PAPER *Special Issue on Biometrics*

Improved Fingerprint Matching by Distortion Removal

Andrew SENIOR[†] ***and*** **Ruud BOLLE**[†], ***Nonmembers***

SUMMARY Fingerprint recognition is a well-researched problem, and there are several highly accurate systems commercially available. However, this biometric technology still suffers from problems with the handling of bad quality prints. Recent research has begun to tackle the problems of poor quality data. This paper takes a new approach to one problem besetting fingerprints — that of distortion. Previous attempts have been made to ensure that acquired prints are not distorted, but the novel approach presented here corrects distortions in fingerprints that have already been acquired. This correction is a completely automatic and unsupervised operation. The distortion modelling and correction are explained, and results are presented demonstrating significant improvements in matching accuracy through the application of the technique.

1. Introduction

Fingerprints have long been used for the identification of individuals for legal purposes. In an increasingly electronic society, automatic fingerprint recognition has become important as a highly accurate method of identification of individuals for a range of commercial as well as civil government purposes. Fingerprint recognition systems have been developed for many years, and perform very well in ideal circumstances, but problems remain with handling exceptional cases — in particular, poor quality prints. Since an authentication system is as weak as its weakest link, it is desirable to handle all prints without exception processing through some other technology. Much can be achieved by controlling the acquisition of the prints, but in some circumstances this is impossible, or inadequate, and poor prints must be handled. This paper deals with the process of removing distortion from fingerprints to achieve the best match quality. The system can remove distortion in prints acquired as rolled or dabbed prints, scanned live or from ink on paper.

1.1 Minutia-based fingerprint matching

A wide range of fingerprint matching algorithms exist, using a number of different techniques. The majority follow the long-established forensic procedures of minutia-matching. Minutia matching recognition algorithms seek to find recurrences of patterns of minutiae. Minutiae are associated with nearby geometric or topological features and a minutia from one print matches one in another print if the associated features are sufficiently similar. If the minutiae of two fingerprints match well enough then the prints are deemed to be from the same finger.

Fingerprint images are usually clean and high-contrast with distinctive features; under good conditions matches can be made with high accuracy. However, a number of effects contribute to the deterioration of the match made between prints from the same finger. These effects include the following: distortion due to elastic deformation of the finger; cuts and abrasions on the finger; dirt, oil or moisture on the finger or scanner; partial imaging of the finger tip; prints imaged with different rotations. Various techniques exist to compensate for these problems, but this paper will concentrate on correcting for elastic distortion of the finger surface.

This problem arises because of the inherent flexibility of the finger tips. Some elastic distortion will necessarily result as the skin of the finger is not a ruled surface. Consequently, pressing or rolling it against a flat surface induces distortions which vary from one impression to another. Such distortions result in relative translations of features when comparing one print with another. In extreme cases, significant distortions can be induced by applying a lateral force or torque while the finger is applied to the sensor or paper. A planar force will tend to compress or stretch the print, while a torque will induce relative rotations in the features.

There are two consequences of poor matching due to elastic deformation, depending on the deployment scenario of the fingerprint recognition system. In a cooperative scenario, where users wish to be identified, such as physical or electronic access control, significant distortion will reduce the matching score and prevent the user from being recognized. Failure to handle distortion will make the system frustrating for the unnecessarily rejected users, or necessitate its operation with such a high false accept rate that security is compromised. However, as users become used to the system, they will learn (or become conditioned) to minimize the distortion forces at capture time and the problem should diminish. Using an unfamiliar scanner or capture configuration may well invalidate this acquired experience.

More serious is the effect on a non-cooperative identification system, such as in driving licence, voter registration or welfare situations. Here malicious users may try to avoid being recognized by the system (for instance when applying for a second drivers licence under a false name). Here the user may deliberately try to distort the fingerprint to prevent a match against their existing record. In this case, distortion is liable to be a significant problem and the user may have every incentive to make it as severe as possible.

[†]IBM T.J.Watson Research Center, Yorktown Heights, NY 10598, USA.`{aws,bolle}@watson.ibm.com`

1.2 Methods for handling distortion

Hitherto, two approaches have been taken to manage the problems of distortion in fingerprints. The simplest is a combination of physical design and operator training. By carefully designing the fingerprint capture setting, for instance by guiding the finger to the capture surface with mouldings around the scanner, the forces produced during capture can be constrained, especially those produced inadvertently. Of the proliferation of fingerprint scanners now commercially available, many have small capture areas and physical guides which combine to limit the distortion produced by cooperative users.

As mentioned before, as cooperative users are rejected by the system they will learn (or can be trained) that minimizing distortion forces will help them be recognized. Providing visual feedback of the current image helps users to see when prints are poor and consequently aids this learning. In the non-cooperative situation, such images might even be counterproductive. In this situation though, the capture can be supervised by a trained person who can observe the acquisition process and the resulting data, to ensure that an undistorted print is acquired. However, distortions are not obvious either in the behaviour of the person giving the print, or in the resultant data. Police officers responsible for fingerprinting suspects receive special training in fingerprinting to enable them to make clean rolled ink prints with minimal distortion, even from non-cooperative subjects.

Other methods to reduce fingerprint distortion have recently been proposed by Ratha *et al.* [1], [2]. The first is to measure the forces and torques on the scanner directly and prevent capture when excessive force is applied. Naturally this requires specialized hardware to measure the forces at capture time. The second method measures distortion in a video sequence of fingerprint images obtained as a finger is presented to the scanner. When excessive distortion is seen, the print can be rejected and a new print requested. Here again there is a hardware requirement in the form of processor power, since the live video feed from the scanner needs to be processed to measure the distortion at the time of capture if distorted prints are to be rejected when there is still an opportunity to capture another print. Both methods can be used to choose the least distorted print from a capture sequence, though there are other criteria affecting the optimal choice of print, including image quality and area.

All of these methods have the limitation that once a print is acquired, nothing can be done about distortions in the data. Large legacy databases are in use, containing distorted prints which cannot benefit from these techniques. None of the techniques avoids distortion completely. Each method just seeks to prevent the acceptance of the most distorted data.

To cope with the residual distortion error present in any fingerprint to some degree, fingerprint matchers have been designed to allow for errors in the location and angle of minutiae. For instance, Ratha *et al.* [3] use tolerance boxes to allow a minutia in one print to match a minutia falling anywhere within a tolerance region in the other print. Similarly Germain *et al.* [4] bin the (s, θ) parameters of their representation, coarsely quantizing to allow robustness to distortion and measurement noise. However, these techniques reduce the accuracy of the match required and consequently have the limitation of increasing the probability of mistakenly believing that prints from two different fingers came from the same finger. Allowing for distortion in this way makes the matcher's false acceptance probability unnecessarily high. The latter augment their representation with the method of *ridge counting* to alleviate problems of distortion. Here, for every pair of minutiae under consideration, a count is made of the number of ridges crossed by a straight line drawn between the minutiae. This measure is sensitive to distortion only when the line passes close to other minutiae.

Kovács-Vajna [5] also proposes a matching method that specifically takes into account distortion, demonstrating graphically the large cumulative effects that can result from small local distortions. The method uses tolerance bounds in inter-minutiae distances and angles to build up corresponding sets of minutiae that have the same geometry to within certain tolerance. Finally Dynamic Time Warping is used to verify grey-level profiles along the inter-minutiae lines — a technique related to ridge counting.

Thebaud [6] has also proposed a method of locally warping fingerprints to make two fingerprints look more similar. In the case where the prints are from the same finger, this will tend to remove distortion, but it has the disadvantage of making fingerprints from different prints look more alike and will consequently increase the false acceptance rate of a matcher. The distortion removal is time-consuming as it requires gradient descent in the space of local warps, with fingerprint correlation scores as the objective function. Furthermore the procedure must be carried out on every pair of prints to be matched, not on every print as the method described here.

1.3 Outline

The following section introduces the topic of distortion mapping, showing how the distortions in the fingerprint can be estimated and represented, and describing the distortion model used in this work. Section 3 describes how the distortion can be removed from a fingerprint through the application of an *Inverse Distortion Transform*. Results for the application of the technique are presented in section 5, showing that it improves matching performance, especially on deliberately distorted prints. Finally conclusions are presented with an outline of further work.

2. Distortion mapping

We describe here a novel approach to the problem of distortion — removing distortion from previously captured fingerprints. The system we describe here attempts to reconstruct a canonical version of the fingerprint without distortion. The fundamental assumption underlying this process is that the ridges in a fingerprint are constantly spaced, and that deviations from constant spacing indicate distortions introduced by elastic deformation of the finger surface. Distortion is removed by enforcing this constraint — mapping the fingerprint image into one in which the assumption is true.

Of course the assumption is not necessarily true for the surfaces of actual fingers, and is less likely to be true for fingerprints. Indeed we would submit that there is no 'true' fingerprint since this planar entity can only be generated by an ill-defined elastic deformation of a time-varying three-dimensional surface. However, ridges and their images are for the most part uniformly spaced, and the method can work even when the assumption is incorrect, since it aims to consistently map the fingerprint image into a canonical representation. The canonical representation can be in any space in which we represent fingerprints, such as the original grey-scale image, ridge locations or minutia coordinates. The canonical representation is not necessarily reality, but the intention is that all prints from the same finger would be transformed to resemble closely the canonical representation, making comparison simple. Further, since the canonical representation is in the same form as a conventional fingerprint representation, the match can continue on the transformed data without any special algorithm. Because there is no loss in discriminative information by applying the transformation (*e.g.* the space of minutiae representations after application of the algorithm is as large as the space with no distortion removal) there is no loss in accuracy, as would be evidenced in Thebaud's method which works by enforcing similarity on two prints. The algorithm could be applied to remove distortion from the original image, or it could be used to translate and rotate the minutiae (as suffices for the matcher used in section 5).

The distortion-estimation algorithm operates as follows:

- Estimate the inter-ridge distance at many points across the fingerprint.
- Estimate the global average ridge spacing for the fingerprint.
- Construct a representation of local dilations representing deviations from uniform ridge spacing.
- Apply local warps to make the local ridge spacings more nearly match the average.

2.1 Ridge spacing estimation

The ridge spacing is estimated here by using the thinned image that is available as a result of the fingerprint image preprocessing carried out to find the minutiae. The preprocessing used is that of [3]. This preprocessing results in an image where the ridges are one pixel wide lines. Minutiae are easily detected in this image as points with only one neighbour (end points) or with three neighbours (bifurcations). For simplicity of subsequent processing, spline curves are fitted to the chains of pixels. A single spline is fitted to each section of curve between two minutiae. The splines used are cubic B-splines, and the number of control points is increased from four until the approximation of the spline to the original pixel locations is sufficiently close. Given the

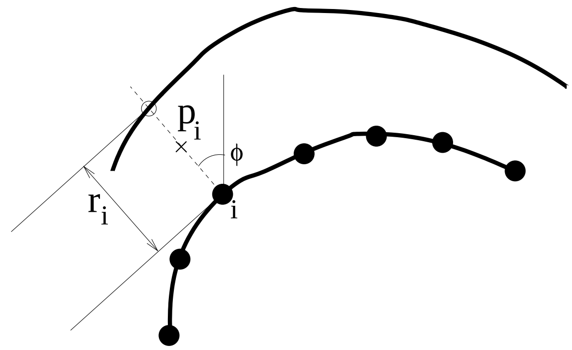

Fig. 1 Two ridges of an enlarged fingerprint, showing the sampling points i (solid circles) and for one such point the search line (dashed) and its intersection with a neighouring ridge (empty circle). The midpoint, $\mathbf{p}_i$, between the search line's two intersections is also shown.

spline representation, we estimate the local ridge separation at regular intervals along each curve. This is done, as shown in Figure 1, by searching perpendicularly to the curve until another ridge is encountered. The distance r_i between the ith initial point, $i = 1 \ldots N$ and the corresponding intersection found is considered as an estimate of the local ridge separation at the midpoint $(x, y) = \mathbf{p}_i$ of the line (at angle ϕ) between the two points. If no ridge is found within a certain number of pixels, then there is deemed to be no matching ridge and no local estimate of the ridge separation is made.

All the local ridge separation estimates are averaged to give an estimate r_{global} for the overall ridge separation, which is deemed to be the 'true' ridge separation for the whole print.

$$r_{global} = \sum_i \frac{r_i}{N} \tag{1}$$

This estimate could be refined by weighting the samples according to their reliability. Factors that can be seen to affect the reliability are the proximity to the edge of the print, or to a minutia. However, for the research carried out here, a simple average was found to suffice.

Deviations of the ridge separation from r_{global} are assumed to be because of distortions in the acquisition process. We now try to model those deviations.

2.2 Ridge distortion model

The distortion is modelled separately in x and y directions. Everywhere on the print axis-parallel relative dilations of the fingerprint are estimated. $D_x(x, y), D_y(x, y)$ contain these

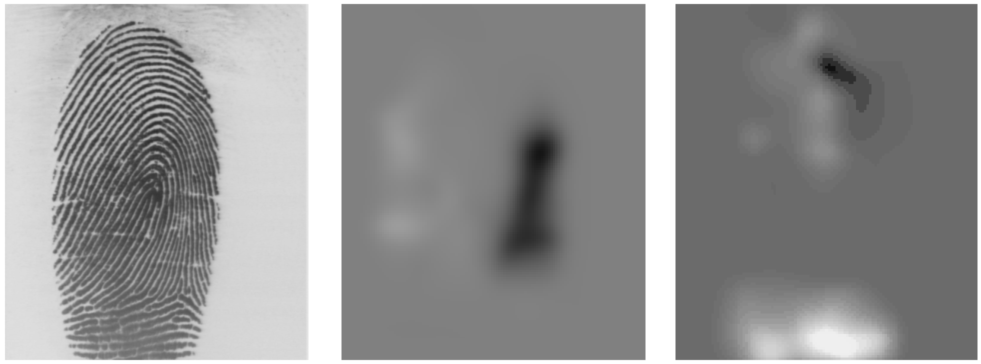

(a) (b) (c)

Fig. 2 A fingerprint distorted by a lateral force. (a) The bitmap captured by a livescan device. (b & c) The estimated dilation maps in x and y directions respectively. Light and dark regions indicate areas where the ridges are estimated to be relatively dilated and compressed respectively.

estimates, termed *Dilation Maps*, with values > 1 for dilation, < 1 for contraction. For estimation purposes, these are stored as arrays, sub-sampled with respect to the original fingerprint resolution. The estimate at (x, y) is contributed to by all the estimates of the local ridge width which fall nearby, and denotes the local dilation relative to the assumed 'true' ridge separation.

For every ridge separation estimate, the local dilation factor d_i is estimated as:

$$d_i = \frac{r_i}{r_{global}}. \tag{2}$$

Since we separate the local dilations into x and y components, the dilation d_i is contributed to by the x dilation and the y dilation except when the ridge separation vector is aligned with the axes. For each direction, we discard ridge separations which are more than an angle ψ from the axis in question (we have used $\psi = 60°$). The dilation estimates are accumulated in a two-dimensional array of bins over all the ridge separation estimates thus:

$$D_x(x, y) = \frac{\sum_{i \in S_x} d_i}{||S_x||} \tag{3}$$

$$\text{where } S_x = \{i : \mathbf{p}_i \approx (x, y), |\cos(\phi)| > \cos(\psi)\}$$

$$D_y(x, y) = \frac{\sum_{i \in S_y} d_i}{||S_y||} \tag{4}$$

$$\text{where } S_y = \{i : \mathbf{p}_i \approx (x, y), |\sin(\phi)| > \cos(\psi)\}$$

where the approximation is the operation of quantizing to the resolution of D. Alternatively, for axis-parallel distortions, we also propose the following 'one-dimensional' model, where dilations are accumulated in one-dimensional bins independent of x and y axes, :

$$D_x(x) = \frac{\sum_{i \in S_x} d_i}{||S_x||} \tag{5}$$

$$\text{where } S_x = \{i : x_i \approx x, |\cos(\phi)| > \cos(\psi)\}$$

$$D_y(y) = \frac{\sum_{i \in S_y} d_i}{||S_y||} \tag{6}$$

$$\text{where } S_y = \{i : y_i \approx y, |\sin(\phi)| > \cos(\psi)\}$$

Since this model has fewer parameters it is more robust, but is also less general. The former two-dimensional models were found to perform better on the larger, less-distorted rolled prints, and the one-dimensional models were found to perform better on distorted dabbed prints.

These accumulated dilation estimates are spatially smoothed, and default to one if no ridge estimate information is available. On a good quality print estimates are available everywhere except at the edges of the print where the information is not used. However, on a poor quality print, where the ridge segments are short, or no ridges are extracted by the ridge-finding algorithm, no ridge-separation estimates are available, and the distortion estimate assumes its prior value of one, leaving the inscrutable areas of the print unaltered.

Now at every point on the print, we have an estimate of the local dilation of the print. Figure 2 shows the distortions estimated automatically in a distorted fingerprint. The estimated dilations can visually be seen to correspond to the change in ridge spacing across the fingerprint (cf. Figure 3a).

Just as with the distortion measures of Ratha *et al.* [1], [2] or indeed any other fingerprint quality measure, the distortion maps can be used to reject bad fingerprints e.g. based upon the maximum distortion or a cumulative measure of the distortion of the whole print, such as:

$$C = \sum_{x,y} |\log D_x(x, y)| + |\log D_y(x, y)|. \tag{7}$$

However, since we know the structure of the dilation, we can attempt to correct it, restoring the fingerprint to a canonical, 'distortion-free' print. (As stated before, this canonical print is not necessarily the distortion-free ground truth, but is the print with constant ridge-width).

3. Distortion correction

Given the distortion maps generated in the previous section, the distortion can be inverted through an *Inverse Distortion Transform* described below. Any fingerprint representation can be generated distortion-free. For the matcher used here, it suffices to apply the Inverse Distortion Transform to the minutiae locations and angles. To generate the distortion-free print shown in figure 3, the Inverse Distortion Transform has been applied to the control points of the ridge splines. The Inverse Distortion Transform could also be applied to the original fingerprint image and any standard fingerprint processing applied to that image to result in any calculable representation of the distortion-free print.

(a) (b)

Fig. 3 The distorted fingerprint of Figure 2 with ridges represented as splines. (a) before and (b) after automatic distortion removal. Note that in (a) the ridges to the right of the loop are closer together than in the rest of the print. At top and bottom they are more widely spaced. In (b) the ridge spacing is uniform.

The Inverse Distortion Transform is a mapping of points in the original image to points in the 'distortion-free' domain. This is simply calculated as the integral of the reciprocals of the dilation map calculated before. To best preserve the most useful, central area of the print, we integrate from its centre, $(0,0)$:

$$(x,y) \mapsto (f_x(x,y), f_y(x,y)) \tag{8}$$

$$f_x(x,y) = \int_0^x \frac{1}{D_x(x',y)}.dx' \tag{9}$$

$$f_y(x,y) = \int_0^y \frac{1}{D_y(x,y')}.dy' \tag{10}$$

Angles (for instance of minutiae) can also be adjusted according to the local torsion induced by the distortion map:

$$\theta \mapsto f_\theta(x,y,\theta) \tag{11}$$

$$f_\theta(x,y,\theta) = \arctan \frac{D_x(x,y)\cos\theta}{D_y(x,y)\sin\theta} \tag{12}$$

4. Computational complexity

The additional time required for distortion removal on a 700 MHz Pentium III processor is about 0.9s for a dabbed print, 1.5s for a rolled print (both at 500dpi). Most of this time is consumed in fitting the splines and estimating the ridge separations, with the creation of the distortion map and the distortion removal being relatively fast operations. The time taken to process a fingerprint is thus roughly proportional to the number of total length of ridges in the print and hence roughly proportional to the print area. The current system has not been optimized for performance; improvements could be made at all stages notably in fitting the splines (or rather working directly on pixel chain representations of the ridges) and optimizing the number of ridge-separation estimates made. Currently about 5000 estimates are made for a typical dabbed print, 6–9000 for rolled prints.

5. Matching results

In this section we present results for matching of fingerprints with and without distortion correction. The matcher used is derived from that mentioned in Ratha *et al.* [3] which gives a similarity score from 0 to 100 for pairs of prints. The experiment uses the image processing described in [3], which extracts ridges and minutiae. The minutiae from two fingerprints are passed to the matcher resulting in a score. In the distortion-removal case, after calculating the distortion maps. The minutiae locations and angles are transformed with the inverse distortion transform. The transformed data are also passed to the matcher to give a score.

To assess the impact of the distortion removal on recognition, we calculate the scores for matching pairs of prints with and without distortion removal. The data used here consists of six sets of images, with each set containing five prints taken from a single finger. All the fingerprints are deliberately distorted by applying a lateral force as the finger comes in contact with the (glass) scanner surface. The distortions are all different. All pair-wise matches are made within each set.

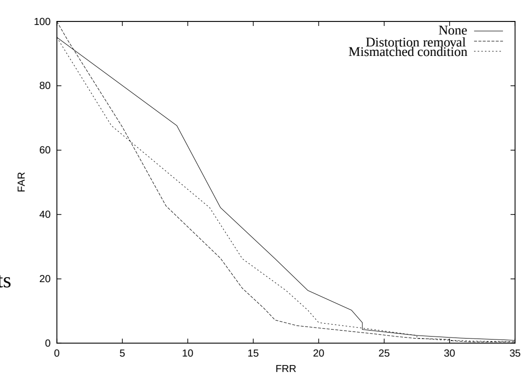

Fig. 4 The Receiver-operating curve for the distorted fingerprint set, with no distortion removal, and with 1-dimensional (x, y, θ) distortion removal.

Table 1 shows the results of the experiments on the small dataset. The fourth column shows the average score

Distortion removal	Angle	Matches	Mean mate score	Significance T(119)	Mean non-mate score	Significance T(749)	Mismatched mate score	Significance T(119)
None	no	120	21.0	–	2.8	–	–	–
2-dimensional	no	120	20.8	-0.44	2.8	0.11	20.2	-1.44
2-dimensional	yes	120	19.9	-1.85	2.7	-0.682	19.7	-2.15
1-dimensional	no	120	22.2	2.21	2.8	0.22	21.2	0.470
1-dimensional	yes	120	22.8	3.26	2.8	0.75	21.2	0.315

Table 1 Changes in matcher performance when correcting for distortion, with and without adjusting minutia angles. Mean scores are presented for mated and non-mated pairs of prints (120 and 750 matches respectively). 'Significance' is the t-statistic for the change in paired matcher scores, for the mated and non-mated pairs. For comparison, the tabulated values are $t_{0.95}(100) = 1.954$ and $t_{0.95}(1000) = 1.646$, showing that the improvement in mated pair scores is significant, but the increase in non-mated pair scores is not. The mismatched column shows the scores when only one print in each match has distortion removal applied.

for mated pairs over 120 trials, with the fifth column showing the significance of the increase of distortion removal over the control. The t-statistic is used calculated on paired matching statistics. Similarly the change in non-mated pair scores is shown to be not significant. The final two columns show that distortion correcting only one print in a match, with the one dimensional model does not harm the match accuracy.

From the table, it can be seen that the 1-dimensional distortion removal helps significantly with matching these distorted prints. Some pairs match less well, but the average match score increases. Figure 4 shows the receiver-operating curve (False accept rate *vs* false reject rate) for the small data set, highlighting the improvement in matching accuracy on distorted prints, and showing that distortion-removal on only one print helps matching accuracy.

PSfrag replacements

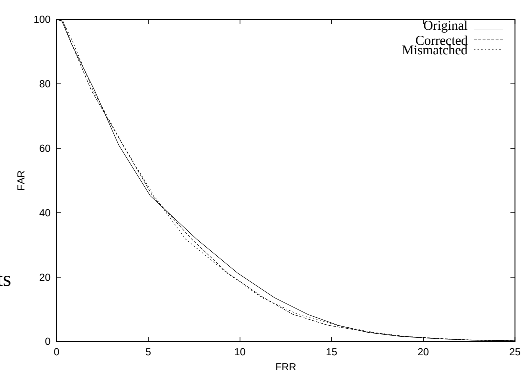

Fig. 5 The Receiver-operating curve for the NIST-4 fingerprint set, with and without distortion removal.

To demonstrate that performance is not harmed when applied to undistorted prints, distortion removal was also carried out on all the pairs of prints from the NIST-4 database. These are ink rolled prints collected professionally, with only a small amount of distortion.

Here an insignificant improvement ($t_{0.25}(\infty) = 0.674$) in the match score is seen by applying distortion removal to the prints. This is reflected in a slight improvement in the ROC as seen in figure 5. Table 2 also shows the insignificant effect of applying the technique to only one print in a pair (which again results in a slight improvement in the ROC). This indicates that the assumption on which the method is based is reasonable — *i.e.* the canonical print matches well with the true print. This has significant implications in practical implementations, particularly for using open fingerprinting standards.

It is true that the features as translated by the distortion removal process are no longer the raw minutiae extracted by the original minutia extraction algorithm on its own and consequently can be considered an algorithm-specific representation. They could thus be deemed not compatible with the public components of currently-proposed common minutiae exchange formats [7],[8, Annex C]. On the other hand it could be argued that distortion removal is simply a further preprocessing step aiming to arrive at the 'true' minutiae locations, though as we have pointed out before these are ill-defined and judging by matching accuracy, the corrected minutiae could be deemed more correct.

Furthermore, table 2 shows that there is no penalty (with this matcher) for presenting 'distortion-free' minutiae and matching against the 'raw' minutiae, as would happen if a distortion-free template were stored in a minutia exchange template, or if this algorithm were always applied when matching against prints from such templates. No penalty is incurred in the mismatch condition, and a significant gain is made in the matched condition when the prints are distorted. Even in a system requiring the original minutiae locations to be stored in a common exchange format, the benefits of the system can be reaped in a number of ways. The 'distortion free' minutiae or sufficient information to derive them could be stored in a 'private', vendor-specific portion of the template.

6. Conclusions and further work

In this work we have proposed a new paradigm for handling distortion in fingerprints. Previous methods of dealing with distortion have sought to prevent distorted fingerprints from being captured or matched, or have allowed for distortion by increasing tolerances which reduce the matcher accuracy. In contrast, we have designed a method which can actually reduce distortion in previously captured fingerprints, in an automatic, unsupervised manner, by exploiting a reasonable assumption about the undistorted fingerprint. The method

Model	Angle	Matched Score	significance	Mismatched score
None	no	26.7	–	4.44
Joint	yes	26.8	0.77	4.45
Cross condition		26.6	-0.57	

Table 2 Changes in matcher performance on the NIST-4 database of rolled prints. With and without correcting for distortion. The 'Cross condition' shows the effect of distortion correcting only one print.

has been demonstrated to improve matching scores significantly, with consequent improvement in accuracy.

Two models have been used. The one-dimensional model appears to work well on badly distorted prints, but these are not handled well by the two-dimensional model, due to lack of constraints. On only lightly distorted prints, the two-dimensional model performs better. Future work work will compare the performance of the two methods with the amount of distortion observed.

Further improvements in matching accuracy can be expected by tuning the matcher to work with the distortion removal system. As noted previously, significant tolerances have hitherto been necessary in minutia pairings to cope with deformation of the fingerprint surface. Since this system reduces the distortion present in the prints, the tolerance boxes could be tightened up, thus reducing the false acceptance rate. Thus far, the system has only been tested on one matcher, which already incorporates some distortion compensation (quantization and ridge-counts). The effect on other matchers is expected to also be beneficial, in inverse proportion to the amount of distortion-invariance already inherent in the matcher, but experimentation is necessary to determine the applicability of the technique.

The system itself could be improved in a number of ways. The assumptions about equal ridge spacing could be relaxed — particularly where it clearly breaks down such as around minutiae or near the edge of the print. More sophisticated models of distortion could be fitted, for instance a thin plate spline, and corresponding Inverse Distortion Transforms applied. While investigating these avenues, we hope to apply the technique to larger databases of fingerprints, including distorted ink rolls and a more general database of prints captured from semiconductor scanners. In the experiments so far, we have not specifically tested distortion due to rotation, though some is certainly present in both databases. It remains to be seen whether this method compensates for this kind of distortion.

References

[1] N. K. Ratha and R. M. Bolle, "Effect of controlled acquisition on fingerprint matching", in *Proceedings of the International Conference on Pattern Recognition*, 1998, vol. II, pp. 1659–1661.

[2] Chitra Dorai, Nalini Ratha, and Ruud Bolle, "Detecting dynamic behavior in compressed fingerprint videos: Distortion", in *Proceedings of Computer Vision and Pattern Recognition*, June 2000, pp. 320–326.

[3] N. K. Ratha, K. Karu, S. Chen, and A. K. Jain, "A real-time matching system for large fingerprint databases", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 8, pp. 799–813, August 1996.

[4] R.S. Germain, A. Califano, and S. Colville, "Fingerprint matching using transformation parameter clustering", *IEEE Computational Science and Engineering*, vol. 4, no. 4, pp. 42–49, Oct-Dec 1997.

[5] Z.M. Kovács-Vajna, "A fingerprint verification system based on triangular matching and dynamic time warping", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 11, pp. 1266–1276, November 2000.

[6] Lawrence R. Thebaud, "Systems and methods with identity verification by comparison and interpretation of skin patterns such as fingerprints", US Patent 5,909,501, June 1999.

[7] R.M. Bolle, N.K. Ratha, A. Senior, and S. Pankanti, "Minutia template exchange format", in *IEEE Workshop on Automatic Identification Advanced Technologies*. October 1999, pp. 74–77, IEEE.

[8] American Association of Motor Vehicles, *AAMVA National Standard for the Driver License/Identification Card*, 2000.