Automatic detection of secondary creases in fingerprints

David Vernon
University of Dublin
Trinity College
Department of Computer Science
Dublin 2, Ireland

Abstract. Human fingerprints comprise a series of whorls or ridges. In some special cases, these whorls are broken by so-called *secondary creases*—colinear breaks across a sequence of adjacent ridges. It is a working hypothesis that the presence of these secondary creases form a physical marker for certain human disorders. A technique to automatically detect such creases in fingerprints is described. This technique utilizes a combination of spatial filtering and region growing to identify the morphology of the locally fragmented fingerprint image. Regions are then thinned to form a skeletal model of the ridge structure. Creases are characterized by colinear terminations on ridges and are isolated by analyzing the Hough transform space derived from the ridge end points. Empirical results using both synthetic and real data are presented and discussed.

Subject terms: fingerprint analysis; secondary creases; Hough transform; virtual lines; Laplacian of Gaussian, spatial scale.

Optical Engineering 32(10), 2616–2623 (October 1993).

1 Introduction

Human fingerprints comprise a series of whorls or ridges. The imprint of these ridges forms a quasi-contiguous locally linear series of ink prints where the surface of the skin has come into contact with the paper on which the fingerprint is formed (see Fig. 1). For most people, these ridges are well formed and unbroken. However, in some people, these ridges are broken by a crease where no skin has been inked and come into contact with the paper. These colinear breaks in the whorls of the fingerprint are called *secondary creases* (See Fig. 2). It is a working hypothesis that such secondary creases form a physical marker for certain human disorders. This paper discusses the research and development of an automated technique to detect and isolate secondary creases so that this working hypothesis can be verified without relying on human interpretation of the morphology of the fingerprint.

2 Segmentation

Before any automatic analysis of the morphology of the fingerprint can be accomplished, it is first necessary to identify and label those parts of the image of the fingerprint that correspond to the ridges or whorls, i.e., the inked part of the paper.

In normal circumstances where the object of interest, i.e., the ridges, are of a distinctly different gray level to the background, i.e., the paper on which the fingerprint is imprinted, this process of segmentation could be effected by "thresholding." In this process, the gray level of each pixel is compared with an appropriate reference level, the threshold, and the pixel is then assigned a value of 255 or 0 depending on whether it is greater or less than the threshold. The resultant binary image comprises pixels of gray level zero, signifying that a pixel represents a ridge, or 255, signifying that a pixel represents the paper. Unfortunately, this straightforward approach is not practicable for two reasons.

First, the level of inking of the finger can vary considerably and consequently the "blackness" of the print varies from fingerprint to fingerprint and from region to region in a given fingerprint. Second, a segmented image of a fingerprint that has been generated by thresholding is extremely fragmented in the sense that a single ridge is broken up into many (of the order of tens or hundreds) isolated, nonadjacent, blobs or regions. This is a natural consequence of the textured nature of the surface of the skin that forms the whorls. While the problem of inhomogeneous inking can be solved through the use of dynamic thresholding[1] wherein the threshold is a function of the image coordinates, the second difficulty of ridge fragmentation is more problematic. An informal investigation of the feasibility of identifying the global structure of each whorl through accepted techniques, such as morphological opening,[2,3] yielded no useful solution.

Paper 19092 received Sep. 21, 1992; revised manuscript received Jan. 25, 1993; accepted for publication Jan. 25, 1993.

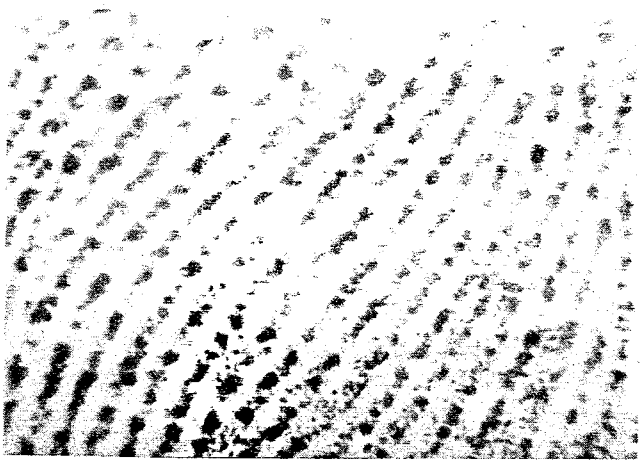

Fig. 1 A normal human fingerprint.

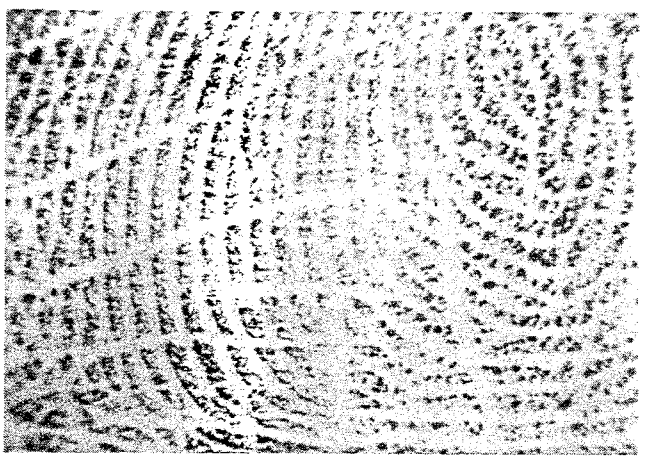

Fig. 2. A human fingerprint with secondary creases.

In the work described here, a robust, if computationally expensive, solution is employed that addresses simultaneously the problems of fragmentation and segmentation. This technique utilizes a combination of spatial filtering and region growing to identify the morphology of the fingerprint image.

Ridges, as entities in themselves, in the fingerprint are substantially larger than the fragments that make up the ridge; that is, they are well represented by the lower spatial frequencies comprising the image. By attenuating the higher spatial frequencies, the fragmented image detail is removed and the global structure of the ridge morphology is retained. This can best be accomplished[4] by convolving the fingerprint image with a 2-D Gaussian function:

$$I(x,y)*G(x,y) \ , \tag{1}$$

where $I(x,y)$ represents the image intensity at a point (x,y) and $G(x,y)$ is the 2-D Gaussian function of a given standard deviation σ defined by:

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

The value of σ governs the spatial scales that are retained: the larger the value of σ, the larger the scale of the objects represented in the filtered image. In all of the results cited in this paper, $\sigma=7.0$ pixels. This value was determined empirically by a calibration procedure, based on the normal distance between the ridges.

The ridges are isolated through the use of a Laplacian second-derivative edge detection filter

$$\nabla^2=\frac{\partial^2}{\partial x^2}+\frac{\partial^2}{\partial y^2} \ , \tag{2}$$

i.e., the sum of second-order, unmixed, partial derivatives.

The evaluation of the Laplacian and the convolution with the Gaussian commute so that the segmentation and selection of the spatial scale can be effected with a single filter: the Laplacian of Gaussian[5]:

$$\nabla^2[I(x,y)*G(x,y)]=\nabla^2G(x,y)*I(x,y) \ . \tag{3}$$

The Laplacian-of-Gaussian operator yields thin continuous closed contours of zero-crossing points that bound regions in the filtered image. These regions are recursively nested, with each region having, alternately, an opposite sign (positive or negative) as one descends through the nesting (see Ref. 6). In this work on the detection of secondary creases, segmentation is achieved by identifying each region by its sign, computing the area of each region, and isolating the region with the largest area. This is the background region, and all other regions are deemed *ridge regions*. This segmentation then is represented as a binary image, for example, see Figs. 5(b) through 7(b).

3 Morphological Processing: Thinning and Isolation of Ridge Termination Points

Once the elongated ridge regions have been isolated, it is necessary to identify locations of their ends. This is accomplished by morphological processing, thinning the region to form "skeletons," which are one pixel wide, and then by identifying the end points of these skeletons. Strictly speaking, a thinned region is a digital approximation of the skeleton and, thus, for the sake of correctness, we refer to it explicitly in the following as a *thinning*. The procedure can be summarized as follows (see Ref. 7 for further details). Let X denote the set of points (pixels) that comprise the ridge regions. The thinning of this set X is accomplished by the morphological filtering of the set X with a sequence of structuring elements L, as follows:

$$X\bigcirc\{L^i\} \ , \tag{4}$$

that is,

$$(((...(X\bigcirc L^1)\bigcirc L^2)\bigcirc L^3)...\bigcirc L^i) \ , \tag{5}$$

where $X\bigcirc L$ is defined as:

$$X\bigcirc L=X\backslash X\circledast L \ . \tag{6}$$

The operator $\circledast$ denotes the *hit-or-miss transformation*. This transformation is defined as:

$$X\circledast L=\{x|L_x^f\subset X;\ L_x^b\subset X^C\} \ , \tag{7}$$

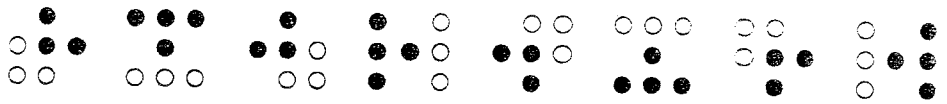

Fig. 3 Sequence of structuring elements L^1 through L^8 used in the thinning operation.

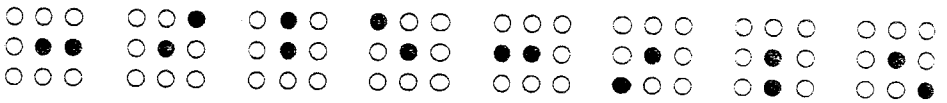

Fig. 4 Structuring elements E^1 through E^8 used to identify end points.

where L_x^f is that subset of L, translated to point x, whose elements belong to the "foreground" (ridge) and L_x^b is the subset of L, translated to point x, whose elements belong to the "background" (i.e., $L^f \cap L^b = \emptyset$). The set complement of X is X^C. A point x belongs to the hit-or-miss transformation if and only if L_x^f is included in X and L_x^b is included in the complement of X. Thus, $X \circledast L$ defines the points where the structuring element L exactly matches the set X, i.e., the ridge pixels in the image. The set $X \bigcirc L$, then, is the set X less the set of points in X that matches L. Thus, if $X \circledast L$ identifies border points, and L is appropriately structured to maintain the connectivity of a set, then repeated application of the thinning process successively removes border points from a set until the thinning is achieved. At this point, further application of the thinning transform yields no change in the thinned set. The sequence $\{L\}$, which is used for thinning, is based on a single-structuring element[7] and is generated by rotating the structuring element (through 360 deg in increments of 45 deg). This sequence $\{L\}$ is shown in Fig. 3. The thinning algorithm then amounts to the repeated transformation of a set $X_i \rightarrow X_{i+1}$ defined:

$$X_{i+1} = (((...(X_i \bigcirc L^1) \bigcirc L^2) \bigcirc L^3) ... \bigcirc L^8) \ . \tag{8}$$

The thinning is achieved when $X_i = X_{i+1}$. Initially, $X_0 = X$, i.e., the original (unthinned) segmented binary image. Examples of thinned ridges can be seen in Figs. 5(d) through 7(d).

Given a thinning X, we then identify the end points, i.e., points that are connected to just one other point, using the hit-or-miss transform and an appropriate set of structuring elements $\{E\}$, shown in Fig. 4. Thus, the end points of the thinning are given by:

$$Y = \bigcup_{i=1}^{8} X \circledast E^i \ , \tag{9}$$

that is, the set of end points is the union of all those points that hit with one of these end-point structuring elements.

4 Hough Transform

4.1 *Computing the Hough Transform*

The technique for the detection of secondary creases described in this paper is based on the assumption that such creases are characterized by colinear ridge ends. Having extracted the morphology of the ridges and having identified the locations of the end points of the segmented ridges, we must now group these end points according to a colinearity criteria: in effect, to find the virtual line formed by the ridge ends. The Hough transform[8] is used to accomplish this. The computation of the Hough transform for the detection of lines is quite straightforward although, as we see in the next section, postprocessing in the Hough transform space is required to effect reliable and robust extraction of lines.

The equation of a straight line is given in parametric form by the equation:

$$x \cos\theta + y \sin\theta = r \ , \tag{10}$$

where r is the length of a normal to the line from the origin and θ is the angle this normal makes with the X axis. For a given line, r and θ are known. In this case, however, r and θ are unknown because we do not yet know which are the crease lines, but we have several specific samples of x and y, namely, x_i and y_i, which are given by the coordinates of the ridge end points. In the Hough transform, the solution to Eq. (10) is computed for each (x_i, y_i) pair, yielding a set of values for r and θ. These values are recorded by incrementing an element of a 2-D array, known as the Hough accumulator, for each (r,θ). From a computational point of view, this is performed quite simply by computing the value of θ from Eq. (10) for all values of r, knowing x_i and y_i. This solution-set is, in effect, a sinusoidal curve in the r-θ space, i.e., in the Hough transform space. The transform is computed for all ridge end points (x_i, y_i), and end points that are colinear will all have a single value of r and θ in common; that is, the solution-set sinusoidal curves will intersect in a single point in the Hough transform space. Such points of intersection are characterized by local maxima in the Hough accumulator; see, for example, Figs. 8(c) and 9(c).

4.2 *Detection of Local Maxima in the Hough Accumulator*

Unfortunately, ridge end points are almost never colinear, even in the most ideal circumstances such as in the synthetic test patterns that have been used to test the technique (see Sec. 5.1). In the case of real fingerprint patterns, this colinearity is, at best, approximate (see Figs. 5 through 7). Consequently, the solution-set curves in r-θ space do not intersect in a single point but in several points that are close together. Before the effective local maximum can be detected by comparing the value of a single accumulator element with a given threshold, it is essential to process the accumulator, i.e., the Hough transform space, so that the accumulator values in a

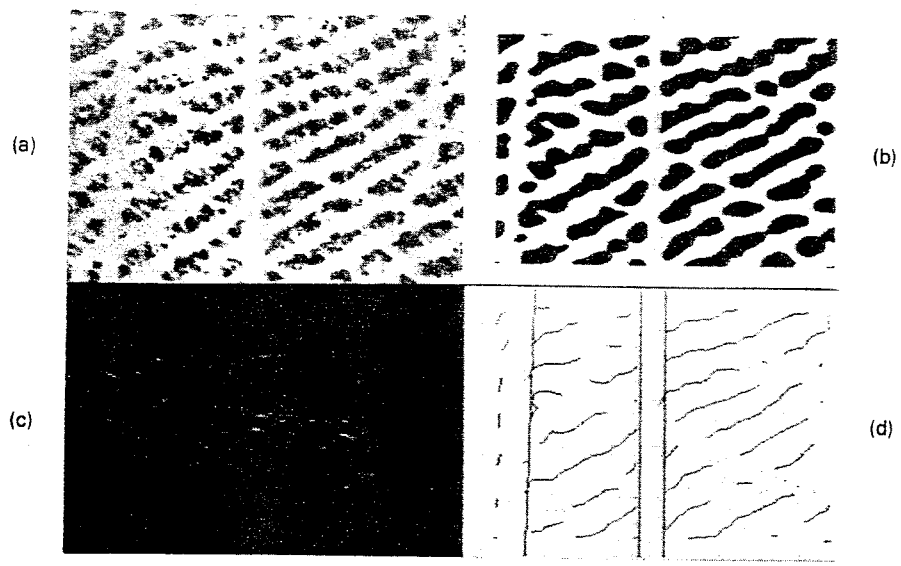

Fig. 5 (a) Gray-level image of a section of a fingerprint; (b) ridge regions derived from Laplacian of Gaussian $\nabla^2 G$ of image, $\sigma = 7.0$ pixels; (c) Hough transform space derived from end points of skeleton of ridge regions; and (d) crease line superimposed on the skeleton of ridge regions.

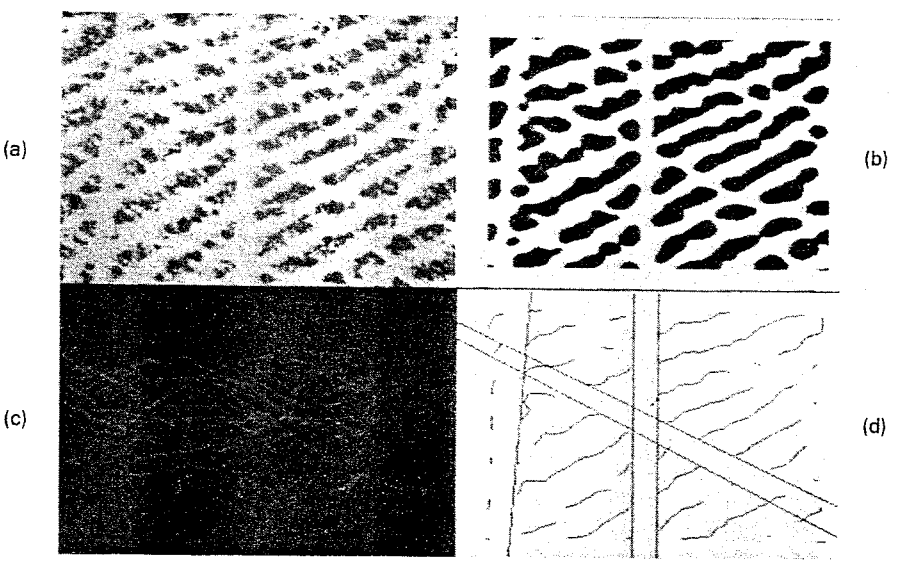

Fig. 6 (a) Gray-level image of a section of a fingerprint; (b) ridge regions derived from Laplacian of Gaussian $\nabla^2 G$ of image, $\sigma = 7.0$ pixels; (c) Hough transform space derived from end points of skeleton of ridge regions, and (d) crease line superimposed on the skeleton of ridge regions (note that the tolerance δ_r^2, which defines the degree to which a crease line can intersect ridges, is equal to 0.2 in this example, whereas in Fig. 5 it is equal to 0.1; refer to Sec. 4.4 for further details).

local region are collected and assigned to a single specific accumulator cell. This is accomplished in the research described in this paper by iteratively reassigning the value of each accumulator element that is not the local maximum of its immediate 3×3 pixel region to one of its eight neighbors. The neighbor to which the value is assigned is required to be the local maximum in that 3×3 region. In this way, the value of a local maximum increases as it is incremented by the value of its neighbors. This iterative process is continued until no more reassignment can be performed, at which point the accumulator comprises a set of isolated points, each of which represents a local maximum, and the point-wise thresholding can proceed.

4.3 *Selection of Threshold for Isolation of Candidate Crease Lines*

Not all of these local maxima correspond to valid lines in the original image and it is necessary to identify a threshold value that accumulator elements must exceed in order to be considered as candidate crease lines. Because the content of the fingerprint images varies considerably and, hence, so too

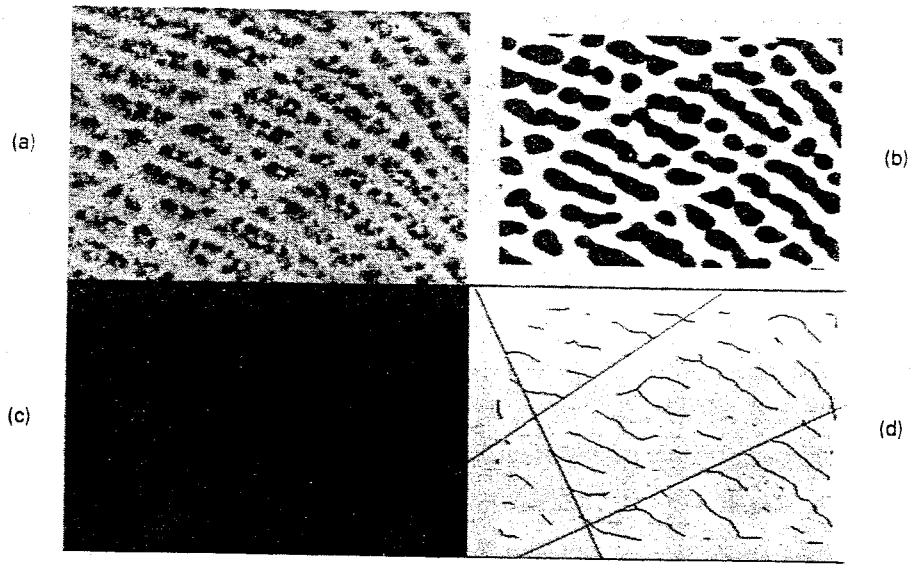

Fig. 7 (a) Gray-level image of a section of a fingerprint; (b) ridge regions derived from Laplacian of Gaussian $\nabla^2 G$ of image, $\sigma = 7.0$ pixels; (c) Hough transform space derived from end points of skeleton of ridge regions; (d) crease line superimposed on the skeleton of ridge regions.

does the resultant form of the Hough accumulator, it is desirable to have this threshold chosen adaptively. This can be accomplished by computing some simple statistics on the distribution of the values in the postprocessed Hough accumulator and by basing the threshold on these statistics. In this implementation, the threshold T is given by:

$$T = \mu + 3\sigma \ , \tag{11}$$

where μ and σ are the mean and standard deviation of the values in the postprocessed accumulator, respectively. This threshold is low enough to ensure all crease lines are included but high enough to remove the majority of the accumulator elements. The choice of the coefficient 3 in Eq. (11) is based on practical experience in the use of the technique in this particular application. A somewhat different value may be required in other applications.

4.4 *Constraints for the Isolation of Crease Lines*

After application of the threshold on the Hough accumulator, a set of candidate crease lines exists. Not all of these lines do, in fact, join the end points of several distinct ridges.

For example, it is common that a single ridge can be fragmented into a number of sections. Each of these sections gives rise to an equal number of roughly colinear skeletal line segments with two end points. Because these end points are colinear, they too will give rise to a valid local maximum in the Hough accumulator. This maximum cannot be distinguished from valid maxima corresponding to creases by analysis in the Hough space. However, it can be identified and removed by reverting to analysis of the structure of the original segmented and thinned images.

A second pathological case also exists. It is possible that lines are formed from several end points of ridges that are not adjacent in the original image and that do not exhibit the spatial relationship required of ridges forming a secondary crease. Again, since all spatial information is lost in the Hough transform, these lines cannot be detected by analysis of the Hough accumulator. And again, they can be isolated by analyzing the structure of the original segmented and thinned images.

Both of these cases can be dealt with by the imposition of two constraints on the lines formed by the ridge end points: these constraints ensure (1) that the crease orientation is significantly different from the average ridge orientation and (2) that the crease line must connect spatially contiguous ridge ends without excessively intersecting unbroken ridges.

The first constraint, which eliminates lines connecting the end points of a single fragmented ridge, is that the orientation of a crease line should be significantly different to the average orientation of the ridges in the region surrounding the crease line. Specifically, candidate crease lines are removed from consideration if they satisfy the following inequality:

$$\theta_r - \delta_r^1 \leq \theta_c \leq \theta_r + \delta_r^1 \ , \tag{12}$$

where θ_r is the average ridge orientation; δ_r^1 is a tolerance, defined to be 20 deg in the implementation described in this paper; and θ_c is the orientation of the crease line. The average ridge orientation is computed as the mean orientation of all adjacent skeleton points in a region around the crease line joining the end points. This region is bounded by two lines, one on either side of and parallel to the crease line. The normal distance from the crease line to the boundary of the region is equal to twice the calibrated ridge width.

The second constraint, which ensures that the crease line connects spatially contiguous ridge end points, concerns the intersection of ridges by the crease line. Similar to the first constraint, candidate crease lines are removed from consideration if a significant proportion of the crease lines intersects ridge points. Thus, lines are removed from consideration if they satisfy the following inequality:

$$\frac{n_r}{l_c} \geq \delta_r^2 \ , \tag{13}$$

where n_r is the number of ridge pixels lying on the crease line, l_c is the length of the crease line, and δ_r^2 is the tolerance. The tolerance δ_r^2 is defined to be 0.1 for all examples in the implementation described in this paper, unless otherwise stated.

Examples of isolated secondary crease lines are shown in Figs. 5(d) through 7(d) and in Figs. 8(d) and 9(d).

5 Verification of Results

To validate the technique, two forms of test have been run. The first uses synthetic test patterns that have been constructed to assess, in a quantitative manner, how the technique performs as the data degrades in a well-understood manner. The second form of test deals with actual fingerprint data.

5.1 *Tests on Synthetic Data*

A series of test patterns, each comprising 10 ridges, were devised and printed on plain white paper using a laser printer (see, for example, Figs. 8 and 9). These patterns were printed at an actual size of 8×6 mm to ensure that the tests on synthetic data operated at the same magnification and field of view as those on the actual fingerprint data. Consequently, there is a natural variation (or noise) in the image of the test pattern due to the limited resolution of the laser printer, the fibrous texture of the paper, as well as imaging noise. This helps ensure that the tests that are carried out on the synthetic data are as realistic and representative as possible.

There are 26 test patterns in total. In the first, each of the ridges has a gap of the same distance as the inter-ridge interval

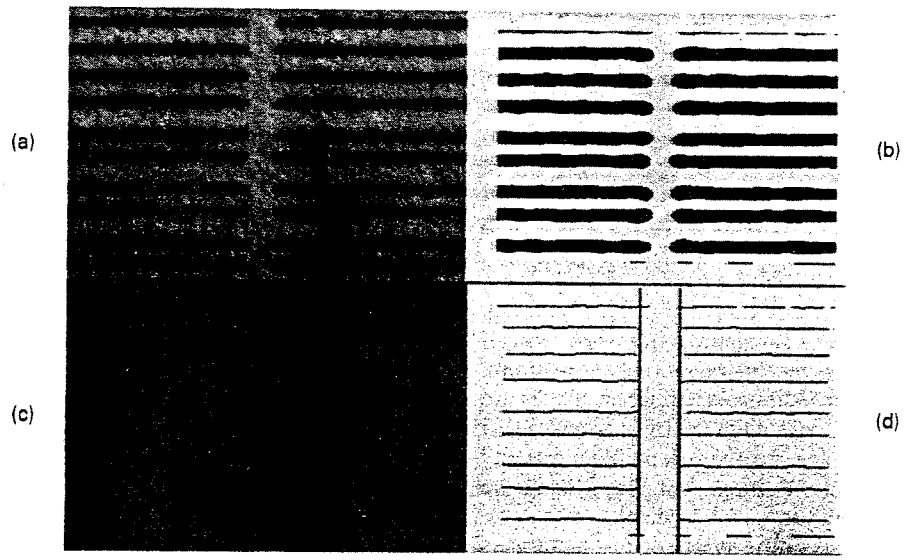

Fig. 8 (a) Gray-level image of a synthetic test pattern; (b) ridge regions derived from Laplacian of Gaussian $\nabla^2 G$ of image, $\sigma = 7.0$ pixels; (c) Hough transform space derived from end points of skeleton of ridge regions; and (d) crease line superimposed on the skeleton of ridge regions.

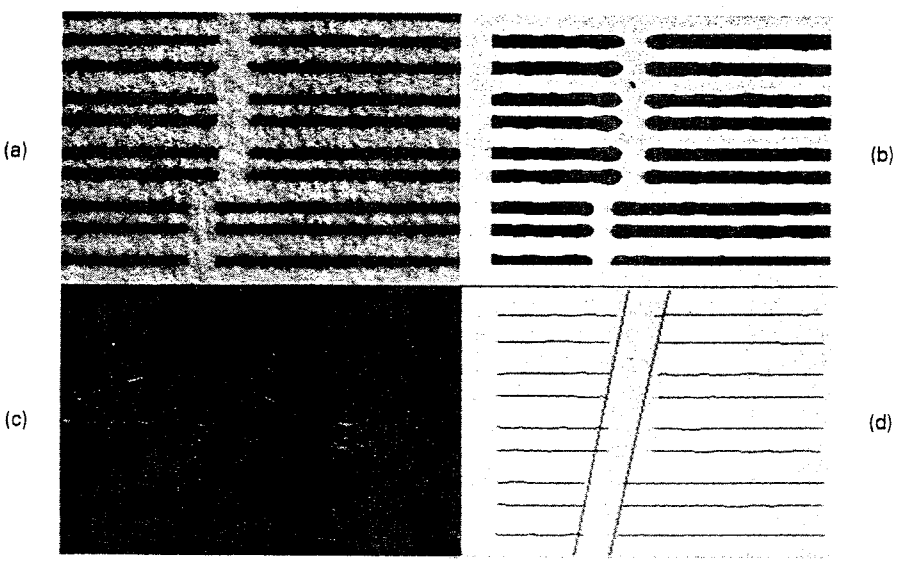

Fig. 9 (a) Gray-level image of a synthetic test pattern; (b) ridge regions derived from Laplacian of Gaussian $\nabla^2 G$ of image, $\sigma = 7.0$ pixels; (c) Hough transform space derived from end points of skeleton of ridge regions; (d) crease line superimposed on the skeleton of ridge regions.

(i.e., the normal distance between the ridges), and this gap is at exactly the same position in each ridge: the ridge end points are all colinear. The remaining 25 patterns are organized in groups of five, with groups 1, 2, 3, 4, and 5 having 10, 20, 30, 40, and 50% of the ridges displaced from the original position, respectively. Within each group, this displacement is varied, with patterns 1, 2, 3, 4, and 5 having the ridge(s) displaced by 50, 100, 150, 200, and 250% of the ridge gap distance.

The technique described in this paper was applied 10 times to each pattern and the number of correct crease detections were recorded. The position and orientation of the pattern in the field of view was altered slightly for each of the 10 applications. Table 1 summarizes the results of this series of tests and details, for each pattern, the rate of correct isolation of the crease. No incorrect crease detections were recorded but some crease lines were not detected. These results demonstrate that the technique is consistent and robust. It fails when 40% or more of the ridges are displaced by 150% or more of the ridge gap. It should be noted that this failure is due, in every case, to the constraint that the detected ridge line cannot contain more than 10% of ridge points along its length (this is the tolerance specified by δ_r^2), and it has been verified that the ridge line would have been detected if this tolerance was altered.

Before proceeding, it is worth considering Fig. 9 again. Notice that a single crease is detected and that the orientation of the crease is approximately 10 degrees from the vertical. Is this a good result? To answer this question, we must first ask several other questions: Is there a crease in the image and how many creases are there? What are their orientations? Are these creases detected and are the orientations correct? Clearly there is a crease in the image but Fig. 9 could be construed to comprise two vertical, but displaced, ridges. However, each of the images of the synthetic test patterns is intended to represent a single crease in which a given, and known, proportion of the ridges are displaced from the true (perfect) position. The orientation of this perfect crease is given by the line joining the end point of the majority of the colinear end points, which, in the case of Fig. 9, is approximately vertical. Nonetheless, Fig. 9 does represent a good result in that the crease has been detected, despite its (intentional) deformation, and there is a quantifiable error in its orientation. However, since in this application it is the detection of the crease and not its orientation which is important, this error in the measurement of orientation is not significant.

5.2 *Tests on Real Data*

The technique described in this paper has also been extensively tested on a somewhat limited data set of fingerprints (in excess of 160 fingerprint regions) and it has proved to be reliable and robust in isolating secondary creases. Representative examples of the results that have been achieved are shown in Figs. 5 through 7.

6 Discussion

The technique described in this paper works well. Nonetheless, a number of issues should be noted.

First, it has been necessary to use a high imaging magnification with an attendantly small field of view in all of the work described. The primary reason for this is to ensure that the features of interest, i.e., the fingerprint ridges, are not undersampled and that they are well represented by the digital images that are the object of the analysis. The current field of view is approximately 8×6 mm with an image resolution of 256×256 pixels and an effective resolution of 32 pixels/mm. (For the sake of comparison, note that the image resolution of fingerprint database at the National Institute of Standards and Technology in the United States is 19.7 pixels/mm.). This means that it would require $4 \times 4 = 16$ images to scan and analyze a single fingerprint of dimensions no greater than 32×24 mm. There is, however, a second reason why so small a field of view is employed. It has been assumed in all of the foregoing that secondary creases are locally linear and that they extend across the greater part of the field of view. If a larger field of view were to be used, these assumptions would no longer be valid.

Second, the spatial filtering of the image to emphasize the morphology of the whorls is a key step in the process. The parameter that governs this spatial scale is σ, the standard deviation in the Laplacian of Gaussian operator. In all of the results cited in this paper, $\sigma = 7.0$ pixels. As we noted at the beginning of the paper, this value is determined empirically through a calibration procedure, based on the normal distance between the ridges. It should also be noted that the technique is not particularly sensitive to the value of this parameter and, for example, a value of $\sigma = 9.0$ pixels has been used successfully in many informal tests.

Third, the extracted ridge end points are almost never exactly colinear and, consequently, the Hough transform space is not as well structured as is often suggested in the literature. Thresholding an unprocessed Hough accumulator gave rise to very unstable results. It is essential to postprocess the Hough accumulator, in the manner described above, prior to the application of a threshold. This threshold should be, and is, adaptively chosen and, importantly, it should be a lower rather than a higher threshold so that candidate lines are not removed from consideration as secondary creases.

Table 1 Results of application of technique on synthetic test patterns: rate of correct isolation of crease (all figures are expressed as percentiles).

Number of Displaced Ridges (% of total)	Displacement of Ridge (% of ridge gap)					
	0	50	100	150	200	250
0	100	-	-	-	-	-
10	-	100	100	100	100	100
20	-	100	100	100	100	100
30	-	100	100	100	100	100
40	-	100	100	0	0	0
50	-	100	100	0	0	0

Invalid lines can then be ignored after subsequent analysis in the original spatial domain of the segmented and skeletonized image, rather than in the Hough transform space.

Two parameters govern the constraints used to validate the secondary crease lines: δ_r^1 and δ_r^2.

The parameter δ_r^1 is a tolerance on the average orientation of the ridges in the vicinity of the crease line. If the crease line is equal to the average orientation, plus or minus this tolerance, then the line is deemed not to be a crease line and is ignored. In all of the examples shown in this paper, δ_r^1 is defined as 20 deg, unless otherwise stated.

The parameter δ_r^2 is a tolerance on the allowable proportion of a crease line that intersects ridges in the original segmented image. That is, a line drawn along the detected crease is likely to cut across parts of a number of ridges; if the ratio of the number of pixels on the line that lie on ridges to the total length of the line is greater than this tolerance δ_r^2, then the line is deemed not to be a crease line and is ignored. In all examples shown in this paper, δ_r^2 is defined to be 0.1 unless it is stated otherwise.

7 Conclusions

The technique for the detection of secondary creases in fingerprints that is described in this paper is robust and works well on all of the data on which it has been tested. This is confirmed by its performance on the synthetic patterns described in the previous section. The clinical usefulness of the technique remains to be assessed by exposing it to a larger data set and by assessing the correlation between the presence of (automatically) detected secondary creases and the incidence of the disease of which secondary creases may be a physical marker.

References

1. J. S. Weszka, "A survey of threshold selection techniques," *Comput. Graph. Image Process.* 7, 259–265 (1978).
2. J. Serra, *Image Analysis and Mathematical Morphology*, Academic Press, London (1982).
3. E. R. Dougherty, *An Introduction to Mathematical Morphology*, Tutorial Text, **TT09**, SPIE Press, Bellingham, WA (1992).
4. D. Marr, *Vision*, W. H. Freeman and Co., San Francisco, CA (1982).
5. D. Marr and E. Hildreth, "Theory of Edge Detection," *Proc. R. Soc. London* **B207**, 187–217 (1980).
6. D. Vernon and G. Sandini, *Parallel Computer Vision—The VIS* $\overrightarrow{a}$ *VIS System*, Ellis Horwood, London (1992).
7. D. Vernon, *Machine Vision*, Prentice-Hall International, London (1991).
8. P. V. C. Hough, "Method and means for recognizing complex patterns," U.S. Patent No. 3,069,654 (1962).

David Vernon was appointed a lecturer in the Department of Computer Science in Trinity College, Dublin, Ireland, in 1983. Prior to this he worked with Westinghouse Electric Inc. as a software engineer from 1979 to 1981 and left Westinghouse to do postgraduate research. He holds the degrees of BA, BAI, MA, and PhD, all awarded by the University of Dublin, Trinity College. His recent research activities have been concerned with the development of robot systems that display true autonomy. His work encompasses computational theories of perception, self-organization, and intelligence. He is the author of one book and a coauthor of another. Dr. Vernon is a fellow of Trinity College, Dublin, Ireland. He is presently on leave at the Commission of the European Communities.