

Histogram Remapping as a Preprocessing Step for Robust Face Recognition

VITOMIR ŠTRUC

University of Ljubljana
Faculty of Electrical Engineering
Tržaška cesta 25, 1000 Ljubljana
SLOVENIA
vitomir.struc@fe.uni-lj.si

JANEZ ŽIBERT

University of Ljubljana
Faculty of Electrical Engineering
Tržaška cesta 25, 1000 Ljubljana
SLOVENIA
janez.zibert@fe.uni-lj.si

NIKOLA PAVEŠIĆ

University of Ljubljana
Faculty of Electrical Engineering
Tržaška cesta 25, 1000 Ljubljana
SLOVENIA
nikola.pavesic@fe.uni-lj.si

Abstract: Image preprocessing techniques represent an essential part of a face recognition systems, which has a great impact on the performance and robustness of the recognition procedure. Amongst the number of techniques already presented in the literature, histogram equalization has emerged as the dominant preprocessing technique and is regularly used for the task of face recognition. With the property of increasing the global contrast of the facial image while simultaneously compensating for the illumination conditions present at the image acquisition stage, it represents a useful preprocessing step, which can ensure enhanced and more robust recognition performance. Even though, more elaborate normalization techniques, such as the multiscale retinex technique, isotropic and anisotropic smoothing, have been introduced to field of face recognition, they have been found to be more of a complement than a real substitute for histogram equalization. However, by closer examining the characteristics of histogram equalization, one can quickly discover that it represents only a specific case of a more general concept of histogram remapping techniques (which may have similar characteristics as histogram equalization does). While histogram equalization remapps the histogram of a given facial image to a uniform distribution, the target distribution could easily be replaced with an arbitrary one. As there is no theoretical justification of why the uniform distribution should be preferred to other target distributions, the question arises: how do other (non-uniform) target distributions influence the face recognition process and are they better suited for the recognition task. To tackle this issues, we present in this paper an empirical assessment of the concept of histogram remapping with the following target distributions: the uniform, the normal, the lognormal and the exponential distribution. We perform comparative experiments on the publicly available XM2VTS and YaleB databases and conclude that similar or even better recognition results that those ensured by histogram equalization can be achieved when other (non-uniform) target distribution are considered for the histogram remapping. This enhanced performance, however, comes at a price, as the nonuniform distributions rely on some parameters which have to be trained or selected appropriately to achieve the optimal performance.

Key-Words: Face recognition, preprocessing techniques, histogram equalization, histogram remapping

1 Introduction

Image preprocessing techniques represent an essential part of a face recognition systems, which has a great impact on the performance and robustness of the recognition procedure. The main objective of these techniques is to enhance the discriminative information contained in the facial images and to ensure that environmental factors, such as the ambient illumination present at the image acquisition stage, do not influence the process of facial-feature-extraction. Especially the latter objective is of major importance, as it is known that the variability induced to a given subject's face images by illumination is often larger than the variability induced to different facial images by the subject's identity [1]. Due to this great sus-

ceptibility of face recognition systems to illumination variations, numerous approaches have been proposed in the literature to achieve illumination invariant face recognition.

Land and McCann [2], for example, tried to achieve illumination invariance through the illumination-reflection model of digital images. The model assumes that each image can be represented as a product of illumination and reflectance, where illumination represents the amount of measured light intensities and reflectance denotes the amount of light reflected by a given object. Under the assumption that the illumination component of the model can be approximated as a blurred, i.e., low pass filtered, version of the original image, the illumination insensitive

reflectance can be computed by dividing the original image with its illumination.

Based on Lands and McCanns work, Jobson et al [3] proposed a multiscale retinex (MSR) technique where the illumination component of the illumination-reflectance model is computed as a weighted sum of the filter outputs of the original image and several gaussian filters of different widths [4]. In comparison to the single filter approach the use of several gaussian filters results in a more stable performance, especially, when high and low pixel intensity regions need to be processed efficiently.

Gross and Brajovic [5] showed that illumination invariance could be achieved with the help of anisotropic smoothing. In their method the illumination component of the illumination-reflectance model is estimated based on the minimization of an energy-based cost function. The method again presumes that the illumination component is a blurred version of the original image, however, the amount of blurring at each pixel location is controlled by the local image contrast, hence, the name anisotropic smoothing.

Du and Ward [6] presented a preprocessing technique based on the wavelet transform. The authors proposed to apply histogram equalization to the sub-band image generated by the so-called approximation-wavelet-coefficients and to emphasize the remaining subbands generated by the detail-wavelet-coefficients. Through the histogram equalization step the image contrast is improved while the second step enhances the edge information. The final illumination compensated image is obtained by simply employing the inverse wavelet transform on the modified coefficient subbands.

While all of the presented techniques represent efficient preprocessing methods, the dominant preprocessing approach most commonly used in face recognition systems today is histogram equalization. Here, the pixel intensity values are mapped from their original distribution to a uniform one, thus improving contrast and simultaneously compensating for the illumination-introduced variations in the appearance of the facial images. Histogram equalization has proven itself to be a powerful preprocessing technique capable of improving both, the recognition performance as well as the robustness of face recognition systems; however, it is usually used solely based on its empirically determined usefulness.

We will show in this paper that histogram equalization represents only a specific case of a more general concept where the pixel values of a given image are altered in such a way that the pixel intensity distribution fits a predefined one. Rather than fitting uniform distributions to images, as it is done in the case of histogram equalization, we propose to use other,

nonuniform distributions potentially more suitable for the task of face recognition. We present experimental results obtained on the XM2VTS and YaleB face databases which show that fitting lognormal distributions to facial images results in similar recognition performance as the use of histogram equalization when identity claims are verified based on images captured in controlled conditions, while it ensures more robust performance on facial images severely affected by illumination variations. We also perform experiments with other distributions and report the comparative results.

The rest of the paper is structured as follows. Section 2 presents the basic structure and operating modes of a face recognition system. In Section 3 the basic principles of histogram equalization are described. Section 4 details the general procedure of fitting arbitrary distributions to pixel intensity values. In Section 5 we present the results of various verification experiments and conclude the paper in Section 6.

2 The structure of face recognition systems

Face recognition systems are basically pattern recognition systems comprising the following four modules [7], [8], [9]:

- *the image acquisition or sensor module* which captures an image of the users face,
- *the preprocessing module* which extracts the facial region from the captured image and normalizes it in respect to size and rotation and ultimately performs a photometric normalization procedure on the facial region,
- *the feature extraction module* which extracts a set of representative features from the normalized facial region, and
- *the matching module* which matches the feature set extracted from the given face image with the templates/models stored in the systems database and based on the outcome of the matching procedure makes a decision regarding the identity of the user.

An example of the structure of a face recognition system is presented in Fig. 1. Here, the upper row shows the enrollment stage, where a user is enrolled into the system. In the enrollment stage a user-model or template is constructed from the feature vectors extracted from a number of (users) enrollment face images and stored in the systems database. In

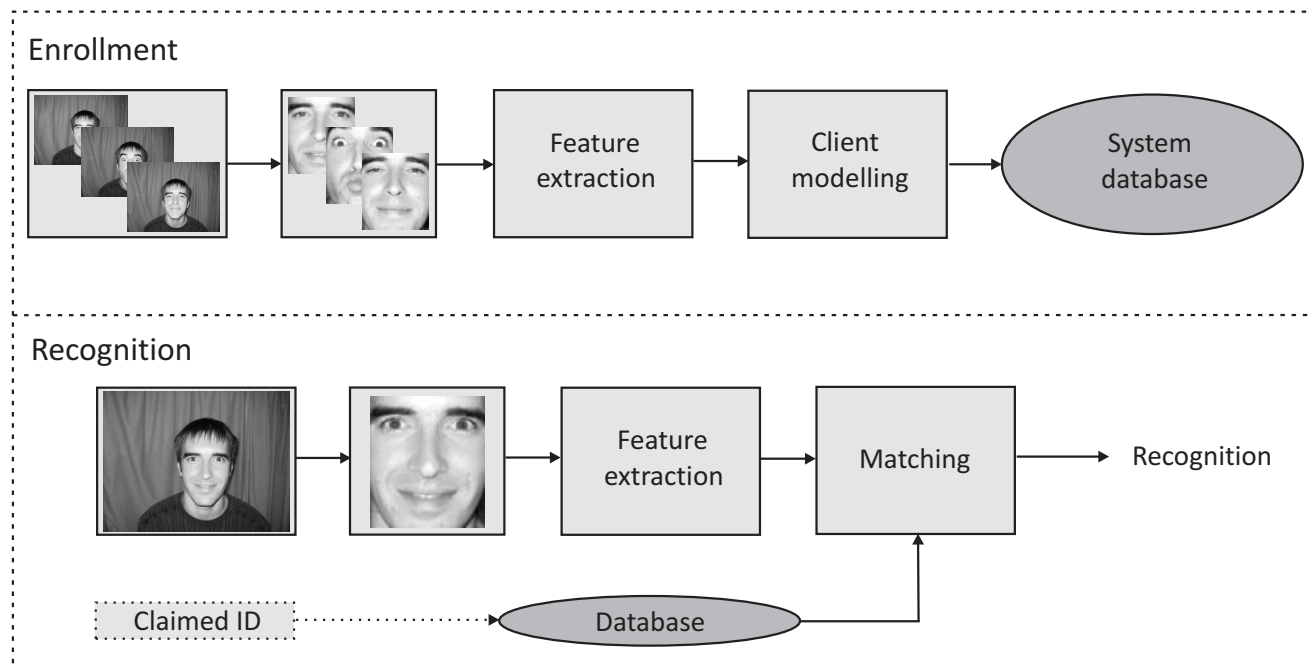


Figure 1: The structure of a face recognition system: upper row - the enrollment stage, lower row - the recognition stage

the experiments presented in Section 5 the mean feature vector of all enrollment images of a given user is computed during the enrollment stage and used as the user-model/template.

The lower row of Fig. 1 shows the recognition stage in a face recognition system. Here, we have to make a distinction between the two operating modes of a face recognition system, i.e., the verification and identification modes. In identification mode, the system outputs the identity associated with the image most similar to the test/query image (in term of feature vector similarity). Thus, to make a decision regarding the identity of the user, the system has to compare the feature vector extracted from the test/query image with all templates/models stored in the systems database. In verification mode, on the other hand, the user has to explicitly claim an identity (dotted part of lower row of Fig. 1). The system then simply compares the feature vector corresponding to the test image with the template associated with the claimed identity and based on this comparison either validates the identity claim or not.

As already indicated in the previous section, we will focus in this paper on the first part of the presented processing chain, namely, the preprocessing module.

3 Histogram equalization

Consider a face image $I(x,y)$ with N pixels and a total number of k grey levels, e.g., 256 grey levels for an 8-bit image. Histogram equalization tries to transform the distribution of the pixel intensity values in the image $I(x,y)$ into a uniform distribution and consequently to improve the images global contrast. It does so by better distributing the pixel intensity values, i.e., by spreading out the most frequent pixel intensity values. Formally, histogram equalization can be defined as follows: given the probability $p(i) = \frac{n_i}{N}$ (i.e., the actual histogram of $I(x,y)$) of an occurrence of a pixel with a grey level of i , where $i \in 0, 1, \dots, k-1$ and n_i denotes the number of pixels in $I(x,y)$ with the grey level value of i , the mapping from a given intensity value i to a new transformed one i_{new} is defined by:

$$i_{new} = \sum_{i=0}^{k-1} \frac{n_i}{N} = \sum_{i=0}^{k-1} p(i). \quad (1)$$

Equation (1) defines a mapping of the pixels' intensity values from their original range (0-255) to the domain of [0,1]. Thus, to obtain pixel values in the original domain, e.g., the 8-bit interval, the values i_{new} have to be rescaled. A visual example of histogram equalization is presented in Fig. 2.

As we can see, the contrast of the processed image is greatly increased - unfortunately, this also ap-




Figure 2: A sample image and its histogram before (upper row) and after (lower row) histogram equalization

plies to the background noise present in the facial image. However, it is this property (i.e., contrast enhancement) which makes histogram equalization one of the most often employed (heuristic) preprocessing techniques in the field of face recognition.

4 From uniform to arbitrary distributions

While transforming the pixel intensity distribution of a facial image to a uniform one (as it is done with histogram equalization) was empirically shown to provide enhanced face recognition performance when compared to unprocessed facial images, it is still just a useful heuristic, which represents a specific case of a more general concept of histogram remapping techniques. In this class of techniques, the target distribution is not limited to a uniform distribution, but can represent an arbitrary one, such as the normal, the lognormal, the exponential or any other distribution. There is no guarantee that the uniform distribution is best suited for the face recognition task, thus, other options of histogram remapping have to be evaluated as well. In the remainder of this section, we will first review the basic concepts governing histogram remapping techniques and then provide examples for three well-known distributions.

4.1 Histogram remapping

The first step common to all histogram remapping techniques is the transformation of the pixel intensity

values of the given image via the rank transform. The rank transform is basically a histogram equalization procedure which renders the histogram of the given image in such a way that the resulting histogram approximates the uniform distribution. Here, each pixel value in a N dimensional image $I(x,y)$ is replaced with the index (or rank) R the pixel would correspond to if the image pixels 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 described procedure is equivalent to the one presented in Section 3, the only difference is in the way the new, mapped pixel intensity values are computed and in the domain they are mapped to.

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 [10],[11]:

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

where the goal is to find t . Obviously, the right hand side represents the target cumulative distribution function (CDF) while the left hand side represents a scalar value. If we denote the CDF with $F(x)$ and the scalar on the left with u , then the mapped value t is found by computing the following expression:

$$t = F^{-1}(u), \quad (3)$$

where F^{-1} denotes the inverse of the CDF, i.e.,

4.2 Normal distribution mapping

The first distribution considered in this paper in the normal distribution. The expression for the normal curve is given with the following expression:

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(\frac{-(x - \mu)^2}{2\sigma^2}\right), \quad (4)$$

where μ denotes the mean value and $\sigma > 0$ represents the standard deviation.

When implementing the histogram mapping technique with the target being the normal distribution, we have to select two parameters, i.e., μ and σ . As we will rescale the mapped pixel values to the 8-bit interval (for visualization purposes), the choice of μ does not influence the outcome of mapping procedure, we do, however set it to 0, for σ , on the other hand, we select the value of 1, so the target distribution used in the experiments presented in Section 5 is the standard normal distribution. A visual example of the performed histogram remapping is shown in Fig. 3. Here




Figure 3: A sample image and its histogram after mapping the histogram to a normal distribution

the same input image was used as for the example in Fig. 2. Note that the pixel values for the mapped image had to be inverted, i.e., subtracted from the value of 255, which is a consequence of ordering the pixel values in an ascending manner. The appearance of the image (inverted or not) has no direct impact on the recognition.

4.3 Lognormal distribution mapping

The second distribution considered in this paper is the lognormal distribution. Here, the expression for the density function is given by:

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} \frac{\exp(-(\ln x - \mu)^2/2\sigma^2)}{x}, \quad (5)$$

where the parameters μ and $\sigma > 0$ again define the

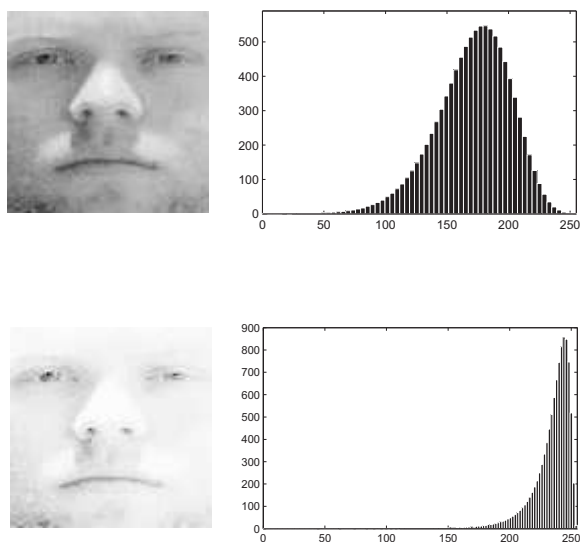


Figure 4: A sample image and its histogram after mapping the histogram to the log-normal distribution: upper row - $\sigma = 0.2$, lower row: $\sigma = 0.7$

shape of the distribution. Note that $x \geq 0$. For our experiments we used two values of σ , i.e., $\sigma = 0.2$ and

$\sigma = 0.7$. Visual examples of the transforms (again inverted) for both values of σ are shown in Fig. 4.

4.4 Exponential distribution mapping

The last distribution considered in this paper is the exponential distribution which for $x \geq 0$ can be defined as follows:

$$f(x) = \lambda \exp(-\lambda x). \quad (6)$$

Here λ denotes the parameter of the distribution often called the rate parameter. For our experiments the rate parameter λ was set to the value of 1. An example of a transformed image using this distribution is presented in Fig. 5




Figure 5: A sample image and its histogram after mapping the histogram to the exponential distribution with $\lambda = 1$

Note that similar to the the previous examples, the pixel values have again been inverted.

5 Experiments and results

This section describes the experiments performed to assess the presented histogram remapping techniques. First, the databases employed in the assessment are briefly introduced and then the actual experiments with the corresponding results are presented.

5.1 Experimental databases

Two popular face databases were used to assess the usefulness of the presented histogram remapping techniques: the XM2VTS [12] and the YaleB [13] databases.

The first database contains 2360 facial images (corresponding to 295 subject), which were captured in controlled conditions, i.e., in frontal view, with a uniform background, with controlled illumination, etc. Each of the 295 subjects of the database is accounted for with 8 facial images recorded during four separate recording session and over a period of approximately five months [14]. Clearly, most of the variability in the appearance of the facial images of



Figure 6: Sample images from the XM2VTS database

a given subject is session induced. Thus, the images differ in the pose of the face, the head-rotation, presence/absence of glasses, mustaches and makeup, different hairstyle, etc. Some examples of the facial images of a subject from the XM2VTS database are shown in Fig. 6.

The second database used in our experiments, i.e., the YaleB database [15], contains images of only 10 subjects. These images, however, exhibit large variations in pose and illumination. The database comprises a total of 5760 grey-scale facial images in 576 viewing conditions (9 poses \times 64 illumination conditions). However, as we are interested only in the impact of the illumination variations on the verification performance, we use only the subset of the facial images from the YaleB database with frontal pose, i.e., a subset of 640 facial images. Some examples of the images from the employed frontal-pose-subset are shown in Fig. 7.

The presented databases were selected for testing purposes to allow us to assess the feasibility of the histogram remapping techniques on facial images captured in controlled as well as uncontrolled conditions. Note that an efficient normalization technique is expected to ensure similar (or better) error rates as those achieved with unprocessed facial images, when identity claims have to be verified based on facial images captured in controlled conditions, and to improve the error rates, when identity claims have to be verified based on facial images captured under severe illumination variations.

In the experiments presented in the next section linear discriminant analysis (LDA) was used as the feature extraction technique and the cosine similarity measure was employed as the scoring function for the nearest neighbor classifier [16].

5.2 Experiments on the XM2VTS database

All experiments on the XM2VTS database were performed in accordance with the first configuration of the established experimental protocol associated with the database [14]. The protocol defines which images should be used for training, which for parameter tuning and which for testing. The protocol represents a closed-set verification protocol¹ and defines that the results of a given series of verification experiments are reported by means of the false rejection error rate (frr), which measures the percentage of falsely rejected clients, the false acceptance error rate (far), which measures the percentage of falsely accepted impostors, and the total error rate (ter), defined as $ter=far+frr$.

In a face recognition system operating in verification mode the claim of the identity of a person presented to the system results in the test/query face image being transformed into the so-called "live" feature vector, and matched against the template associated with the claimed identity. If the live feature vector and the template display a degree of similarity that is higher than a predefined decision threshold, the claim of identity is verified, otherwise the claim of identity is rejected [17]. Unfortunately, the two error rates FAR and FRR both depend on the value of the decision threshold. Thus, selecting a decision threshold which ensures a low FAR, result in an increase of the FRR and vice versa, a decision threshold that ensures a low FRR results in an increase of the FAR. Obviously, a decision threshold which ensures predefined values of the FAR and FRR has to be chosen to allow for a fair comparison of different face veri-

¹In a closed-set protocol the images used for training ones feature extraction technique and the images used for enrolling a user, i.e., building a user model, coincide.



Figure 7: Sample images from the YaleB database

fication systems. This threshold defines an operating point at which different face verification systems can be compared. The equal error rate (EER) operating point, i.e., the threshold which ensures that the FAR equals FRR, is a popular operating point commonly used in comparative assessments and will, therefore, also be employed in our experiments. Of course, an evaluation image set is needed to determine the value of the decision threshold which ensures equal error rates. This threshold is then simply used on the test images set to simulate real operating conditions.

Prior to the experiments, the facial images were first aligned and then cropped to a standard size of 100×100 pixels in accordance with the manually marked eye coordinates. This way we ensured that our results are comparable with other results presented in the literature. Some examples of the preprocessed images are shown in Fig. 8. Note that assessing the performance of the preprocessing techniques with automatically localized facial regions is beyond the scope of this paper. The reader should, however, refer to [18], [19], [20], [21], [22] and [23] to get an overview of existing face detection/localization techniques.



Figure 8: Examples of preprocessed images from the XM2VTS database

After the alignment and cropping procedures, five sets of images were produced, which formed the basis for our experiments. The first set contained images which have not been further processed (denoted as UP in Tables 1, 2, and 3), the second contained images processed with histogram equalization (HQ), the third set comprised images with a fitted normal distribution (NM), the fourth set featured images with a lognormal distribution (LN) and the fifth set contained images with an exponential distribution (EX). The results of the verification experiments conducted on all five image sets for the evaluation set are presented in Table 1 and for the test set in Table 2. The total error rates for the experiments on the two images sets are also presented in Fig. 9.

Table 1: The error rates at the equal error operating point for different preprocessing techniques - evaluation set

<i>Method</i>	<i>UP</i>	<i>HQ</i>	<i>NM</i>	<i>LN</i>	<i>EX</i>
FRR(%)	3.50	3.00	2.83	3.00	4.33
FAR(%)	3.55	2.97	2.84	3.02	4.08
TER(%)	7.05	5.97	5.67	6.02	8.41

Table 2: The error rates at the equal error operating point for different preprocessing techniques - test set

<i>Method</i>	<i>UP</i>	<i>HQ</i>	<i>NM</i>	<i>LN</i>	<i>EX</i>
FRR(%)	3.25	2.50	2.50	2.50	3.25
FAR(%)	3.73	3.39	2.96	3.19	4.18
TER(%)	6.98	5.89	5.44	5.69	7.43

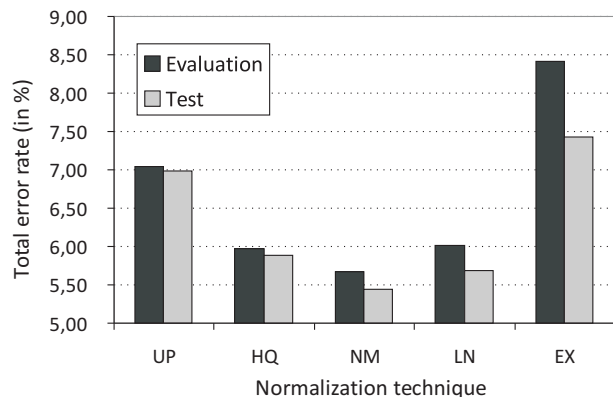


Figure 9: The values of the TER for the evaluation and test images sets of the XM2VTS database

From the results we can see that histogram equalization, lognormal distribution and normal distribution fitting provided similar verification errors, all outperforming the unprocessed images². The only image set resulting in higher error rates than the original images was the set with the exponential distribution. Based on these first results we can conclude that other distribution besides the uniform also ensure lower error rates (when fitted to facial images) than those achievable with unprocessed images. It is also interesting to see that three out of the four tested preprocessing techniques actually improved the verification performance when compared to the unprocessed images. This finding suggest that the assessed techniques are useful not only for image normalization (as will be shown in the next series of our experiments), but for image enhancement as well. This finding is supported by the fact that the images from the XM2VTS database have already been free of illumination induced appearance changes.

5.3 Experiments on the YaleB database

As already indicated in Section 5.1, we used the 640 images with frontal pose for our experiments on the YaleB database [15]. Hence, for every subject in the database, we made use of 64 images - each captured in different illumination conditions. The 640 images were distributed into five subsets depending on the extremity in illumination and were preprocessed in a similar fashion as the images from the XM2VTS database. Some examples of the preprocessed facial images from the five subsets are presented in Fig. 10.

In our experiments the first subset (with the most

²Note that the results for the lognormal distribution correspond to $\sigma = 0.2$, which performed better on the XM2VTS database than the variant with $\sigma = 0.7$.



Figure 10: Examples of preprocessed images from the five subsets of the YaleB database (from top- to bottom-row): subset 1, subset 2, subset 3, subset 4 and subset 5

controlled-like conditions) was used for training and enrollment, while the remaining subsets were employed for testing. Such an experimental setup resulted in highly miss-matched conditions for the verification procedure and posed a great challenge to the preprocessing techniques. On the other hand, it ensured real-life conditions, as the enrollment process is usually supervised and, hence, the training/enrollment images are always of good quality.

The five techniques, i.e., UP, HQ, NM, LN and EX, already used in the experiments on the XM2VTS database were employed to produce five image sets. Based on these sets again a series of identification experiments was performed. The results of these experiments in terms of the rank one recognition rate are presented in Table 3 and Fig. 11. Note that the rank one recognition rate is defined as the percentage of tested images that were correctly identified.

Histogram equalization and normal distribution fitting again resulted in a similar performance with significantly higher recognition rates than those achieved with the unprocessed images. However, even better results were achieved when the histogram of the images was mapped to the lognormal (here a value of

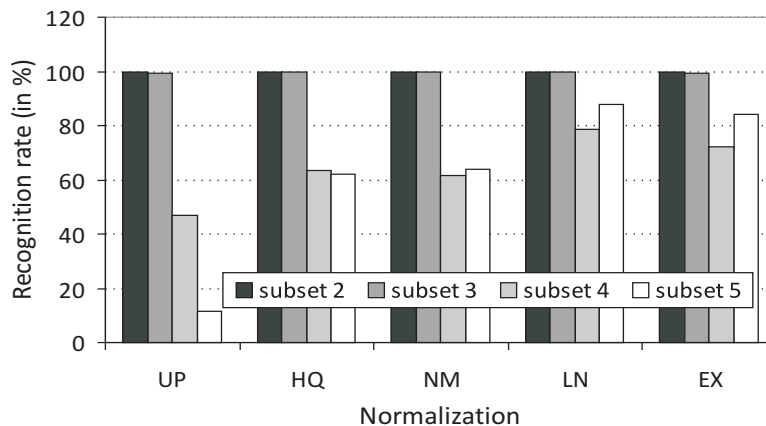


Figure 11: Rank one recognition rates for the five subsets of the YaleB database

Table 3: The rank one recognition rates for the Yale database

<i>Method</i>	<i>UP</i>	<i>HQ</i>	<i>NM</i>	<i>LN</i>	<i>EX</i>
subset 2	100	100	100	100	100
subset 3	99.2	100	100	100	99.2
subset 4	47.1	63.6	61.4	78.6	72.1
subset 5	11.6	62.1	63.7	87.9	84.2

$\sigma = 0.7$ was used) and exponential distributions.

Based on the results obtained on both databases we can conclude that mapping the histograms of facial images to lognormal distributions ensures the most consistent recognition performance. It achieves similar results as histogram equalization when facial images captured in controlled conditions are subjected to the recognition process and outperforms histogram equalization when images with illumination induced appearance variations are used for the recognition. However, even though distribution fitting may result in better recognition performance than histogram equalization, there is also a disadvantage to it, as each distribution relies on one or more parameters which have to be determined in advance and can have a great impact on the face recognition performance.

6 Conclusion

We have presented an empirical assessment of the impact histogram remapping has on the performance of a face recognition system. In a comparative evaluation where in the first step four target distributions were mapped to the facial images and in the second step the mapped images were employed for assess-

ing the performance of a LDA-based face recognition system, the best and most consistent results were achieved when a lognormal distribution was used as the target distribution for the histogram remapping. Our results suggest that other distributions besides the uniform can also ensure enhanced and robust recognition performance, in some cases even surpassing histogram equalization, but at the expense that one or two parameters determining the shape of the predefined distribution have to be set in advance or through some additional training procedure. Our future work in respect to histogram remapping will, therefore, be focused on:

- assessing the performance of face recognition systems with other target distributions (which have not been considered in this paper),
- finding a way of automatically determine the optimal parameters for the given target distribution, and
- combining histogram remapping techniques with other preprocessing approaches such as the multiscale retinex technique or anisotropic smoothing.

Acknowledgements: The research was partially supported by the national research program P2-0250(C) Metrology and Biometric Systems, the bilateral project with the People's Republic of China BICN/07-09-019, the bilateral project with the Bulgarian Academy of Sciences - Face and Signature Biometrics, the national project AvID M2-0210, the COST Action 2101 Biometrics for Identity Documents and Smart Cards and the EU-FP7 project 217762 Homeland security, biometric Identification and personal Detection Ethics (HIDE).

References:

- [1] Y. Adini, Y. Moses and S. Ullman, Face Recognition: The Problem of Compensating for Illumination Changes, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 19, 1997, pp. 721–732.
- [2] E. Land and J. McCann, Lightness and Retinex Theory, *Journal of the Optical Society of America* 61, 1971, pp. 1–11.
- [3] D. Jabson, Z. Rahmann and G. Woodell, A Multiscale retinex for Bridging the Gap Between Color Images and the human Observations of Scenes, *IEEE Transactions on Image Processing* 6, 1997, pp. 897–1056.
- [4] J. Short, J. Kittler and K. Messer, A Comparison of Photometric Normalisation Algorithms for Face Verification, *Proceedings of the International Conference on Automatic Face and Gesture Recognition, AFGR'04*, 2004, pp. 254–259.
- [5] R. Gross and V. Brajovic, An Image Preprocessing Algorithm for Illumination Invariant Face Recognition, *Proceedings of the International Conference on Audio- and Video-Based Biometric Person Authentication, AVBPA'03*, 2003, pp. 10–18.
- [6] S. Du and R. Ward, Wavelet-based Illumination Normalization for Face Recognition, *Proceedings of the IEEE International Conference on Image Processing, ICIP'05*, 2005, pp. 954–957.
- [7] A.K. Jain, A. Ross and S. Prabhakar, An introduction to biometric recognition, *IEEE Transactions on Circuits and Systems for Video Technology* 14, 2004, pp. 4–20.
- [8] V. Štruc and N. Pavešić, Phase Congruency Features for Palm-print Verification, *IET Signal Processing*, to be published.
- [9] T. Savič and N. Pavešić, Personal Recognition Based on an Image of the Palmar Surface of the Hand, *Pattern Recognition* 40, 2007, pp. 3152–3163.
- [10] J. Pelecanos and S. Sridharam, Feature Warping for Robust Speaker Verification, *Proceedings of the Speaker Recognition Workshop*, 2001, pp. 213–218.
- [11] V. Štruc and N. Pavešić, A comparison of feature normalization techniques for PCA-based palm-print recognition. *Proceedings of the Conference on Mathematical Modelling, MATHMOD'09*, 2009, pp. 2450–2453.
- [12] Homepage of the XM2VTS face database, <http://www.ee.surrey.ac.uk/CVSSP/xm2vtsdb/>, accessed March 2009.
- [13] Homepage of the YaleB face database, <http://cvc.yale.edu/projects/yalefacesB/yalefacesB.html>, accessed March 2009.
- [14] K. Messer, J. Matas, J. Kittler and J. Luettin, XM2VTSDB: the Extended M2VTS Database, *Proceedings of the International Conference on Audio- and Video-Based Biometric Person Authentication, AVBPA'99*, 1999, pp. 72–77.
- [15] A.S. Georhiades, P.N. Belhumeur and D.J. Kriegman, From Few to Many: Illumination Cone Models for Face Recognition Under Variable Lighting and Pose, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23, 2001, pp. 643–660.
- [16] P.N. Belhumeur, J.P. Hespanha and D.J. Kriegman, Eigenfaces vs. Fisherfaces: Recognition Using Class Specific Linear Projection, *Proceedings of the 4th European Conference on Computer Vision, ECCV'96*, 1996, pp. 45–58.
- [17] V. Štruc, F. Mihelič and N. Pavešić, Face authentication using a hybrid approach. *Journal of Electronic Imaging* 17, 2008, pp. 1–11.
- [18] K. Sundaraj, Real-Time Face Detection using Dynamic Background Substraction, *WSEAS Transactions on Information Science and Applications* 5, 2008, pp. 1531–1540.
- [19] Y. Gizatdinova, J. Erola and V. Surakka, Automatic Real Time Localization of Frowning and Smiling Faces under Controlled Head Rotations, *WSEAS Transactions on Signal Processing* 4, 2008, pp. 463–473.
- [20] A. Bottino and S. Cumani, A Fast and Robust Method for the Identification of Face Landmarks in Profile Images, *WSEAS Transactions on Computers* 7, 2008, pp. 1250–1259.
- [21] O. Jesorsky, K.J. Kirchberg and R.W. Frischholz, Robust face detection using the Hausdorff distance, *Proceedings of the International Conference on Audio- and Video-based Biometric Person Authentication, AVBPA'01*, 2001, pp. 90–95.
- [22] P. Viola and M. Jones, Robust Real-Time Object Detection, *Proceedings of the Second International Workshop on statistical and Computational Theories of Vision - Modelling, Learning, Computing and Sampling*, 2001, pp. 1–25.
- [23] E. Hjelmås and B.K. Low, Face Detection: A Survey, *Computer Vision and Image Understanding* 83, 2001, pp. 236–274.