2):85–88.

Lamdan, Y., Schwartz, J.T., and Wolfson, H., 1988. On recognition of 3-d objects from 2-d images, *Proc. IEEE Intern. Conf. Robot. Autom. Philadelphia*, 1407–1413.

Locher, P.J. and Nodine, C.F., 1986. Symmetry catches the eye. In A. Levy Schoeh (ed.), *Eye Movements: From Physiology to Cognition*, Elsevier: North Holland, pp. 353–361.

Loftus, G. and Mackworth, N., 1978. Cognitive determinants of fixation location during picture viewing, *Human Perception and Performance* 4:565–572.

Marola, G., 1989. On the detection of the axis of symmetry of symmetric and almost symmetric plannar images, *IEEE Trans. Patt. Anal. Mach. Intell.*, 11(1):104–108.

Moravec, H.P., 1977. Towards automatic visual obstacle avoidance, *5th Intern. Joint Conf. Artif. Intell.* Cambridge, MA, pp. 584–590.

Nevatia, R. and Binford, T.O. 1977. Description and recognition of curved objects. *Artificial Intelligence*, 8:77–98.

Posner, M.L. and Peterson, S.E., 1990. The attention system of the human brain., *Annu. Rev. Neurosci.* 13:25–42.

Reisfeld, D., 1994. *Generalized symmetry transforms: attentional mechanisms and face recognition*, Ph.D. thesis, Computer Science Department, Tel-Aviv University, January.

Reisfeld, D., Wolfson, H., and Yeshurun, Y., 1990. Detection of interest points using symmetry, *Proc. 3rd Intern. Conf. Comput. Vis.*, Osaka, Japan, December, pp. 62–65.

Reisfeld, D. and Yeshurun, Y., 1992. Robust detection of facial features by generalized symmetry, *Proc. 11th Intern. Conf. Image Anal. Patt. Recog.*, The Hague, The Netherlands, August, pp. 117–120.

Rimey, R.D. and Brown, C.M., 1992. Where to look next using a bayes net: incorporating geometric relations, *Proc. 2nd Europ. Conf. Comput. Vis.*, S. Margherita, Ligure, Italy, May, pp. 542–550.

Rojer, A. and Schwartz, E., 1990. Design considerations for a space-variant visual sensor with complex logarithmic geometry, *Proc. 10th Intern. Conf. Patt. Recog.*, pp. 278–285.

Salapatek, P. and Kessen, W., 1973. Prolonged investigation of a plane geometric triangle by the human newborn, *J. Exper. Child Psychol.* 15:22–29.

Tistarelli, M. and Sandini, G., 1990. Estimation of depth from motion using an anthropomorphic visual sensor, *Image Vis. Comput.*, 8(4):271–278.

Ullman, S., 1984. Visual routines, *Cognition*, 18:97–159.

Xia, Y., 1989. Skeletonization via the realization of the fire front's propagation and extinction in digital binary shapes. *IEEE Trans. Patt. Anal. Mach. Intell.*, 11(10):1076–1089.

Yeshurun, Y. and Schwartz, E.L., 1989. Shape description with a space-variant sensor: Algorithm for scan-path, fusion, and convergence over multiple scans, *IEEE Trans. Patt. Anal. Mach. Intell.*, 11(11):1217–1222.

Yuille, A. and Leyton, M., 1990. 3d symmetry-curvature duality theorems, *J. Comput. Vis. Graphics, Image Process.* 52:124–140.

Zabrodsky, H., Peleg, S., and Avnir, D., 1992. A measure of symmetry based on shape similarity, *Proc. Conf. Comput. Vis. Patt. Recog.*, Champaign, IL, June.

Zucker, S.W., Dobbins, A., and Iverson, L., 1989. Two stages of curve detection suggest two styles of visual computation, *Neural Computation* 1:68–81.