ELSEVIER

Medical Image Analysis 4 (2000) 269–282

MEDICAL IMAGE ANALYSIS

www.elsevier.com/locate/media

Symmetry axis computation for almost-symmetrical and asymmetrical objects: Application to pigmented skin lesions

Phillippe Schmid-Saugeon*

EPFL, Signal Processing Laboratory, Department of Elecrical Engineering, EL-Ecublens, 1015 Lausanne, Switzerland

Received 27 January 1999; received in revised form 27 December 1999; accepted 31 March 2000

Abstract

The detection of symmetry axes through the optimization of a given symmetry measure, computed as a function of the mean-square error between the original and reflected images, is investigated in this paper. A genetic algorithm and an optimization scheme derived from the self-organizing maps theory are presented. The notion of symmetry map is then introduced. This transform allows us to map an object into a symmetry space where its symmetry properties can be analyzed. The locations of the different axes that globally and locally maximize the symmetry value can be obtained. The input data are assumed to be vector-valued, which allow to focus on either shape, color or texture information. Finally, the application to skin cancer diagnosis is illustrated and discussed.

Keywords: Symmetry, Optimization, Multi-resolution, Genetic algorithm, Self-organizing maps, Dermatoscopy

1. Introduction

The goal of this research is to develop robust feature extraction methods for the computer-aided diagnosis and follow up of pigmented skin lesions. Digital methods have shown to be reliable tools for the early diagnosis of skin cancer (Green et al., 1994). The computer-aided evaluation of diagnostic features offers empirical measures, in contrast to the more qualitative evaluation done by physicians. In the clinical diagnosis of pigmented skin lesions, one of the main feature is the lesion asymmetry, which is evaluated by drawing two orthogonal axes that maximize the perceived symmetry (Stolz et al., 1994). The evaluation is then binary: the lesion is either symmetrical or asymmetrical. In addition to the small number of possible outcomes, the evaluation is highly subjective and depends on the physicians experience. Therefore, the development of automated techniques for the quantification of symmetry and the detection of symmetry axes is necessary. The approach developed throughout this paper provides the physician with a complete analysis tool for the symmetry criterion.

In this study we have used digitized *epiluminescence microscopy* (ELM) images of pigmented skin lesions (Binder et al., 1995). Epiluminescence microscopy, also called *dermatoscopy* or *dermoscopy*, is a non-invasive technique that uses an oil immersion to render the skin translucent and make pigmented structures visible. There is more information in these images than in clinical images and their use should therefore lead to a significant improvement of the diagnosis accuracy. A review of the technique can be found in (Argenyi, 1997). A set of 50 benign melanocytic lesions and 50 malignant melanoma, along with their histology, has been used to evaluate the different techniques. Photographic pictures of the lesions have been taken and later digitized at a size of 768×512 pixels. To localize the lesions, two different approaches have been developed and presented in previous publications to obtain a binary mask of the lesion: a segmentation technique using color clustering (Schmid, 1999a) and a contour detection technique using nonlinear isotropic

*Corresponding author. Tel.: +41-21-693-4622; fax: +41-21-693-7600. Now with the Video and Image Processing Lab at the University of California at Berkeley.

E-mail address: philippe.schmid@ieee.org (P. Schmid-Saugeon).

1361-8415/00/$ – see front matter
PII: S1361-8415(00)00019-0

diffusion and morphological flooding (Schmid, 1999a). Both techniques have been used in our investigations in order to evaluate the sensitivity of our system to possible changes in the lesion boundary location.

The symmetry of a lesion should be evaluated according to its shape, color and texture (Stolz et al., 1994). While it is straightforward that symmetrical objects show the same symmetry axes for all these components, it is not necessarily the case when objects tend to be asymmetrical. In that case, the axis that minimizes the visual difference between an object and its reflected version may be different when shape, color and texture are considered separately. Because the lesions are almost-symmetrical or asymmetrical objects, a candidate symmetry axis is an axis that locally maximizes a given symmetry measure. This has led our investigations to the general problem of symmetry quantification and optimization. This paper proposes two approaches to this optimization problem: the computation of the axis that maximizes a given symmetry measure, and the computation of a symmetry map that completely describes the symmetry properties of objects. In the former case two techniques are proposed, a genetic algorithm and a self-organizing map, which are intended to provide a unique axis as quickly as possible. In the latter case, a kind of 'symmetry transform' is computed, that maps any object into a symmetry space. Two multiresolution techniques are proposed to compute a symmetry map.

In Section 2, a brief review of the existing literature is given for application to almost-symmetrical and asymmetrical objects. In Section 3, the different parameters used in the optimization process are presented. Sections 4 and 5 are concerned with the heuristic search methods and the symmetry maps, respectively. The application to pigmented skin lesions is illustrated and discussed in Section 7, and finally the techniques are summarized, and conclusions are drawn in Section 8.

2. Symmetry axis detection

2.1. What is a symmetry axis?

In geometry, symmetry is *an exact correspondence in position or form about a given point, line or plane*. From this definition it follows that the extraction of a symmetry axis is possible only for intrinsically symmetrical objects. In general, this is not the case for natural images such as those of pigmented skin lesions. Even for symmetrical objects, perfect symmetry is impossible to obtain in digital images due to imperfect lighting, digitization or occlusion. The term of symmetry axis is often used for almost-symmetrical images or objects, and corresponds to the axis that minimizes the visual difference between the original and its reflection.

The different approaches to symmetry axis computation are:

- based on the human perception of symmetry,
- model based,
- measure based.

The former approach is based on the human visual system and the knowledge we have about how the human brain perceives and analyses symmetrical patterns. However, it appears that this research field has often led to contradictory results (Tyler et al., 1996). The second approach takes into account a model of the observed objects to compute the symmetry axes. In the case of pigmented skin lesions, for example, benign melanocytic skin lesions usually tend to have an elliptic shape. This characteristic may be used to facilitate the axes computation (Schmid, 1997). However there is not always an a priori knowledge of what will be analyzed. The last approach, which is investigated in this study, defines a symmetry measure and attempts to find the axis that maximizes this function.

2.2. Symmetry in computer vision

Several methods have been proposed to compute the *orientation* of two-dimensional discrete objects. A widely used method is the discrete Karhunen-Loève transform (KL), also called principal component decomposition. This approach may be used for binary images (Gonzalez and Wintz, 1992) or gray level images, where the pixel coordinates are weighted by the gray level values or the spatial frequencies (Bigün, 1991, 1994). The analogy between orientation and symmetry axis may be possible only for intrinsically symmetrical objects. It is however not guaranteed. A simple example is to consider an equilateral triangle with a very large base: the principal component will be parallel to the base while the symmetry axis is perpendicular to the base. In addition, the definition of a symmetry value can only be a function of the data variance from the axis, which is actually the minimization criteria. Any other measure would not lead to optimal results.

A method to compute the best symmetry axis called *n-transform* was first proposed in (Marola, 1989). The *n*-transform F_n of an image or object is obtained by multiplying each coordinate vector $\mathbf{x}_i$ n times by itself, the image being centered at the center of mass. It is shown in (Marola, 1989) that both the number and the location of all the axes of symmetry can be computed from the center of mass of F_{n_i} and $F_{n_{i+1}}$, where $n_1, n_2, \ldots$ are the values of n for which the center of mass is not located at the origin. A coefficient of symmetry β is also introduced. The exact location of the symmetry axis may be found by maximizing β, starting from the axis provided by the *n*-transform. For this purpose, β is decomposed into a power series. The complete derivation can be found in (Marola, 1989). This method is mainly devoted to symmetrical and almost-symmetrical objects and has been shown to require long processing time. Especially when the objects tend to be asymmetrical, the result obtained after optimization of β

can strongly diverge from the result initially obtained with the *n*-transform.

A method to extract rotational and reflectional symmetry, that is symmetry about a point and an axis respectively, is proposed in (Masuda et al., 1993). It is based on successive matching of transformed versions of the image to obtain the best match. The directional correlation of edge features is used to evaluate the matching. This method may require a large number of iterations to find the symmetry axis location. However, the advantage of this method is that it does not require any knowledge about the centroid, and that it also allows for the detection of symmetry properties in almost-symmetrical images.

The automatic detection of rotational symmetry is presented in (Leou and Tsai, 1987). This method, devoted to shape analysis, is based on the determination of the number of intersections between a closed curve *S* and the circle centered at the centroid of *S* with a radius equal to the mean distance to *S*. However, this method works only for symmetrical objects.

In (Zielke, 1993) a method for detecting mirror symmetry using gray level information and local orientation applied to the recognition and tracking of cars is presented. This method, however, assumes that only vertical or near vertical symmetry are of interest. In (Sun, 1995) the orientation of the gradient vectors in the gray scale image is used to compute the symmetry axis orientation. This method, like all gradient based methods, has shown to be excessively noise sensitive. Symmetry has been treated as a continuous feature in (Zabrodsky et al., 1995). They define a symmetry distance that allows for reconstructing symmetry of occluded shapes. Symmetry in planar shapes is discussed in (Van Gool et al., 1995). Other approaches based on shape information and polygonal representation can be found in (Davis, 1977) and (Parui and Dutta Majumder, 1983). A more general approach to finding symmetry axes of skewed objects can be found in (Friedberg, 1986).

Most of the studies do not make use of a symmetry measure for the symmetry axis computation. The introduction of such a measure becomes necessary when dealing with images that might be clearly asymmetrical but where symmetry must be quantified. This approach generates new constraints that must be taken into account. For example, the segmentation problem, which should separate the object from the background, must be solved in advance. The reason is that we use a symmetry measure that *must not* be influenced by the image frame or any smooth window.

2.3. Symmetry and automated skin cancer diagnosis

The symmetry measure applied to digital images of pigmented skin lesions has been presented in previous papers (Stoecker et al., 1992; Gutkowicz-Krusin et al., 1997). Usually, the extraction of this feature is part of a complete classification scheme, including other features like color and border regularity. In (Stoecker et al., 1992), classification results using the principal component method are shown. The classification uses a coefficient of asymmetry, although this notion is not uniquely defined, and a threshold to separate symmetrical and asymmetrical lesions. The results are then compared with those given by dermatologists. An interesting extension would be to compute the statistical distribution of the symmetry values according to the type of lesion to discriminate benign from malignant melanocytic lesions. One may wonder how accurate is the evaluation of a subjective parameter like the degree of asymmetry done by a physician, and if he is able to reproduce this measure.

In (Gutkowicz-Krusin et al., 1997) the principal component approach has also been applied to dermatoscopic images. However, the authors point out that computing axes from a binary image is only shape dependent, and that the inner structure of the lesion should be included to account for the inhomogeneous distribution of the pigmentation.

The methods used up to now are of low complexity and allow a quick computation but they have a number of serious shortcomings. First, the symmetry measure and the symmetry axis computation are defined separately, which means that the former is not optimized. Secondly, the computation is mostly based on shape information, sometimes on luminance information, although clinical research has shown that color and texture are important in the evaluation of symmetry (Stolz et al., 1994). New approaches which intend to fill these gaps are presented in the following sections.

3. Optimization parameters and data representation

3.1. Symmetry measure

Ideally, the maximum symmetry axis is the one that minimizes the *visual difference* between the original and reflected objects. The problem of quantifying the difference between two objects, and more generally two images, is non trivial because the limit between similar and different objects is very subjective. In addition, the symmetry quantification should be a uniform measure, which means that a constant change in the visual difference should result in a constant change in the symmetry value. This notion of uniformity has previously been introduced for color spaces (Wyszecki and Stiles, 1982) and for video sequence coding (Van den Branden-Lambrecht and Verscheure, 1996). The correspondence between the latter and symmetry measure is immediate, since measuring the distortion between an image and a coded/decoded version is a similarity measure. Hence, these methods may be used to evaluate the degree of symmetry, but they are unfortunately too complex and need too much computation time to be used in optimization problems.

The symmetry measure used in this study is based on the *mean square error* (MSE) between the original and reflected image for a given axis:

$$\mathrm{MSE} = \mathrm{E}\left[\|\boldsymbol{\Gamma}(x,y) - \boldsymbol{\Gamma}(\tilde{x},\tilde{y})\|^2\right], \tag{1}$$

where $\boldsymbol{\Gamma}(x,y)$: $[0,c-1]\times[0,r-1]\mapsto\mathbb{R}^n$ is the vector-valued input image and n the vector dimension, (x,y) the pixel coordinate and $(\tilde{x},\tilde{y})$ the symmetric coordinate. The image dimensions are r rows and c columns. Knowing that any object can be decomposed into a symmetric component $\boldsymbol{\Gamma}_s(x,y)$ and an asymmetric component $\boldsymbol{\Gamma}_a(x,y)$ using the following relationships:

$$\boldsymbol{\Gamma}_s(x,y) = \frac{\boldsymbol{\Gamma}(x,y) + \boldsymbol{\Gamma}(\tilde{x},\tilde{y})}{2}, \tag{2}$$

$$\boldsymbol{\Gamma}_a(x,y) = \frac{\boldsymbol{\Gamma}(x,y) - \boldsymbol{\Gamma}(\tilde{x},\tilde{y})}{2}, \tag{3}$$

with $\boldsymbol{\Gamma}(x,y) = \boldsymbol{\Gamma}_s(x,y) + \boldsymbol{\Gamma}_a(x,y)$, then the asymmetric component can be considered as *symmetry noise* and the MSE given by Eq. (1) is proportional to its energy. As for signal compression applications for example, the distortion due to noise can then be measured through the *peak signal-to-noise ratio* (PSNR):

$$\mathrm{PSNR} = 10\ \log_{10}\left(\frac{(N_q-1)^2}{\mathrm{MSE}}\right), \tag{4}$$

where N_q is the number of quantization levels in the image. A symmetry coefficient can finally be defined as follows:

$$\psi = \psi(\varphi,\rho,\Gamma) = 1 - \frac{1}{1+\mathrm{PSNR}}, \tag{5}$$

where φ and ρ are two axis parameters (see Section 3.2), and $\psi\in[0,1]$. This coefficient is equal to 1 only if symmetry is perfect and decreases down to zero for increasing asymmetry. Using the PSNR is not an optimal way of quantifying the *visual difference* between two images but this kind of measure is easy to implement and reduces significantly the processing time.

The techniques developed later in this paper are not bound to the above defined symmetry coefficient. Any other measure, depending on the application, can be used instead. Therefore, in the following text, the symmetry measure will be called *fitting function*.

3.2. Parameter space

Before discussing the optimization of the fitting function, one has to define the two parameters needed to describe an axis. The usual way of describing a line with a mathematical equation is $y = ax + b$, where a is the slope and b is an offset in the y direction. To have finite values and uniform scales, the two following parameters are defined:

$$\varphi = \arctan(a), \tag{6}$$

$$\rho = -b\cos(\varphi), \tag{7}$$

where $\varphi\in[0;\pi)$ is the angle with the x axis (clock-wise) and ρ is a radial offset, i.e. the normal distance from the origin to the axis. The x axis is assumed to be left–right oriented and the y axis top–bottom oriented. The line equation becomes

$$y = \tan(\varphi)x - \frac{\rho}{\cos(\varphi)}, \tag{8}$$

when $\varphi \neq \pi/2$ and

$$y = \rho, \tag{9}$$

when $\varphi = \pi/2$. A uniform variation of the parameters φ and ρ corresponds to a uniform rotation and a uniform lateral shift respectively.

These two parameters must be normalized to be insensitive to any linear geometrical transform of the object of interest, such as rotation, scaling and shifting. A binary mask of the object, here the lesion, is obtained by segmentation (Schmid, 1999a,b). The image origin is set to the center of gravity (c_x,c_y) of this mask. The angle φ_0 of its principal component (Gonzalez and Wintz, 1992) and the largest distance $\rho_{\max}$ from the center to the object (here lesion) border are used for normalization. Finally the two following parameters are obtained:

$$\tilde{\varphi} = \varphi - \varphi_0, \tag{10}$$

$$\tilde{\rho} = \rho/\rho_{\max}. \tag{11}$$

The use of $\tilde{\varphi}$ and $\tilde{\rho}$, as well as taking (c_x,c_y) as origin, renders the processing insensitive to scale change, shifting or rotation of the object of interest. Fig. 1 illustrates the two axis parameters and the normalization components. A circular *search area* is defined, centered at the center of gravity and with radius R equal to $\rho_{\max}/2$. The reason is that searching for symmetry axes close to the object boundary is not sensible. The limit above has shown to be sufficient in our experiments, even for completely asymmetrical lesions.

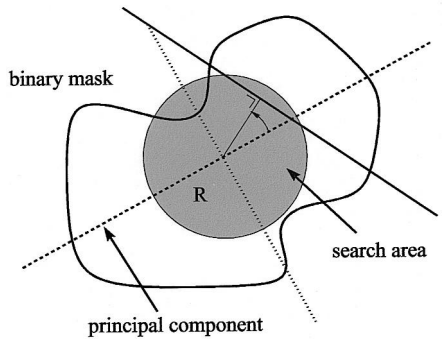

Fig. 1. Axis parameters used to construct a symmetry map: angle and offset. Both parameters are normalized as described in Section 3.2.

3.3. Feature vectors for color and texture description

Since there is no proof that symmetry in shape, color and texture are strongly correlated, especially in almost-symmetrical and asymmetrical objects, three symmetry values can be computed using the corresponding information. Therefore feature vectors that describe local color and structure must be used. The goal here is not to investigate the description of texture and color, but to integrate these components into symmetry measurements.

Starting with *RGB* color images, one has to compute the frames which will be used for extracting color and texture information. The latter may be simply the combination of the three color input components. From a spectral point of view, however, these components also contain luminance information. It seems that the luminance component strongly influences our color perception, however it should be established in future studies if luminance can be skipped from the color representation of pigmented skin lesions. In this study we have used texture descriptors extracted from the luminance component and two filtered chrominance components as color descriptors. The light absorption in pigmented skin lesions is mainly due the melanin, which is actually the pigment. The density of melanin will control the light absorption by the skin, and therefore the luminance component will contain most of the structural information. This is the reason we have chosen the luminance component in our study. It can be very well obtained from a *uniform color space* like the *CIE* $L^*u^*v^*$ color space (Wyszecki and Stiles, 1982). While the luminance component L^* may not render pigmented structures in an optimal way, we have noticed that all visible structures are present in this component. Uniform color spaces have the advantage of defining color distance as the Euclidean distance between components, which is in agreement with the MSE computation. Low-pass versions of the u^* and v^* components are used as color descriptors.

The use of Gabor filters for producing texture features has shown to be very powerful for texture discrimination (Turner, 1986). However, it is not suited in this case, since the filters are directional. Thus, symmetrical features must have symmetrical directions, which depends on the symmetry axis orientation. Isotropic filters are needed to avoid this dependency. In this study a filter bank built of isotropic Gaussian filters has been used. The filter equations are

$$G_i(\omega_r) = \exp\left(-\frac{(\omega_r - \omega_{r_i}^0)^2}{2\sigma_{r_i}^2} \right), \tag{12}$$

$$\omega_{r_i}^0 = \omega_{\min} + \sigma\{1 + 3(2^{i-1} - 1)\}, \tag{13}$$

$$\sigma = (\omega_{\max} - \omega_{\min})/2(2^M - 1), \tag{14}$$

$$\sigma_{r_i} = \sigma 2^{i-1}, \tag{15}$$

with M the number of filters, $\omega_{\max}$ and $\omega_{\min}$ the frequency limits and $1 \leqslant i \leqslant M$. ω_r is the radial spatial frequency. These equations have been derived from Gabor filters (Bigün, 1994), which have shown to be one of the most efficient methods for texture description.

3.4. Windowing

The fitting function must be computed from the pixels which constitute the object of interest. If all the image pixels are considered, the system will detect the image frame symmetry axes. In addition, the same object at different scales would give different symmetry values. In our approach only pixels within the lesion are used while the reflected pixels may be anywhere within the image frame. Pixels reflected outside the image are considered to have a symmetric value of zero.

Fig. 18 shows a pigmented skin lesion that will be used throughout this paper to illustrate different techniques. Binary masks are obtained from two different techniques, a segmentation and a contour detection technique, as illustrated in Fig. 19. In the former case only the outer boundary has been kept.

4. Optimization schemes

4.1. Genetic algorithm

We will first investigate optimization techniques that might be used to obtain the best fitting symmetry axis as quickly as possible. To this end, two techniques are proposed, one based on a *genetic algorithm* (GA), and the other derived from the *self-organizing maps* (SOM) theory. The former is presented in this paragraph while the latter is presented in Section 4.2.

A GA is an active search method that uses operators inspired by evolutionary processes. The parameters used for the optimization are seen as *alleles* and are part of a *chromosome*. The goal of this technique is to find the best chromosome, i.e. the set of parameters that maximizes a given function. The ‘quality’ of the chromosome should improve through the different generations. A new generation is obtained from the current set of chromosomes, all encoding a different combination of parameters, using different operators, typically crossover and mutation. The reproduction stage usually mixes part of two randomly selected chromosomes. The complete scheme used in this study is shown in Fig. 2.

Most of the developments in GA deal with binary encoded chromosomes. This means that parameters can only take a limited number of values and with a predefined precision. When using real parameters, one possibility is to sample the parameter space and to define the chromosome length according to the required precision (Wright, 1991) (the variables are ‘encoded’). The direct encoding of real

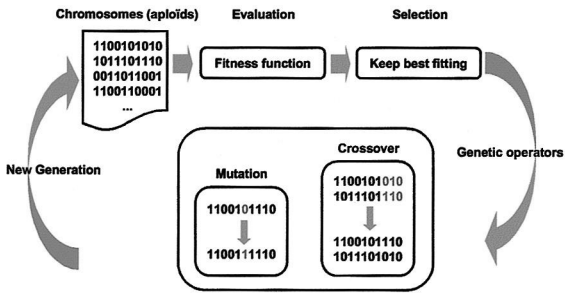

Fig. 2. Complete scheme of the genetic algorithm used in this study.

parameters has been also investigated, leading for example to *blend crossover* (BLX-α), which uniformly picks values that lie on a line passing through two parents (Eshelman and Schaffer, 1993). *Simulated binary crossover* (SBX) has been introduced in (Deb and Kumar, 1995a,b). Its search power is similar to that of single-point crossover used for binary GAs, where the search power is defined in terms of the probability of creating an arbitrary child from two parents. In this study we use a modified version of this technique.

Crossover is used for producing offsprings from a couple of parent chromosomes. In SBX the following spread factor is first defined:

$$\beta = \left| \left| \frac{c_1 - c_2}{p_1 - p_2} \right| \right|, \tag{16}$$

$$p(\beta) = \begin{cases} 0.5(n+1)\beta^n & \text{if } \beta \leqslant 1\,, \\ 0.5(n+1)/\beta^{n+2} & \text{otherwise}\,, \end{cases} \tag{17}$$

where p_1, p_2 are the two parent's real variable values and c_1, c_2 the produced offsprings. The spread factor is randomly chosen according to the above distribution (see (Deb and Kumar, 1995a) for the detailed procedure) and used to compute c_1 and c_2:

$$c_1 = 0.5[(p_1 + p_2) + \beta(p_1 - p_2)]\,, \tag{18}$$

$$c_2 = 0.5[(p_1 + p_2) - \beta(p_1 - p_2)]\,. \tag{19}$$

We propose a different probability distribution for the spread factor β, which is the combination between an *attraction* term and a *rejection* term:

$$P(\beta) = \frac{1}{\sqrt{2\pi}(\sigma_a - \sigma_r)} \left\{ \exp\left(-\frac{(\beta-1)^2}{2\sigma_a^2} \right) - \exp\left(-\frac{(\beta-1)^2}{2\sigma_r^2} \right) \right\}, \tag{20}$$

with $\sigma_r \ll \sigma_a$. The shape of this distribution is close to that obtained with the SBX and has the advantage of producing children different from their parents, which is due to the rejection term in Eq. (20). When using the SBX, the probability that the children are exactly equal to the parent is the highest. However, if one assumes that the best fitting parents are kept for the next generation, it is not desirable to duplicate them. By keeping the best fitting chromosomes, one ensures the convergence to the optimal solution. The normalizing term is not useful for a numerical implementation, which is very easy when using normal distributions (Press et al., 1994). The distribution may also be bounded without modifying the shape (we must have $\beta \in [0;\infty)$). Fig. 3 plots the original and modified distributions.

To speed up the computation, a different subset of the input data is used at each generation. The consequence is that the less symmetrical the object, the higher the introduced uncertainty. Depending on the subset, an axis may show a better symmetry than the true symmetry axis. This effect is especially strong when using active search methods like the GA, where a well fitting chromosome may be suppressed when the next generation is created. However this simplification is necessary to speed up the processing.

The optimization techniques have been evaluated in terms of robustness, that we define as the ability to reproduce results. We have reiterated the processing 50 times for every lesion in order to visualize the convergence of the algorithms. Figs. 4–6 plot the different results obtained for the lesion shown in Fig. 18. Here we use texture (red 'x') and color (green '+') feature vectors as input. Two symmetry axes are revealed, which is due to their close symmetry values. The main drawback of using GAs is that the optimal solution may need a very long convergence time. GAs are powerful for finding an approximate solution in very complex problems. Our results have shown that the successive runs lead to close but not equal results (within a given precision). Figs. 5 and 6 show results obtained with our spread factor. While Fig. 5 shows a clear improvement over the results displayed in Fig. 4,

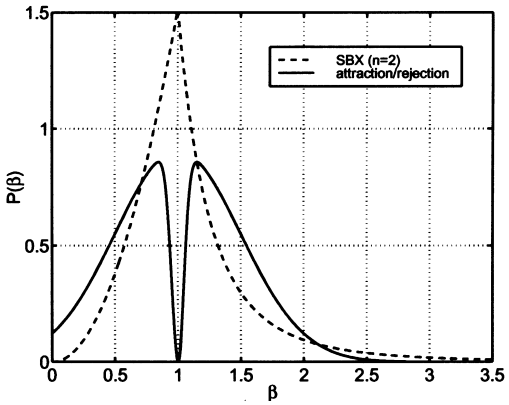

Fig. 3. SBX distribution and attraction/rejection Gaussian distribution used to produce offsprings.

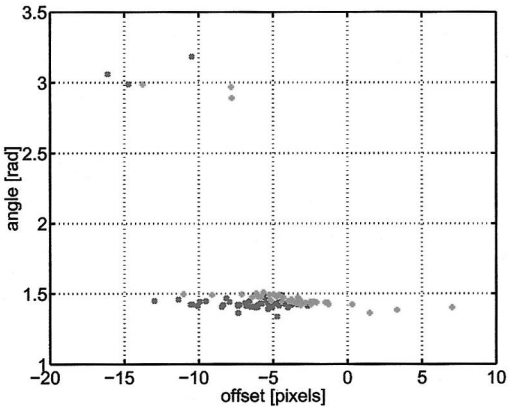

Fig. 4. The genetic algorithm has been applied 50 times on the lesion shown in Fig. 18. The SBX spread factor has been used, with 50 chromosomes, 1000 generations and 1024 input samples. The red 'x' correspond to *texture symmetry*, while the green '+' correspond to *color symmetry*.

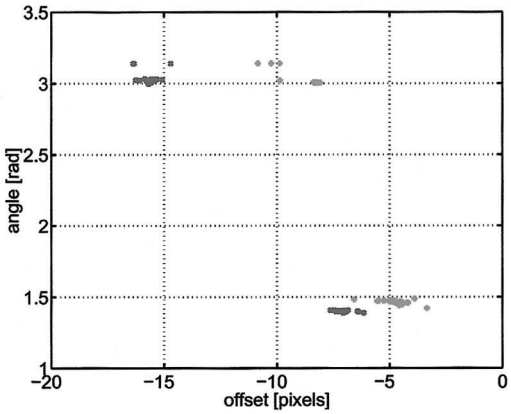

Fig. 5. Results obtained using the same approach as in Fig. 4 but with the attraction/rejection spread factor.

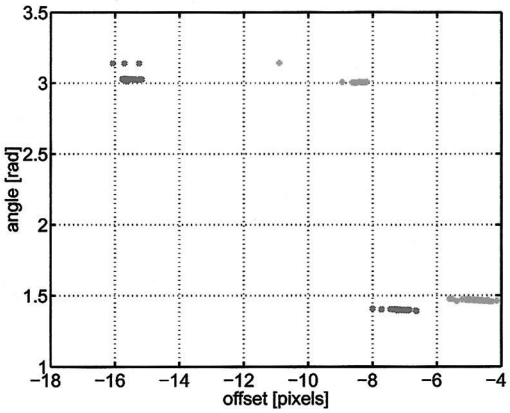

Fig. 6. Results obtained with the genetic algorithm using 50 chromosomes, 1000 generations and 4096 input samples. The attraction/rejection spread factor has been used.

increasing the data size does not improve significantly the convergence, as shown in Fig. 6.

4.2. Self-organizing maps

Using *self-organizing maps* (SOM) is a way of mapping an N-dimensional data distribution on a generally two-dimensional array. The two-dimensional map is an array of nodes to each of which is associated an N-dimensional vector. Organizing this map is an iterative procedure based on *competition* between the different nodes: at each iteration an input vector is randomly selected and associated to the best fitting node (the reference node) using a specific metric. The vectors associated to the different nodes are updated using the following relationship (Kohonen, 1995):

$$\boldsymbol{m}_i(t+1) = \boldsymbol{m}_i(t) + h_{ci}(t)[\boldsymbol{x}(t) - \boldsymbol{m}_i(t)]\,, \tag{21}$$

where $t = 0,1,2,\ldots$ is the iteration number (time variable), $\boldsymbol{x}(t)$ is the randomly selected input vector and $h_{ci}(t)$ is called the *neighborhood function*. Subscript i indicates the node number and c the reference node number.

The neighborhood function weights the update term $\boldsymbol{x}(t) - \boldsymbol{m}_i(t)$ according to the distance of node i to the reference node c within the map. Different functions may be defined under the condition that $h_{ci} \to 0$ with increasing $\|\boldsymbol{r}_c - \boldsymbol{r}_i\|$, where $\boldsymbol{r}_i \in \mathbb{R}^2$ is the position of node i. It is also necessary for convergence that $\lim_{t\to\infty} h_{ci}(t) = 0$, i.e. the vector update decreases in time.

One may wonder how this technique can be used to solve an optimization problem. Let the N-dimensional feature space be the parameter space (in our case $n = 2$). A random set of input vectors is selected and the fitting function is approximated. The closest vector in the map is the reference vector. The nodes are updated using a modified version of Eq. (21):

$$\boldsymbol{m}_i(t+1) = \boldsymbol{m}_i(t) + f(\psi)h_{ci}(t)[\boldsymbol{x}(t) - \boldsymbol{m}_i(t)]\,. \tag{22}$$

The function $f(\psi)$ (for example $f(\psi) = \psi$) is introduced to weight the update according to the fitting function value.

Since we are seeking a unique set of parameters, that is the axis parameters which maximize the fitting function, an alternative approach is to ensure the convergence of all the nodes to a unique solution. This 'moving map' allows for searching actively the parameter space with the advantage of improving the search precision when the nodes are close to the optimal parameter vector. To achieve this result, the fitting function is evaluated for every vector in the map. The best fitting vector becomes the reference vector.

In the case of the SOM-based method, the slowly moving map guarantees optimal results at the end. The same experiment that we did in Section 4.1 for the genetic algorithm has been repeated for the SOM technique. Fig. 7 shows the result we obtained using an 11×11 grid and three processing steps of respectively 7000, 2000 and 500

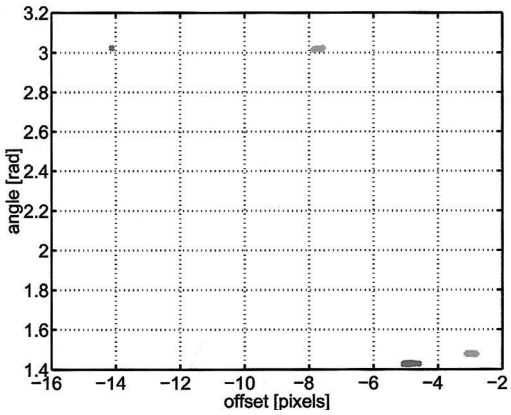

Fig. 7. Results obtained with the modified SOM technique. In the above example, 32 input samples have been used and three steps of 7000, 2000 and 500 iterations have been necessary.

iterations. At every step the neighborhood size and amplitude have been decreased, and only 32 input samples have been used to reduce the computation time. Both symmetry axes are still present but even if they have very close values the system is able, most of the time, to find the best fitting axis.

5. The symmetry map

5.1. Raw processing

Another approach is to compute the fitting function ψ for $0 \leqslant \tilde{\varphi} \leqslant \pi$ and $-0.5 \leqslant \tilde{\rho} \leqslant 0.5$. The result is a *symmetry map* (SM) where all the maxima correspond to local optima of the fitting function and where the global maximum corresponds to the best fitting symmetry axis (i.e. the axis that maximizes the fitting function, as previously defined). For an SM of size $m \times n$, $\psi(i,j) = \psi(\varphi,\rho,\Gamma)$ with

$$\varphi = i\,\frac{\pi}{m} + \varphi_0\,, \tag{23}$$

$$\rho = \left(j\,\frac{1}{n-1} - 0.5\right)\rho_{\max}\,, \tag{24}$$

and $0 \leqslant i \leqslant m-1$ and $0 \leqslant j \leqslant n-1$.

Fig. 8 shows a typical result obtained for a melanocytic skin lesion[1]. The main drawback of this approach is the processing time (about 1 h on a Silicon Onyx computer for a 256×256 SM).

To reduce the processing time, the sampling precision in the parameter space is reduced. Fig. 9 shows the result obtained when reducing the precision by a factor 5. The full resolution is obtained by interpolation. The effect is a small shift of the maxima location and therefore the procedure requires the use of local optimization methods (Press et al., 1995) to find the exact maxima location. Fig. 10 shows the result obtained by using a reduced subset (5%) of input samples. The fitting function is only approximated and the result is a noisy SM. A (very) low-pass filter can be used to smooth out the map, as shown in Fig. 11, but it introduces a strong border effect. The result is a strong shift of the maxima location. Because the processing time is still long for a reduced precision, more elaborated methods are needed to construct an SM.

[1] A unique lesion has been used for the different examples given through this paper to allow the comparison, and the different SMs are scaled and displayed using a non-uniform lookup table.

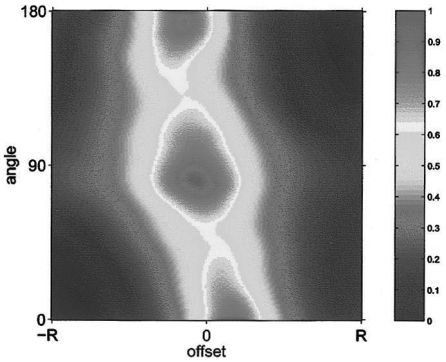

Fig. 8. Symmetry map obtained by direct computation.

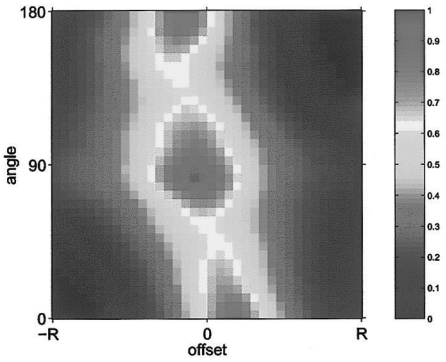

Fig. 9. The same symmetry map as that shown in Fig. 8 has been recomputed with a larger sampling step.

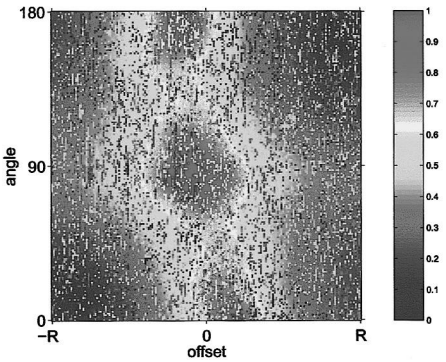

Fig. 10. Symmetry map computed with a randomly chosen subset of image pixels.

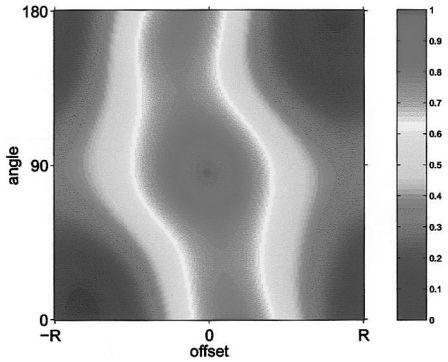

Fig. 11. The symmetry map shown in Fig. 10 has been smoothed out with a Gaussian low-pass filter.

In the next paragraphs, two multi-resolution (MR) approaches for computing an SM are presented.

5.2. Multi-resolution parameter space

The first multi-resolution approach proposed in this paper uses successive refinements of the parameter space to obtain the final SM. For an SM of size $m \times n$, the fitting function is approximated by

$$\psi(i,j) \simeq \begin{cases} \sum_{x \in S} \psi_x(i+\hat{i}, j+\hat{j}) \\ \quad \text{if } i \bmod 2^r \equiv 0 \text{ and } j \bmod 2^r \equiv 0, \\ \psi(i-(i \bmod 2^s), j-\bmod 2^s) \\ \quad \text{otherwise}, \end{cases} \tag{25}$$

where s is the scale ($s = 0$ is full resolution), $0 \leqslant i \leqslant m-1$ and $0 \leqslant j \leqslant n-1$. To reduce the computation time, only a subset S of size N of the image samples is used ($\boldsymbol{x} \in S \subset [0,r-1] \times [0,c-1]$, where r and c are respectively the number of rows and columns in the image). $\hat{i}$ and $\hat{j}$ are two random numbers having a normal distribution $N(\mu = (2^s-1)/2, \sigma^2 = \mu^2)$. The introduction of $\hat{i}$ and $\hat{j}$ is intended to give at every sampling location a local value instead of a punctual value. Fig. 12 shows how the sampling grid is organized.

φ and ρ are obtained using Eqs. (22) and (23), respectively. After each iteration, only regions having a value higher than a given threshold are kept for further processing.

The second multi-resolution approach presented in the next paragraph uses the true fitting function computed for the input image at different resolutions.

5.3. Pyramidal image representation

Here we use a pyramidal representation of the input image (Burt and Adelson, 1983; Burt, 1981): the image size is successively reduced by a factor of two to generate a down-sampled version of the previous image. At scale s the down-sampled image is given by

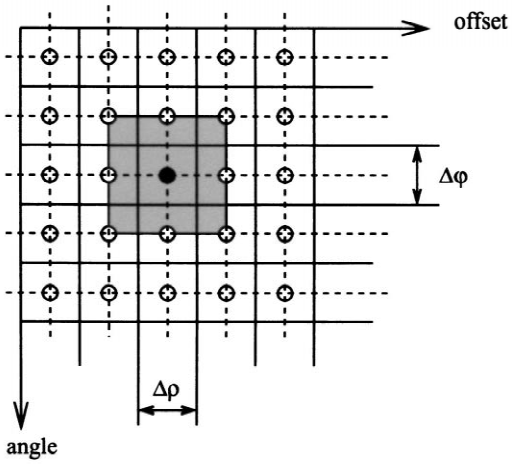

Fig. 12. Organization of the sampling grid used to compute a multi-resolution symmetry map.

$$\boldsymbol{\Gamma}^{(s)} = \boldsymbol{\Gamma}^{(0)}(2^s(\boldsymbol{x}+\boldsymbol{1})-\boldsymbol{1}), \tag{26}$$

where $\boldsymbol{x} = (x,y)$ is the pixel coordinate. The image origin is at (0,0).

To avoid aliasing, a *bicubic interpolation kernel* called Mexican hat is used, given by (Sonka et al., 1993)

$$h_3 = \begin{cases} 1-2|x|^2+|x|^3 & \text{for } 0 \leqslant |x| < 1, \\ 4-8|x|+5|x|^2-|x|^3 & \text{for } 1 \leqslant |x| < 2, \\ 0 & \text{otherwise}, \end{cases} \tag{27}$$

and

$$\int_{-\infty}^{\infty} h_3(x)\,\mathrm{d}x = 1. \tag{28}$$

A typical filter size for a 0.25 cut-off frequency is $N = 7$.

The SM is initially computed using a small version of the image, in our case 1/16 of the initial size (512×512). The obtained SM is a low-pass version of the true SM, and only values above a given threshold are recomputed at a finer scale. The procedure is repeated until the initial image size has been reached.

Fig. 13 shows an SM obtained with the multi-resolution approach described in the previous paragraph. The main drawback of this approach is the amount of noise produced by the fitting function approximation. Even though it can be removed using morphological filtering, one looses precision in the maxima location. The real improvement compared to the direct computation is the processing time (up to a factor 60). On a 400 MHz PC, the 256×256 SM computation takes less than 1 minute. Best results are obtained using a pyramidal image representation, as shown in Fig. 14. The SM is smooth and the maxima are precisely located (up to the grid resolution).

To obtain the different maxima in an SM, we used a simple but efficient procedure. A morphological opening

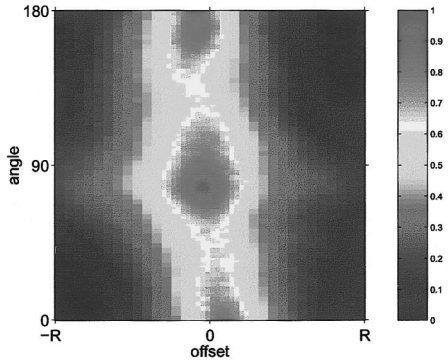

Fig. 13. Symmetry map obtained using a multi-scale approach.

operator (Gonzalez and Wintz, 1992) with a cylindrical structuring element is applied to the SM, followed by a gradient computation. We assume that regions of null gradient contain a peak or maximum. In all our experiments, this simple procedure allowed us to extract the full list of local maxima in the SMs.

6. Effect of noise

This section gives the analytical value of the expected MSE between a *noisy* image and its reflected version.

The MSE is given by $\mathrm{E}[(I + n - \tilde{I} - \tilde{n})^2]$, where $I(x,y)$ is the image intensity function and $n(x,y)$ the added noise. The tilde denotes the symmetrical value. While $\mathrm{E}[I]$ and $\mathrm{E}[n]$ remain constant, the values of $\mathrm{E}[\tilde{I}]$ and $\mathrm{E}[\tilde{n}]$ vary with the symmetry axis. However, for large objects and assuming that most of the reflected object is inside the image, we can write $\mathrm{E}[n] = \mu \cong \mathrm{E}[\tilde{n}]$. Since the image and the noise are uncorrelated, and the noise has no spatial correlation, we can write

$$\mathrm{E}[n\tilde{n}] = \mathrm{E}[n]\mathrm{E}[\tilde{n}] \cong \mu^2 , \tag{29}$$

$$\mathrm{E}[In] = \mathrm{E}[I]\mathrm{E}[n] = \mathrm{E}[I]\mu , \tag{30}$$

$$\mathrm{E}[\tilde{I}\tilde{n}] = \mathrm{E}[\tilde{I}]\mathrm{E}[\tilde{n}] \cong \mathrm{E}[\tilde{I}]\mu . \tag{31}$$

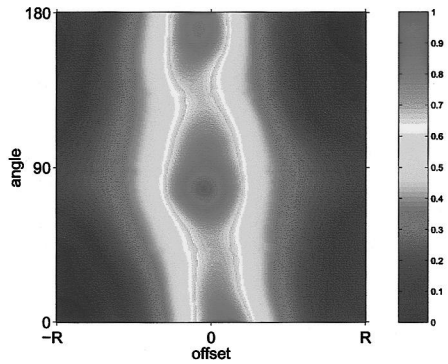

Fig. 14. Symmetry map obtained using a pyramidal image representation.

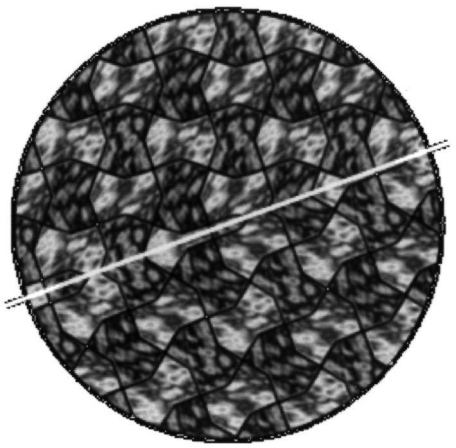

Fig. 15. Symmetry axis of synthetic object with circular shape and symmetrical pattern. Both the gray-level principal component approach and the optimization approach give the true axis.

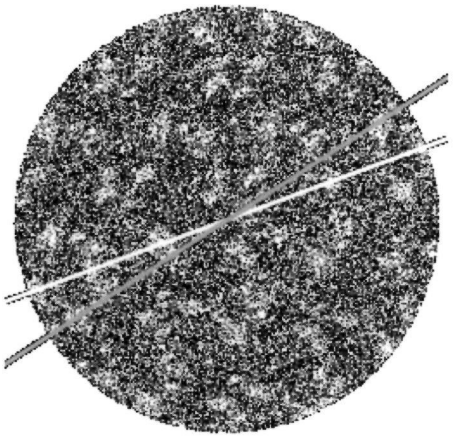

Fig. 16. Synthetic object corrupted with zero mean Gaussian noise. The result deviates from the true axis for the KL approach (gray axis), while the optimization approach is successful.

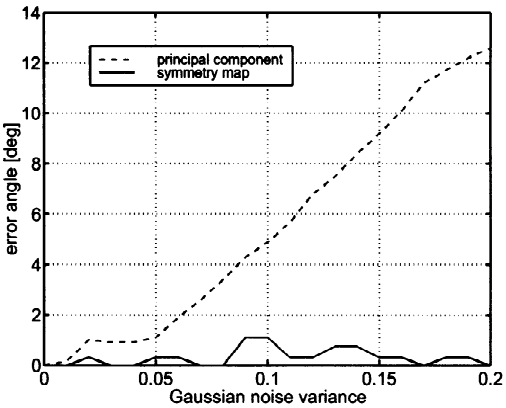

Fig. 17. The two curves show the angular difference between the true symmetry axis and the axis obtained when the image is corrupted by noise, for the synthetic image shown in Fig. 15 and for both the KL and symmetry optimization approaches.

The MSE becomes

$$\begin{aligned}\mathrm{E}[(I+n-\tilde{I}-\tilde{n})^2] &\cong \mathrm{E}[(I-\tilde{I})^2]+\mathrm{E}[(n-\tilde{n})^2]\\ &\cong \mathrm{E}[(I-\tilde{I})^2]+2\mathrm{E}[n^2]-2\mu^2\\ &= \mathrm{E}[(I-\tilde{I})^2]+2\sigma^2 . \end{aligned} \quad (32)$$

The added noise affects *globally* the SM, which means that existing maxima do not move and new maxima do not appear. This result can be easily extended to vector-valued images.

Fig. 15 shows a circular object with symmetrical pattern. Since the symmetry is perfect, both the gray-level principal component and the symmetry optimization approaches give the exact result. When symmetry is imperfect, as illustrated with the noisy version shown in Fig. 16, only the symmetry optimization approach is robust and still provides the exact result. Fig. 17 plots the angular deviation from the true symmetry axis versus the Gaussian noise variance for both the KL and symmetry optimization approaches. While the former shows increasing sensitivity, the latter performs well up to the numerical approximation errors.

7. Application to digital dermatoscopy

7.1. Material and methods

A set of 50 malignant melanoma (MM) and 50 benign melanocytic lesions, along with their histology, has been used to illustrate and evaluate the proposed technique. For better comparison with the approach used by physicians and other automated techniques, only the two axes with highest local symmetry value have been used for every lesion. A symmetry feature vector of dimension 6 has finally been obtained by separating texture, color and shape information.

A linear classifier with training by epoch (Schalkoff, 1992) has been used to classify the lesions. This method uses a gradient descent approach to determine the coefficients of the hyperplane that minimizes the misclassification error. Since we evaluate a single feature, the use of a linear classifier is acceptable.

A number of measurements can be used to quantify the classification accuracy. In our investigation we have evaluated the *sensitivity* (SE), which is the proportion of MM that have been classified as MM, the *specificity* (SP), which is the proportion of lesions that are not MM and that have been classified accurately, and finally the *diagnosis accuracy* (DA), which is the proportion of cases in which the classification fits the diagnosis. These three measures are respectively defined by the following relationships:

$$\text{sensitivity} = \frac{\mathrm{TP}}{\mathrm{TP}+\mathrm{FN}}, \quad (33)$$

$$\text{specificity} = \frac{\mathrm{TN}}{\mathrm{TN}+\mathrm{FP}}, \quad (34)$$

$$\text{diagnostic accuracy} = \frac{\mathrm{TP}}{\mathrm{TP}+\mathrm{FP}+\mathrm{FN}}, \quad (35)$$

where TP is the number of true positives, FN is the number of false negatives, TN is the number of true negatives, and FP is the number of false positives. A ‘positive’ is a malignant melanoma.

The Karhunen-Loève transform (KL), which is widely used to quickly compute the direction of maximum data variance in binary or gray-level objects, has been also evaluated with the same set of images. The orientation of a 2D object obtained through this technique is often associated with the symmetry axis, even if the relationship is not straightforward, except for intrinsically symmetrical objects. For the binary KL transform, a binary mask of every lesion is needed. An asymmetry index is computed for both (orthogonal) axes using the following definition (Stoecker et al., 1992):

$$\text{Asymmetry index} = \frac{\Delta A}{A}, \quad (36)$$

where ΔA is the non-overlapping area between the original and reflected mask, and A is the area of the original mask. This index is computed for both axes. In (Stoecker et al., 1992), only the minimum value is kept. Here we will keep both values. For the gray level KL transform, the L^* component has been used to compute the following asymmetry index (Gutkowicz-Krusin et al., 1997):

$$\text{Asymmetry index} = \frac{\sum |\Gamma(\mathbf{x}) - \Gamma(\tilde{\mathbf{x}})|}{\sum \Gamma(\mathbf{x})}, \quad (37)$$

where $\Gamma(\mathbf{x})$ is the original gray level image and $\Gamma(\tilde{\mathbf{x}})$ the reflected version. The index is computed for both axes and the image is masked with a binary mask of the lesion before processing.

Different results are shown in the following paragraph and discussed in Section 7.3.

7.2. Results

Figs. 20 and 22 give two examples of texture and color SMs obtained for pigmented skin lesions, and the corresponding axes are superimposed on the lesions in Figs. 21 and 23, respectively. In the former case, the lesion is benign and symmetrical, while in the second case it is malignant and asymmetrical. In any case, at least one axis is found, namely the axis that maximizes the symmetry measure. Therefore even asymmetrical objects have a symmetry axis, but with a low symmetry value. It is interesting to note that when lesions are asymmetrical they have more than two local maxima, as can be seen from Fig. 23. This characteristic may be exploited in future studies.

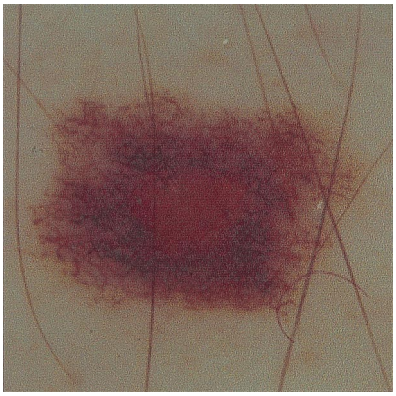

Fig. 18. Melanocytic skin lesion used throughout this paper to illustrate the different optimization schemes.

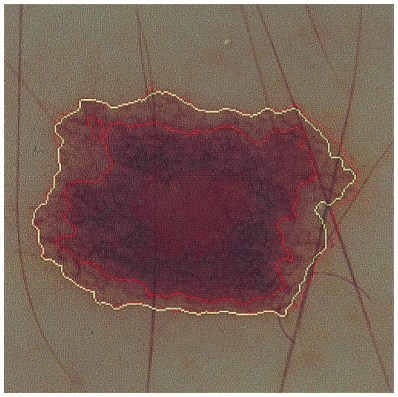

Fig. 19. Lesion boundaries obtained through color clustering (Schmid, 1999) (red) and nonlinear isotropic diffusion and morphological flooding (Schmid, 1999) (green).

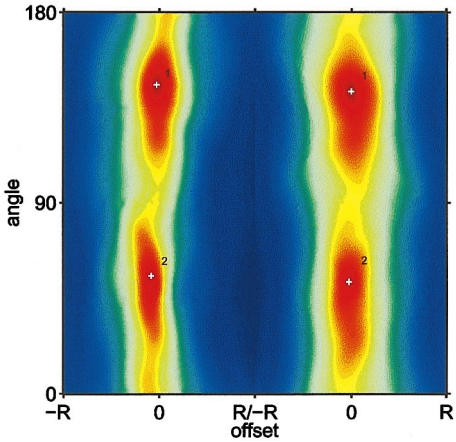

Fig. 20. Texture (left) and color (right) symmetry maps of the lesion shown in Fig. 18.

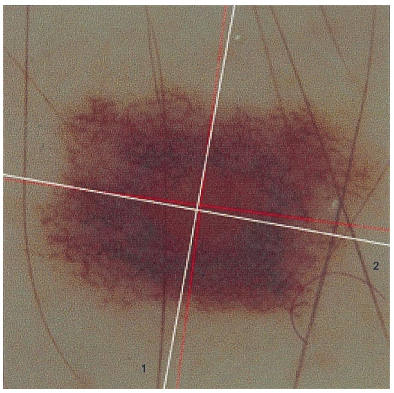

Fig. 21. Axes obtained from the map of Fig. 20 (red for color, white for texture).

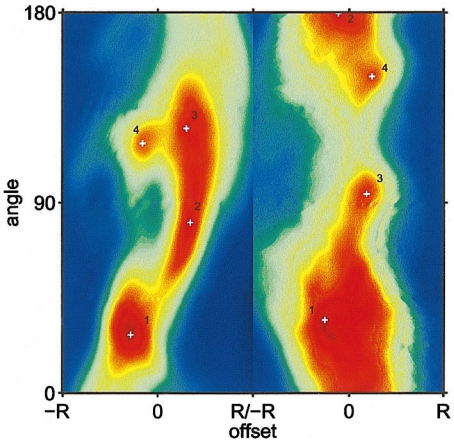

Fig. 22. Texture (left) and color (right) symmetry maps of asymmetric lesion.

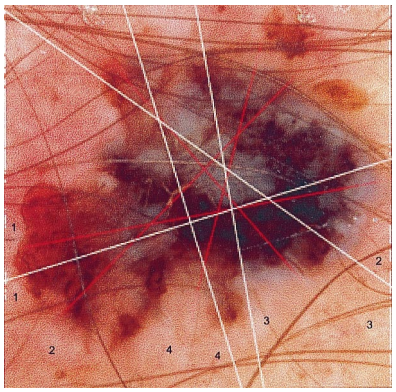

Fig. 23. Axes obtained from the map of Fig. 22 (red for color, white for texture).

Table 1
Linear classification of symmetry values for benign and malignant lesions

	Optimization	Binary KL	Gray-level KL
TP	39	30	35
FP	4	9	4
TN	45	41	46
FN	11	20	15
SE	78%	60%	70%
SP	90%	82%	92%
DA	72.2%	50.8%	64.8%

Table 1 gives the different parameter values for different symmetry features: the 6-D vector obtained from the shape, color, and texture SMs, a 2-D vector obtained with the binary KL transform, and a 2-D vector obtained with the gray-level KL transform.

7.3. Discussion

Table 1 shows that the optimization approach combining shape, color and texture information improves the separation between benign and malignant lesions. Especially, shape and texture are relevant features, while the color symmetry does not significantly improve the results. Malignant melanoma show to be rather asymmetrical, but this criteria is not sufficient to separate malignant from benign lesions. Different measurements have been repeated with the masks obtained through both the color segmentation and contour detection techniques. While our approach did not provide different classification results, this was not the case for the binary KL approach, which is intuitively appealing. In general, the binary approach gives weak results while the gray-scale approach improves significantly the classification.

It is interesting to note from the above results that a very simple approach such as the gray-level principal component decomposition performs close to the more elaborated optimization approach developed in this paper. One possible reason is that this method uses asymmetry indexes computed from two orthogonal axes, thus evaluating the asymmetry criteria in two orthogonal directions. For the different SMs the two largest values have been used, which do not always correspond to orthogonal axes, especially for malignant melanoma. It is therefore important to extend this work to the extraction of all the local maxima revealed by the SMs as well as different criteria such as their respective location. According to our observations, it may be a very promising way for answering the symmetry problem in dermatoscopic images. A symmetry map contains much more information than only the parameters of the axis that minimizes the difference between the original and reflected images. Local maxima are also detected, as well as their respective location in the symmetry map. This distribution of maxima can be analyzed to improve the classification based on symmetry measures. The more information we use, the better we can describe a single object. This kind of analysis is left for future studies.

Another important study that must follow these initial results is the evaluation of different texture and color descriptors. The features used in this study have shown to work well on synthetic images and in noisy environments but they may not be optimal for dermatoscopic images.

8. Summary and conclusions

The problem of finding the axis that maximizes a given symmetry measure has been investigated in this paper. An improved version of the SBX genetic algorithm has been proposed, as well as the use of self-organizing maps for optimization applications. The notion of symmetry map (SM) has been introduced, which provides a 2-D mapping of any object in a symmetry space. The symmetry properties can be completely analyzed from this representation. Two multi-resolution techniques are provided to improve the processing time of SMs. A coarse-to-fine technique works on the sampling grid of the parameter space, while a pyramidal approach is used to select regions of interest in the SM that will need better resolution. The use of three SMs, for shape, color and texture information, is illustrated for the detection of malignant melanoma. Initial results have shown that our approach performs better than classical principal component approaches used so far for this application.

Beside their applications to computer-aided diagnosis systems, as we use them for skin cancer detection, the applications of SMs are various:

- object recognition,
- database retrieval,
- medical diagnosis.

The SM are especially interesting because they provide a signature that may be compared with that of other objects, and lead to a classification. The two search methods, GA and SOM, and especially the latter one are very fast when only the detection of the axis with highest symmetry is needed. In conclusion, our future work will focus on the extraction of optimal color and texture descriptors, as well as the analysis of SMs for recognition tasks.

Acknowledgements

The author would like to thank Dr. Joël Guillod from the CHUV-DHURDV, Lausanne, Switzerland for his help and for providing the images, and the anonymous reviewers for their fruitful comments. Finally, the author would like to thank Prof. Dr. Murat Kunt for his support. This research is funded by the Swiss National Science Foundation, fund no. 3252-053175.

References

Argenyi, Z.B., 1997. Dermatoscopy (epiluminescence microscopy) of pigmented skin lesions. Dermatol. Clin. 15 (1), 79–95.

Bigün, J., 1991. Multidimensional orientation estimation with application to texture analysis and optical flow. IEEE Trans. Pattern Anal. Machine Intelligence 13 (8), 775–790.

Bigün, J., 1994. N-folded symmetries by complex moments in Gabor space and their application to unsupervised texture segmentation. IEEE Trans. Pattern Anal. Machine Intelligence 16 (1), 80–87.

Binder, M., Schwarz, M., Winkler, A., Steiner, A., Kaider, A., Wolff, P., Pehamberger, 1995. Epiluminescence microscopy; a useful tool for the diagnosis of pigmented skin lesions for formally trained dermatologists. Arch. Dermatol. 131, 286–291.

Burt, P.J., 1981. Fast filter transforms for image processing. Computer Graphics and Image Processing 16, 20–51.

Burt, P.J., Adelson, E.H., 1983. The Laplacian pyramid as a compact image code. IEEE Trans. Commun. COM-31 (4), 532–540.

Davis, L.S., 1977. Understanding shape: II. symmetry. IEEE Trans. on Systems, Man, and Cybernetics SMC-7 (3), 204–212.

Deb, K., Kumar, A., 1995a. Real-coded genetic algorithms with simulated binary crossover: Studies on multimodal and multiobjective problems. Complex Systems 9 (6), 431–454.

Deb, K., Kumar, A., 1995b. Simulated binary crossover for continuous search space. Complex Systems 9 (1), 115–148.

Eshelman, L.J., Schaffer, J.D., 1993. Real-coded genetic algorithms and interval-schemata. In: Whitley, L.D. (Ed.), Foundations of genetic algorithms. Morgan Kaufmann, Vail, Colorado, Vol. 2, pp. 187–202.

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

Gonzalez, R.C., Wintz, P.A., 1992. Digital Image Processing. Addison-Wesley, Reading, MA.

Green, A., Martin, N., Pfitzner, J., O'Rourke, M., Knight, N., 1994. Computer image analysis in the diagnosis of melanoma. J. Am. Acad. Dermatol. 31, 958–964.

Gutkowicz-Krusin, D., Elbaum, M., Szwaykowski, P., Kopf, A.W., 1997. Can early malignant melanoma be differentiated from atypical melanocytic nevus by in vivo techniques? Part II. automatic machine vision classification. Skin Res. Technol. 3, 15–22.

Kohonen, T., 1995. Self-Organizing Maps. Springer Series in Information Sciences. Springer.

Leou, J.-J., Tsai, W.-H., 1987. Automatic rotational symmetry determination for shape analysis. Pattern Recognition 20 (6), 571–582.

Marola, G., 1989. On the detection of the axes of symmetry of symmetric and almost symmetric planar images. IEEE Trans. Pattern Anal. Machine Intelligence 11 (1), 104–108.

Masuda, T., Yamamoto, K., Yamada, H., 1993. Detection of partial symmetry using correlation with rotated-reflected images. Pattern Recognition 26 (8), 1245–1253.

Parui, S.K., Dutta Majumder, D., 1983. Symmetry analysis by computer. Pattern Recognition 16 (1), 63–67.

Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P., 1994. Numerical recipes in C: The Art of Scientific Computing, 2nd Edition. Cambridge University Press, Cambridge.

Schalkoff, R., 1992. Pattern Recognition: Statistical, Structural and Neural Approaches. Wiley.

Schmid, P., 1997. Symmetry measure for the computer aided diagnosis of pigmented skin lesions. In: Lemke, H.U., Vannier, M.W., Inamura, K. (Eds.), Proceedings of the 11th International Symposium in Computer Assisted Radiology and Surgery. Berlin, Germany, 25–28 June 1997. Elsevier Science, p. 998.

Schmid, P., 1999a. Lesion detection in dermatoscopic images using anisotropic diffusion and morphological flooding. In: Proceedings of the International Conference on Image Processing (ICIP-99), Kobe, Japan, 25–28 October 1999. Vol. 3, pp. 449–453.

Schmid, P., 1999b. Segmentation of dermatoscopic images by 2D color clustering. IEEE Trans. Med. Imag. 18 (2), 164–171.

Sonka, M., Hlavac, V., Boyle, R., 1993. Image Processing, Analysis and Machine Vision, 1st Edition. Chapman & Hall Computing, London, UK.

Stoecker, W.V., Li, W.W., Moss, R.H., 1992. Automatic detection of asymmetry in skin tumors. Computerized Med. Imag. Graphics 16 (3), 191–197.

Stolz, W., Braun-Falco, O., Landthaler, M., Bilek, P., Cognetta, A.B., 1994. Color Atlas of Dermatoscopy. Blackwell Science.

Sun, C., 1995. Symmetry detection using gradient information. Pattern Recognition Lett. 16, 987–996.

Turner, M.R., 1986. Texture discrimination by Gabor functions. Biol. Cybernetics 55, 71–82.

Tyler, C.W. (Ed.), 1996. Human Symmetry Perception and its Computational Analysis. VSP, Utrecht, The Netherlands.

Van den Branden-Lambrecht, C.J., Verscheure, O., 1996. Perceptual quality measure using a spatio-temporal model of the human visual system. In: Proc. SPIE, San Jose, CA, 28 January–2 February 1996. Vol. 2668, pp. 450–461.

Van Gool, L., Moons, T., Ungureanu, D., Pauwels, E., 1995. Symmetry from shape and shape from symmetry. Int. J. Robotics Res. 15 (1), 407–424.

Wright, A.H., 1991. Genetic algorithms for real parameter optimization. In: Rawlins, G.J.E. (Ed.), Foundations of Genetic Algorithms. Morgan Kaufmann, San Francisco, California, pp. 205–218.

Wyszecki, G., Stiles, W.S., 1982. Color Science: Concepts and Methods, Quantitative Data and Formulae. Wiley, New York.

Zabrodsky, H., Peleg, S., Avnir, D., 1995. Symmetry as a continuous feature. IEEE Trans. Pattern Anal. Machine Intelligence 17 (12), 1154–1166.

Zielke, T., 1993. Intensity and edge based symmetry detection with application to car following. CVGIP: Image Understanding 58 (1), 177–190.