Detecting Symmetry in Grey Level Images: The Global Optimization Approach

Nahum Kiryati and Yossi Gofman
Department of Electrical Engineering
Technion - Israel Institute of Technology
Haifa 32000, ISRAEL

Abstract

The detection of significant local reflectional symmetry in grey level images is considered. Prior segmentation is not assumed, and it is intended that the results could be used for guiding visual attention and for providing side information to segmentation algorithms. A local measure of reflectional symmetry that transforms the symmetry detection problem to a global optimization problem is defined. Reflectional symmetry detection becomes equivalent to finding the global maximum of a complicated multimodal function parameterized by the location of the center of the supporting region, its size, and the orientation of the symmetry axis. Unlike previous approaches, time consuming exhaustive search is avoided. A global optimization algorithm for solving the problem is presented. It is related to genetic algorithms and to adaptive random search techniques. The efficiency of the suggested algorithm is experimentally demonstrated. Just one thousand evaluations of the local symmetry measure are typically needed in order to locate the dominant symmetry in natural test images.

1 Introduction

Symmetry appears in man made objects and is also common in nature. Classical references on symmetry are Weil [42] and Gardner [17]. The omnipresence of symmetry has motivated many studies on various aspects of symmetry in images. One line of research has addressed the following class of problems: *Given a shape, detect, quantify and describe its symmetry.* The results could be applied in object recognition, visual inspection, shape representation, etc. These studies have initially focused on the basic reflectional and rotational symmetry types, then on generalized and refined concepts such as the medial axis transform, skeletons [6, 8, 10, 11, 16, 29, 33, 34], ribbons and skewed symmetry [9, 15, 36, 44].

The input for algorithms that quantify or describe symmetry in shapes has normally been assumed to be the output of a successful segmentation procedure. However, a general purpose segmentation algorithm that consistently provides reliable results when applied to a wide variety of images is still to be developed, and the segmentation problem now seems to be fundamentally difficult. The applicability of shape-based symmetry analysis algorithms is thus limited.

There is a great computational advantage in performing symmetry analysis using edges and lines rather than grey level data, and indeed, most of the studies on symmetry have so far taken that approach [1, 13, 35, 31, 41, 45]. It can also be argued that edge symmetry is more significant than region symmetry due to the importance of edges in human vision. Yet, extracting clean, well separated contour representations from unconstrained real images is as difficult as segmentation, and is beyond the ability of current edge detection and linking algorithms. Thus, regardless of whether contour or region input is used, as long as the input of a symmetry analysis algorithm is assumed to come from a single shape, its usefulness is restricted to special cases.

Toward overcoming these limitations, the developers of symmetry analysis algorithms have gradually begun to try and relax the shape-input requirement. A computationally attractive possibility is to use the edge detector output as input to symmetry analysis algorithms, without assuming that the various objects in the edge image are readily separated [12, 25, 32, 43]. However, valuable information is often lost in edge detection, and noise, shadows and specularity effects are accentuated. In many cases symmetry that is detectable in the grey level image, or in some functional of the grey level image, such as the (unthresholded) gradient image, might be obscured after edge detection. It is thus clear that in demanding applications the grey level data should be available to the symmetry analysis module.

Marola [26] developed an algorithm for detecting the reflectional symmetry axis of almost symmetric images. He defined a measure of symmetry in grey level images and, relying on the assumption that the image is almost symmetric, described a way to find with respect to which axis the maximum of the measure is attained. Kuehnle [24] and Zielke *et al* [46] used contour, grey level and local orientation symmetry with respect to nearly vertical axes for the detection of vehicle rears. Masuda *et al* [27] suggested an algorithm for the detection of partial reflectional as well as rotational symmetry in unsegmented images by computationally expensive exhaustive search for maximal correlation of directional edge features. Sun [39] developed a symmetry detection algorithm that uses the unsegmented grey level image and gradient information. The direction of the symmetry axis is obtained from the gradient orientation histogram and its position according to the center of gravity or from the projection of the image in the direction of the symmetry axis. The algorithm is thus applicable when the symmetry in the image is a global feature. Minovic *et al* [28] proposed a method for identifying symmetry of a 3-D object represented by an octree.

Pizer and his coworkers [10, 16, 29, 34] and Kelly and Levine [22] suggested the insightful and interesting concept of performing segmentation *in conjunction* with symmetry analysis based on edge images. Reisfeld *et al* [37] presented a symmetry operator that can be applied to gradient images without prior segmentation, and can be used as an interest operator to guide visual attention. Di Gesu and Valenti [14] suggested an alternative moment-based symmetry operator that can be applied to the grey level image itself and discussed its utility as a cue for segmentation.

An important issue in the design of a local symmetry measure is scale. In [37] a scale parameter determines the effective radius of the support for symmetry measurement at each pixel. By varying the scale parameter, multiresultion schemes can be implemented. In [14] the symmetry operator is applied at various scales, and the final result is obtained by merging significant symmetry zones obtained at all scales. In practice, both approaches have been aimed primarily at small scale features, such as eyes in face images. See also [7]. While this is very reasonable (and compatible with human vision) for guiding foveation, the ability to identify large symmetric or nearly symmetric objects in an image is crucial for segmentation or recognition purposes. The bottleneck in identifying large symmetric features using symmetry operators is computational. The time needed for evaluating a symmetry measure necessarily depends on the area of the support. Suppose that a symmetry operator that measures the reflectional symmetry as a function of the location of center of the support, the orientation of the axis and the scale is given. Applying the operator at all locations, scales and orientations is a formidable computational task[1].

The key observation in this research is that evaluating the symmetry measure at all locations, scales and orientations is not necessary. The symmetry detection problem is regarded as a global optimization problem, in which the location, scale and orientation parameters that maximize the

[1]Since the the operator suggested by Reisfeld *et al* [37] cleverly uses directional information from the gradient image, it does not have to be computed separately for each orientation. Computing it at all locations and at numerous large scales would still be very expensive.

measure are to be determined. By applying suitable global optimization algorithms, the number of symmetry measure evaluations needed for finding the dominant symmetry can be greatly reduced with respect to existing exhaustive search approaches.

In this paper the global optimization approach to symmetry detection in grey level images is demonstrated in the context of *reflectional* symmetry. In section 2 a measure of local reflectional symmetry as a function of the location of the center of the supporting region, its size, and the orientation of the symmetry axis is presented. The suggested approach is modular in the sense that other symmetry types can be detected if a suitable local symmetry operator is plugged in, see e.g. [4, 20]. For any non-trivial image, local symmetry is a highly complex, multimodal function of four parameters (or more, depending on support parameterization). Standard maximization procedures designed for convex functions, such as steepest descent, are inapplicable. It is well known that finding the global maximum of a general function is hard. However, if the function to be maximized is not too badly-behaved, and in particular, if the regions of attraction [40] associated with the global maximum are not too small, various global optimization techniques can be applied. In section 3 we present an algorithm for determining the global maximum of the reflectional symmetry measure. It is related to genetic algorithms [21] and to adaptive random search techniques [40]. In section 4 the performance of the suggested algorithm is demonstrated on various test images. Just one thousand evaluations of the (four parameter) local symmetry measure are typically needed in order to locate the dominant reflectional symmetry in these images. This requires just two minutes on a standard SGI Indigo workstation under (interpreted) MATLAB^{TM}, and is likely to require much less with optimized code.

2 Measure of Local Reflectional Symmetry

The suggested approach for the detection of reflectional symmetry in grey level images requires a measure of local 2-D symmetry. Generally, the measure must be a function of the position and orientation of the symmetry axis, and of the local measurement support. The local symmetry measure that we eventually recommend is related to the measure suggested by Marola [26], extended to deal with finite supports and varying scales. We present the development of the measure from principles, beginning with 1-D functions.

2.1 Symmetry of 1-D Functions

Consider a 1-D real function $y = f(x)$, defined on $x \in [-L, L]$. The function is said to be reflectionally symmetric with respect to the origin if $f(x) = f(-x)\ \forall x$ and antisymmetric if $f(x) = -f(-x)\ \forall x$. It is easy to show that any function $f(x)$ can be uniquely represented as a sum of a symmetric function and an antisymmetric function $f(x) = f_s(x) + f_{as}(x)$. Based on this decomposition a measure of the symmetry of $f(x)$ with respect to the origin can be defined as[2]

$$S\{f(x)\} = \frac{\|f_s(x)\|^2}{\|f(x)\|^2} = \frac{\|f_s(x)\|^2}{\|f_s(x)\|^2 + \|f_{as}(x)\|^2} \quad . \tag{1}$$

This measure satisfies

- $S\{f(x)\} \in [0, 1]$
- $S\{f(x)\} = 1$ if $f(x)$ is purely symmetric
- $S\{f(x)\} = 0$ if $f(x)$ is purely antisymmetric.

The decomposition of $f(x)$ into symmetric and antisymmetric components can easily be obtained:

$$f_s(x) = (f(x) + f(-x))/2 \tag{2}$$

$$f_{as}(x) = (f(x) - f(-x))/2 \tag{3}$$

Consider the reflectional correlation coefficient $C\{f(x)\}$ defined as

$$C\{f(x)\} = \frac{\int_{-L}^{L} f(x)f(-x)\,dx}{\int_{-L}^{L} f^2(x)dx} \quad . \tag{4}$$

By direct substitution it is easy to show that

$$C\{f(x)\} = \frac{\|f_s(x)\|^2 - \|f_{as}(x)\|^2}{\|f_s(x)\|^2 + \|f_{as}(x)\|^2} \tag{5}$$

[2] A similar measure is used in [46].

which is the measure used by Zielke *et al*[3] [46]. Clearly,

$$C\{f(x)\} \;=\; 2S\{f(x)\} - 1 \;. \tag{6}$$

$S\{f(x)\} = 1$ (pure symmetry) thus corresponds to a correlation coefficient of 1, and $S\{f(x)\} = 0$ corresponds to pure antisymmetry and a correlation coefficient of -1.

2.2 Reflectional Symmetry of 2-D Functions

Let $f(x, y)$ be a 2-D real function that can be considered to be zero everywhere except within a circle of radius L centered at the origin[4]. Suppose that the reflectional symmetry of $f(x, y)$ should be measured with respect to an axis t that passes through the origin and makes an angle θ with the positive x axis as shown in Figure 1.

$$\begin{aligned} t &= x\cos\theta + y\sin\theta \\ s &= x\sin\theta - y\cos\theta \end{aligned} \tag{7}$$

In terms of the t, s coordinate system the function is reflectionally symmetric if $f(t, s) = f(t, -s)$ for all t, s in the domain. It is said to be antisymmetric if $f(t, s) = -f(t, -s)$. 2-D reflectional symmetry with respect to the t axis thus requires $\forall t$ 1-D symmetry on line segments parallel to the s axis.

A measure of 2-D reflectional symmetry with respect to the t axis is defined as follows :

$$S_\theta\{f\} = \frac{\int \|f_s(t,s)\|^2 dt}{\int \|f(t,s)\|^2 dt} = \frac{\int \|f_s(t,s)\|^2 dt}{\int \|f_s(t,s)\|^2 dt + \int \|f_{as}(t,s)\|^2 dt} \tag{8}$$

where for each t the norms are of 1-D functions of s along a line segment extending from $-L$ to L in parallel to the s axis. It satisfies

[3] Zielke *et al* [46] proposed to subtract the average of $f_s(x)$. In this research, DC removal within *2-D* supports will be suggested.

[4] As described in the sequel, this condition is realized by multiplying the image with a 2-D Gaussian window that also ensures that the transition is smooth.

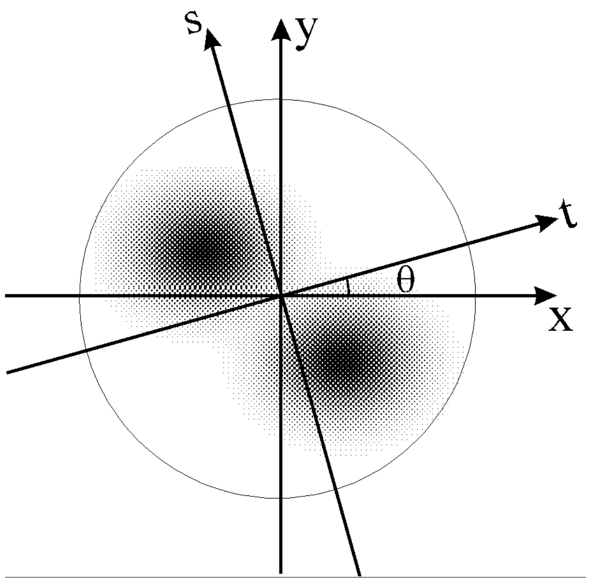

Figure 1: The coordinate system for 2-D symmetry measurement

- $S_\theta\{f\} \in [0, 1]$
- $S_\theta\{f\} = 1 \quad$ if f is purely symmetric
- $S_\theta\{f\} = 0 \quad$ if f is purely antisymmetric.

This 2-D symmetry measure can be expressed as a weighted average of 1-D (line) symmetry values:

$$S_\theta\{f\} = \frac{\int S\{f(t,s)\} \|f(t,s)\|^2 dt}{\int \|f(t,s)\|^2 dt}, \tag{9}$$

where, for each t, $S\{f(t,s)\}$ and $\|f(t,s)\|$ are respectively the symmetry value and the norm of a 1-D function of s. It is also easy to show that

$$S_\theta\{f\} = \frac{1}{2} \frac{\int\int f(t,s) f(t,-s) ds dt}{\int f^2(t,s) ds dt} + \frac{1}{2}, \tag{10}$$

hence, at this stage the measure is similar to the measure used by Marola [26]. It can be extended to support skewed symmetry and medial axis symmetry.

2.3 Finite Supports, Implementation and Complexity

The measurement of the local symmetry in specified regions in an image requires windowing. The "hard" circular window is undesirable and can be substituted by a "soft" Gaussian window:

$$G(x, y, r) = \frac{1}{2\pi r^2} \exp\left[-\left(\frac{x^2 + y^2}{2r^2}\right)\right] \tag{11}$$

where r is referred to as the effective radius of the support. In practice, the value of the Gaussian window can be regarded as zero from a certain distance to its center and on. The window is applied to the region in the image in which the local symmetry should be measured, i.e., $f(t, s)$ is multiplied by $G(t, s, r)$ [5].

The computation of the symmetry measure can be carried out in various ways in the spatial or in the frequency domains. In the spatial domain, the image is clipped by a $2L \times 2L$ square window centered at the coordinates of the center of the support in which symmetry is to be measured. L is set so that a 2-D Gaussian window with effective radius r (Eq. 11) is nearly zero at distance L from its center[6]. The average of the clipped image is subtracted, the result is multiplied by the Gaussian window and rotated according to Eq. 7. The 2-D symmetry measure is now calculated via discrete approximation of Eq. 10. The overall complexity is linear in the area of the clipped image, i.e., $O(L^2)$.

2.4 Examples

The suggested symmetry measure is a function of four parameters: the coordinates of the center of the support, the effective radius of the support and the orientation of the symmetry axis. For visualization only two dimensional slices through the symmetry function can be shown. Two parameters must be held constant, and the symmetry measure is shown as a function of the other two parameters. Figures 2 and 3 exemplify the suggested symmetry measure.

Note that in general the support does not have to be circular. As shown in Fig. 4, by increasing

[5] In many visual tasks, the reference intensity level in the examined part of the image should be ignored. This can be accomplished by subtracting the local average of $f(t, s)$.

[6] In our implementation L was set to be $3.03r$, hence at distance L from its center the Gaussian window reached $1/100$ of its maximal value.

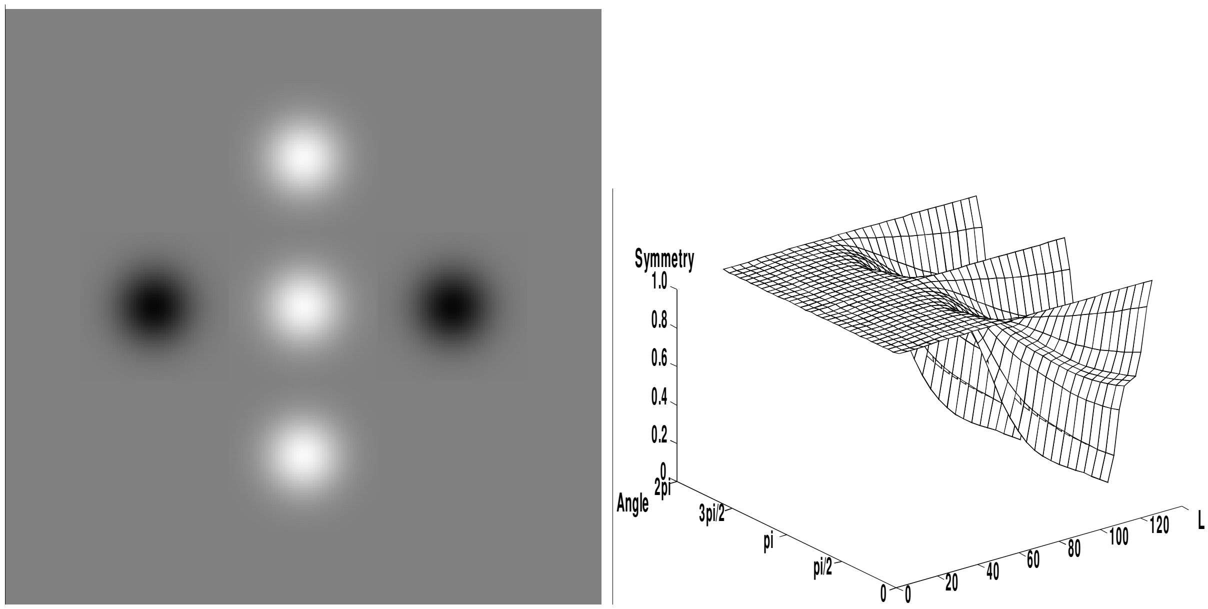

Figure 2: The left image is a 256×256 image in which symmetry is measured. The right image is a two dimensional slice through the symmetry function in which the coordinates of the center of the support are held constant at the center of the image, and the symmetry measure is shown as a function of the orientation of the symmetry axis and the size of the support.

the number of parameters that describe the support, better fits to symmetric regions can be obtained. On the other hand, by adding parameters the dimension of the symmetry search space increases beyond the current value of four. This carries a significant price tag in terms of computational complexity, but might be desirable in certain applications, especially if symmetry is being used as a cue for segmentation.

2.5 Introducing Scale Dependence

As defined above, the measure of local symmetry is scale invariant. Therefore, the symmetry measure associated with a tiny symmetric region will be higher than that associated with a much larger support in which symmetry is slightly imperfect. This is an undesirable state of affairs, since in all images there are many unimportant tiny symmetric regions, and at the limit each single pixel is always perfectly symmetric. Regions of interest for image understanding are usually much larger.

Consider the image of the monkey (Figure 3). One tends to say that the significant symmetry

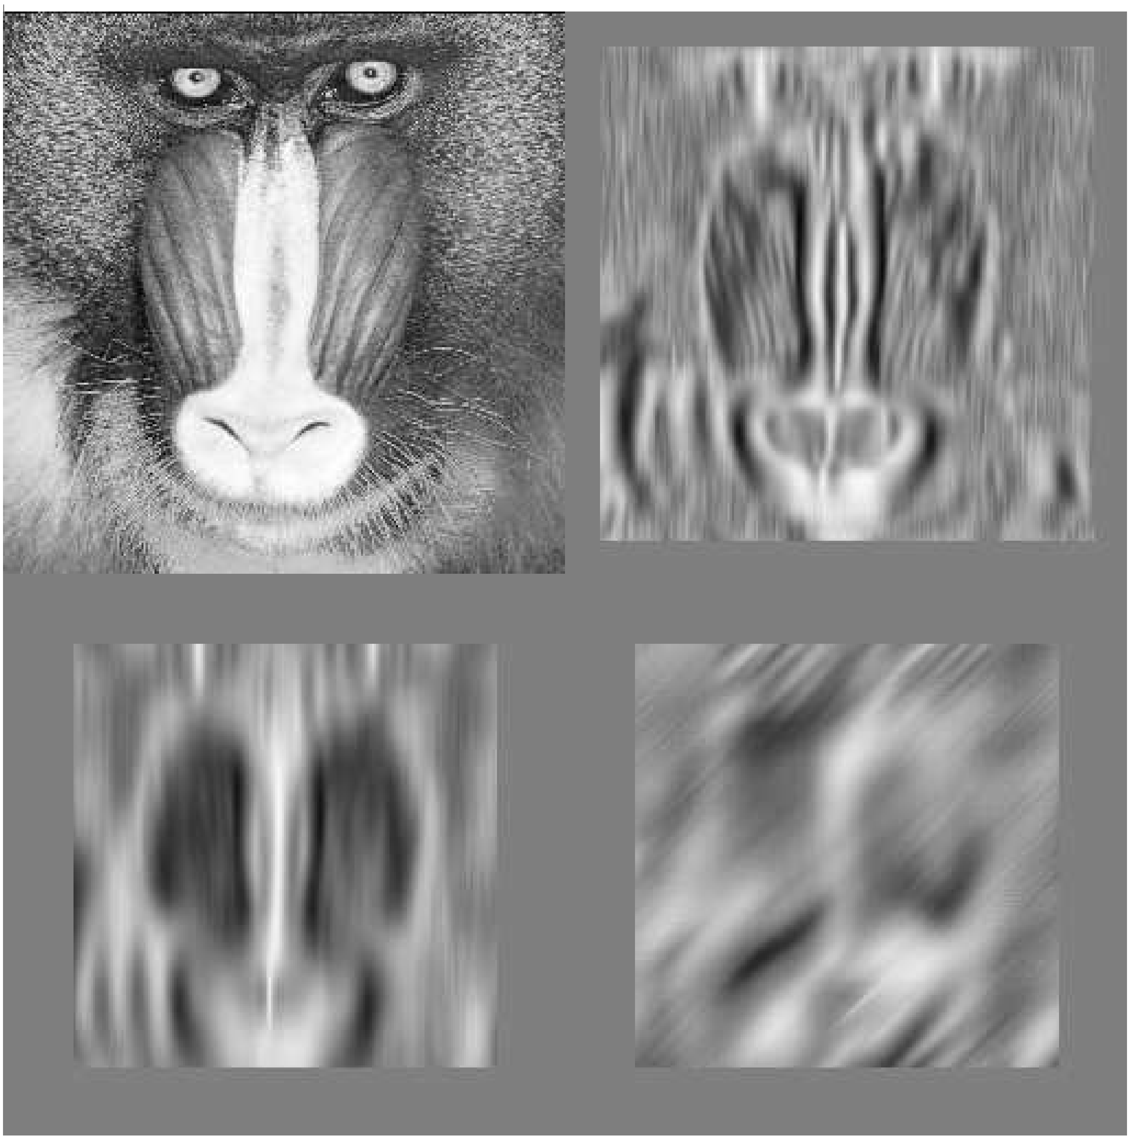

Figure 3: The upper left image is a 256×256 image in which symmetry is measured. The other three images are two dimensional slices through the symmetry function in which the size of the support and the orientation of the symmetry axis are held constant, and the symmetry measure is shown as a function of location. Thus, the brightness of each pixel represents the symmetry measure for a fixed size support around that pixel with respect to an axis at a given constant orientation. The lower left symmetry map corresponds to a vertical axis and 64×64 ($L = 32$) window. The symmetry of the nose and eyes is apparent. In the upper right symmetry map the axis is still vertical, but the window is reduced to 32×32 ($L = 16$). Smaller scale symmetry features are obtained. The lower right symmetry map corresponds to 64×64 ($L = 32$) window but a slanted ($\theta = \pi/4$) axis. It is evident that no dominant symmetric features are detected at that orientation.

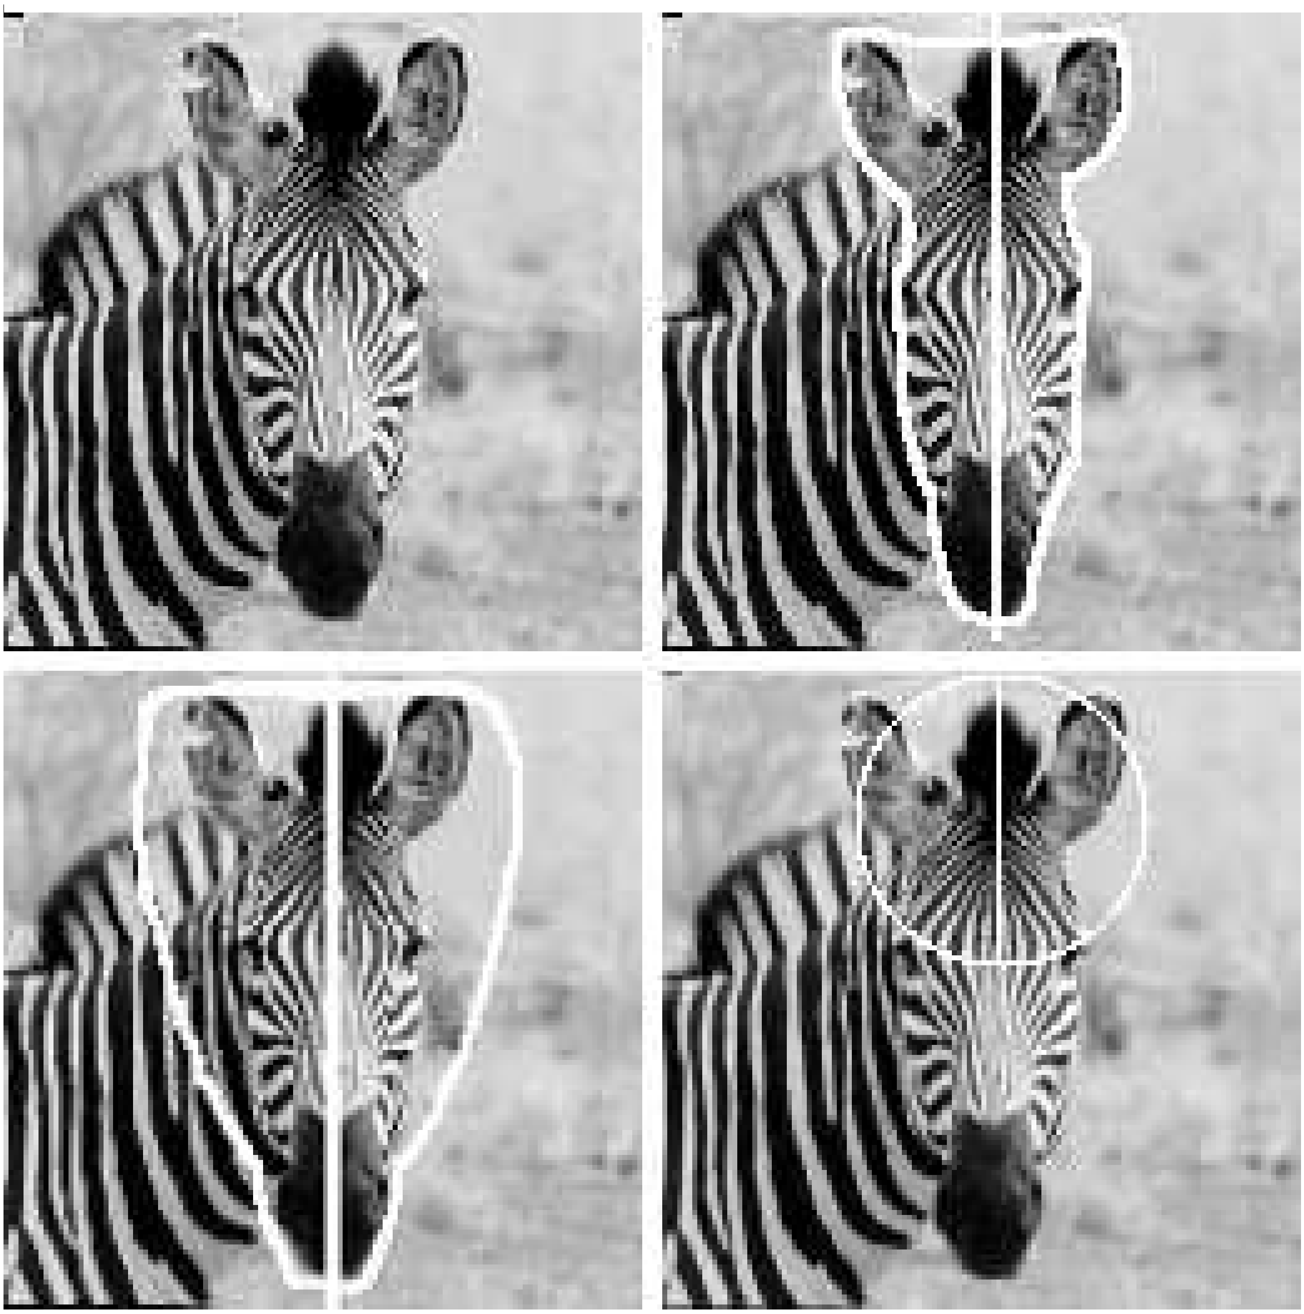

Figure 4: *Top left* - Original Image, *Top right* - A support in which symmetry is high. *Bottom left* - Modified support in which symmetry is lower. *Bottom right* - Circular support in which symmetry is high.

in that image is the overall symmetry of the face. However, several regions in the image are locally symmetric, notably the eyes. In most applications the symmetry of the eye can be argued to be less important than the overall facial symmetry. On the other hand, in Figure 11 in the sequel, the symmetry of each flower could be of greater interest than the overall symmetry of the flower arrangement.

For meaningful symmetry detection, it is thus necessary to introduce scale dependence in the measure, in a task oriented way. A reasonable general scale-dependent symmetry measure is of the form

$$S^{SD}(x, y, \theta, r) = f(r) \cdot S^{SI}(x, y, \theta, r) \tag{12}$$

where S^{SD} is the scale-dependent symmetry measure, S^{SI} is the scale invariant symmetry measure and $f(r)$ is some function of the effective radius of the support.

In many applications, such as in robotics and manufacturing, in face image analysis and in certain medical imaging tasks, reliable a-priori estimates of the scale R of the features of interest are available. $f(r)$ can then be selected to be a smooth, unimodal function, scaled such that its maximum corresponds to R. Reasonable choices are the Rayleigh function

$$f_R(r) = \frac{r}{R^2} e^{-\frac{r^2}{2R^2}} \mu(r) \tag{13}$$

where $\mu(r)$ is the unit step function, or a truncated Gaussian

$$f_{R,W}(r) = \frac{1}{\sqrt{2\pi} W} e^{-\frac{(r-R)^2}{2W^2}} \mu(r). \tag{14}$$

The latter has the advantage that the level of certainty in the a-priori scale estimate can be expressed in $f(r)$ via the width parameter W.

In various applications good scale estimates are not available, but a minimal scale of interest R_{min} can be defined. $f(r)$ can then be specified to be a saturation function such as the scaled and truncated error function

$$f_{R_{min},W}(r) = \frac{1}{\sqrt{2\pi} W} \mu(r) \int_{-\infty}^{\infty} e^{-\frac{(t-R_{min})^2}{2W^2}} dt. \tag{15}$$

If significant symmetric regions are to be detected in the absence of a-priori scale estimates, a reasonable strategy could be to search for the largest region in which the scale-independent symmetry level exceeds a certain threshold. This can be accomplished by

$$S^{SD}(x, y, \theta, r) = r^2 \cdot S^{SI}(x, y, \theta, r) \cdot \mu(S^{SI}(x, y, \theta, r). - S_{th}) \tag{16}$$

Fig. 5 shows a test image and the region in which $S^{SD}(x, y, \theta, r)$ is maximal for various values of the threshold level S_{th}. The radius of the drawn circles is L. The effective radius r is 3 times smaller. As expected, the size of the highlighted region decreases as S_{th} is increased. In the experimental section of this paper, scale dependence is introduced via Eq. 16.

3 Symmetry and Global Optimization

For any non-trivial image, the local symmetry measure is a highly complex, multimodal function of four parameters: the x and y coordinates of the center of the supporting region, its effective radius, and the orientation of the symmetry axis. Finding the dominant local symmetry in the image is thus an elaborate global optimization problem.

It is known that the global maximum of a general function above a bounded continuous domain cannot be found in finite time. However, if the function to be maximized is not too badly-behaved, various approaches to global optimization can be taken [40]. Global optimization algorithms carry out a search, guided by some strategy, in the domain. Their outcome is considered to be acceptable if it sufficiently close, in some sense, to the global maximum. A trade-off exists between the quality of the outcome and the time (i.e., the number of function evaluations) spent in the search.

Mathematical performance analysis of global optimization algorithms is usually very difficult. The applicability of a certain global optimization technique to a given problem must be carefully verified by application specific experimental benchmarks. We developed a global optimization algorithm and experimentally verified its applicability to the detection of symmetry in grey level images. The suggested algorithm is related to genetic algorithms on one hand, and to adaptive random search techniques on the other.

In genetic algorithms [21] the parameters of the search space are binary encoded over a string

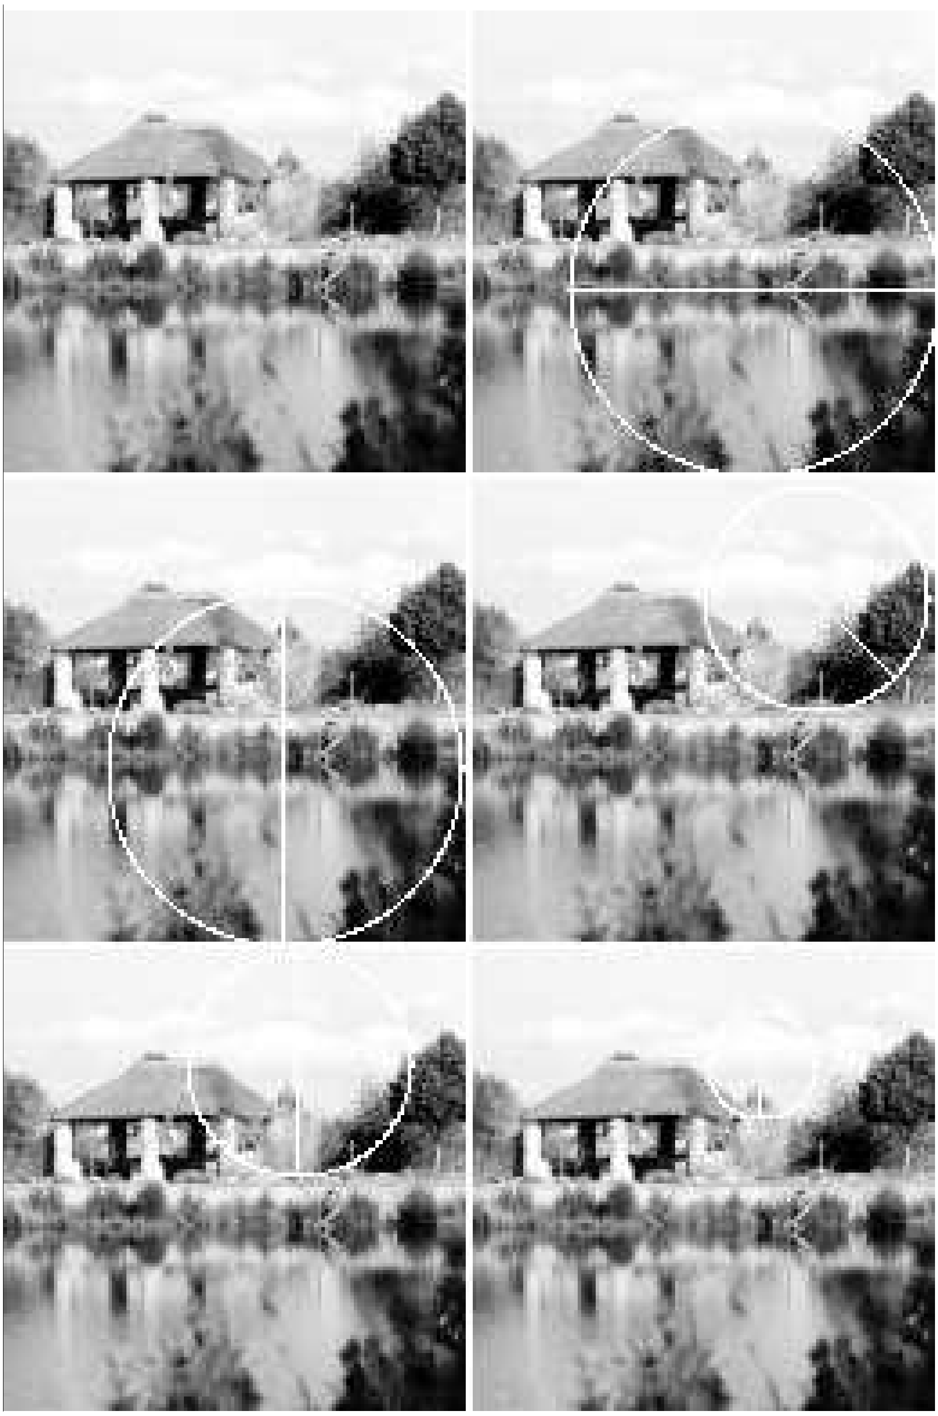

Figure 5: Introducing scale dependence via Eq. 16. The regions in which the scale dependent symmetry measure is maximal are shown for various threshold levels. The radius of the drawn circles is L. The effective radius r is 3 times smaller. *Upper left:* Original test image. *Upper right:* $S_{th} = 0.5$. *Middle left:* $S_{th} = 0.6$. *Middle right:* $S_{th} = 0.7$. *Lower left:* $S_{th} = 0.8$. *Lower right:* $S_{th} = 0.9$.

of predefined length. Each bit is referred to as a "gene", and the string that represents a point in the search space as a "chromosome". An initial set of chromosomes is created at random. This initial "population" is allowed to "breed", by performing crossovers between random chromosome pairs. "Mutation" which is a random, low probability alteration of genes, is applied to the new population and the goal function is evaluated for each new member of the population. The best chromosomes in the initial and new population, i.e., the best points in the search space in terms of the value of the goal function, are kept for the next iteration. The procedure is repeated until a stopping condition is satisfied.

In standard genetic algorithms, a crossover between two chromosomes is carried out by selecting at random a crosspoint location, and exchanging one part of each chromosome with the respective part in the other chromosome, leaving the remaining part in each chromosome unchanged. This implies that the course and efficiency of the search depends on the encoding of the parameters, in particular on the position of the "genes" within the "chromosomes". Since in most applications it is difficult to select a "good" encoding strategy, the significance of the encoding should be regarded as a drawback.

Without mutations, the search points specified by a genetic algorithm usually converge to the neighborhood of a certain local maximum. Mutations are introduced to prevent this deadlock. In the classical genetic algorithms the "mutation rate", i.e., the probability that a gene (bit) is changed is constant in time and identical for all genes. Note however that at the initial stages of the search the diversity in the population is large. At that stage mutation is not necessary. After many iterations the "genepool" is very depleted and a high mutation rate is needed in order to escape from a local maximum. A fixed mutation rate may be effective in certain stages of the search, but is likely to have a superfluous effect in the initial stages and an insufficient effect at later stages.

The suggested approach is to modify the standard genetic algorithm (GA) and create a probabilistic genetic algorithm (PGA) that alleviates these drawbacks of the GA. Briefly, as in the GA, in the PGA an initial population is randomly selected and iterations start. In GA the next generation is obtained by performing crossovers between pairs of randomly selected chromosomes. In the PGA each member of the next generation can be regarded as the offspring of the entire current generation. Each gene in every chromosome of the next generation is obtained by randomly drawing '0' or '1' according to the distribution of the values of that gene in the current population, weighted by the values of the goal function. Thus, in the PGA there are no crosspoints and the binary encoding has no significance.

In addition, in the PGA a different mutation probability is assigned to each gene in each generation based on its diversity. This is accomplished by calculating the standard deviation σ of the values of each gene. A depleted gene (either "1" or "0" for all chromosomes) yields $\sigma = 0$ and a diverse gene (equal number of "1" and "0") yields $\sigma = 0.5$. The mutation probability of a gene is then made proportional to $0.5 - \sigma$. This has the effect of increasing the mutation rate as a gene becomes depleted. Finally, as in GA, the new candidate chromosomes are evaluated and the best among the combined pool of new and old ones are kept for the next generation.

We have independently developed the PGA, but its principles are not necessarily new. We show that the algorithm can be successfully applied to the problem of symmetry detection, and outperform the standard GA. For a comprehensive, interesting discussion on related extensions to Genetic Algorithms the reader is referred to [2]. A detailed description of the PGA follows.

3.1 The Probabilistic Genetic Algorithm

As in the standard genetic algorithm, in the PGA each search parameter is encoded as a binary string. The strings are tied together to form a larger string, of total length K, referred to as a chromosome, that represents a point in the search space. As will shortly be seen, in the PGA the position of each bit in the chromosome can be arbitrary. A chromosome $\vec{C_i}$ has the following form:

$$\vec{C_i} = [c_{i,1}, c_{i,2}, ..., c_{i,K-1}, c_{i,K}], \quad c_{i,j} \in \{0, 1\} \tag{17}$$

As in the GA, each iteration of the PGA yields a "population", or "generation" of chromosomes. Let N denote the number of chromosomes in each generation. A generation can be represented as a matrix:

$$\mathbf{C} = \begin{bmatrix} \vec{C_1} \\ \vdots \\ \vec{C_N} \end{bmatrix} \tag{18}$$

The initial population is generated at random. For each chromosome the goal function is evaluated. These values can be represented as a vector

$$\vec{v} = [v_1, \cdots, v_N]^T. \tag{19}$$

Then iterations begin.

Unlike the GA, in the PGA each member of the next generation can be regarded as the offspring of the entire current generation. Each gene in every new chromosome is obtained by randomly drawing '0' or '1' according to the distribution of the values of that gene in the current population, weighted by the values of the goal function.

Formally, the values of the goal function v_i in the current population are first normalized:

$$V_i = \frac{v_i - \min_k(v_k)}{\max_k(v_k) - \min_k(v_k)}, \quad i = 1...K. \tag{20}$$

These values satisfy $V_i \in [0, 1]$, with the value "1" assigned to the "best" chromosome, and can be arranged as a vector $\vec{V} = \{V_i\}$. Now the probability with which the value "1" should be assigned to each of the bits in the new chromosomes is calculated. It is the weighted average of each bit (gene) in columns of $\mathbf{C}$ with $\vec{V}$ taken as a vector of weights:

$$P_j = \frac{\sum_{i=1}^{N} c_{i,j} V_i}{\sum_{i=1}^{N} V_i}, \quad j = 1...K. \tag{21}$$

P_j is used as the probability to draw the value "1" for the bit (gene) in position j in each new chromosome (the number of new chromosomes created N_n is a parameter). Thus, in the PGA there are no crosspoints and the binary encoding has no significance.

To these new chromosomes mutation is applied. Unlike the GA, in the PGA a different mutation probability is assigned to each gene based on the current diversity of that gene. This is accomplished by calculating the standard deviation of the values of each gene in the current generation:

$$\sigma_j = \sqrt{\frac{\sum_{i=1}^{N}(c_{i,j} - \frac{1}{N}\sum_{k=1}^{N} c_{k,j})^2}{N}}, \quad j = 1...K. \tag{22}$$

Clearly $\sigma_j \in [0, 0.5]$ $j = 1...K$. A fully depleted gene (either "1" or "0" for all chromosomes) corresponds to $\sigma_j = 0$ and a maximally diverse gene (equal number of "1" and "0") yields $\sigma_j = 0.5$. The mutation probability of a gene is increased as its diversity decreases:

$$M_j = m \cdot (0.5 - \sigma_j) \;, \quad j = 1...K \tag{23}$$

where m is a parameter. Mutation occurs in the bit in position j in each of the new chromosomes with probability M_j. Finally, as in the GA, the goal function is evaluated at points represented by the new chromosomes and the best chromosomes in the combined pool of current and new chromosomes are kept for the next generation.

4 Experiments

We implemented the suggested algorithm within the MATLAB^{TM} environment and applied it to several 128×128 test images. Scale dependence was introduced via Eq. 16, using a default threshold $S_{th} = 0.65$ except where otherwise indicated. In the PGA the location parameters x, y of the center of the support were encoded by 7 bits each, the orientation of the symmetry axis θ by 3 bits and the effective radius r by 6 bits, hence the total length of the "chromosomes" was $K = 24$. The number of chromosomes in the beginning and end of each generation was $N = 50$ and the number of new chromosomes created in each generation was $N_n = 100$. The mutation parameter in the PGA was set to $m = 0.05$. About 1000 evaluations of the local symmetry measure were typically needed to reach convergence in these examples. This corresponds to about two minutes in (interpreted) MATLAB^{TM} on a standard SGI Indigo workstation.

Symmetry detection results are shown in Figs. 6-14. In these figures the radius of the circle is L, and indicates the size of the symmetry measurement support. On the circle the value of the Gaussian window is just 1/100 of its maximal value. Outside these circles its value is considered to be negligible. The effective radius r, where the Gaussian window reaches about 0.6 of its

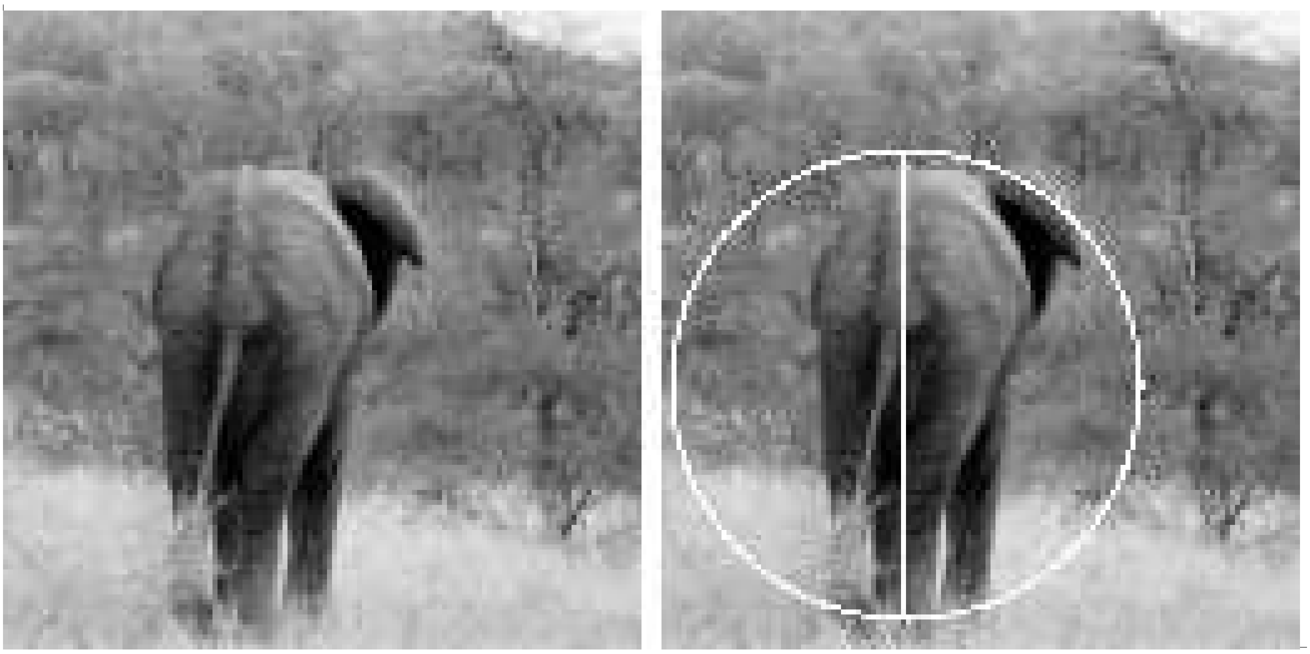

Figure 6: Elephant. *Left:* original image. *Right:* detected symmetry ($S_{th} = 0.65$). In this and the following figures, the circles indicate the size of the symmetry measurement support. The radius of the circle is L, so on the circle the value of the Gaussian window is just $1/100$ of its maximal value. The effective radius r, where the Gaussian window reaches about 0.6 of its maximal value, is 3 times smaller.

maximal value, is 3 times smaller.

The results shown in Figs. 6-9 are generally pleasing and conform well to the intuition of a human observer. Figs. 10-12 exemplify the possibility of incorporating task dependent information via tuning of the threshold S_{th}. In Fig. 10 reducing the threshold to $S_{th} = 0.55$ allowed for the slight imperfections due to partial occlusion, the door and the window. With a higher threshold, a symmetry axis perpendicular to the road at the lower-right part of the image would have been found. In Fig. 11 reducing the threshold to $S_{th} = 0.55$ led to the detection of the symmetry of the bouquet, while with a larger threshold the symmetry of the individual flower would have been found. In Fig. 12 a higher threshold of $S_{th} = 0.85$ focuses the search on the small but highly symmetric face of the lion.

For comparison we have also implemented a standard Genetic Algorithm with similar bit allocation parameters and a mutation probability of 0.05 per bit. Table 1 shows the average number of goal function evaluations over 10 runs needed for the GA and for the PGA to converge to the global maximum in the symmetry detection examples shown. The PGA is significantly

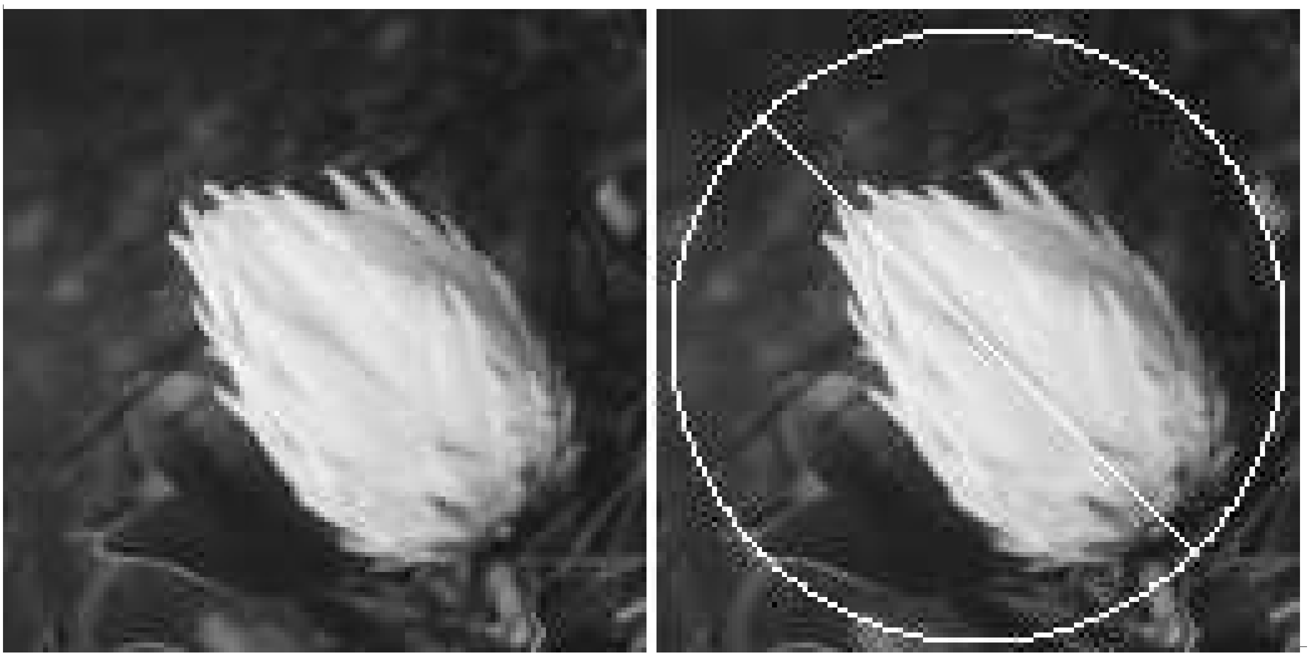

Figure 7: Flower. *Left:* original image. *Right:* detected symmetry ($S_{th} = 0.65$).

Figure 8: Roadsign. *Left:* original image. *Right:* detected symmetry ($S_{th} = 0.65$).

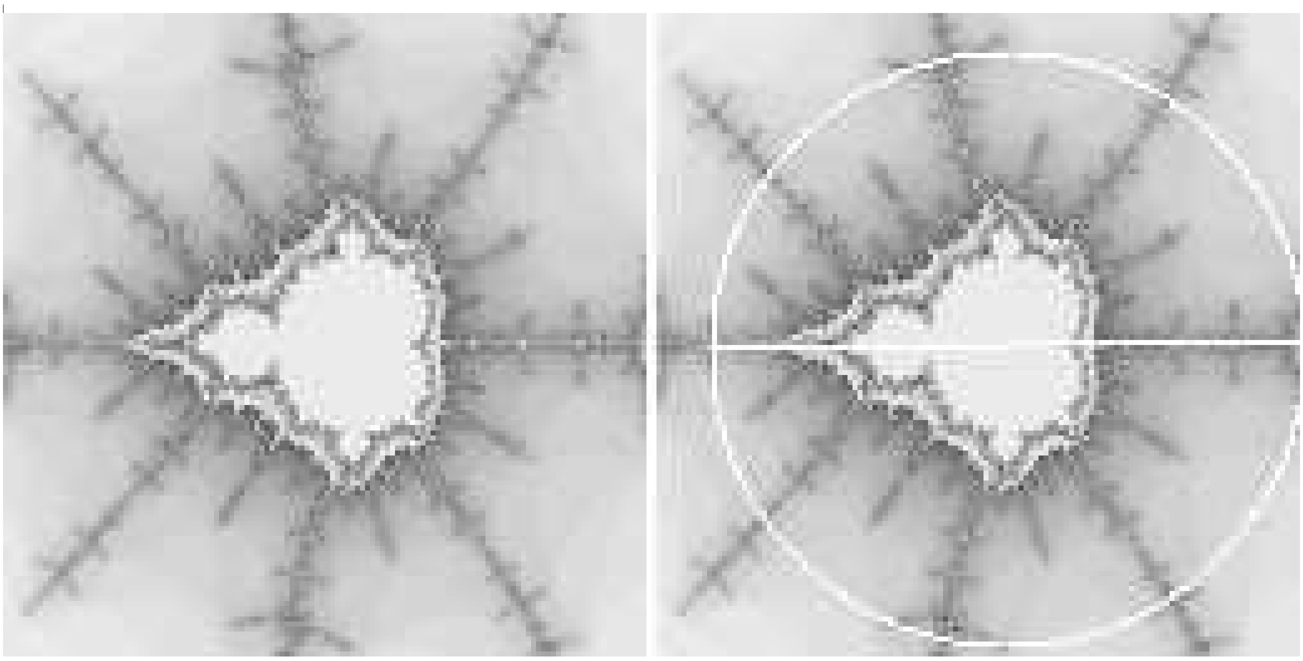

Figure 9: Fractal. *Left:* original image. *Right:* detected symmetry ($S_{th} = 0.65$).

Figure 10: House. *Left:* original image. *Right:* detected symmetry ($S_{th} = 0.55$).

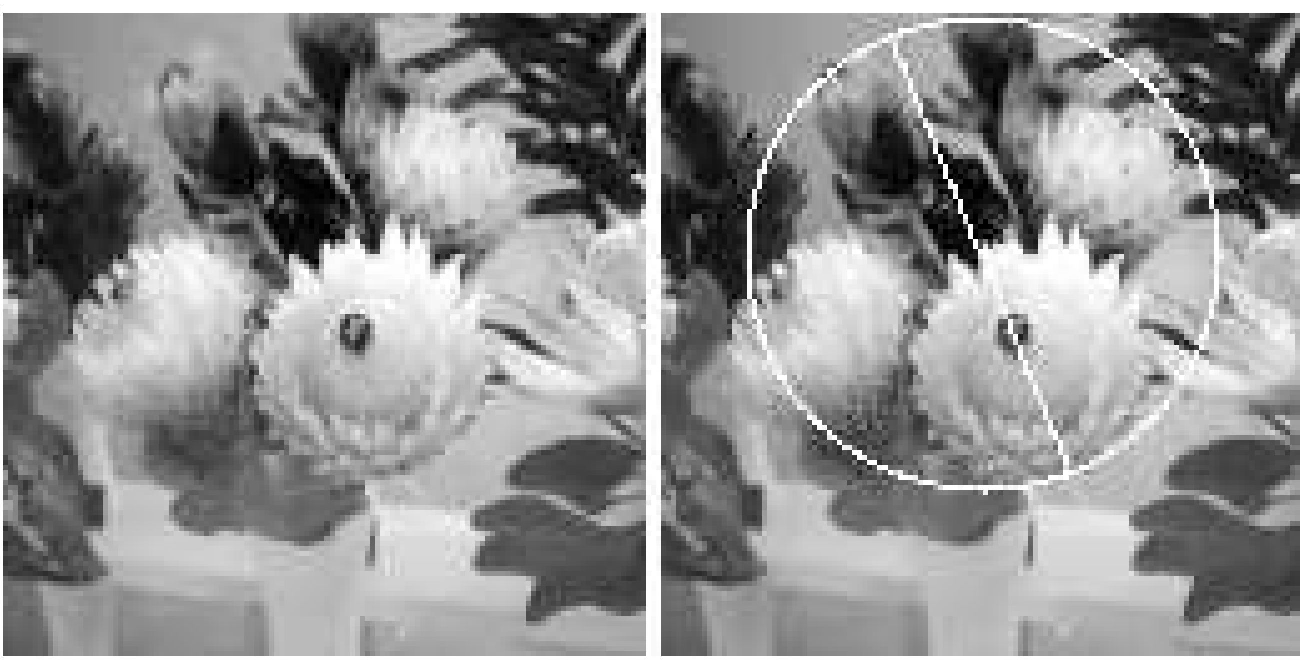

Figure 11: Bouquet. *Left:* original image. *Right:* detected symmetry ($S_{th} = 0.55$).

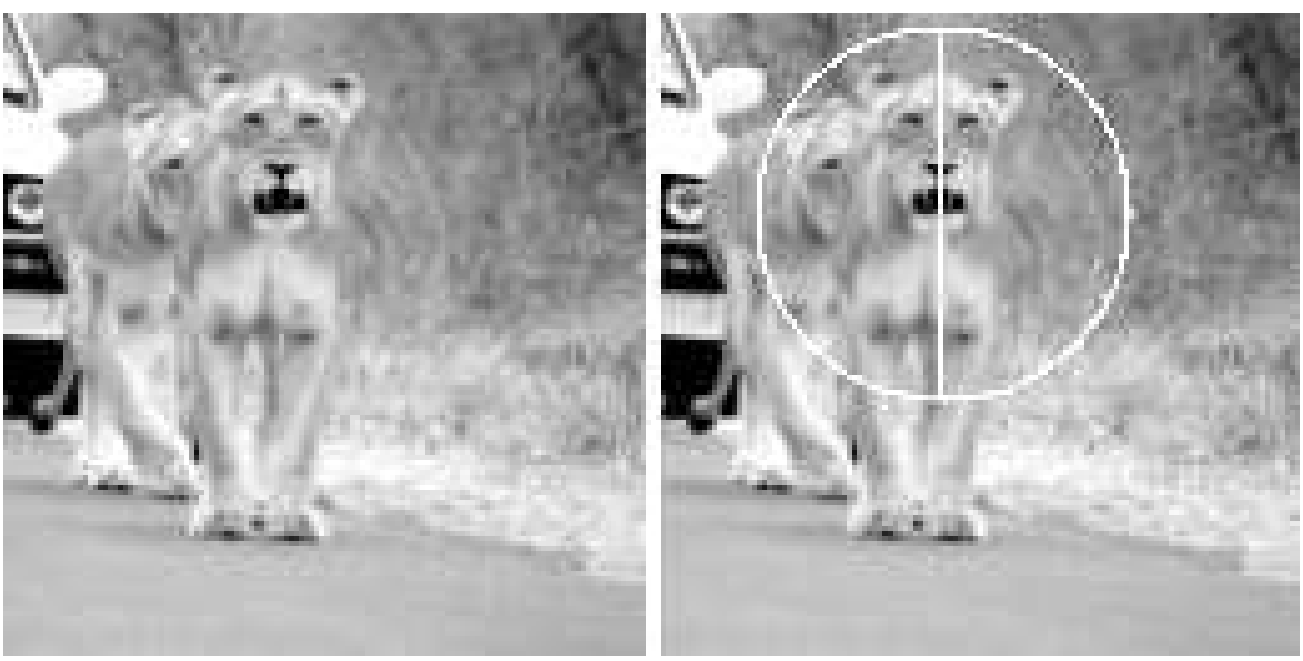

Figure 12: Lion. *Left:* original image. *Right:* detected symmetry ($S_{th} = 0.85$).

faster, and its performance generally more consistent than that of the standard GA.

Image	Number of goal function evaluations GA	PGA
Elephant	1627	593
House	1648	667
Flower	1137	687
Lion	1124	423
Roadsign	1431	631
Fractal	957	532

Table 1: The average number of goal function evaluations over 10 runs needed for the GA and for the PGA to converge to the global maximum in the symmetry detection examples shown. The PGA is faster and its performance more consistent than that of the GA.

5 Discussion

In this study the detection of the dominant local symmetry in grey level images is treated as a global optimization problem. A measure of local grey level symmetry and an efficient algorithm for solving the resulting global optimization problem have been developed. The suggested measure of local symmetry has been designed to be mathematically well behaved and visually relevant, but keeping in mind the subsequent global optimization stage, computationally economical as well.

It is important to realize that a mathematical definition of symmetry may not always coincide with human perception. In the tower image (Fig. 13) a human initially sees a vertical symmetry axis, but due to perspective distortion the actual reflectional symmetry with respect to the horizontal axis is higher. In Fig. 14 the tree itself is slanted, hence the surprising symmetry with the shadow is detected. Allocating more bits to the representation of θ would allow finer fits in these two test images, but the problem seems to be fundamental and related to the wonderful human ability to unconsciously compensate for various imaging distortions. An attempt to deal with the problem by allowing additional degrees of freedom in the search to compensate for, say, perspective, would be computationally expensive but also likely to lead to the detection of visually insignificant random arrangements in the image. Trying to *simultaneously* detect other symmetry types can be expected to lead to similar difficulties. Still, that might be an interesting topic for future research.

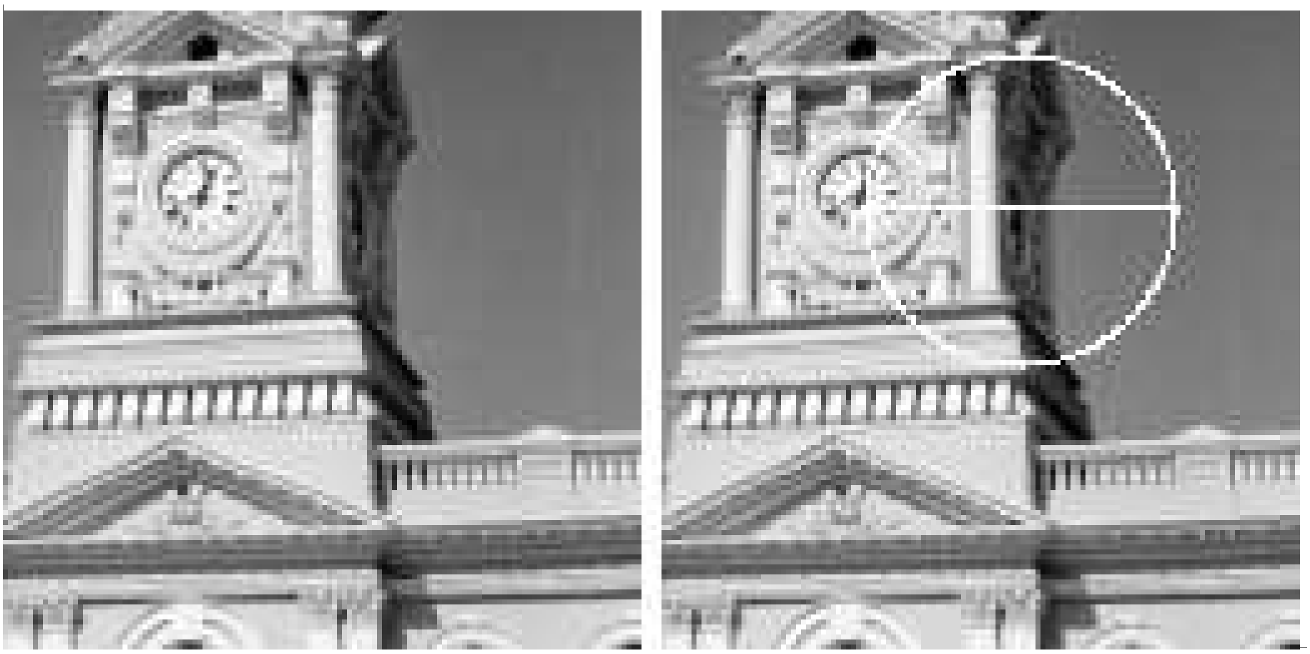

Figure 13: Tower. *Left:* original image. *Right:* detected symmetry ($S_{th} = 0.65$).

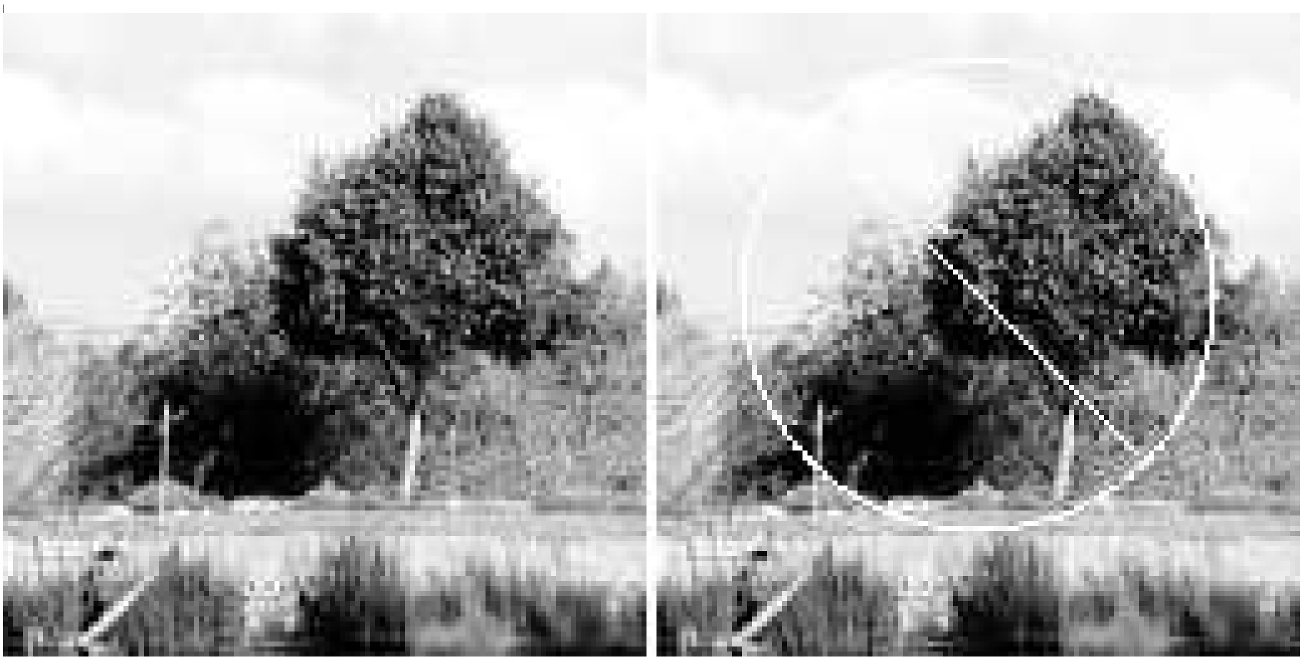

Figure 14: Tree. *Left:* original image. *Right:* detected symmetry ($S_{th} = 0.65$).

The difficulty of global optimization problems depends on the smoothness of the objective function, the number of local maxima, the dimension of the search space and the cost of a single objective function evaluation. Grey level symmetry analysis leads to difficult maximization of a multidimensional and multimodal objective function. The suggested Probabilistic Genetic Algorithm seems to be well suited to this problem. Convergence is typically obtained after just about 1000 evaluations of the symmetry measure in the 2^{24} quantized search space. This takes roughly two minutes in a MATLAB^{TM} environment on a standard SGI Indigo workstation. The main computational load is objective function evaluations. Therefore, computing time can be shortened by implementing an efficient cache of symmetry measure values that will prevent repeated evaluation of dominant hypotheses. Also, since within each generation all function evaluations can be performed in parallel, the suggested algorithm can efficiently utilize parallel architectures of various grain sizes.

As defined in this paper, the goal of the optimization algorithm is to detect the global maximum of the local symmetry measure, i.e., the dominant symmetry in the image. In many applications it would be interesting to also detect other significant maxima, in addition to the global one. Fortunately, strategies for adapting genetic algorithms to detect multiple maxima have been developed. Goldberg and Richardson [19] described the method of *sharing functions* to permit the formation of several different stable subpopulations within the genetic algorithm. Beasley *et al* [3] presented a sequential niche technique in which the genetic algorithm is iterated using knowledge gained during one iteration to avoid re-searching regions where solutions have already been found. It is interesting to apply these techniques to multiple symmetry detection.

It is well known that the general continuous global optimization problem cannot be solved in finite time. Global Optimization algorithms thus come without guarantees. If, however, the objective function satisfies certain smoothness constraints, it might be possible to upper bound the number of function evaluations needed in order to locate the global maximum to a given accuracy. This requires careful analysis of the a-priori knowledge on the structure of the goal function. In [38] probabilistic convergence guarantees for the multiresolution smooth-kernel Hough transform have been obtained. It is theoretically important to analyze the symmetry detection problem in the same spirit, to develop a covering global optimization algorithm with convergence guarantees and compare its performance to the PGA.

How can the symmetry detection results be used in conjunction with a segmentation algorithm? The symmetry detector provides a circular contour that roughly encompasses the dominant symmetrical support in the image. One approach is to use that contour as an initial condition for a curve evolution [23] procedure that would modify the contour to obtain better fits according to certain segmentation criteria. The outcome could be evaluated and refined according to shape-based symmetry criteria, and vice-versa until a stopping condition is satisfied. Another approach would be to attach a symmetry term to an energy functional used for segmentation [30]. Using an active surface functional, Gauch and Pizer [18] have already demonstrated that grey-level skeletonization and segmentation can be successfully tied together. We share the belief that segmentation should be performed in conjunction with the detection of visual features such as symmetry [10, 22], and that the combined result of this bootstrap process would be better than possible with independent procedures.

Acknowledgments

We are grateful to Professors E. Granum, H.H. Nagel and H. Wechsler and to the anonymous referees for bringing important references to our attention. This research has been supported in part by the Israeli Ministry of Science, by the R. and M. Rochlin Research Fund and by the Ollendorff Center of the Department of Electrical Engineering.

References

[1] M.J. Attalah, "On Symmetry Detection", *IEEE Trans. Comput.*, Vol. 34, pp. 663-666, 1985.

[2] S. Baluja, *Population-Based Incremental Learning: A Method for Integrating Genetic Search Based Function Optimization and Competitive Learning*, Technical Report CMU-CS-94-163, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania.

[3] D. Beasley, D.R. Bull and R.R. Martin, "A Sequential Niche Technique for Multimodal Function Optimization", *Evolutionary Computation*, Vol. 1, pp. 101-125, 1993.

[4] J. Bigün, "Recognition of Local Symmetries in Gray Value Images by Harmonic Functions", *Proc. Int. Conf. on Pattern Recognition (ICPR)*, pp. 345-347, Rome, 1988.

[5] J. Bigün, "A Structure Feature for Some Image Processing Applications Based on Spiral Functions", *Comput. Vision Graph. Image Process.*, Vol. 51, pp. 166-194, 1990.

[6] H. Blum and R.N. Nagel, "Shape Description Using Weighted Symmetric Axis Features", *Pattern Recognition*, Vol. 10, pp. 167-180, 1978.

[7] Y. Bonneh, D. Reisfeld and Y. Yeshurun, "Quantification of Local Symmetry: Application for Texture Discrimination", *Spatial Vision*, to appear.

[8] M. Brady and H. Asada, "Smoothed Local Symmetries and Their Implementation", *Int. J. Robotics Research*, Vol. 3, pp. 36-61, 1984.

[9] A.M. Bruckstein and D. Shaked, "Skew Symmetry Detection via Invariant Signatures", *Proc. 6th Int. Conf. on Comput. Analysis of Images and Patterns (CAIP)*, pp. 17-24, Prague, 1995.

[10] C.A. Burbeck and S.M. Pizer, "Object Representation by Cores: Identifying and Representing Primitive Spatial Regions", *Vision Research*, Vol 35, pp. 1917-1930, 1995.

[11] C.A. Burbeck, S.M. Pizer, B.S. Morse, D. Ariely, G.S. Zauberman and J.P. Rolland, "Linking Object Boundaries at Scale: a Common Mechanism for Size and Shape Judgments", *Vision Research*, Vol 36, pp. 361-372, 1996.

[12] T.J. Cham and R. Cipolla, "Symmetry Detection Through Local Skewed Symmetries", *Image and Vision Computing*, Vol. 13, pp. 439-450, 1995.

[13] L.S. Davis, "Understanding Shape: II. Symmetry", *IEEE Trans. Syst. Man Cybernet.*, pp. 204-212, 1977.

[14] V. Di Gesù and C. Valenti, "Symmetry Operators in Computer Vision", preprint, 1995.

[15] S.A. Friedberg, "Finding axes of Skewed Symmetry", *Comput. Vision Graph. Image Process.*, Vol. 34, pp. 138-155, 1986.

[16] D.S. Fritsch, S.M. Pizer, B.S. Morse, D.H. Eberly and A. Liu, "The Multiscale Medial Axis and its Applications in Image Registration", *Pattern Recognition Letters*, Vol. 15, pp. 445-452, 1994.

[17] M. Gardner, *The New Ambidextrous Universe - Symmetry and Asymmetry from Mirror Reflections to Superstrings*, Freeman, New York, 1979.

[18] J.M. Gauch and S.M. Pizer, "The Intensity Axis of Symmetry and its Applications to Image Segmentation", *IEEE Trans. Pattern Anal. Machine Intell.*, Vol. 15, pp. 753-770, 1993.

[19] D.E. Goldberg and J. Richardson, "Genetic Algorithms with Sharing for Multimodal Function Optimization", *Proc. 2nd Int. Conf. on Genetic Algorithms*, pp. 41-49, Cambridge, Mass., 1987.

[20] O. Hansen and J. Bigün, "Local Symmetry Modeling in Multi-dimensional Images", *Pattern Recognition Letters*, Vol. 13, pp. 253-262, 1992.

[21] J. H. Holland, "Genetic Algorithms", *Scientific American*, pp. 44–50, July 1992.

[22] M.F. Kelly and M.D. Levine, "Annular Symmetry Operators: A Method For Locating and Describing Objects", *Proc. Int. Conf. on Computer Vision (ICCV)*, pp. 1016-1021, Cambridge, Mass., 1995.

[23] R. Kimmel, N. Kiryati and A.M. Bruckstein, "Analyzing and Synthesizing Images by Evolving Curves with the Osher-Sethian Method", *Int. J. Computer Vision*, to appear.

[24] A. Kuehnle, "Symmetry Based Recognition of Vehicle Rears", *Pattern Recognition Letters*, Vol. 12, pp. 249-258, 1991.

[25] T.S. Levitt, "Domain Independent Object Description and Decomposition", *Proc. AAAI*, pp. 207-211, 1984.

[26] G. Marola, "On The Detection of the Axes of Symmetry of Symmetric and Almost Symmetric Planar Images", *IEEE Trans. Pattern Anal. Machine Intell.*, Vol. 11, pp. 104–108, 1989.

[27] T. Masuda, K. Yamamoto and H. Yamada, "Detection of Partial Symmetry Using Correlation with Rotated-Reflected Images", *Pattern Recognition*, Vol. 26, pp. 1245-1253, 1993.

[28] P. Minovic, S. Ishikawa and K. Kato, "Symmetry Identification of a 3-D Object Represented by Octree", *IEEE Trans. Pattern Anal. Machine Intell.*, Vol. 15, pp. 507-514, 1993.

[29] B.S. Morse, S.M. Pizer and A. Liu, "Multiscale Medila Analysis of Medical Images", *Image and Vision Computing*, Vol. 12, pp. 327-338, 1994.

[30] D. Mumford and J. Shah, "Boundary Detection by Minimizing Functionals", *Proc. IEEE Computer Society Conf. on Computer Vision and Pattern Recognition (CVPR)*, pp. 22-26, San Francisco, 1985.

[31] V.S. Nalwa, "Line-Drawing Interpretation: Bilateral Symmetry", *IEEE Trans. Pattern Anal. Machine Intell.*, Vol. 11, pp. 1117-1120, 1989.

[32] H. Ogawa, "Symmetry Analysis of Line Drawings Using the Hough Transform", *Pattern Recognition Letters*, Vol. 12, pp. 9-12, 1991.

[33] S.M. Pizer, W.R. Oliver and S.H. Bloomberg, "Hierarchical Shape Description Via the Multiresolution Symmetric Axis Transform", *IEEE Trans. Pattern Anal. Machine Intell.*, Vol. 9, pp. 505-511, 1987.

[34] S.M. Pizer, C.A. Burbeck, J.M. Coggins, D.S. Fritsch and B.S. Morse, "Object Shape before Boundary Shape: Scale-Space Medial Axes", *Journal of Mathematical Imaging and Vision*, Vol. 4, pp. 303-313, 1994.

[35] S.K. Parui and D. Dutta Majumder, "Symmetry Analysis by Computer", *Pattern Recognition*, Vol. 16, pp. 63-67, 1983.

[36] J. Ponce, "On Characterizing Ribbons and Finding Skewed Symmetries", *Comput. Vision Graph. Image Process.*, Vol. 52, pp. 328-340, 1990.

[37] D. Reisfeld, H. Wolfson and Y. Yeshurun, "Context Free Attentional Operators: The Generalized Symmetry Transform", *Int. J. Computer Vision*, Vol. 14, pp. 119-130, 1995.

[38] M. Soffer and N. Kiryati, "Guaranteed Convergence of the Hough Transform", *Computer Vision and Image Understanding (CVIU)*, to appear.

[39] C. Sun, "Symmetry Detection Using Gradient Information", *Pattern Recognition Letters*, Vol. 16, pp. 987-996, 1995.

[40] A. Törn and A. Žilinskas, *Global Optimization*, Lecture Notes in Computer Science #350, Springer-Verlag, 1989.

[41] L. Van Gool, T. Moons, D. Ungureanu and E. Pauwels, "Symmetry from Shape and Shape from Symmetry", *Int. J. Robotics Research*, Vol. 14, pp. 407-424, 1995.

[42] H. Weil, *Symmetry*, Princeton University Press, 1952.

[43] A. Ylä-Jääski and F. Ade, "Grouping Symmetrical Structures for Object Segmentation and Description", *Computer Vision and Image Understanding*, Vol. 63, pp. 399-417, 1996.

[44] K.S.Y. Yuen and W.W. Chan, "Two Methods for Detecting Symmetries", *Pattern Recognition Letters*, Vol. 15, pp. 279-286, 1994.

[45] H. Zabrodsky, S. Peleg and D. Avnir, "Symmetry as a Continuous Feature", *IEEE Trans. Pattern Anal. Machine Intell.*, Vol. 17, pp. 1154-1166, 1995.

[46] T. Zielke, M. Brauckmann and W. Von Seelen, "Intensity and Edge Based Symmetry Detection with Application to Car Following", *CVGIP: Image Understanding*, Vol. 58, pp. 177-190, 1993.