

Recognition of live-scan fingerprints with elastic distortions using correlation filters

Craig Watson

National Institute of Standards and
Technology MS8940
Gaithersburg, Maryland 20899

David Casasent, FELLOW SPIE

Carnegie Mellon University
Department of Electrical and Computer
Engineering
Pittsburgh, Pennsylvania 15213

Abstract. A special NIST database of live-scan fingerprints with elastic distortion was prepared. It is used to evaluate the effect of elastic and other distortions on correlation filters. The need for normalized and finely or coarsely rotationally aligned data is addressed, and performance gains for various cases noted. © 2004 Society of Photo-Optical Instrumentation Engineers. [DOI: 10.1117/1.1783278]

Subject terms: biometric recognition; distortion-invariant filters; elastic distortions; fingerprint recognition; pattern recognition.

Paper 431006 received Aug. 15, 2003; revised manuscript received Sep. 19, 2003; accepted for publication Oct. 21, 2003.

1 Introduction

Fingerprints (FPs) are a well-established biometric for personal identification with proven use and performance. Live-scan FP sensors have become the main FP acquisition sensors for limited access, the U.S. Immigration and Naturalization Service, etc., in the face of recent terrorism. These new sensors produce FP data that are quite different from the classic rolled FPs. Specifically, live-scan FP images exhibit elastic distortions, as shown in Figures 1 and 2. Figure 1 shows two FP images taken at two times in succession. They appear very similar. However, upon examination of the location of specific minutiae (denoted by crosses in Fig. 2), we see that the locations of the minutiae are shifted differently in the two images and that the shifts are different for different regions of the FP. This phenomenon is referred to as elastic distortion and creates recognition problems for live-scan FP data. We have produced a database using such variations and filters to achieve FP recognition in the face of such distortions and others.

Section 2 presents the modified version of NIST Special Database 24¹ that we use. It is the only extensive database of elastic distortions for live-scan data. That section also notes numerous other expected FP distortions (oily, dry, partial, etc., FPs). (Section 6 shows example imagery.) Section 2 also describes the finely and coarsely aligned versions of the database we used (different degrees of preprocessing are needed for each case) and the normalized data we prepared. Section 3 notes the different correlation filters considered; several are novel with respect to their data preprocessing. Each of these FP recognition filters is a sum of training-set images for one FP. Section 4 notes the test and evaluation procedures developed and used to evaluate different algorithms for FP verification and identification. This is a major new result. Section 5 presents initial test results. Analysis of the results follows in Sec. 6.

Much prior work exists in FP recognition. Most uses image processing based on minutia matching;²⁻⁴ neural net classification methods have been found attractive.^{4,5} These methods all require extensive image preprocessing,^{3,6,7} and they require good FPs, since tracking FP ridges to locate minutia regions requires continuous high-quality ridge

lines. These standard methods are thus not of use for about 10% or more of available FPs. Use of correlation filters is thus attractive.⁸⁻¹⁶ However, no prior work we have seen has addressed FP recognition in the face of the realistic elastic distortions expected. Several well-engineered optical correlations have been suggested for FP recognition,⁸⁻¹¹ however, and can benefit greatly from the filter results we present.

2 Modified NIST Special Database 24

The second portion of this database¹ was used, consisting of separate dab FP images obtained over 10 s (the user dabs his finger on the reader a number of times), since these data represent the elastic distortions we wish to consider. There were 55 cases (out of 200) in which the test subjects were able to provide at least nine FP samples. These were the 55 FPs used in our more extensive recognition tests on the ability of various algorithms to perform recognition of FPs with elastic distortions present. We felt that nine versions of each FP should be enough samples to represent most elastically distorted representations. Section 6 investigates the number of samples needed. Our database is not a large one, but it is the largest available livescan one with a number of multiple samples.

For each FP, one image (the one with the least rotation) was selected as the test image; the other eight or more FP images were used to form filters. For each set of data on a given FP, we used the NIST centering algorithm¹⁷⁻¹⁹ to locate the center of the FP and to thus initially align all versions of each FP. We then used NIST software to locate two minutiae close to the FP center (core). These locations were then used to rotate each FP to produce a set of rotationally aligned images for each FP. This yielded a set of *coarsely* centered and rotationally aligned images of each FP. We then iterated shifts and rotations on this set of images and used correlation of the different images to produce a set of *finely* aligned FP images for each FP. Both alignment methods were automated. These represent our coarsely and finely aligned FP data, with elastic distortions left as the major difference between them. The central 350-pixel-diameter area of each centered and rotationally

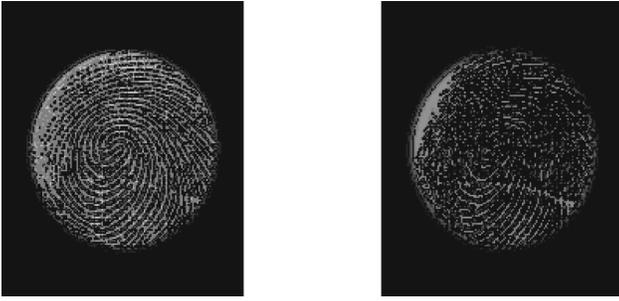


Fig. 1 Same fingerprint with different elastic distortions.

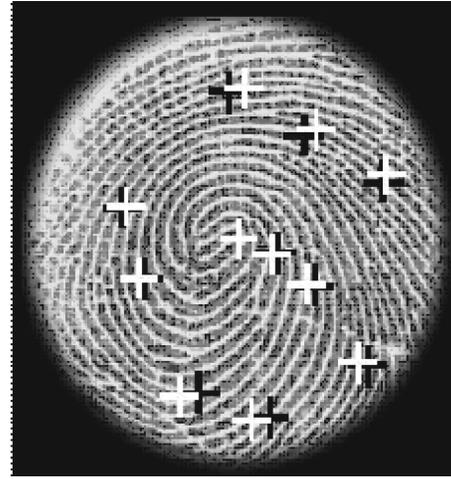


Fig. 2 Locations of minutia shifts in Fig. 1 due to elastic distortions.

aligned FP image was used in filter synthesis and testing. Figure 3(a) shows the original image, Fig. 3(b) shows its rotated version, and Fig. 3(c) shows the 350-pixel-diameter final image used. A set of normalized and nonnormalized images were obtained for each FP and for each finely and coarsely normalized case. Normalized images have the energy within the 350-pixel radius normalized. Normalized data were initially considered for handling the presence of partial FPs, but were eventually found to also be useful for handling dry and oily FP variations.

Thus, the modified NIST 24 database used contained primarily elastically distorted data, but it also included oily, dry, scarred, etc., FP variations. It represents the most extensive database of elastic distortions. Coarse FP alignment is easier to achieve, and good filter performance on such data is thus preferable to making use of fine-aligned data. Distortion-invariant filters (DIFs) were found to be very good for handling such poor FPs, including FPs with scars. Conventional minutia-matching methods often cannot handle such FPs and refuse decision on many of them, while our proposed DIF filters easily handle such cases. Section 5.4 notes this, and Sec. 6 shows examples of such FP images successfully recognized.

3 Distortion-Invariant Filters

Many different types of DIFs exist.²⁰ We considered the synthetic discriminant function (SDF)²¹ and the minimum average noise and correlation discriminant function (MINACE)²² filters as well as a simple averaging filter. The averaging filter, \mathbf{H}_{avg} , gives the sum of all N training-set

images, $\mathbf{F}(I_n)$, divided by N , with filter values scaled so that the maximum training-set peak was 1.0. The scaling factor c is $1/(\text{maximum correlation peak of the training images})$. Thus we have

$$\mathbf{H}_{\text{avg}} = c \frac{\sum \mathbf{F}(I_n)}{N}. \quad (1)$$

The SDF filter \mathbf{H}_{sdf} is also a linear combination of training-set images, but with an unequal set of combination coefficients \mathbf{a} determined by the elements of the vector inner-product matrix \mathbf{V} , since they show the cross-correlation (similarity) of the different training set images. The coefficients are selected to yield equal correlation peak values of 1.0 (the elements of \mathbf{u}) for all training-set images included in the filter. The coefficients are given by

$$\mathbf{a} = \mathbf{V}^{-1} \mathbf{u}, \quad (2)$$

and the filter is the sum

$$\mathbf{H}_{\text{sdf}} = \sum a(n) \mathbf{F}(I_n). \quad (3)$$

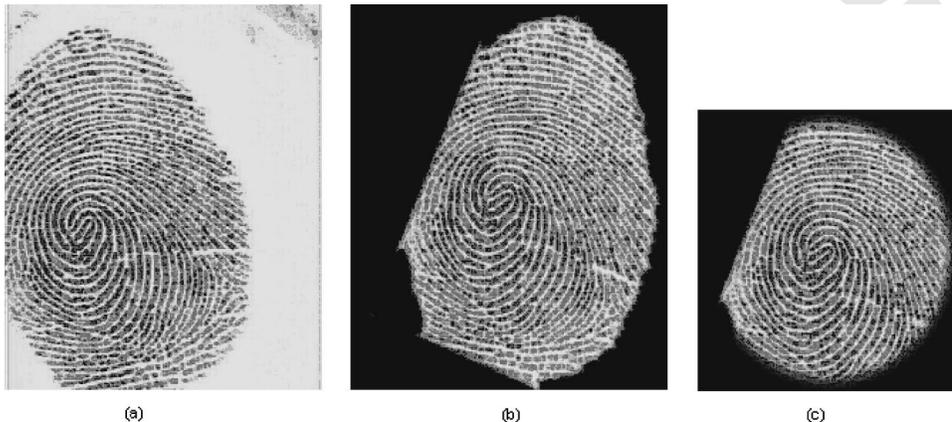


Fig. 3 Processing of NIST Special Database 24.

The MINACE filter \mathbf{H}_{\min} also specifies a correlation peak value of 1.0 (the elements of \mathbf{u}) for each training-set image, but it also minimizes a combination of the effect of distortions: $\mathbf{T}(u, v)$, modeled by white Gaussian noise \mathbf{S}_n (this improves distortion tolerance), and the correlation-plane signal energy for different training-set images $\mathbf{S}_1, \dots, \mathbf{S}_N$ (this reduces sidelobes and produces a sharp correlation peak). The parameter c determines how much each term is minimized. The rows of the transpose conjugate data matrix \mathbf{X}^H are the conjugate Fourier transforms of the training-set images that are included in the filter. The filter must satisfy peak constraints

$$\mathbf{X}^H \mathbf{H}_{\min} = \mathbf{u} \quad (4)$$

and minimize an objective function. This is achieved by the preprocessing function

$$\mathbf{T}(u, v) = \max[\mathbf{S}_1(u, v), \dots, \mathbf{S}_N(u, v), c\mathbf{S}_n(0, 0)]. \quad (5)$$

The MINACE filter is a linear combination of nonlinearly preprocessed images described by

$$\mathbf{H}_{\min} = \mathbf{T}^{-1} \mathbf{X} (\mathbf{X}^H \mathbf{T}^{-1} \mathbf{X})^{-1} \mathbf{u}. \quad (6)$$

In addition, the training-set images used in the averaging and SDF filters are highpass-filtered; this is a simple approximation of the preprocessing provided by advanced filters such as MINACE. Thus, these are not conventional averaging and SDF filters; they are actually much closer to MINACE filters. This highpass preprocessing was done to improve discrimination between similar FPs (by suppressing dc and low-spatial-frequency data). In this work, all training-set images are included in each filter for consistency; this is not generally done.

4 Test Procedure

This represents the first formulation of the test procedure to be used for evaluating candidate filters for FP verification or identification. By *FP verification* we mean the case when a person enters his FP and some PIN indicator of his identity. In use, only the single filter corresponding to this FP will be accessed and compared with the input test FP. If the agreement is above some threshold, the system will verify his identity, else access will be denied. In evaluating different verification algorithms, all filter outputs are considered, as we now discuss. We consider the case of a 50-person (50-FP) database. For each test FP input (50 test inputs, one per FP), there is one filter; we consider all 50 filter outputs for the test input FP. We determine the correlation-peak output for each filter and produce a 50×50 array of correlation-peak values (Fig. 4). If any wrong filter output is above T (the threshold), it is a false alarm (there are thus a maximum of $50^2 - 50 = 2450$ false alarms), and P_{FA} is the percentage of this outcome. If the correct filter output is above T , it contributes to P_C . If no filter output is above T for a given input, it is rejected and contributes to P_R . To produce receiver operating characteristic (ROC) data (P_C versus P_{FA}), we vary T ; thus for a given test input, the

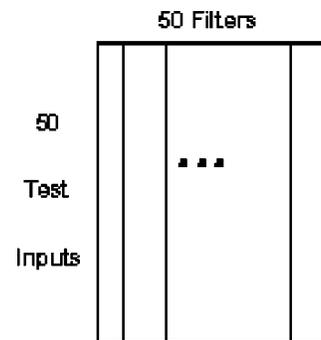


Fig. 4 Verification test procedure for the case of 50 fingers.

correct filter can give the largest output (above T); but as T is lowered, many other filters may give outputs above T and thus false alarms.

In *identification*, the user enters a test FP, it is checked against all filters in the database; the largest filter output yields the class. Only the largest filter output is considered. If the correct filter output is largest, it contributes to P_C ; if any wrong filter is largest ($>T$) for a test input, it is a false alarm. If no filter output is above T , it contributes to P_R . There are a maximum of only 50 false alarms in identification.

Rejection rates (P_R) are also present and of concern. For a given test input, if the true filter output is not above T , that test FP is considered to be rejected so that a false alarm will not occur. We vary T to obtain ROC data. For the identification case, $P_C + P_{FA} + P_R = 100\%$ (and P_{FA} is a percentage out of 50 for our example). For verification, P_{FA} is a percentage of 2450 for our example. This evaluation procedure is equivalent to testing the system on nondata-base test inputs.

5 Initial Test Results

We present test results for the 55 FPs for which eight or more training-set images were available. As noted earlier (Sec. 2), only one test image per FP is available. We note that $\approx 70\%$ of the database consists of partial FPs with a large portion ($>30\%$) missing (in the 350-pixel-diameter image after our centering steps) and that many of the FPs are oily, dry, or scarred. Thus, this is a formidable pattern recognition problem. The fundamental questions to be addressed were: Can DIFs handle elastic distortion? Can DIFs recognize oily or dry or scarred FPs (standard minutia-matching methods have problems with such FPs)? What type of preprocessing (rotation, normalization) is needed, and why? Which type of DIF is best?

5.1 Highpass-Filter Preprocessing

The MINACE and advanced DIFs tend to highpass-filter the training-set (and test-set) data, emphasizing higher spatial frequencies to improve discrimination and reduce false alarms (P_{FA}). We also highpass-filtered all training-set (and test-set) data for the other filter types. All pixels within a radius R of dc in the frequency domain were set to 0, and the response beyond R was tapered over 10 pixels to avoid ringing effects. A validation set was used to select R ; it consisted of a number of distorted versions of one print

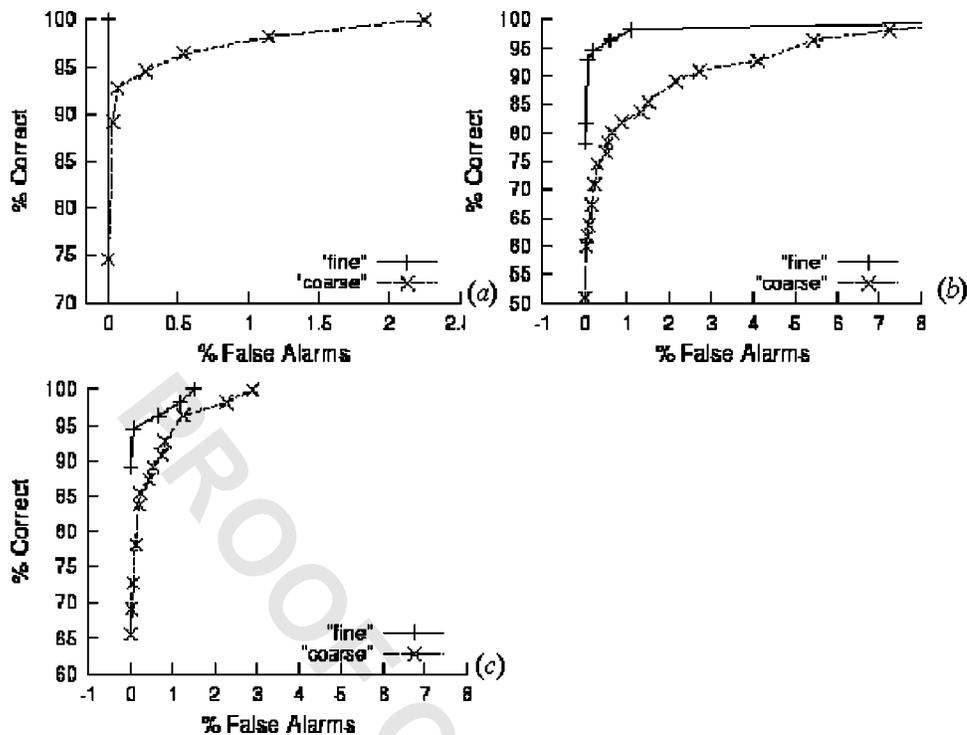


Fig. 5 Fine versus coarse alignment (shift and rotation) improvement using normalized FP data. ROC verification results for the MINACE (a), SDF (b), and average (c) filters.

(true class) and a number of other FPs of the same type (false class). After tests, we chose $R=30$, since it gave a large true-to-false separation; a wide range of R choices performed comparably. Without blocking R pixels, many false FPs were present with larger outputs than the distorted true FPs gave.

The MINACE filter requires selection of one parameter c that determines whether to emphasize detection of distorted FPs or false-alarm reduction. A validation set was used to select c , as suggested elsewhere.²³ Three filters were made from each of these sets of FPs. Tests were against three true FPs and 50 others. All FPs were of the same type. As c for the filter was varied, the lowest true correlation-peak value minus the largest false correlation-peak value was recorded. The c value (5×10^{-5}) with the largest difference was chosen. The exact c value was not critical. Normalization is done after application of c preprocessing (MINACE filter) or after highpass filtering (with SDF and averaging filters).

5.2 Verification Results

We show P_C (the percentage of all 55 test FPs correctly recognized) versus P_{FA} (the percentage of false alarms out of 2970 possible errors—false inputs), and we record the rejection rate P_R (the percentage of the 55 test inputs with correct filter outputs below the threshold T). For verification, $P_C + P_R = 100\%$; $P_{FA} = 0.034\%$ corresponds to one error. At a given T , if any wrong filter gives an output $\geq T$, it contributes to P_{FA} , even if the correct filter output is largest.

5.2.1 Need for finely aligned data

Figure 5 shows test results using normalized data for both coarsely and finely aligned data for the different filters. We expect better results (larger true correlation peaks) with better-aligned FP data, due to the rotation sensitivity. Filters formed from aligned data also have more structure and less blur, due to the summation of shifted training-set FPs; this also yields lower false correlation peaks. All the ROC data in Fig. 5 are better when finely aligned data are used, as expected. We see that the improvement in P_C is at least 5% and can be as much as 25% when finely rather than coarsely aligned data are used.

We now discuss results [Fig. 5(a)] for the MINACE filter. When finely aligned data are used, the test results are perfect ($P_C = 100\%$, $P_{FA} = P_R = 0\%$). These results occur for a threshold $T \geq 0.48$ up to 0.56 (the minimum true correlation test peak is 0.56, and the maximum false peak is 0.47). With coarsely aligned data, the performance is worse ($P_C = 74.5\%$ at $P_{FA} = 0\%$ and $P_R = 25.5\%$ for $T = 0.55$; at a lower $T = 0.41$, we can achieve $P_C = 100\%$ but with $P_{FA} = 2.26\%$, which is very large). Thus, finely aligned data improve performance by 25%, they also provide a larger minimum true correlation peak (0.56 versus 0.41) and a smaller maximum false-alarm rate (0.47 versus 0.546). This result is attributed to the better structure of the filter (less blur) when training data are more finely aligned; true correlation peaks are always larger when rotation alignment of test inputs is better. It is important to note that the correct MINACE filter output is always the largest of those of all 55 filters; this occurred for all cases (coarsely or finely aligned, and normalized or not). This fact is needed in identification

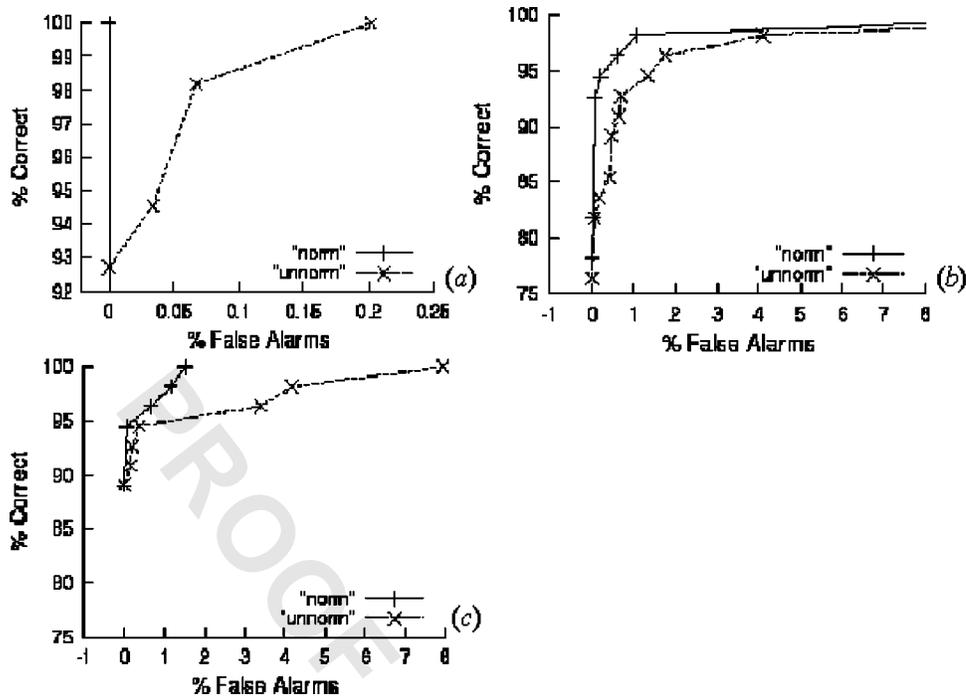


Fig. 6 Normalized versus unnormalized improvement (with fine alignment). ROC verification results for fine-shifted data for the MINACE (a), SDF (b), and average (c) filters.

tests; thus, fine alignment is less needed for identification than for verification. In our verification tests, the same T is used for all test inputs; this is why finely aligned data are needed to increase all true test peaks above some T . Use of different T for known problematic FPs could allow relaxation of the need for fine alignment.

We now discuss results for the SDF [Fig. 5(b)] and averaging [Fig. 5(c)] filters. The trends are similar to those in Fig. 5(a). We first discuss SDF filter results. For the low $P_{FA} < 1\%$ region of interest, P_C is better by 9%–27% when finely aligned data are used. With such data, SDF filters give good performance: $P_C = 92.7\%$ with $P_{FA} = 0.067\%$ (only two false alarms). The problem is that with coarsely aligned data, true correlation peaks are too low (minimum 0.25) and false peaks are too high (maximum 0.55), and use of low T values such as 0.25 is not realistic. Even with finely aligned data, very low T values are needed for high P_C , too low to be realistic: at $T = 0.395$, we have $P_C = 94.5\%$ (3 of 55 FPs missed) and $P_{FA} = 0.2\%$ (6 false alarms).

Averaging-filter results [Fig. 5(c)] are now discussed. Trends are again similar. The value of P_C improves by 23.5% (from 65.4% to 89.1%) with no false alarms, and by 22% (from 72.7% to 94.5%) with $P_{FA} = 0.067\%$ (2 false alarms) when finely rather than coarsely aligned data are used. As with SDF filters, the T levels are too low to be of use if high P_C is needed. In general, MINACE-filter results are much better than those using other filters. SDF- and averaging-filter results are not appreciably different; we attribute this to the highpass preprocessing used, which improves discrimination and makes all filters more similar.

5.2.2 Need for normalized data

Our original motivation for the use of normalized data was for the cases of partial input test FPs (Fig. 11). This is still valid. However, only one of our present 55 test inputs is a partial FP (intentionally). In retrospect, normalization also greatly aids recognition of oily and dry FPs, as we now discuss. Dry FPs have more white areas (and hence higher energy) and lower contrast. Oily FPs have more dark area (and hence lower energy). See Fig. 9(c) and 9(d) in Sec. 6. Dry FPs produce larger correlation peaks (and possibly false alarms); oily FPs produce lower true correlation peaks (and thus require low T , causing problems). Thus, use of normalized data (in training and in testing) will aid in such realistic FP cases. Figure 6 shows test results for all three filters with finely aligned data and with and without use of normalized data. For the MINACE filter, at a low $P_{FA} = 0.03\%$ (one error), use of normalized data improves the results by 5.5%. For the SDF filter, the improvement is $\approx 7\%$ at $P_{FA} = 0.03\%$; for the averaging filter, it is 5% at $P_{FA} = 0.17\%$ (5 errors). Thus, use of normalized data is clearly of help. It increases the lowest true peak (from 0.52 to 0.56 for the MINACE filter), and it significantly reduces the maximum false peak (from 0.67 to 0.48 for the MINACE filter). These lower false maxima allow use of a lower T and result in better ROC data.

5.2.3 Verification test summary

Figure 7 shows ROC (P_C versus P_{FA}) results for the low- P_{FA} region of interest for the case of finely aligned and normalized data. We also include test results for a one-to-one filter case, in which one training image was used as the

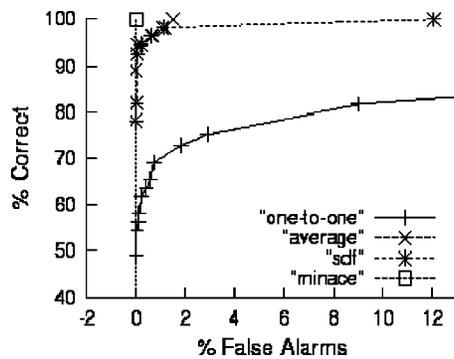


Fig. 7 Verification test results.

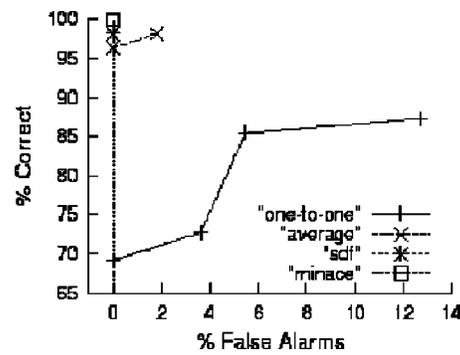


Fig. 8 Identification test results.

filter and another as the test input. This last filter case performs worst; thus clearly DIF methods are needed to handle elastic distortions, as expected. The MINACE filter performs perfectly ($P_C=100\%$, $P_{FA}=P_R=0\%$) and better than all other filters. Thus, its combination weights (used to assemble the filter from the training sets) and its highpass filtering are preferable to those of the other filters (averaging filters have equal weights for all training-set images). Recall that with the highpass filter preprocessing used, the SDF and averaging filters are not the conventional ones.

5.3 Identification Test Results

Recall that in these tests, the user does not state his identity, and each test input is correlated with all 55 filters. If any filter output is $\geq T$, the input is accepted. Thus only the filter with the largest output is considered. If it is the correct one, P_C is incremented; if it is the wrong class, P_{FA} is incremented; if no filter output is above T , the input is rejected. Thus, $P_C + P_{FA} + P_R = 100\%$. The quantity P_{FA} is now a percentage of 55 (not 2970); and the percentages P_{FA} are now much larger than in verification. We use the same T for all FPs in the present test procedures. The only way to achieve $P_C=100\%$ is with no errors ($P_{FA}=0\%$).

Identification test results are now analyzed. The MINACE filter gave perfect results ($P_C=100\%$, $P_{FA}=0\%$, and $P_R=0\%$) over a large range ($T=0.48$ to 0.56). The correct MINACE filter output is always the largest, and the largest false peak is below the minimum true peak for all $55^2=3025$ filter cases. This occurs for all fine versus coarse and normalized versus unnormalized data cases and for all partial, oily, dry, and elastically distorted test FPs. Thus, the MINACE filter is very robust and for identification does not require any preprocessing beyond coarse rotation alignment, which can easily be automated.

SDF filters using normalized and finely aligned data performed much better for identification than for verification ($P_C=98.2\%$ versus $P_C=78.2\%$ with $P_{FA}=0\%$). The correct filter output is largest for 54 of the 55 cases ($P_C=98.2\%$). Thus, SDF filters are very attractive; however, a very low $T=0.315$ was needed. In practice, test prints with low score would be rejected, as the confidence of such classifications is low.

Averaging filters also perform well (a bit worse than the SDF filter). Their identification performance is better than their verification performance, as expected. We obtain P_C

$=98.4\%$ versus 89.1% with $P_{FA}=0\%$, but this occurs at a low $T=0.245$. The problem is again that these averaging filters (like the SDFs ones) give low true correlation peaks (even though the largest peak is the correct one), compared to the MINACE filters. In practice, most errors in the present identification filter tests would be handled by rejecting the input FP and requiring the user to reenter his FP. Figure 8 shows identification test results; only finely aligned and normalized data results are included, as they provide the best scores.

5.4 Minutia-Matching Tests

The performance of minutia-matching software at NIST on this database of 55 fingerprints was also obtained. Since minutia-matching methods require only one reference FP, only the first four elastically distorted versions of each FP were analyzed. This provides useful minutia-matching data for the case of multiple test inputs; how to use these data is the subject of ongoing work. Other minutia-matching methods may perform better, but the one used is the benchmark NIST employs. It does not need rotationally aligned data. For each FP, there are six possible matching combinations. A minutia-matching score of 40 is generally considered acceptable for the matcher used. At this threshold, twelve of the 55 FPs had at least one miss (a true FP version that was not recognized); this is more than 20% of the FPs. Six of the FPs had three or more misses (at least half of the six possible cases), and for one FP all six combinations were missed (regardless of which of the four FP versions was used as the reference and which was the test input, none matched). At this threshold level of 40, eight of the 55 prints, or $\approx 15\%$, gave false alarms (some other false FPs also gave a score above 40); for three of the cases, there were three false alarms for the given FP. These tests do not allow direct comparison of the performance of minutia-matching filters versus DIFs (we could have used different FP versions as the test input for our DIFs; a much larger database is obviously needed; etc.). In addition, DIFs need at least seven or eight training samples, while minutia matchers do not. However, MINACE filters gave no errors in initial tests. Thus, these initial results seem to indicate more misses and more false alarms using standard minutia-matching methods. A comparison of minutia matching and DIFs is not our purpose. These initial results are, however, of use and merit further attention on a larger database with attention to how to perform valid comparisons.

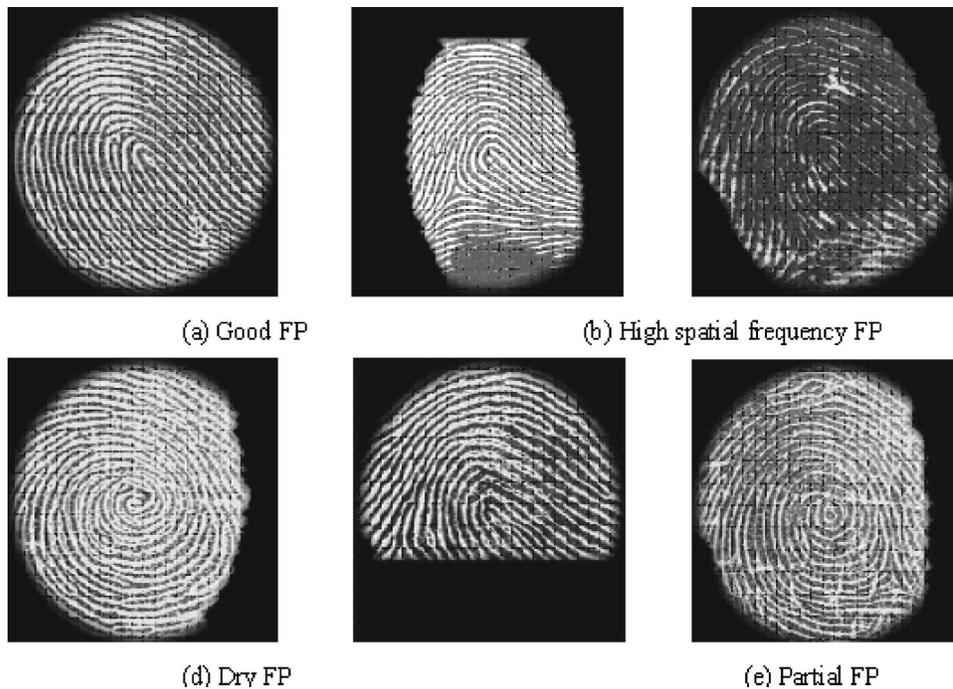


Fig. 9 Representative examples of different test FP variations to be expected.

6 Analysis of Results and FP Imagery

Elastic distortions are the primary distortions we intended to address in this study. DIFs are able to handle such distortions with proper preprocessing. However, many other FP variations arise in practice and in our present database (Fig. 9); all examples shown are test-set FPs, which are (ideally) the best for a given finger. Figure 9(a) shows a good FP (it occupies the entire 350-pixel diameter). Figure 9(b) shows an FP with noticeably higher ridge spatial frequency than other FPs (the person was female, and the smallness of her fingers seems to be the cause of this). Her test errors seem to be due to the need for more accurate rotational alignment, even better than our fine rotational alignment. Figure 9(c) shows an oily FP; as seen, it is very dark and has less energy (thus, it requires use of normalized data for good performance); its ridges and valleys are less clear, being broken up and having low contrast. Figure 9(d) shows a dry FP; it is whiter and has more energy; its ridges are also broken up and less clear, and its contrast is low (dry FPs do not make good contact with the FP scanner).

Figure 9(e) shows a partial FP; Fig. 9(b) is also a partial FP (normalization should help such cases, especially when different parts of the FP are present in training and in testing). Figure 9(f) shows a scarred test FP. Automated minutia systems have problems with scarred FPs, and some of the other cases shown and often reject such FPs, since a sufficient number of minutiae may not be located (e.g., for partial FPs). These examples in Fig. 9 are typical, not isolated, cases. With finely aligned and normalized data, all three filter types successfully performed verification and identification of all cases in Fig. 9.

Some oily FPs [Fig. 10(a)] can cause problems. The correct averaging filter for this test case (with finely aligned and normalized data) gave only a modest (but reasonable) correlation peak of 0.49, and no other filter output was above 0.18. With finely aligned but unnormalized data, a lower correct peak (0.37) occurred, due to the low energy of the test FP; this peak was still the largest among all 55 filters. It gave no identification errors; however, the peak is too low, as other test inputs have false filter peaks above

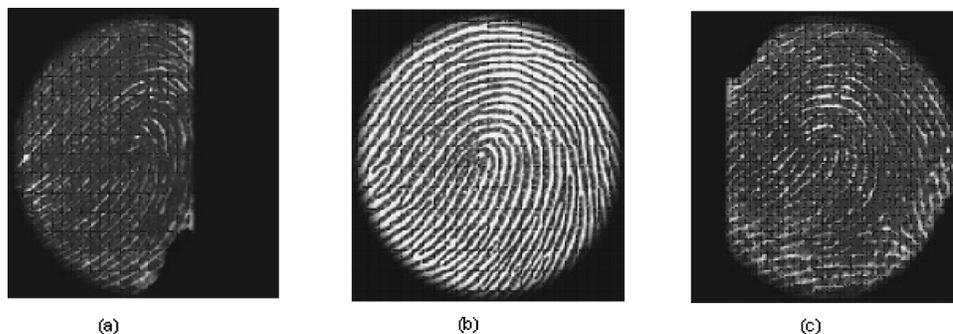


Fig. 10 Images of similar oily partial test fingerprints.

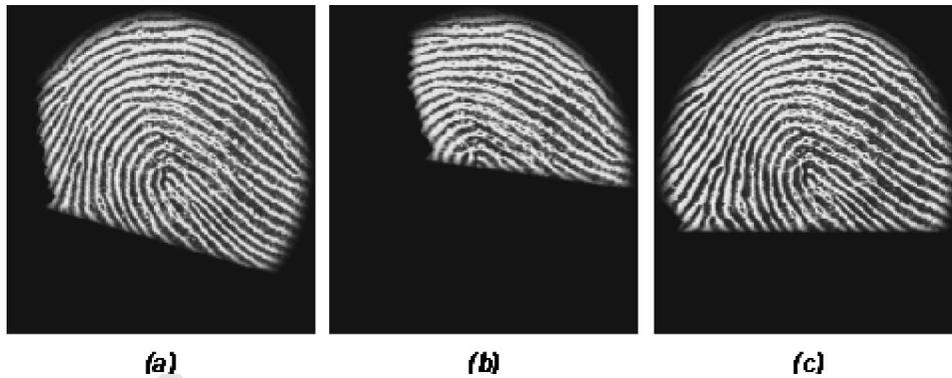


Fig. 11 Images of a typical partial fingerprint from the test (c) and training (a,b) sets.

0.37, and thus errors occur in verification tests. For the test FP in Fig. 10(a), the MINACE filter (with finely aligned and normalized data) give a large peak (0.64); no other filter output was above 0.34. With finely aligned but unnormalized data, a lower but still large peak (0.52) resulted; however, other filters for other FPs [such as in Fig. 10(b)] of a similar class (loops) gave larger outputs (0.67, 0.56) and thus errors. That this was a partial as well as an oily FP contributed to these effects. However, the training set is also an obvious factor. The oily FP in Fig. 10(c) was not a problem. The MINACE filter (with finely aligned data) give large correct peaks of 0.97 and 1.1 for unnormalized and normalized data. The SDF (0.95 and 0.99) and the averaging filter (0.79 and 0.9) also performed well. This is attributed to the fact that the training set captured the oily nature of the test FP (as well as its elastic distortions). This may not always occur when the test FP is taken at a very different time from the training set used to synthesize the filter. Thus, training-set effects are also of concern; variations must be captured by the training set.

Figure 11 shows several of the training-set images and the test-set FP for a case in which all images were of severe partial FPs. Standard minutia matching systems would reject such inputs, as an insufficient number of minutiae would be located. Surprisingly, all filters performed well on this test input (using finely aligned and normalized data). With finely aligned and either normalized or unnormalized data, the correct MINACE filter gave a large correlation output of 0.84 or 0.78. The correct averaging filter (0.84 and 0.80) and the correct SDF filter (0.96 and 0.84) also did

well. We attribute these very high test-set peak values to the fact that the training set is very typical of the test set and captured this expected partial nature of the test FP. If the test FP had been taken at a much later time, then normalization might be expected to have been of more help.

Figure 12 shows ROC MINACE-filter verification results for persons with fewer than eight training-set FPs. For all cases there are some errors. This seems to indicate that having at least eight training-set images is a good choice; much more extensive tests are needed, of course, to confirm this. There are 145 fingers with less than eight training samples: 42 had seven samples, 34 had six samples, etc., and seven did not even produce four samples. As expected, performance degrades with fewer training samples.

7 Summary

A database of FPs with elastic distortions was assembled. Various distortion-invariant filters were considered for classification in the face of elastic and other distortions. The MINACE filter was found to be best among all filters tested. It gave perfect verification and identification results.

The need for finely aligned and normalized databases was investigated. They were found to be needed for verification tests, but not for identification. The database also contained oily, dry, partial, and scarred FPs. These represent other distortion cases that the filters handled successfully. Filters seem tolerant of poor FPs and can achieve correct recognition of them; conventional minutia-matching FP systems seem to have large rejection rates for such cases. Much more extensive tests are needed to confirm these results.

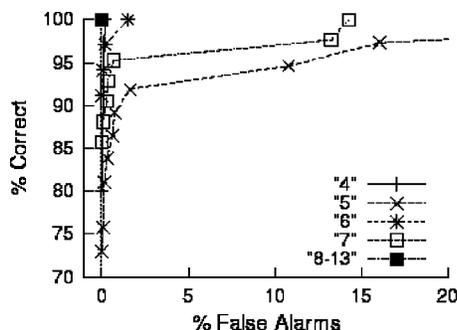


Fig. 12 Verification tests of MINACE filter using 4–7 training-set images per FP.

References

1. C. I. Watson, "NIST Special Database 24: digital video of live-scan fingerprint data," www.itl.nist.gov/iad/894.03/databases/defs/dbases.html#finglist (1998).
2. K. Karu and A. Jain, "Fingerprint classification," *Pattern Recogn.* **29**, 389–404 (Mar. 1996).
3. N. Ratha, K. Karu, and A. Jain, "A real-time matching system for large fingerprint databases," *IEEE Trans. Pattern Anal. Mach. Intell.* **18**(8), 799–813 (1996).
4. M. D. Garris, C. I. Watson, C. L. Wilson, and A. Hicklin, "Studies of fingerprint matching using the NIST verification test bed (VTB)," NISTIR 7020, National Institute of Standards and Technology (July 2003).
5. C. Wilson, C. Watson, and E. Paek, "Combined optical and neural network fingerprint matching," *Proc. SPIE* **3073**, 373–382 (Apr. 1997).
6. D. Maio and D. Maltoni, "Direct gray-scale minutiae detection in

- fingerprints," *IEEE Trans. Pattern Anal. Mach. Intell.* **19**(1), 27–40 (1997).
7. A. K. Jain, L. Hong, S. Pankanti, and R. Bolle, "An identity-authentication system using fingerprints," *Proc. IEEE* **85**(9), 1365–1388 (1997).
 8. A. Stoianov, C. Soutar, and A. Graham, "High-speed fingerprint verification using an optical correlator," *Proc. SPIE* **3386**, 242–252 (1998).
 9. H. Rajbenbach, C. Touret, J. Huignard, M. Curon, and C. Bricot, "Fingerprint database search by optical correlation," *Proc. SPIE* **2752**, 214–223 (Apr. 1996).
 10. F. Gamble, L. M. Frye, and D. R. Grieser, "Real-time fingerprint verification system," *Appl. Opt.* **31**, 652–655 (Feb. 1992).
 11. D. Casasent and C. Wilson, "Optical metrology for industrialization of optical information processing," *Proc. SPIE* **3386**, 2–13 (Apr. 1998).
 12. K. Fielding and J. L. Horner, "Optical fingerprint identification by joint transform correlation," *Opt. Eng.* **30**, 1958 (1991).
 13. T. Grycewicz, "Fingerprint identification with the joint transform correlator using multiple reference fingerprints," *Proc. SPIE* **2490**, 249–254 (1995).
 14. T. Grycewicz, "Fingerprint recognition using binary nonlinear joint transform correlators," *Proc. SPIE* **65**, 57–77 (1997).
 15. D. Roberge, C. Soutar, and B. Kumar, "Optimal correlation filter for fingerprint verification," *Proc. SPIE* **3386**, 123–133 (1998).
 16. C. I. Watson, P. Grother, E. G. Paek, and C. L. Wilson, "Composite filter for VanderLugt correlator," in *Optical Pattern Recognition X*, *Proc. SPIE* **3715**, 53–59 (1999).
 17. G. T. Candela, P. J. Grother, C. I. Watson, R. A. Wilkinson, and C. L. Wilson, "PCASYS—a pattern-level classification automation system for fingerprints," NISTIR 5647, National Institute of Standards and Technology (Aug. 1995).
 18. J. H. Wegstein, "An automated fingerprint identification system," National Institute of Standards and Technology, NBS Special Publication 500-89 (Feb. 1982).
 19. M. D. Garris, C. I. Watson, R. M. McCabe, and C. L. Wilson, "User's guide to NIST fingerprint image software (NFIS)," NISTIR 6813, National Institute of Standards and Technology (2001).
 20. B. V. K. Vijaya Kumar, "Tutorial survey of composite filter designs for optical correlators," *Appl. Opt.* **31**, 4773–4801 (1992).
 21. C. F. Hester and D. Casasent, "Multivariate techniques for multiclass pattern recognition," *Appl. Opt.* **19**, 1758–1761 (1980).
 22. G. Ravichandran and D. Casasent, "Minimum noise and correlation energy (MINACE) optical correlation filter," *Appl. Opt.* **31**, 1823–1833 (10 Apr. 1992).
 23. D. Casasent and S. Ashizawa, "Synthetic aperture radar detection, recognition and clutter rejection with new minimum and correlation energy filters," *Opt. Eng.* **36**(10), 2729–2736 (1997).
 24. H. Zhou, T. H. Chao, and G. Reyes, "Practical filter dynamic range compression for grayscale optical correlator using bipolar-amplitude SLM," in *Optical Pattern Recognition XI*, *Proc. SPIE* **4043**, 90–95 (2000).
 25. R. Gonzalez and P. Wintz, *Digital Image Processing*, Addison-Wesley, ■ (1997).



2003 for work related to measurements and standards in biometrics.

Craig Watson received his BS in electrical engineering technology from the University of Pittsburgh at Johnstown in 1991, and his MS in electrical engineering from Johns Hopkins University in 1997. He has worked with the Image Group and the National Institute of Standards and Technology for the past 12 years. His work has included image compression, automated fingerprint classification, pattern matching, and biometrics. He received a DOC gold medal in



70 journal and conference volumes, a contributor to chapters in 20 books, and an author of over 700 technical publications, on various aspects of optical data processing, image pattern recognition, and product inspection.

David Casasent is a professor at Carnegie Mellon University, Pittsburgh, Pennsylvania, in the Department of Electrical and Computer Engineering, where he holds the George Westinghouse Chair. He is a Fellow of the IEEE, OSA, and SPIE and has received various best-paper awards and other honors. He is a past president of SPIE and of the International Neural Net Society (INNS). He is the author of two books, the editor of one text, the editor of