ARTICLE IN PRESS

ELSEVIER

Image and Vision Computing 21:221 (2003)

image AND vision COMPUTING

www.elsevier.com/locate/imavis

Detecting ellipses of limited eccentricity in images with high noise levels

K.-U. Kasemir, K. Betzler*

Universität Osnabrück, Fachbereich Physik, D-49069, Osnabrück, Germany

Received 10 September 2002; accepted 12 October 2002

Abstract

We present a newly developed algorithm for the detection of elliptical shapes in images in the presence of high noise levels. The algorithm combines a modified version of the Hough transform with a genetic algorithm, namely Differential Evolution. Suggestions for a parallel implementation are given. In our implementation the algorithm is restricted due to the technical problem to be solved, yet it can be easily generalized to arbitrary ellipse detection.

Keywords: Ellipse detection; Hough transform; Differential evolution

1. Introduction

Originally developed for the recognition of simple shapes in images like e.g. lines, the Hough transform [1] may be applied to detect any sort of curve which can be appropriately parameterized. However, the size of the parameter space grows tremendously with the number of parameters needed to represent the curve. While a straight line is described by two parameters, a circle uses three, and for an ellipse generally five are necessary, yielding a five-dimensional parameter space. As the detection of circles or ellipses is a standard problem in computer vision, various approaches have been developed to reduce time and space complexity.

The solution we propose, a combination of the Hough transform with a genetic algorithm, is a new attempt in this direction. The presented algorithm was designed to detect ring patterns resulting from experiments with Spontaneous Non-Colinear Frequency Doubling (SNCFD) in order to characterize crystals [2]. As shown in Fig. 1, the ideal SNCFD ring image consists of a nearly circular ellipse with an additional spot close to the ring center. In real SNCFD images, portions of the ellipse are usually missing, and most images contain a high level of noise. A comparable test image may be represented by Fig. 2 where 40% of the original pixels are replaced by random noise.

The proposed algorithm for ellipse detection is, however, not limited to SNCFD rings but of general use in computer vision: Many pieces of machinery as well as e.g. bottles do have circular features. In production control systems, cameras taking images from a conveyor often cannot be placed perpendicular to these features because tools limit the field of vision, consequently circles are transformed into ellipses whose eccentricity is limited by the camera position. Moreover, the combination of the Hough transform with a genetic optimization algorithm might be of general interest in various fields of pattern recognition.

2. The technological problem

Using high laser intensities, in crystals with appropriate point symmetry various nonlinear effects can be produced. One of these effects is the generation of harmonic waves due to nonlinearities in the polarization induced by the fundamental laser beam. These harmonic waves are amplified when the so-called phase-matching condition is met: Fundamental and harmonic wave propagate through the crystal at equal velocities. Usually all light beams involved are mutually parallel to each other, a *colinear* geometry.

* Corresponding author. Tel.: +49-541-969-2636; fax: +49-541-969-2670.

E-mail address: klaus.betzler@physik.uni-osnabrueck.de (K. Betzler).

0262-8856/02/$ - see front matter
PII: S0262-8856(02)00155-5

IMAVIS 1939—28/11/2002—18:06—UMASHANKAR—58463— MODEL 5

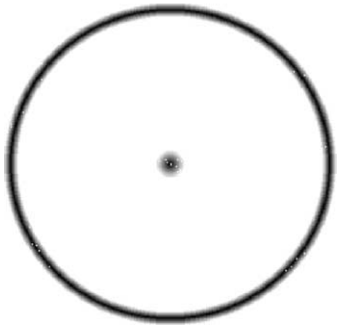

Fig. 1. Test image showing ideal SNCFD-Ring with central spot.

Yet for the generation of the second harmonic wave (twice the fundamental frequency) the fundamental beam can also interact with stray light in the crystal producing second harmonic light in arbitrary directions. In this *noncolinear* geometry only for certain combinations of directions phase matching is achieved. One can show that the locus of directions, where intense second harmonic light is produced, is an elliptic cone.

The velocity of light inside a crystal is defined by the refractive indices of the material. Thus the exact shape of the phase-matching cone depends very sensitively on the refractive indices. The indices are influenced in a finger-print-like way by nearly all other material parameters as for instance composition, density, or defect concentration. Thus in turn one can use the measurement of the cone angle to determine or monitor those other parameters. The technique is completely non-destructive and contactless, thus well suited for production monitoring. Yet it is limited to crystal of certain symmetry classes. Fortunately the applicability of crystals in nonlinear optics and electrooptics is governed by similar symmetry restrictions, thus nearly all of the materials interesting in these fields can be investigated using this technique.

The application of the method –Spontaneous SNCFD-is fairly simple: A slightly focused laser beam is directed onto the sample, the cone of second harmonic light is measured using a CCD camera. A two-dimensional topographical inspection is possible when the sample is moved by means of a suitable translation stage. The only severe problem is the automatic and quick extraction of the ellipse parameters from the measured set of video pictures. The problem is illustrated by Fig. 3 where some typical pictures are shown. All pictures are characterized by a comparably high noise level, which is due to stray light in the samples. Due to the parameters of the crystals under inspection, the eccentricity of the ellipses is always limited to a certain range. Furthermore, the crystals under inspection always can be oriented such that the principal axes of the ellipse are parallel to fixed, known directions. For simplicity one can chose a geometry where the axes are parallel to the x- and y-direction, respectively.

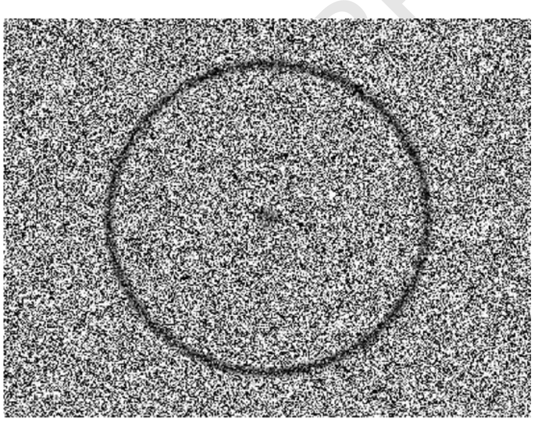

Fig. 2. Test image generated from fig. 1 by substituting 40% of pixels by random intensity.

3. Fundamentals of the implementation

To apply the technique as a standard characterization and monitoring method it was necessary to develop an efficient algorithm capable of extracting the ellipse parameters from pictures of that sort. The algorithm was tailored to fit our application problem albeit can be easily generalized. As we deal with ellipses at a fixed orientation, we neglect the respective parameter and use only four parameters for the description instead of the full set of five.

3.1. Hough transform

The Hough transform (HT) [1] is known to be a highly robust detection method for objects which can be parameterized. It is especially tolerant to pixel noise and shapes different from the object sought, e.g. the central spot in SNCFD images. Usual parameters for an ellipse with horizontal resp. vertical axes are the center position (x_0, y_0) and the radii r_x and r_y :

$$\left(\frac{x - x_0}{r_x}\right)^2 + \left(\frac{y - y_0}{r_y}\right)^2 = 1. \tag{1}$$

After extracting the edge information from the original image, the parameters for each edge point (x, y) are calculated from Eq. 1 and registered as votes in a discrete parameter space called accumulator. Finally each accumulator entry indicates the likelihood of the corresponding ellipse parameters, for details including references see [3].

To detect ellipses with an accuracy of one pixel given a 800×600 size image, the parameter space $\{x_0, y_0, r_x, r_y\}$ must contain $800 \times 600 \times 400 \times 300 \approx 10^{10}$ entries. The requirements for memory space and computational time would render the standard HT impractical for ellipse detection.

Most of the specialized HTs try to use gradient information in order to improve on memory and speed [4]. In addition, the Randomized HT samples only a random subset of the original image [5] while novel approaches try

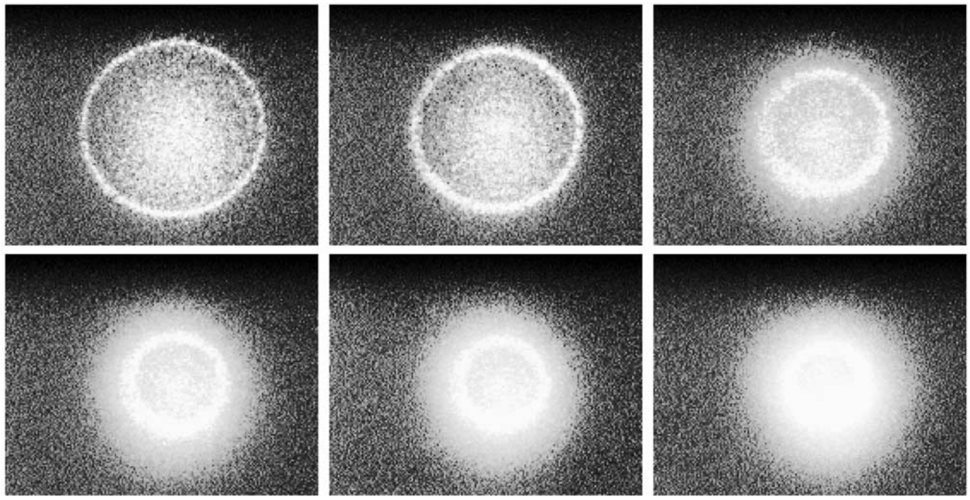

Fig. 3. Real images of SNCFD characterization measurements. Most images show a broadened elliptic ring, high noise levels, and an intense central spot.

to find the optimal subset [6]. Focusing algorithms can be used to provide Fast Hough Transforms [7]. Another way to reduce the time and space complexity is to decompose the parameter space [8]. A comparison [9], however, has shown that HTs utilizing gradient information are only suited for images containing less that 10% of speckle noise. When applied to our test images, these special HTs as well as similar methods based on local features failed because the gradient information is degenerated by noise. Even the initial step of edge extraction poses a problem with our images: Common examples of HT application ideally use high contrast disk images from which circular edges result. Filtering a blurred ring image with usual edge operators will at best yield two concentric circles. A Laplacian would generate a single ellipse for each optimal ring but is also much more sensitive to noise because of the second derivative. For the algorithm presented in the following sections we obtained the best results when applying the HT to the original, unfiltered image, weighting the votes into the Hough accumulator by the intensity of the image pixel.

3.2. Elliptic circle parameterization

The parameter space can be reduced in size by acknowledging the limited eccentricity of the ellipses under investigation. In a new parameter space $\{x_0, y_0, r := r_x, e := r_y/r_x\}$ Eq. 1 becomes:

$$(x - x_0)^2 + 1/e^2(y - y_0)^2 = r^2. \tag{2}$$

For SNCFD-rings a suitable range would be $e = 0.92\ldots1.08$ with 20 discrete steps. To further shorten the accumulator, only a two-dimensional HT into $\{r, e\}$ is calculated, requiring merely 400×20 entries.

As depicted in Fig. 4, the accumulators for both ideal and noisy ring images exhibit a well defined maximum corresponding to the parameters r and e sought, as long as the center coordinates (x_0, y_0) are chosen correctly. For positions (x_0, y_0) increasingly different from the ring center, the accumulator peak flattens up to a degree where it cannot be found.

3.3. Determination of ellipse center

If a peak is found in the accumulator, its height divided by its width is defined as the peak value v, otherwise v is set to 0. As shown in Fig. 5, this value v has a well defined maximum where (x_0, y_0) matches the ellipse center. Conventional methods for multidimensional maximization of $v(x_0, y_0)$ like the Simplex search [10,10.4] will nonetheless often fail to find the optimal (x_0, y_0) because of local side maxima. They might also stop without any valid solution as they depend on gradients in $v(x_0, y_0)$. Many positions, though, yield an accumulator without an observable peak. While a complete test of all possible (x_0, y_0) will finally detect the correct ring center, the computational effort would be too high.

3.4. Differential evolution

Genetic algorithms (GA) provide an elegant way for optimizing $v(x_0, y_0)$ as they combine means of directed search with random modifications of (x_0, y_0). Consequently they do not stop in regions of v with vanishing gradient. Starting with an initial set (population) of coordinates, GAs generate new trial coordinates via a combination of calculations named mutation and crossover. Finally selection decides which trial coordinates will replace members of the previous population in order to form a new generation. Differential Evolution (DE) [11] is a powerful yet exceptionally simple to implement GA.

Usually evolutionary strategies mutate vectors by simply adding zero-mean Gaussian noise to them. To be efficient, this noise distribution has to be adapted dynamically to the population distribution itself. This can be a complicated process, especially when the parameters behave differently. We used a newer, more convenient approach, mutation with

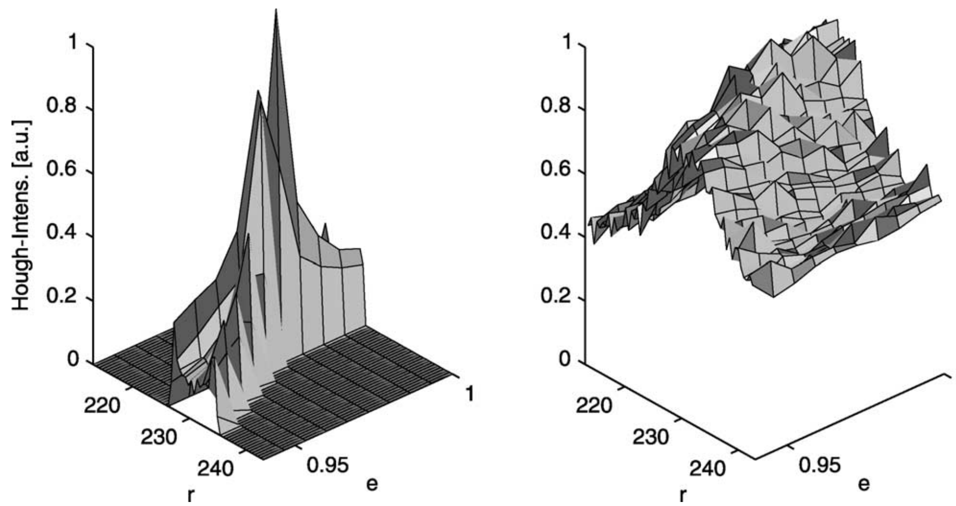

Fig. 4. {r,e}-subspace of Hough accumulator for Fig. 1 (left) and Fig. 2 (right) test images, (x_0, y_0) set to match correct ellipse center.

vector differentials. In this mutation scheme the population itself is used as source of appropriately scaled perturbations. Every pair of vectors, (z_a, z_b), defines a vector differential $z_a - z_b$. When both vectors are chosen randomly, their weighted difference can be used in place of Gaussian noise to perform the mutation step.

For the task of optimizing $v(x_0, y_0)$, an initial population of test centers (x_i, y_i), $i = 1 \ldots N$, is chosen, the HTs are calculated and the corresponding values $v(x_i, y_i)$ are determined. The behavior of DE is fully controlled by the population size N, usually set to 10, the scaling factor for mutation $F = 0\ldots1$, and the crossover rate $CR = 0\ldots1$. New generations are calculated by genetically evolving each coordinate according to the following loop:

for $i = 1$ to N do
  repeat
  choose random integers $1 <= \{a, b, c\} <= N$.
until i, a, b, c different.
choose random integer $1 <= k <= 2$.
choose random float $0 <= r <= 1$.
if $k = 1$
  $x_t = x_i$.
  if $r < CR$
    $y_t = y_c + F \cdot (y_a - y_b)$.
  else
    $y_t = y_i$.
else
  $y_t = y_i$.
  if $r < CR$
    $x_t = x_c + F \cdot (x_a - x_b)$.
  else
    $x_t = x_i$.
calculate $v(x_t, y_t)$, i.e. perform HT into $\{r, e\}$ for center (x_t, y_t) and evaluate accumulator.
if $v(x_t, y_t) > v(x_i, y_i)$
  replace (x_i, y_i) by (x_t, y_t).

This evolution loop is repeated until all coordinates are nominally equivalent, e.g. all inter-coordinate distances fall below one pixel. We tried several values for the scaling factor F and the crossover rate CR and found good, dependable operation—for our problem—setting both to 0.2. Higher values do accelerate the algorithm by increasing the mutational change (F) or diversity (CR) but might also cause it to miss the best solution.

3.5. Parallel computation of Hough Transform

Each iteration of the DE main loop as described in section 3.4 requires N Hough Transforms. While the test center coordinates (x_t, y_t) for each HT are originally float numbers, integer pixel coordinates are a more sensible choice for pixel images. To avoid—after rounding (x_t, y_t)—recalculation of equivalent HTs, a cache for $v(x_t, y_t)$ is used. Because of the inherent parallelism of the DE main loop, all N iterations needed to calculate a new generation can be performed in parallel.

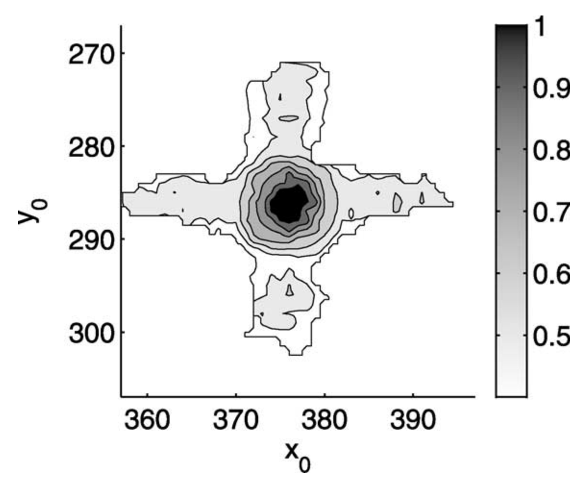

Fig. 5. Value v of peak in Hough accumulator for Fig. 2 test image for various center coordinates (x_0, y_0).

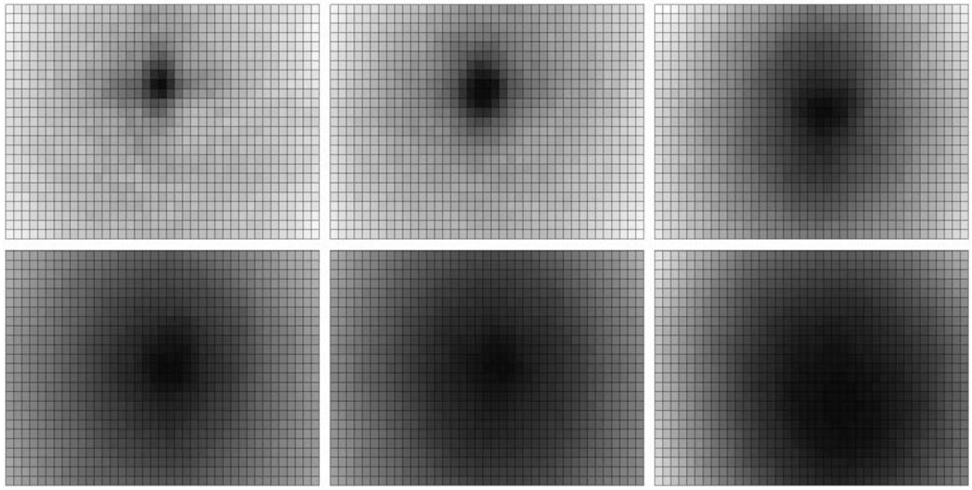

Fig. 6. Peak values in the *{r,e}*-subspaces of the Hough accumulators as a function of the ellipse center coordinates (x_0, y_0) for the six real images of Fig. 3. $x_0(80\ldots120)$ from left to right, $y_0(65\ldots90)$ from top to bottom in each of the graphs.

The major portion of CPU cycles is consumed by the HTs. Therefore in our implementation up to N Hough servers are launched on different CPUs. The main loop was modified to calculate all the N test centers (x_t, y_t) first. If a coordinate has already been used, the associated peak value $v(x_t, y_t)$ is taken from the cache, otherwise the HT is enqueued to be performed by the next idle Hough server. Because of the high rate of cache hits when the DE is converging, experiments have shown that the number of Hough servers can be reduced to about $N/3$ without increasing overall search time.

4. Results

4.1. Real images

To illustrate the technique, the evaluation of the real images shown in Fig. 3 is demonstrated. We use low resolution pictures (183 × 141 pixels) for simplification. Fig. 6 shows the height of the maximum values of the *{r,e}*-subspaces of the Hough accumulators as a function of the ellipse center (x_0, y_0) for each of the real images. The coordinate range for the ellipse center is restricted to $80 \le x_0 \le 120$ and $65 \le y_0 \le 90$, respectively. In each case a maximum value (darkest spot) of these $\{r, e\}$ maxima is clearly detectable, indicating the coordinates of the optimum ellipse center. Yet the distribution of the maxima gets considerably flattened out when noise and background intensity increase in the corresponding original pictures.

The distributions in the *{r,e}*-subspaces of the Hough accumulators for the optimized ellipse centers are shown in Fig. 7 ($10 \le r \le 70$, $0.8 \le e \le 1.2$.) As can be seen, the maxima (dark regions) are less expressed yet can be well discriminated. Again, corresponding to the broadening of the ellipses and to the increasing background intensity in the

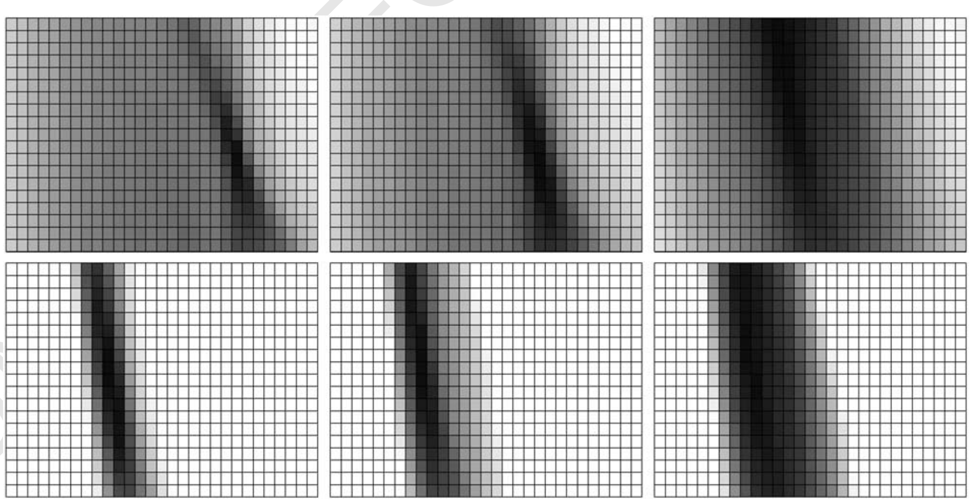

Fig. 7. *{r,e}*-subspaces of the Hough accumulators for optimized center coordinates x_0, y_0. $r(10\ldots70)$ from left to right, $e(0.8\ldots1.2)$ from bottom to top in each of the graphs. For better visibility of the maxima, in the lower three graphs the gray scaling is expanded by a factor of 10, all values less than 90% of the maximum thus are suppressed.

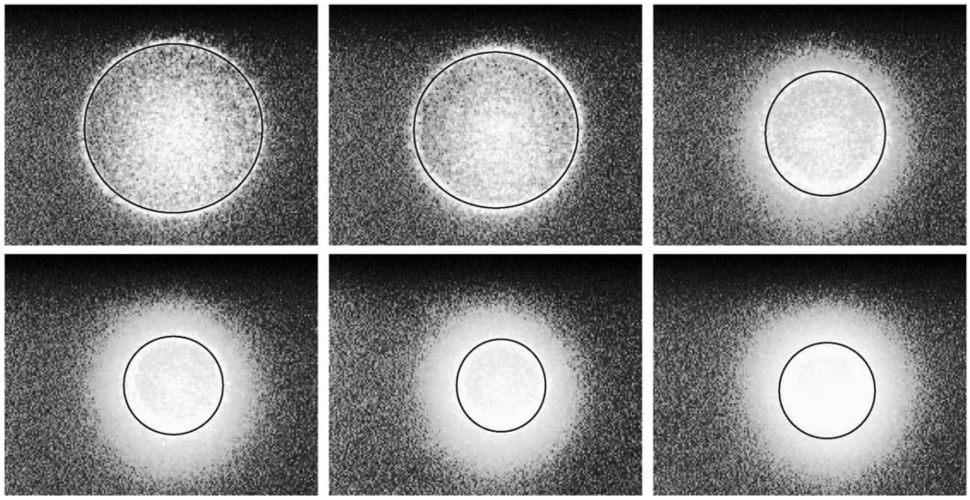

Fig. 8. Ellipses detected by the algorithm. The black lines indicating the ellipses are superimposed to the images of Fig. 3.

series of the six pictures, the distributions flatten out accordingly.

Fig. 8 shows the final results, the ellipses found by the described algorithm. They are indicated as solid black lines superimposed to the original pictures.

4.2. Bias problems

In Figs. 6 and 7, besides the main maxima also less expressed minor maxima can be detected. These probably could cause bias problems when using gradient algorithms for maximum detection. Depending on the starting point, such an algorithm could get stuck in one of the minor maxima. For the DE algorithm used here we did not experience such a behavior when choosing the relevant parameters (scaling factor F and crossover rate CR as described in section 3.4 appropriate for a dependable operation.

4.3. Sample application

Using the described algorithm for ellipse detection, the technique is now readily applicable for the characterization of crystals. An example is shown in Fig. 9. A lithium niobate crystal is topographically inspected to test homogeneity and composition. SNCFD rings are measured in a two-dimensional scan all over the test crystal. The ellipse parameters are automatically extracted using the described evaluation scheme; from the parameters, the crystal composition is derived. As can be seen, there is a slight compositional variance in the growth direction of the crystal. To our knowledge no other technique available in this field is capable to detect such small variations [12].

The calculation time needed for the extraction of the ellipse parameters for one image was typically 100 s using four Pentium 133 in parallel. In comparison, a full Hough transform for the same problem lasted about 1 h.

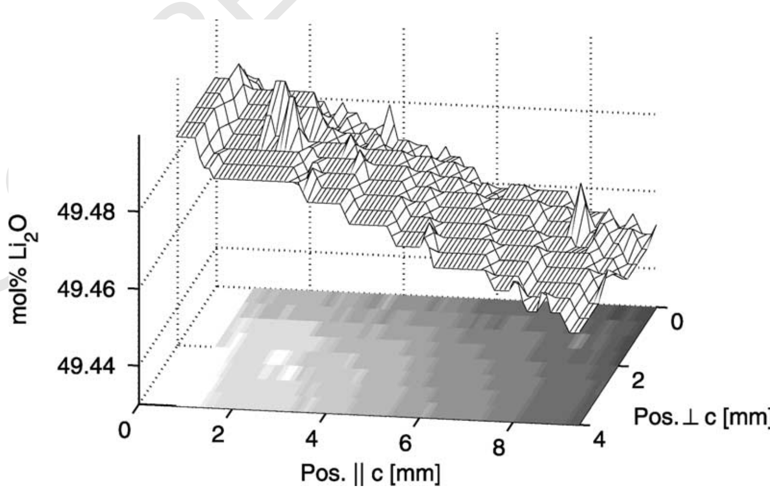

Fig. 9. Result of a typical characterization measurement: Two-dimensional composition topography of a lithium niobate crystal. The achieved composition resolution is in the order of 0.01%, a value not reached by any other technique.

5. Conclusion

A specialized Hough transform has been presented which is optimized for detecting ellipses with limited eccentricity like nearly circular rings even in the presence of high noise levels. As common simplifications of the standard HT like utilizing edge and gradient information were not successful, we fell back on transforming the full original image. The Hough accumulator was reduced in size by optimizing the remaining parameters with DE. A parallel implementation of the genetic algorithm was used to reduce the overall detection time.

Acknowledgements

We are greatly indebted to the reviewer for pointing out various important aspects. Support from the Deutsche Forschungsgemeinschaft (grant SFB 225, A10) is gratefully acknowledged.

References

1 P.V.C. Hough, Methods and means for recognizing complex patterns. U.S. Patent 3069654, 1962.

2 A. Reichert, K.-U. Kasemir, K. Betzler, Crystal characterization by noncolinear frequency doubling, Ferroelectrics 184 (2) (1996) 21–30.

3 D.H. Ballard, C.M. Brown, Computer Vision, Prentice Hall, Englewood Cliffs NJ, 1982.

4 H.K. Yuen, J. Illingworth, J. Kittler, Detecting partially occluded ellipses using the Hough Transform, Image and Vision Computing 7 (1) (1989).

5 H. Kalviainen, P. Hirvonen, L. Xu, E. Oja, Probabilistic and non-probabilistic Hough Transforms: Overview and comparisons, Image and Vision Computing 13 (4) (1995) 239–252.

6 J.Y. Goulermas, P. Liatsis, Genetically fine-tuning the Hough Transform feature space for the detection of circular objects, Image and Vision Computing 16 (1998) 615–625.

7 N. Guil, E.L. Zapata, Lower order circle and ellipse Hough transform, Pattern Recognition 30 (1997) 1729–1744.

8 P.S. Nair, A.T. Saunders Jr., Hough transform based ellipse detection algorithm, Pattern Recognition Letters 17 (1996) 777–784.

9 Robert A. McLaughlin. Technical report –Randomized Hough Transform: Improved ellipse detection with comparison. Tech.Rep.TR97-01, Univ. of Western Australia, CIIPS, Dept of Electrical and Electronic Eng., Nedlands, Australia, 1997.

10 W.H. Press, W.H. Vetterling, S.A. Teukolsky, B.P. Flannery, Numerical Recipes in C:The Art of Scientific Computing, Cambridge University Press, Cambridge, 1992.

11 K. Price, R. Storn, Differential Evolution, Dr Dobb's Journal 4 (1997) 18–24.

12 M. Wöhlecke, G. Corradi, K. Betzler, Optical methods to characterize the composition and homogeneity of lithium niobate single crystals, Applied Physics B 63 (1996) 323.