Elektrotehniški vestnik 76(5): 245–250, 2009
Electrotechnical Review, Ljubljana, Slovenija

Gaussianization of image patches for efficient palmprint recognition

Vitomir Štruc, Nikola Pavešić

University of Ljubljana, Faculty of Electrical Engineering, Tržaška 25, 1001 Ljubljana, Slovenija
E-mail: vitomir.struc@fe.uni-lj.si, nikola.pavesic@fe.uni-lj.si

Abstract. In this paper we present a comparison of the two dominant image preprocessing techniques for palmprint recognition, namely, histogram equalization and mean-variance normalization. We show that both techniques pursue a similar goal and that the difference in recognition efficiency stems from the fact that not all assumptions underlying the mean-variance normalization approach are always met. We present an alternative justification of why histogram equalization ensures enhanced verification performance, and, based on the findings, propose two novel preprocessing techniques: gaussianization of the palmprint images and gaussianization of image patches. We present comparative results obtained on the PolyU database and show that the patch-based normalization technique ensures stat-of-the-art recognition results with a simple feature extraction method and the nearest neighbor classifier.

Key words: Image enhancement, pattern recognition, palmprint recognition, Gaussianization, the PolyU database.

Normiranje slik za učinkovito razpoznavanje dlani

Povzetek. V članku predstavljamo primerjavo dveh uveljavljenih postopkov normiranja slik za samodejeno razpoznavanje dlani, t. j. izravnave histograma in normiranja srednje vrednosti ter variance. Oba postopka spadata v skupino pogosteje uporabljenih načinov normiranja slik dlani, pri čemer izravnava histograma na splošno zagotavlja boljše rezultate razpoznavanja. Učinkovitost normiranja z izravnavo histograma pogosto pripišemo izboljšanju kontrastnih razmer v sliki, medtem ko z normiranjem srednje vrednosti in variance svetilnosti slikovnih elementov slike zagotovimo podobna območja vrednosti, s katerimi je opisana slika (zagotovimo podobne vrednosti svetilnosti slikovnih elementov). Ob predpostavki, da so začetne porazdelitve svetilnosti slik dlani po obliki podobne (in se razlikujejo kvečjemu po vrednosti parametrov, ki porazdelitev določajo), lahko oba postopka normiranja obravnavamo kot transfromaciji vhodne porazdelitve svetilnosti slike. Pri tem z izravnavo histograma (poljubno) vhodno porazdelitev preslikamo v enakomerno porazdelitev, z normiranjem srednje vrednosti in variance pa preslikamo zgolj parametre vhodne porazdelitve. Na podlagi povedanega lahko sklepamo, da je razlika v učinkovitosti obeh postopkov normiranja posledica neprimernih predpostavk postopka normiranja srednje vrednosti in variance, ki zahtevajo zgolj transformacijo parametrov vhodne porazdelitve in ne predvidevajo preslikave celotne porazdelitve svetilnosti slike dlani. V članku tako predstavljamo nova postopka noramiranja, ki vhodno porazdelitev svetilnosti slike (v celoti) preslikata v standardno normalno (t. j., porazdelitev s srednjo vrednostjo nič in varianco ena) in zagotavljata podobne rezultate razpoznavnja kot normiranje z izravnavo histograma. Učinkovitost predlaganih opstopkov smo preverili na javno dostopni podatkovni zbirki PolyU in pri tem dosegli spodbudne rezultate.

Ključne besede: normiranje slik, razpoznavanje vzorcev, razpoznavanje dlani, podatkovna zbirka PolyU

Received 14 July 2009
Accepted 6 September 2009

1 Introduction

Biometric recognition systems represent an emerging technology which holds the potential of substituting (or at least complementing) the classical token- and knowledge-based security schemes. Many such systems exploiting biometric traits such as fingerprints, the face, the iris, the voice, palmprints, etc. have already been presented in the literature. While each of these biometric characteristics has its own strengths and weaknesses, palmprints have an advantage over other modalities, as recognition systems based on palmprint images are considered both user-friendly as well as sufficiently accurate [1].

Palmprint recognition systems have received a great deal of attention from the scientific community in recent years. However, most of the research is focused on the recognition step (i.e., the feature extraction and matching-and-decision procedures) rather than on preprocessing (i.e., enhancement) of palmprint images, which has, without a doubt, a great impact on the final verification performance. Up until the present, two techniques have commonly been employed for the enhancement of palmprints: histogram equalization (HQ) and the transformation of the pixel intensity distribution to a specific mean and variance (MV) [1]. While the employment of HQ is usually justified with its contrast enhancing property, MV is used with the goal of transforming the pixel intensity distribution of a specific palmprint to a common distribution and, hence, making it easier to recognize

palmprints of different classes.

In this paper we will show that both HQ as well as MV pursue the same goal and that the difference in the verification performance stems from the fact that not all assumptions underlying the MV technique are met in the task of palmprint recognition. We present an alternative justification of why histogram equalization ensures enhanced verification performance and, based on the findings, propose two novel preprocessing techniques, namely, Gaussianization (GS) of palmprints and Gaussianization of image patches (GP). Both techniques were successfully evaluated on the PolyU database.

2 Global preprocessing

Consider a palmprint ROI $I(x, y)$ of the size $a \times b$ pixels. The goal of preprocessing (or image enhancement) techniques in palmprint verification systems is to transform the pixels (or their distributions) in $I(x, y)$ in such a way as to enhance the similarity of palmprints belonging to the same class, i.e., person, while simultaneously decreasing the similarity of palmprints belonging to different classes.

2.1 Zero-mean and unit-variance normalization

Zero-mean and unit-variance (ZMUV) normalization is a specific case of MV normalization where the target mean and variance correspond to the values of 0 and 1, respectively. Without the loss of generality we can focus our reasoning to ZMUV normalization rather than MV, as the choice of the values 0 and 1 is as valid as any other.

To normalize a palmprint ROI to ZMUV, each pixel intensity value in $I(x, y)$ is transformed in accordance with the following expression:

$$I^*(x, y) = \frac{I(x, y) - \mu}{\sigma}, \qquad (1)$$

where $I^*(x, y)$ stands for the normalized image and μ and σ denote the mean value and standard deviation of the pixels in $I(x, y)$, respectively.

From the above equation it is clear that ZMUV normalization performs a linear transform to a common, predefined mean and variance. By doing so, it ensures a similar intensity range of the palmprint-pixels of different images and compensates for possible illumination changes present during the image acquisition stage. An example of the deployment of ZMUV normalization on the histogram of a sample palmprint image is shown in Fig. 1.

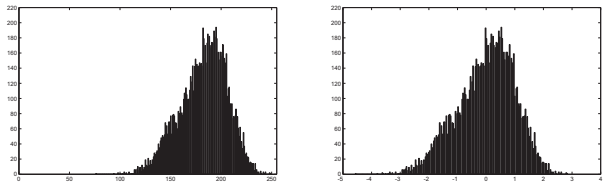

Figure 1. Initial histogram (left), the histogram after ZMUV processing (right)
Slika 1. Izvirni histogram slike (levo), histogram po normiranju s postopkom ZMUV (desno)

The normalization technique is often used in the field of face recognition, where illumination variations represent one of the biggest problems; however, its use in palmprint recognition systems has (to the best of our knowledge) not been theoretically justified, since illumination changes during the acquisition stage do not represent a significant obstacle for palmprint recognition systems and, therefore, scaling the pixel intensities to a similar range does not justify its deployment. The most obvious explanation for its deployment is a common distribution to which the palmprint images are mapped to. This distribution then serves as a common ground to compare different palmprints.

Unfortunately, this justification implies that the initial distributions of the pixel intensities are identical for all palmprint images (regardless of the identity they correspond to) and furthermore that these distributions are symmetric*. If the images deviated from the presented assumptions, the normalization procedure would have no beneficial effect on the recognition rates. A similar finding can also be made for the case where the initial distribution of the pixel intensities for different palmprint images is very similar. In such situations the normalization procedure would perform a scaling operation with similar estimates of μ and σ on all palmprint images and again having only a minimal or no effect at all on the recognition performance. To effectively improve the performance with a normalization technique, the entire distribution of the images should be rendered to a common one (not just its parameters), as it is done with the procedures presented in the remainder of this section.

2.2 Histogram equalization

Histogram equalization (HQ) tries to improve the palmprint ROIs contrast by transforming the distribution of the pixel intensity values in $I(x, y)$ into a common, uniform distribution. The main difference between ZMUV and HQ lies in the fact that ZMUV presumes a certain shape of the initial distribution of the pixel intensities and tries to adjust the mean and variance, while HQ performs a remapping of the pixel intensity values to render the whole distribution of $I(x, y)$ and not just the mean and variance.

In practice, HQ is implemented via the rank transform, in which each pixel value in the $N = ab$ dimensional im-

*Note that if the distribution of the pixel intensities is not symmetric, μ and σ in (1) represent poor estimates of the first and second order statistical moments and the resulting distribution is not the same for all images - especially if the pixel intensity distribution varies in shape from image to image.

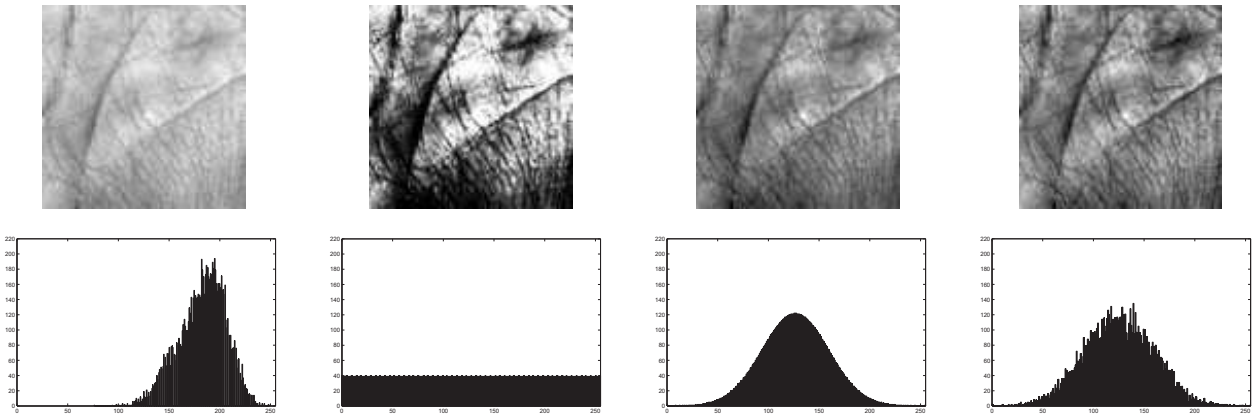

Figure 2. Visual examples of the deployment of different histogram remapping techniques (upper row from left to right): original palmprint image, histogram equalized image, image after gaussianization and image after gaussianization of image patches. The lower row shows the corresponding histograms.
Slika 2. Primeri slik dlani po transformaciji porazdelitve svetilnosti slike (zgornja vrstica od leve proti desni): izvirna slika dlani, slika dlani z izravnanim histogramom, slika dlani po transformaciji porazdelitve svetilnosti slike na standardno normalno, slika dlani po transformaciji porazdelitve svetilnosti delov slike na standardno normalno. V spodnji vrstici so prikazani ustrezni histogrami.

age $I(x, y)$ is replaced with the index (or rank) R to which the pixel intensity would correspond if the pixels intensities were ordered in an ascending manner. For example, the most negative pixel value is assigned a ranking of 1, while the most positive value is assigned a ranking of N. The result is then simply rescaled to the 8-bit interval. An example of the deployment of histogram equalized and its effect on the histogram is shown in the second row of Fig. 2, where the first row corresponds to the original, unprocessed palm-print image and its histogram.

While the use of HQ as a preprocessing technique for palmprint verification systems is commonly justified by its contrast enhancing property, there is another aspect to it as well. We argue that the main reason that HQ improves the verification performance of a palmprint recognition system is the common distribution that the pixel intensities of the palmprint ROIs are mapped to. By remapping the pixel intensities to an uniform distribution, local characteristics potentially intrinsic to a specific palmprint class are revealed and, hence, the discriminative information in the images is enhanced. However, these characteristics are not necessarily linked to the image contrast. To examine this assumption, we propose to map a nonuniform distribution to the palmprint ROIs, as will be done in the next section.

2.3 Gaussianization

Let us again focus on the ZMUV normalization technique presented in Section 2.1. We can generalize this technique to a histogram remapping approach, where the whole distribution of the palmprint ROIs is rendered to the normal distribution with zero-mean and unit-variance. Such an approach is insensitive to the initial distribution of the pixel values in the palmprint ROIs and works without any prior assumptions.

When mapping a target distribution to an image, the first step is rank normalization (as described in the previous section). Once the rank R of each image pixel is determined, the general mapping function to match the target distribution $f(x)$ may be calculated from [2]:

$$\frac{N - R + 0.5}{N} = \int_{x=-\infty}^{t} f(x)dx, \qquad (2)$$

where N denotes the number of pixels in $I(x, y)$ and the goal is to find t via the inverse of the cumulative distribution function, $CDF = \int_{x=-\infty}^{t} f(x)dx$. Obviously, $f(x)$ represents the probability density function for the normal distribution with $\mu = 0$ and $\sigma = 1$. An example of the deployment of the presented gaussianization procedure is shown in the third row of Fig. 2. Here the image in the upper row depicts the transformed image, while the image in the lower row shows its corresponding histogram.

3 Local preprocessing

Considering our assumption that a common distribution can help the verification task by emphasizing local characteristics of the palmprints, we can apply the same reasoning to smaller image patches. As long as these patches are not too small, mapping a common distribution to these subimages should further enhance the performance of palmprint verification systems. It is clear that if the image patches correspond to a predefined distribution, i.e., the pixels intensities of each patch are drawn from the predefined distribution, the global distribution also corresponds to the predefined distribution.

3.1 Gaussianization of image patches

Following the procedure presented in Section 2.3, we propose to map the Gaussian distribution with zero-mean and unit-variance to image patches of the palmprints. Considering the size of the palmprint ROIs to be 100×100 pixels as used in our experiments, we apply the GP technique to image patches of 50×50 pixels, while drawing the image patches with a 50% overlap in the horizontal and vertical directions from the images. After the normalization we combine the patches using a weighting window to ensure that there are no sharp transitions between the patches. Of course, this process affects the global distribution; however, its impact is minimal. Performing local histogram remapping without patch overlapping would induce distortions at patch borders and would limit the choice of the feature extractor. An example of the deployment of image patch gaussianization together with the corresponding histogram is presented in the fourth row of Fig. 2.

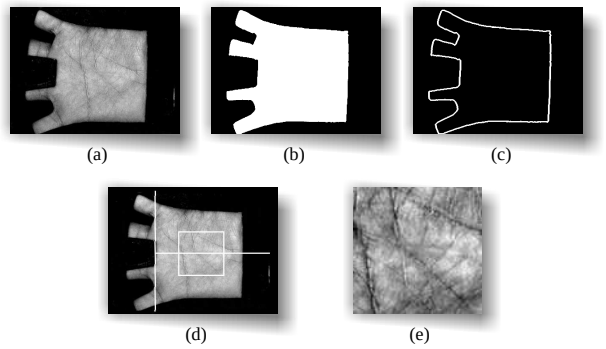

Figure 3. ROI extraction procedure: (a) a hand-image from the PolyU database, (b) the binary image of the hand region, (c) the image of the hand region contour, (d) extraction of the palmprint ROI, (e) the extracted palmprint ROI
Slika 3. Postopek določanja področja zanimanja (ROI): (a) primer slike dlani iz podatkovne zbirke PolyU, (b) binarna slika področja roke, (c) obris področja roke, (d) določitev področja zanimanja, (e) izrezano področje zanimanja

4 Experiments and results

In this section, we present comparative palmprint-verification experiments using the preprocessing techniques introduced in the previous sections, namely, histogram equalization (HQ), zero-mean and unit-variance (ZMUV) normalization, gaussianization (GS) of the pixel intensity distribution and the local preprocessing technique which performs gaussianization on patches (GP) of the palmprint image. As a reference we also provide results obtained in verification experiments with unprocessed (UP) palmprint images.

4.1 PolyU database

To assess the performance of the preprocessing techniques for palmprint verification, we make use of the publicly available PolyU palmprint database [3]. The database was recorded at the Hong Kong Polytechnic University and contains 7752 grey-scale images that correspond to 386 subjects. Each subject in the database is accounted for with approximately 20 palmprints, the exact number, however, varies from case to case.

To extract the region-of-interest (ROI), i.e., the palmprint region, from the images contained in the database, we employed the segmentation procedure proposed by Zhang et al. in [4]. The procedure, which is shown in Fig. 3, is tailored to the image characteristics induced by the recording setup and comprises the following steps: *(i)* the conversion of the original grey-scale palmprint image (Fig. 3a) into its binary form (Fig. 3b), *(ii)* the extraction of the hand region contour from the binary image (Fig. 3c), *(iii)* the detection of the reference points between the little and ring fingers and the middle and index fingers, followed by a subsequent alignment (i.e., rotation) of the image based on these points (Fig. 3d), and *(iv)* the final extraction of the palmprint ROI (Fig. 3e).

The segmented palmprint ROI is rescaled to a standard size of 100×100 pixels and forms the foundation for the assessment of the preprocessing techniques.

To measure the efficiency of the tested preprocessing techniques, we use the three standard error rates commonly employed for assessing the performance of biometric verification systems, i.e., the false rejection error rate (FRR) defined as the frequency with which a client (authorized user) is falsely rejected, the false acceptance error rate (FAR) defined as the frequency with which an impostor (unauthorized user) is falsely accepted and the half total error rate (HTER) defined as the mean of the FAR and FRR. As both, the FAR as well as the FRR depend on the value of the decision threshold t (selecting a value of t which ensures a small FAR results in an increase of the FRR and vice versa), an operating point is usually chosen to ensure predefined values of the FAR and FRR. In our experiments we selected the equal error operating point, where the FAR equals the FRR.

In addition to the presented error rates, we also provide graphical results in the form of receiver operating characteristic (ROC) curves, which plot the FAR against the FRR at various decision thresholds, and the expected performance curves (EPC) [5], which plot the HTER against the parameter α, which controls the relative importance of the two error rates in the expression: $\alpha FAR + (1 - \alpha)FRR$. To produce the EPC curves, an evaluation image set and a test image set are required. For each α a decision threshold t is computed on the evaluation image set which minimizes the weighted sum of the FAR and FRR; this threshold is then used on the test images to determine the value of the HTER employed for constructing the EPCs.

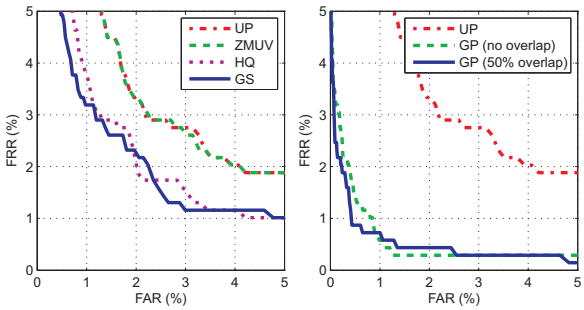

Figure 4. ROC curves of the experiments
Slika 4. ROC krivulje poizkusov

4.2 Experiments

For the experiments presented in the remainder of this section, the PolyU database was partitioned into two groups in the ratio of approximately 6:4 (in terms of subjects). The first group was considered as the client group and the second as the impostor group. Images from these two groups were then assigned to image sets used for training, evaluation and testing. The training images (three per subject) were employed for constructing the principal component (PCA) transformation matrix and building the client templates, i.e., the mean of the feature vectors corresponding to the training images of a given client; the evaluation images were used for determining the decision threshold t, which was then employed in the final performance assessment on the test image set. The presented experimental protocol, which is presented in more detail in [6], resulted in 690 genuine and 279910 impostor verification attempts during the evaluation stage and 3230 genuine and 442750 impostor verification attempts during the test stage. As already indicated above, PCA was used for the feature extraction in our experiments and the nearest neighbor classifier in conjunction with the cosine similarity measure was employed for the matching score calculation. The length of the feature vectors was set to its maximum value. For the experiments we used a HP desktop computer (HP Desktop dc7600 Convertible) with an Intel Pentium 3.2GHz duo-core processor running Windows XP. All techniques were implemented with Matlab R2007b.

In our first series of verification experiments we compared the performance of the global preprocessing techniques HQ, GS and ZMUV to that of unprocessed palmprints (UP). At this stage only images from the evaluation image set were used. From Fig. 4 (left), where the results of the experiments are presented, we can see that ZMUV normalization did not improve the verification performance when compared to unprocessed palmprints. Gaussianization, on the other hand, resulted in a similar performance as histogram equalization clearly improving the verification efficiency upon ZMUV normalization. The presented results confirm our hypothesis that the common distribution rather than contrast enhancement is the main reason for achieving enhanced recognition results. To achieve results comparable to HQ, the whole pixel intensity distribution of the palmprints had to be rendered to $\mathcal{N}(0,1)$.

The second series of our experiments focused on the local preprocessing techniques. To this end, the proposed GP approach as described in Section 3, i.e., with a 50% overlap of image patches, was implemented and compared to the GP approach without patch overlapping. Again, only images from the evaluation images set were used in the experiments. Fig. 4 (right) shows the ROC curves generated during the assessment. Here, the ROC curve obtained with unprocessed palmprints is also given as a reference. While both tested implementations of the GP technique resulted in a similar performance, the variant with a 50% overlap of image patches is still preferable, as it does not induce sharp borders between individual patches and is, therefore, suitable for employment with other feature extraction techniques like, for example, Gabor-wavelet-based methods, which might be susceptible to spurious lines in the palmprints.

In the third series of experiments we aimed at assessing the computational complexity of the normalization techniques. To this end, we measured the time needed to normalize 2000 palmprint images and calculated the average time for a single image. The results of this series of experiments are presented in Table 1. Recall that all of our experiments were conducted on a HP desktop computer (HP Desktop dc7600 Convertible) with a Intel Pentium 3.2GHz duo-core processor and that all the assessed techniques were implemented with Matlab R2007b.

Method	ZMUV	HQ	GS	$GP_{0\%}$	$GP_{50\%}$
$\bar{t}$ (ms)	0.7	2.9	3.4	10.6	23.9

Table 1. Average time needed to normalize a single palmprint image (estimated over 2000 samples).
Tabela 1. Povprečni čas potreben za normalizacijo ene slike dlani (ocenjen na populaciji 2000 slik)

In the fourth series of our verification experiments we made use of the test image set. By using the test images, we produced a number of EPC curves shown in Fig. 5. In addition, the three standard error rates for the decision threshold that ensured an equal value of the FAR and FRR on the evaluation image set are presented in Table 2.

We can see that the local preprocessing techniques clearly outperformed the global ones. The GS approach performed similarly as the HQ technique, both improving upon the error rates obtained with UP and ZMUV. The proposed GP technique performed best and resulted in reduction of the HTER (with the threshold set at the equal error operating point) of more than 65% when compared to the unprocessed palmprints and more than 50% when compared to normalized palmprint images using the HQ

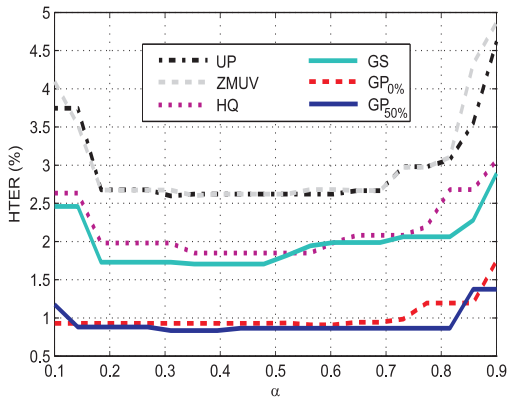

Figure 5. EPC curves obtained on the test image set
Slika 5. EPC krivulje poizkusov na testnem delu zbirke PolyU

Method	UP	HQ	GS	ZMUV	$GP_{0\%}$	$GP_{50\%}$
FAR	2.78	2.02	2.17	2.78	0.97	0.81
FRR	3.53	2.41	2.17	3.53	1.30	1.24
HTER	3.16	2.22	2.17	3.16	1.14	**1.03**

Table 2. Error rates FAR, FRR and HTER (in %) obtained on the test image set
Tabela 2. Napake FAR, FRR in HTER (v %) za poizkuse na testnem delu zbirke PolyU

and GS techniques. It is also interesting to note that the ZMUV-processed images resulted in the same recognition rates as the unprocessed images confirming our assumption that when the pixel intensity distributions of the images are very similar, ZMUV has no (or only a minor) effect on the recognition performance.

As last we compared the performance of the simple PCA technique applied on our preprocessed images (i.e., processed with $GP_{50\%}$) to that of several state-of-the-art methods found in the literature. A detailed description of the assessed methods can be found in the corresponding references.

Method	$GP_{50\%}$	From [7]	From [1]	From [6]
FAR	0.81	1.15	2.08	1.58
FRR	1.24	1.21	1.98	1.76
HTER	**1.03**	1.18	2.03	1.67

Table 3. Comparison with the state-of-the-art methods
Tabela 3. Primerjava z uspešnostjo verifikacije uveljaljenih postopkov razpoznavanja

From Table 3, where the results of the comparison are presented, we can see that the proposed preprocessing technique ensured the best performance for the PCA recognition approach. While these results may or may not be statistically significant, it is far more important that they were achieved using a simple preprocessing technique which can be combined with any of the assessed recognition methods.

5 Conclusion

In this paper we presented two novel preprocessing techniques for palmprint recognition. The techniques are based on the gaussianization of the pixel intensity distributions and serve as means to enhance the discriminative information contained in the palmprints. In verification experiments (with a PCA-based feature extractor and the nearest neighbor classifier) performed on the PolyU database, the techniques were shown to significantly reduce the half total error rate when compared to unprocessed palmprint images.

6 References

[1] A. Kumar, D.C.M. Wong, H.C. Shen, and A.K. Jain, "Personal authentication using hand images," *Pattern Recognition Letters*, vol. 27, no. 13, pp. 1478–1486, 2006.

[2] J. Pelecanos and S. Sridharam, "Feature warping for robust speaker verification," in *Proceedings of the Speaker Recognition Workshop*, 2001, pp. 213–218.

[3] The PolyU palmprint database (accessed January 2009), *http://www4.comp.polyu.edu.hk/∼biometrics/*.

[4] D. Zhang, W.K. Kong, and M. Wong, "Online palmprint identification," *IEEE TPAMI*, vol. 25, pp. 1041–1050, September 2003.

[5] S. Bengio and J. Marithoz, "The expected performance curve: a new assessment measure for person authentication," in *Proceedings of the speaker and Language recognition workshop (Odyssey)*, 2004, pp. 279–284.

[6] V. Štruc and N. Pavešić, "Phase-congruency features for palm-print verification," *IET Signal Processing*, vol. 3, no. 4, pp. 258–268, 2009.

[7] X. Wu, D. Zhang, and K. Wang, "Fisherpalms based palmprint recognition," *Pattern Recognition Letters*, vol. 24, no. 15, pp. 2829–2838, 2003.

Vitomir Štruc received his B.Sc. degree in electrical engineering from the University of Ljubljana in 2005. He is currently working as a researcher at the Laboratory of Artificial Perception, Systems and Cybernetics at the Faculty of Electrical Engineering of the University in Ljubljana. His research interests include pattern recognition, machine learning and biometrics.

Nikola Pavešić received his B.Sc. degree in electronics, M.Sc. degree in automatics, and Ph.D. degree in electrical engineering from the University of Ljubljana, Slovenia, in 1970, 1973 and 1976, respectively. Since 1970 he has been a staff member at the Faculty of Electrical Engineering in Ljubljana, where he is currently head of the Laboratory of Artificial Perception, Systems and Cybernetics. His research interests include pattern recognition, neural networks, image processing, speech processing, and information theory.