

Research Article

The Complete Gabor-Fisher Classifier for Robust Face Recognition

Vitomir Štruc and Nikola Pavešić

*Laboratory of Artificial Perception, Systems and Cybernetics, Faculty of Electrical Engineering,
University of Ljubljana, SI-1000 Ljubljana, Slovenia*

Correspondence should be addressed to Vitomir Štruc, vitos@luks.fe.uni-lj.si

Received 2 December 2009; Revised 15 April 2010; Accepted 20 April 2010

Academic Editor: Robert W. Ives

Copyright © 2010 V. Štruc and N. Pavešić. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

This paper develops a novel face recognition technique called Complete Gabor Fisher Classifier (CGFC). Different from existing techniques that use Gabor filters for deriving the Gabor face representation, the proposed approach does not rely solely on Gabor magnitude information but effectively uses features computed based on Gabor phase information as well. It represents one of the few successful attempts found in the literature of combining Gabor magnitude and phase information for robust face recognition. The novelty of the proposed CGFC technique comes from (1) the introduction of a Gabor phase-based face representation and (2) the combination of the recognition technique using the proposed representation with classical Gabor magnitude-based methods into a unified framework. The proposed face recognition framework is assessed in a series of face verification and identification experiments performed on the XM2VTS, Extended YaleB, FERET, and AR databases. The results of the assessment suggest that the proposed technique clearly outperforms state-of-the-art face recognition techniques from the literature and that its performance is almost unaffected by the presence of partial occlusions of the facial area, changes in facial expression, or severe illumination changes.

1. Introduction

Biometrics is a scientific discipline that uses unique and measurable physical, biological, or/and behavioral human characteristics that can be processed to establish identity, to perform identity verification, or to recognize a person through automation [1–3]. Among the different characteristics suitable for biometric recognition, the human face and the associated face recognition technology bear the most potential. This potential is fueled by the countless application possibilities of face recognition technology in the private as well as the public sector. Examples of potential application domains range from entertainment, human-machine interaction, homeland security, smart surveillance, access and border control to user authentication schemes in e-commerce, e-health, and e-government services [1, 2, 4].

While, for example, access control applications can often ensure stable and controlled external conditions for the

image acquisition procedure, the majority of applications (especially those linked to unconstrained face recognition, e.g., surveillance) cannot. In such cases, image characteristics, such as changes in illumination, partial occlusions of the facial area, or different facial expressions, heavily influence the appearance of the face in the acquired image and render much of the existing face recognition technology useless [5]. To prove useful in unconstrained environments, the deployed face recognition system has to utilize recognition techniques capable of providing reliable recognition results regardless of the (variable) characteristics of the acquired images.

Many researchers have tackled the problem of robust face recognition in uncontrolled (out-door) environments by trying to develop face recognition techniques insensitive to image degradations caused by various external factors. Sanderson and Paliwal, for example, proposed a feature extraction technique called DCT-mod2. The DCT-mod2

technique first applies the discrete cosine transform (DCT) to subregions (or image blocks) of facial images to extract a number of DCT coefficients. Next, it compensates for any potential illumination changes, by replacing the coefficients most affected by illumination with their corresponding horizontal and vertical delta coefficients. By doing so, the technique derives a face representation (partially) insensitive to external lighting changes. The authors assessed their approach on various databases comprised of images with illumination-induced variability. On all databases the DCT-mod2 technique resulted in promising results [6].

Gao and Leung [7] went a different way and proposed a face representation called Line Edge Map (LEM). Here, a given face image is first processed with the Sobel operator to extract edge pixels, which are then combined into line segments that constitute the LEMs. While the authors suggest that the LEMs ensure illumination and expression invariant face recognition, the developed face representation, nevertheless, inherits the shortcomings of the gradient-based Sobel operator, which is known to struggle with its performance under severe lighting variations.

Fidler et al. [8] tried to achieve robust face recognition by exploiting an elaborate subsampling procedure. The subsampling procedure first detects image pixels representing statistical outliers in each of the facial images and then derives a low-dimensional representation of each facial image by considering only valid image pixels (i.e., based on statistical “inliers”). The authors show that their procedure ensures (partial) robustness to facial occlusions as well as to extreme facial expression changes.

More recently, Wright et al. [9] introduced a novel method for robust face recognition exploiting recent advances from the field of compressed sensing. Their method, called the Sparse Representation Classifier (SRC), derives a sparse representation from the given face image and simultaneously assumes that the image is contaminated with a spatially sparse error. Under the assumption of the sparse error, the authors are able to construct robust classifiers capable of performing well under a variety of image degradations caused, for example, by illumination changes, noise, or facial occlusions.

One of the most popular solutions to the problem of robust face recognition was, however, proposed by Liu and Wechsler in [10]. Here, the authors proposed to adopt a filter bank of forty Gabor filters to derive an augmented feature vector of Gabor magnitude features, and then apply a variant of the multiclass linear discriminant analysis to the constructed Gabor feature vector to improve the vector’s compactness. The efficiency of the technique, named the Gabor-Fisher Classifier (GFC), was determined on a large and challenging database and is, furthermore, evidenced by the large number of papers following up on the work in [10], for example, [11–14].

It should be noted that the Gabor face representation (as proposed in [10]) exhibits (partial) robustness to changing facial expressions as well as illumination variations. The former is a consequence of the local nature of the Gabor feature vector, while the latter is linked to the properties of the Gabor filter bank (as Gabor filters represent band-limited

filters, the filter bank can be constructed in such a way that it excludes the frequency bands most affected by lighting variations, resulting in robustness to lighting changes).

While the existing Gabor-based methods are among the most successful face recognition techniques, one could still voice some misgivings, as they rely only on the Gabor magnitude information and discard the potentially useful Gabor phase information. In this paper, we tackle this issue and propose a novel face representation called oriented Gabor phase congruency image, which, as the name suggests, is derived from the Gabor phase congruency model [15]. The proposed face representation is based on the phase responses of the Gabor filter bank rather than the Gabor magnitude responses and as such offers an alternative (or complement) to the established Gabor magnitude methods. We show that the face representation derived from the oriented Gabor phase congruency images is more compact than the commonly used Gabor magnitude representation of face images and that it also exhibits an inherent robustness to illumination changes.

The novel representation is combined with the multiclass linear discriminant analysis to obtain the so-called phase-based Gabor-Fisher classifier (PBGFC). The developed PBGFC method is ultimately fused with the GFC technique to result in the complete Gabor-Fisher classifier (CGFC), which effectively uses Gabor magnitude as well as Gabor phase information for robust face recognition. The feasibility of the proposed techniques is assessed in a series of face recognition experiments performed on the popular XM2VTS, FERET, AR, and Extended YaleB databases. The results of the assessments show that the proposed technique compare favorably with face recognition methods from the literature in terms of robustness as well as face recognition performance.

The rest of the paper is structured as follows. In Section 2 a brief review of Gabor filters and Gabor filter base face recognition techniques is given. In Section 3, the novel face representation in form of oriented Gabor phase congruency images is introduced. Sections 4 and 5 develop the phase-based and complete Gabor-Fisher classifiers for face recognition. Section 6 presents the classification rules, while Section 7 describes the employed experimental databases. The feasibility of the proposed technique is assessed in Section 8. The paper concludes with some final remarks in Section 9.

2. Review of Gabor Filters for Face Recognition

This section briefly reviews the use of Gabor filters for face recognition. It commences with the introduction of Gabor filters and the basic concepts of feature extraction using the Gabor filter bank and proceeds with the presentation of the Gabor (magnitude) face representation, which forms the foundation for many popular face recognition techniques, including the prominent Gabor-Fisher Classifier [10].

2.1. Gabor Filter Construction. Gabor filters (also called Gabor wavelets or kernels) have proven themselves to be




FIGURE 1: An example of the real (a) and imaginary (b) part of a Gabor filter.

a powerful tool for facial feature extraction and robust face recognition. They represent complex band-limited filters with an optimal localization in both the spatial as well as the frequency domain. Thus, when employed for facial feature extraction, they extract multiresolutional, spatially local features of a confined frequency band [6]. Like all filters operating in the scale-space, Gabor filters also relate to the simple cells of the mammalian visual cortex and are, hence, relevant from the biological point of view as well.

In general, the family of 2D Gabor filters can be defined in the spatial domain as follows [10, 11, 13, 14, 16–20]:

$$\psi_{u,v}(x, y) = \frac{f_u^2}{\pi\kappa\eta} e^{-((f_u^2/\kappa^2)x'^2 + (f_u^2/\eta^2)y'^2)} e^{j2\pi f_u x'}, \quad (1)$$

where $x' = x \cos \theta_v + y \sin \theta_v$, $y' = -x \sin \theta_v + y \cos \theta_v$, $f_u = f_{\max}/2^{(u/2)}$, and $\theta_v = v\pi/8$. As can be seen from the filters definition, each Gabor filter represents a Gaussian kernel function modulated by a complex plane wave whose center frequency and orientation are given by f_u and θ_v , respectively. The parameters κ and η determine the ratio between the center frequency and the size of the Gaussian envelope and, when set to a fixed value, ensure that Gabor filters of different scales behave as scaled versions of each other [6]. It should also be noted that with fixed values of the parameters κ and η , the scale of the given Gabor filter is uniquely defined by the value of its center frequency f_u .

While different choices of the parameters determining the shape and characteristics of the filters define different families of Gabor filters, the most common parameters used for face recognition are $\kappa = \eta = \sqrt{2}$ and $f_{\max} = 0.25$ [6, 10, 11, 13, 19, 20]. When using the Gabor filters for facial feature extraction, researchers typically construct a filter bank featuring filters of five scales and eight orientations, that is, $u = 0, 1, \dots, p-1$ and $v = 0, 1, \dots, r-1$, where $p = 5$ and $r = 8$. An example of the real and imaginary parts of a Gabor filter is presented in Figure 1, while the real parts of the entire filter bank commonly used for facial feature extraction (comprised of 40 filters) are shown in Figure 2.

2.2. Feature Extraction Using Gabor Filters. Let $I(x, y)$ stand for a grey-scale face image of size $a \times b$ pixels and, moreover, let $\psi_{u,v}(x, y)$ denote a Gabor filter given by its center frequency f_u and orientation θ_v . The feature extraction procedure can then be defined as a filtering operation of the




FIGURE 2: The real parts of the Gabor filter bank commonly used for feature extraction in the field of face recognition.

given face image $I(x, y)$ with the Gabor filter $\psi_{u,v}(x, y)$ of size u and orientation v [10, 11, 13, 14, 17, 19, 20], that is

$$G_{u,v}(x, y) = I(x, y) * \psi_{u,v}(x, y), \quad (2)$$

where $G_{u,v}(x, y)$ denotes the complex filtering output that can be decomposed into its real ($E_{u,v}(x, y)$) and imaginary ($O_{u,v}(x, y)$) parts:

$$\begin{aligned} E_{u,v}(x, y) &= \text{Re}[G_{u,v}(x, y)], \\ O_{u,v}(x, y) &= \text{Im}[G_{u,v}(x, y)]. \end{aligned} \quad (3)$$

Based on these results, the magnitude ($A_{u,v}(x, y)$) and phase ($\phi_{u,v}(x, y)$) responses of the filtering operation can be computed as follows:

$$\begin{aligned} A_{u,v}(x, y) &= \sqrt{E_{u,v}^2(x, y) + O_{u,v}^2(x, y)}, \\ \phi_{u,v}(x, y) &= \arctan\left(\frac{O_{u,v}(x, y)}{E_{u,v}(x, y)}\right). \end{aligned} \quad (4)$$

The majority of Gabor-based face recognition techniques found in the literature discard the phase information of the filtering output and rely solely on the magnitude information when constructing the Gabor face representation. By doing so, they discard potentially valuable discriminative information that could prove useful for the recognition task. The magnitude responses usually retained by the Gabor-based




FIGURE 3: An example of the Gabor magnitude output: a sample image (a) and the magnitude output of the filtering operation with the entire Gabor filter bank of 40 Gabor filters (b).

recognition techniques are presented in Figure 3 for a sample face image.

2.3. The Gabor (Magnitude) Face Representation. When deriving the Gabor (magnitude) face representation from a given facial image, the first step is the construction of the Gabor filter bank. As we have pointed out already, most of the existing techniques in the literature adopt a filter bank comprising Gabor filters of five scales ($u = 0, 1, \dots, 4$) and eight orientations ($v = 0, 1, \dots, 7$).

Next, the given face image is filtered with all 40 filters from the filter bank resulting in an inflation of data dimensionality to 40 times its initial size. Even for a small face image of, for example, 128×128 pixels, the 40 magnitude responses reside in a 655360 ($128 \times 128 \times 40$) dimensional space, which is far too extensive for efficient processing and storage. Thus, to overcome this dimensionality issue, downsampling strategies are normally exploited. The downsampling techniques reduce the dimensionality of the Gabor magnitude responses, unfortunately often at the expense of valuable discriminatory information. One of the most popular downsampling strategies relies on a rectangular sampling grid (as shown in Figure 4) superimposed over the image to be sampled. In the downsampled image only the values located under the sampling grid's nodes are retained, while the rest is discarded. The downsampling procedure is applied to all magnitude responses, which are ultimately normalized using a properly selected normalization procedure and then concatenated into the final Gabor (magnitude) face representation or, as named by Liu and Wechsler [10], into the augmented Gabor feature vector. (Note that typically zero-mean and unit variance normalization is applied at this step. However, as other normalization techniques might be superior to the zero-mean and unit-variance scheme, the issue of selecting the most appropriate normalization procedure will empirically be investigated in the experimental section.)

If we denote the downsampled Gabor magnitude responses in vector form at the u th filter scale and v th orientation by $\mathbf{g}_{u,v}$, then the augmented Gabor (magnitude)

feature vector \mathbf{x} can be defined as follows [10, 11, 13, 14, 19, 20]:

$$\mathbf{x} = (\mathbf{g}_{0,0}^T, \mathbf{g}_{0,1}^T, \mathbf{g}_{0,2}^T, \dots, \mathbf{g}_{4,7}^T)^T. \quad (5)$$

It should be noted that in the experimental section, we use images of size 128×128 pixels and a rectangular sampling grid with 16 horizontal and 16 vertical lines, which corresponds to a downsampling factor of $\rho = 64$. The resulting feature vector, or, in other words, the resulting Gabor (magnitude) face representation forms the foundation for the Gabor-Fisher classifier, which will be presented in Section 4 in more detail.

3. The Oriented Gabor Phase Congruency Face Representation

Up until now we have been concerned with Gabor magnitude responses and face representations derived from them. In this section, however, we will focus on face representations derived from Gabor phase responses and their usefulness for face recognition. The section commences by reviewing the existing attempts at incorporating the Gabor phase information into the face recognition procedure and, thus, the attempts at further improving the recognition performance of Gabor-based recognition techniques. Next, it presents a novel representation of face images called the oriented Gabor phase congruency image, and, finally, it develops the oriented Gabor phase congruency face representation, which forms the basis for the phase-based Gabor-Fisher classifier presented in Section 4.

3.1. Background. Before we turn our attention to the novel representation of face images, let us take a closer look at why the Gabor phase information is commonly discarded when using Gabor filters for face recognition.

Unlike the (Gabor) magnitude, which is known to vary slowly with the spatial position, the (Gabor) phase takes very different values even if it sampled at image locations only a few pixels apart [6, 21, 22]. This inherent instability




FIGURE 4: Down-sampling of a magnitude filter response with a rectangular sampling grid (from left to right): (a) an example of a magnitude response, (b) an example of the magnitude response with a superimposed rectangular sampling grid, and (c) a downsampled magnitude response.

of the Gabor phase makes it difficult to extract stable and discriminative features from the phase responses of (4) and is also the primary reason that most of the existing methods rely solely on the Gabor magnitude when constructing the Gabor feature vector.

To the best of our knowledge, there are only a few studies in the literature that successfully derive useful features from Gabor phase responses for the task of face recognition, that is, [6, 21–28]. A common characteristic of these methods is the fact that they do not rely on face representations constructed directly from the Gabor phase responses; rather they use features derived from the “raw” Gabor phase responses or combine the phase information with other descriptors to compensate for the instability of the Gabor phase.

Zhang et al. [21, 22], for example, adopt local histograms of the phase responses encoded via the local binary patterns (LBPs) [29, 30] as facial descriptors and consequently show that, over small image areas, the Gabor phase patterns exhibit regularity that can be exploited for face recognition [6].

A similar procedure is introduced in [27] by Guo and Xu and later extended by Guo et al. in [26]. Different from the procedure of Zhang et al. presented in [21], Guo et al. rely on Gabor phase differences instead of the “raw” phase values to compute the binary patterns. In the second step, the computed patterns corresponding to the phase response of the Gabor filter of a given scale and orientation are grouped to form local (subregion-based) histograms and ultimately concatenated into extended histograms, which encode local as well as global aspects of the given phase response.

Other authors (e.g., [23–25]) incorporate the Gabor phase information by employing the so-called phase congruency model (developed by Kovess [15]) for edge detection in the facial image and then deploy the “edge” image for detection of interest points that are used with other image descriptors, such as Gabor magnitude features.

The face representation developed in this paper differs greatly from the existing Gabor phase-based approaches presented above. It is related to work presented in [23–25] only as far as it uses the concept of phase congruency for encoding the Gabor phase information. However, unlike previous work on this subject, it derives a face representation based on a modified model of phase congruency and employs

the resulting representation directly for recognition rather than solely for feature selection.

The difference to the existing face recognition methods using Gabor phase information is even more pronounced if we consider only techniques adopting histograms of Gabor phase (or phase difference) patterns. (Note that in the remainder of the paper we will use the term *Gabor phase patterns* for all Gabor phase-based patterns whether they were computed from actual phase responses or phase-differences due to the similarity of the descriptors.) These methods alleviate the problem of phase instability by observing local histograms of the phase patterns, or in other words, by adopting histograms of the Gabor phase patterns computed from image subblocks as the basic image descriptors. In a sense, they (e.g., [21, 22, 26, 27]) assume that despite the irregular changes of the Gabor phase from one spatial position to a neighboring one, the distribution of the Gabor phase values over a small spatial area is consistent and, thus, useful for recognition. Furthermore, to reduce the variability of the Gabor phase responses prior to histogram construction (it should be noted that the computed histograms serve as non-parametric estimates of the true Gabor phase distribution), they encode the phase using different versions of binary patterns, for example, [29]. Unlike the presented Gabor-phase-pattern-based methods, which exploit regularities of the Gabor phase in the spatial domain, the face representation presented in this paper relies on regularities in the scale-space domain (or frequency domain—see Section 2.1). Hence, it exploits a completely different approach to overcome the instability problems related to the Gabor phase.

As will be shown in the remainder of this section, the proposed representation does not offer only an efficient way of overcoming the Gabor phase instability, but also exhibits several desirable properties for the task of face recognition.

3.2. The 2D Phase Congruency Model. The original 2D phase congruency model as proposed by Kovess in [15] was developed with the goal of robust edge and corner detection in digital images. Unlike classical gradient-based edge detectors, which search for image points of maximum intensity

gradients and are known to be susceptible to image contrast and illumination conditions, the phase congruency model searches for points of order in the frequency spectrum, and provides an illumination invariant model of edge detection [15, 24].

For 1D signals, the phase congruency $PC(x)$ is defined implicitly by the relation of the energy at a given point in the signal $E(x)$ and the sum of the Fourier amplitudes A_n as shown by Venkatesh and Owens [31]:

$$E(x) = PC(x) \sum_n A_n, \quad (6)$$

where n denotes the number of Fourier components. Thus, phase congruency at a given location of the signal x is defined as the ratio of the local energy at this location and the sum of Fourier amplitudes.

Kovesi extended the above concept to 2D signals by computing the phase congruency with logarithmic Gabor filters using the following expression:

$$PC_{2D}(x, y) = \frac{\sum_{v=0}^{r-1} \sum_{u=0}^{p-1} A_{u,v}(x, y) \Delta\Phi_{u,v}(x, y)}{\sum_{v=0}^{r-1} \sum_{u=0}^{p-1} A_{u,v}(x, y) + \varepsilon}, \quad (7)$$

where $A_{u,v}(x, y)$ denotes the magnitude response of the logarithmic Gabor filter at scale u and orientation v , ε represents a small constant that prevents divisions with zero, and $\Delta\Phi_{u,v}(x, y)$ stands for a phase deviation measure defined as

$$\Delta\Phi_{u,v}(x, y) = \cos(\phi_{u,v}(x, y) - \bar{\phi}_v(x, y)) - \left| \sin(\phi_{u,v}(x, y) - \bar{\phi}_v(x, y)) \right|. \quad (8)$$

Here $\phi_{u,v}(x, y)$ denotes the phase angle of the logarithmic Gabor filters at the u th scale and v th orientation, while $\bar{\phi}_v(z)$ represents the mean phase angle at the v th orientation.

Clearly, the expression $\sum_{v=0}^{r-1} \sum_{u=0}^{p-1} A_{u,v}(x, y) \Delta\Phi_{u,v}(x, y)$ approximates the local energy at the spatial location (x, y) , while the denominator of (7) represents the sum of the (logarithmic) Gabor amplitudes over all orientations and scales. Obviously, the phase congruency as defined by (7) represents a quantity that is independent of the overall magnitude of its underlying signal and is, hence, invariant to changes in contrast and illumination [6, 15, 24, 25]. The model detects points in an image where the logarithmic Gabor filter responses are maximally in phase, or in other words, scans the logarithmic Gabor phase responses for regularities in the scale-space.

At closer examination of the 2D phase congruency model we can notice that it first computes the phase congruency for each of the employed filter orientations and subsequently combines the results to form the final output. Some examples of a facial image subjected to logarithmic Gabor filter banks with different numbers of scales p and orientations r are shown in Figure 5. We can see that both parameters effect the appearance of the resulting phase congruency image (PCI).

3.3. The Oriented Gabor Phase Congruency Model. While the 2D phase congruency model given by (7) is suitable

for robust edge and corner detection, its usefulness for face recognition is at least questionable. First of all, the edges detected by the model are highly localized, suggesting that even small variation in facial expression or misalignment would drastically change the appearance of the PCI of a given subject; and, second of all, the phase congruency representation does not make use of multiorientational information, which can provide important clues for the recognition task.

To overcome the presented shortcomings, we propose in this paper a novel face representation, called the oriented Gabor phase congruency image (OGPCI). Rather than combining phase congruency information computed over several orientations, and using the result for construction of the facial feature vector, we compute the oriented form of phase congruency for each of the employed filter orientations and construct an augmented Gabor phase congruency feature vector based on the results [6]. Note that differently from the original model of Kovesi [15], we deploy conventional Gabor filter as given by (1) rather than logarithmic Gabor filters.

Taking into account the original definition of phase congruency, we derive an oriented form of phase congruency, which, when presented in image form, reveals the OGPCI for the v th orientation:

$$OGPCI(x, y) = \frac{\sum_{u=0}^{p-1} A_{u,v}(x, y) \Delta\Phi_{u,v}(x, y)}{\sum_{u=0}^{p-1} A_{u,v}(x, y) + \varepsilon}. \quad (9)$$

Some examples of the OGPCIs for different number of employed filter scales p and a fixed orientation of $\theta_v = 0^\circ$ are shown in Figure 6. We can see that the choice of the number of filter scales p influences the appearance of the OGPCIs and for optimal face recognition performance should be set based on preliminary results on some development data (see Section 8.4 for more details).

3.4. The Oriented Gabor Phase Congruency Face Representation. The OGPCIs introduced in the previous section form the foundation for the derivation of the oriented Gabor phase congruency face representation or, in accordance with the notation used by Liu and Wechsler in [10], the augmented Gabor phase congruency vector, which is computed by taking the following steps.

- (i) For a given face image the OGPCIs are computed for all r orientations (an example of all OGPCIs for a sample image with $r = 8$ and $p = 2$ is presented in Figure 7).
- (ii) The computed OGPCIs are downsampled by a down-sampling factor ρ (similar as depicted in Figure 4).
- (iii) The downsampled OGPCIs are normalized using an appropriate normalization procedure.
- (iv) The downsampled and normalized OGPCIs in vector form (denoted by \mathbf{D}_v) are concatenated to form the augmented Gabor phase congruency feature vector \mathbf{x} .

Formally, the augmented Gabor phase congruency feature vector can be defined as follows:

$$\mathbf{x} = (\mathbf{D}_0^T, \mathbf{D}_1^T, \mathbf{D}_2^T, \dots, \mathbf{D}_{r-1}^T)^T, \quad (10)$$




FIGURE 5: Examples of phase congruency images (from left to right): (a) the original face image, (b) the PCI for $p = 3$ and $r = 6$, (c) the PCI for $p = 5$ and $r = 6$, (d) and the PCI for $p = 3$ and $r = 8$, (e) the PCI for $p = 5$ and $r = 8$.




FIGURE 6: Examples of OGPCIs (from left to right): (a) the original face image, (b) the OGPCI for $p = 2$, (c) the OGPCI for $p = 3$, (d) and the OGPCI for $p = 4$, (e) the OGPCI for $p = 5$.

where T denotes the transform operator and \mathbf{D}_v , for $v = 0, 1, \dots, r - 1$, represents the vector form of the OGPCI at the v th orientation.

Note that in the experiments presented in Section 8 the augmented Gabor phase congruency feature vector was constructed using a downsampling factor of $\rho = 16$, as opposed to the augmented Gabor magnitude vector, where a downsampling factor of $\rho = 64$ was employed. This setup led to similar lengths of the two augmented feature vectors allowing for a fair comparison of their usefulness for face recognition [6].

4. The Gabor-Fisher and Phase-Based Gabor-Fisher Classifiers

We have already emphasized that both the augmented Gabor magnitude feature vector the augmented Gabor phase congruency feature vector, despite the downsampling procedure, still reside in a very high-dimensional space. The Gabor-Fisher Classifier presented by Liu and Wechsler [10] and the Phase-based Gabor-Fisher Classifier introduced here overcome this dimensionality issue by subjecting the augmented feature vectors to Fisher's Discriminant Analysis (also known as Linear Discriminant Analysis). The subspace projection reduces the size of the augmented feature vectors and allows for an efficient implementation of the matching procedure.

The employed dimensionality reduction technique (i.e., LDA) derives a transformation matrix (i.e., the projection basis) which is used to project the augmented feature vectors into a subspace where between-class variations of the projected patterns are maximized while within-class variations are minimized [32].

Given a set of n d -dimensional training patterns (i.e., augmented feature vectors) \mathbf{x}_i arranged into a $d \times n$ data matrix $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n]$, each belonging to one of N

classes C_1, C_2, \dots, C_N , one first computes the between-class and the within-class scatter matrices Σ_B and Σ_W :

$$\begin{aligned} \Sigma_B &= \sum_{i=1}^N n_i (\mu_i - \mu) (\mu_i - \mu)^T, \\ \Sigma_W &= \sum_{i=1}^N \sum_{\mathbf{x}_j \in C_i} (\mathbf{x}_j - \mu_i) (\mathbf{x}_j - \mu_i)^T, \end{aligned} \quad (11)$$

and then one derives the LDA transformation matrix \mathbf{W} which maximizes Fisher's discriminant criterion [32, 33]:

$$J(\mathbf{W}) = \arg \max_{\mathbf{W}} \frac{|\mathbf{W}^T \Sigma_B \mathbf{W}|}{|\mathbf{W}^T \Sigma_W \mathbf{W}|}, \quad (12)$$

where n_i denotes the number of samples in the i th class, μ_i stands for the class conditional mean and μ represents the global mean of all training samples [4].

Fisher's discriminant criterion is maximized when \mathbf{W} is constructed by a simple concatenation of the $d' \leq N - 1$ leading eigenvectors of the following eigenproblem:

$$\Sigma_W^{-1} \Sigma_B \mathbf{w}_i = \lambda_i \mathbf{w}_i, \quad i = 1, 2, \dots, d', \quad (13)$$

that is, $\mathbf{W} = [\mathbf{w}_1 \mathbf{w}_2 \dots \mathbf{w}_{d'}]$.

Once the transformation matrix \mathbf{W} is calculated, it can be used to project a test pattern (i.e., an arbitrary augmented feature vector) \mathbf{x} into the LDA subspace, thus reducing the pattern's dimension from d to d' :

$$\mathbf{y} = \mathbf{W}^T (\mathbf{x} - \mu), \quad (14)$$

where \mathbf{y} represents the d' -dimensional projection of the centered pattern \mathbf{x} [4]. To avoid singularity issues, when computing the inverse of the within-class scatter matrix Σ_W , LDA is implemented in the PCA subspace as suggested by Belhumeur et al. in [33].




FIGURE 7: An example of all OGPCIs: the original face image (a), the OGPCIs (for $r = 8$) (b), which in their downsampled and normalized form constitute the augmented Gabor phase congruency vector.

When the presented technique is applied to the augmented Gabor magnitude feature vectors, we obtain the Gabor-Fisher Classifier and, similarly, if the underlying feature vectors take the form of the augmented Gabor phase congruency vectors, we obtain the Phase-based Gabor-Fisher Classifier.

5. The Complete Gabor-Fisher Classifier

The presented Gabor-Fisher and phase-based Gabor-Fisher classifiers (GFC and PBGFC, resp.) operate on different feature types derived from the Gabor filter responses. Since the first relies on Gabor magnitude information and the second encodes Gabor phase information, we combine both classifiers into the Complete Gabor-Fisher Classifier (CGFC), which should exhibit enhanced face recognition performance when compared to either of the classifiers on their own.

The fusion of the classifiers is implemented at the matching score level using the fusion scheme shown in Figure 8. Here, the final matching score of the CGFC method δ_{CGFC} is computed using the following expression [34]:

$$\delta_{CGFC} = (1 - \gamma)\delta_{GFC} + \gamma\delta_{PBGFC}, \quad (15)$$

where δ_{GFC} denotes the matching score obtained with the GFC technique, δ_{PBGFC} denotes the matching score obtained with the PBGFC approach and $\gamma \in [0, 1]$ denotes the fusion parameter that controls the relative importance of the two matching scores. (Note that the matching scores for the individual classifiers are computed based on the procedure described in Section 6.) When set to $\gamma = 0$, the CGFC method turns into the GFC method, when set to $\gamma = 1$, the CGFC technique turns into the PBGFC technique, while for any other value of γ the CGFC technique considers both feature types. It has to be noted that the value of the fusion parameter γ should be optimized for the best possible performance (see Section 8.5).

6. The Classification Rule

In general, a face recognition system can operate in one of two modes, either in verification or in identification mode

[35]. In verification mode, the goal of the system is to determine the validity of the identity claim made by the user currently presented to the system. This is achieved by comparing the so-called “live” feature vector \mathbf{y} extracted from the given face image of the user with the template corresponding to the claimed identity. Based on the outcome of this comparison, the identity claim is either rejected or accepted.

Formally this can be written as follows: given the live feature vector \mathbf{y} and a claimed identity C_i associated with a user-template \mathbf{y}_i , where $i \in 1, 2, \dots, N$ and N represents the number of enrolled users, determine the validity of the identity claim by classifying the pair (\mathbf{y}, C_i) into one of two classes w_1 or w_2 [2, 36]:

$$(\mathbf{y}, C_i) \in \begin{cases} w_1, & \text{if } \delta(\mathbf{y}, \mathbf{y}_i) \geq \Delta, \quad i = 1, 2, \dots, N, \\ w_2, & \text{otherwise,} \end{cases} \quad (16)$$

where w_1 denotes the class of genuine identity claims, w_2 stands for the class of illegitimate identity claims, $\delta(\cdot, \cdot)$ represents a function measuring the similarity of its arguments, which in our case takes the form of the cosine similarity measure, that is,

$$\delta(\mathbf{y}, \mathbf{y}_i) = \frac{\mathbf{y}^T \mathbf{y}_i}{\sqrt{\mathbf{y}^T \mathbf{y} \mathbf{y}_i^T \mathbf{y}_i}}, \quad (17)$$

and Δ stands for a predefined decision threshold.

In a face recognition system operating in the identification mode the problem statement is different from that presented above. In case of the identification task we are not interested whether the similarity of the live feature vector with a specific user-template is high enough; rather, we are looking for the template in the database that best matches the live feature vector. This can be formalized as follows: given a live feature vector \mathbf{y} and a database containing N user-templates $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N$ of the enrolled users (or identities) C_1, C_2, \dots, C_N , determine the most suitable identity [2], that is,

$$\mathbf{y} \in \begin{cases} C_i, & \text{if } \delta(\mathbf{y}, \mathbf{y}_i) = \max_{j=1}^N \delta(\mathbf{y}, \mathbf{y}_j) \geq \Delta, \\ C_{N+1}, & \text{otherwise,} \end{cases} \quad (18)$$

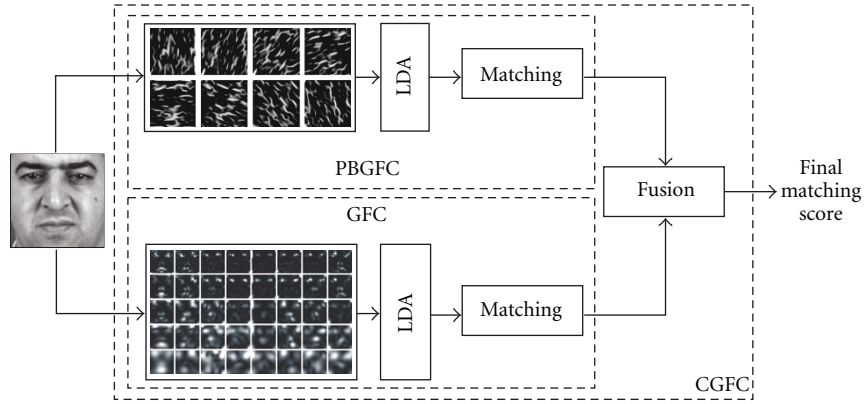


FIGURE 8: Block scheme of the Complete Gabor-Fisher Classifier.

where $\delta(\mathbf{y}, \mathbf{y}_i)$ again denotes the cosine similarity measure and C_{N+1} stands for the case, where no appropriate identity from the database can be assigned to the live feature vector \mathbf{y} . The presented expression postulates that, if the similarity of the live feature vector and the template associated with the i th identity is the highest among the similarities with all user-templates in the system, then the i th identity is assigned to the live feature vector \mathbf{y} .

It should be noted that, in the experiments presented in the remainder of this paper, the user-templates are constructed as the mean vectors of the feature vectors extracted from the enrollment images of the users.

7. The Databases and Experimental Setups

This section presents the experimental databases, setups and performance measures used to assess the feasibility of the Complete Gabor-Fisher Classifier (CGFC) for face recognition. Four popular face databases are selected for the experiments presented in the next section, namely, the XM2VTS database [37], the Extended YaleB database [38, 39], the FERET database [40], and the AR database [41, 42]. These databases are employed either in face verification or face identification experiments to demonstrate the effectiveness and robustness of the proposed CGFC framework.

7.1. The XM2VTS Database. The XM2VTS database is a large multimodal database featuring image, video and speech data of 295 subjects [37]. For the experiments presented in Section 8 we adopt only the (face) image part of the database—the datasets labeled as CD001 and CD006. These two datasets contain a total of 2360 images that were captured in four separate recording sessions. The recording sessions were distributed evenly over a period of approximately five months and at each session the external conditions were controlled. This means that all images were taken against a more or less uniform background, that good illumination conditions were present during the recording, and that only small tilts and in-plane rotations were allowed. The described recording setup resulted in the facial images exhibiting variations mainly induced by the temporal factor




FIGURE 9: Sample images from the image part of the XM2VTS database.

(as shown in Figure 9). Thus, images of the same subject differ in terms of hair-style, presence or absence of make-up and glasses, pose, expression, and so forth. Since two images were taken at each of the four recording sessions, 8 facial images are featured in the database for each of the 295 subjects.

To ensure that our results are comparable to other results obtained on the XM2VTS database and published in the literature, we follow the first configuration of the experimental protocol (for face verification) associated with the database, also known as the Lausanne protocol [37]. The first configuration of the protocol was chosen for the experiments, since it is considered to be the most difficult of the different experimental configurations defined by the Lausanne protocol. As stated by the protocol, we split the subjects of the database into two disjoint groups of 200 clients and 95 impostors (25 evaluation impostors and 70 test impostors). (Note that the term client refers to a subject making a legitimate identity claim, while the term impostor refers to an user making an illegitimate identity claim.) These two groups are then further partitioned into image sets employed for training, evaluation and testing. Specifically, the first configuration of the protocol results in the following experimental setup [35]:

- (i) number of training images: 3 per client (600 in total),
- (ii) number of client verification attempts on the evaluation image sets: $n_{ce} = 600$ (3×200),



FIGURE 10: Sample images from the Extended YaleB database.

- (iii) number of impostor verification attempts in the evaluation image sets: $n_{ie} = 40000$ ($25 \times 8 \times 200$),
- (iv) number of client verification attempts in the test image sets: $n_{ct} = 400$ (2×200), and
- (v) number of impostor verification attempts in the test image sets: $n_{it} = 112000$ ($70 \times 8 \times 200$).

The above numbers are obtained when matching each client image from the evaluation or test image set against the corresponding client template and all impostor images from the evaluation or test image sets against all client templates stored in the systems database.

The training set is used to train the system, that is, to generate the face space where facial images are compared, and to build the client templates (or models). The evaluation image set is employed to tune any potential parameters of the face space and adjust the decision threshold, while the test set is used exclusively for the final performance assessment with predefined system parameters.

7.2. The Extended YaleB Database. The second database used in our experiments is the Extended YaleB (EYB) database [38, 39]. Different from the XM2VTS database, the EYB database is used in our experiments to assess the relative usefulness of the CGFC method for face identification.

The EYB database was recorded at the Yale University and features 2415 frontal face images of 38 subjects. Different from the XM2VTS database, images of the EYB were captured at a single recording session in a relatively short time. Hence, the images are free from severe expression-changes and session-induced variability, but exhibit large variations in illumination, as shown in Figure 10.

To make the experimental protocol as challenging as possible, we partition the EYB database into five image subsets based on the extremity in illumination, as suggested by the authors of the database [38, 39], and use the first subset (the subset with images captured in “good” illumination conditions) for training and the remaining four subsets for testing. This setup results in highly mismatched conditions between the training and test images and represents quite a challenge to the recognition techniques used. Furthermore, it is also in accordance with real-life settings, where the enrollment stage can typically be supervised, while the operation environment is unknown in advance and can feature arbitrary conditions. Specifically, the presented partitioning of the database results in the following experimental setup:

- (i) number of training images: 7 per client (265 in total),
- (ii) number of identification experiments with images from subset 2: $n_{s2} = 456$,
- (iii) number of identification experiments with images from subset 3: $n_{s3} = 455$,
- (iv) number of identification experiments with images from subset 4: $n_{s4} = 525$,
- (v) number of identification experiments with images from subset 5: $n_{s5} = 714$.

It should be noted that not all subjects from the database are represented with an equal number of images due to difficulties during the image acquisition stage. The corrupted images were excluded from the database prior to our experiments. This exclusion resulted in less than the initial 64 images being available for each of the 38 subjects and in the image subset sizes presented above.

7.3. The FERET Database. The third database chosen for the evaluation of the Complete Gabor Fisher Classifier is the FERET database. The database has long been the standard database to assess new face identification techniques, not only due to its size but also due to the great challenges that it poses to the existing face recognition technology. The images in the database differ in terms of facial expression, illumination, age and ethnicity.

For our experiments we adopt the standard FERET evaluation protocol, where 1196 frontal face images of 1196 subjects are defined as gallery (target) images, and four different probe (query/test) sets are employed for determining the recognition rates of the assessed techniques, that is, [40, 43]:

- (i) the Fb probe set, which contains 1195 images exhibiting different facial variations in comparison to the gallery images ($n_{sfb} = 1195$),
- (ii) the Fc probe set, which contains 194 images exhibiting different illumination conditions in comparison to the gallery images ($n_{sfc} = 194$),
- (iii) the Dup I probe set, which contains 722 images acquired between one minute and 1031 days after the corresponding gallery images ($n_{sDupI} = 722$),
- (iv) the Dup II probe set, which contains 234 images acquired at least 18 months after the corresponding gallery images ($n_{sDupII} = 234$).

Some examples of these images are shown in Figure 11. It should be noted that the standard experimental protocol associated with the FERET database does not define a fixed set of training images. Therefore, we select the most commonly adopted training set of 1002 images (corresponding to 428 subjects) for our experiments. (Please visit <http://luks.fe.uni-lj.si/en/staff/vitimir/index.html> for the list of training images used in our experiments.)

7.4. The AR Database. The last database employed in the experimental section is the AR database [41]. The database



FIGURE 11: Sample images from the FERET database.

contains more than 4000 color images of 126 subjects taken during two separate recording sessions. While the variability in the face images of the AR database is caused, for example, by different illumination and facial expressions, the main characteristic that has popularized the database are occlusions of the facial area due to the presence of scarves and glasses.

Each subject in the AR database is accounted for with 26 images taken under different conditions. Figure 12 shows all 13 images of one subject from the AR database acquired in the first recording session. The remaining 13 images were recorded under the same conditions during the second recording session and are, hence, similar in appearance.

Following the experimental setup presented in [9], we select a subset of 100 subjects (50 males and 50 females) for our experiments. We choose the first three images from both recording sessions for training, that is, images denoted by N, H, and A in Figure 12 (6 images per subject, 600 in total), and group the remaining images into a number of probe (or test) sets. Here, each of the probe sets is designed in such a way that only a predefined type of image variability (or combination of specific variability types) and its (their) impact on the CGFC technique is assessed at a time, that is:

- (i) the *scarves* probe set is designed to assess the impact of lower face occlusions on the recognition accuracy of the CGFC technique, and features a total of $n_{s_s} = 600$ images (6 per subject)—denoted by S, SL, and SR in Figure 12,
- (ii) the *glasses* probe set is designed to assess the impact of upper face occlusions on the recognition accuracy of the CGFC technique, and features a total of $n_{s_g} = 600$ images (6 per subject)—denoted by G, GL, and GR in Figure 12,
- (iii) the *scream* probe set is designed to assess the impact of extreme facial expression variations on the recognition accuracy of the CGFC technique, and features a total of $n_{s_{sc}} = 200$ images (2 per subject)—denoted by SC in Figure 12,
- (iv) the *lighting* probe set is designed to assess the impact of illumination variations on the recognition accuracy of the CGFC technique, and features a total of $n_{s_l} = 600$ images (6 per subject)—denoted by L, R, and F in Figure 12,
- (v) the *all* probe set is designed to assess the robustness of the CGFC technique to various types of image variability, and features a total of $n_{s_A} = 2000$ images

(20 per subject)—denoted by SC, L, R, F, G, GL, GR, S, SL, and SR in Figure 12.

7.5. Performance Measures. The recognition performance of the techniques assessed in the next section is measured by the standard error and recognition rates commonly used in the field of face recognition.

For the verification experiments the false acceptance error and false rejection error rates (FAR and FRR, resp.) as well as the half total error rate (HTER) are used. The FAR and FRR are defined as follows:

$$\text{FRR} = \frac{n_{rc}}{n_c} 100\%, \quad \text{FAR} = \frac{n_{ai}}{n_i} 100\%, \quad (19)$$

while the HTER is given by

$$\text{HTER} = 0.5(\text{FAR} + \text{FRR}). \quad (20)$$

In the above equations n_{rc} denotes the number of rejected legitimate identity claims, n_c stands for the number of all legitimate identity claims made, n_{ai} denotes the number of accepted illegitimate identity claims, and n_i represents the number of all illegitimate identity claims made.

Note that both the FAR and the FRR depend on the value of the decision threshold Δ (see (15)). Selecting a threshold that ensures a small value of the FAR inevitably results in a large value of the FRR and vice versa, a threshold that ensures a small FRR results in a large value of the FAR. Thus, to fairly compare the different recognition techniques the decision threshold has to be set in such a way that it ensures some predefined ratio of the FAR and FRR on some evaluation dataset or, alternatively, the two error rates have to be plotted against all possible values of the decision threshold, resulting in the so-called performance curves. For our assessment we chose the latter approach and represent the results in the form of Detection Error Trade-off (DET) curves, which plot the FAR against the FRR at different values of Δ on a scale defined by the inverse of a cumulative Gaussian density function.

For the performance evaluation on the test image sets we again use performance curves, which this time take the form of Expected Performance Curves (EPCs) [44]. To generate an EPC two separate image sets are needed. The first image set, that is, the evaluation image set, is used to find a threshold that minimizes the following weighted error function (WER) for different values of α :

$$\text{WER}(\Delta, \alpha) = \alpha \text{FAR}(\Delta) + (1 - \alpha) \text{FRR}(\Delta), \quad (21)$$

where α denotes a weighting factor that controls the relative importance of the FAR and FRR in the above expression. Next, the second image set, that is, the test image set, is employed to estimate the value of the HTER at the given α and with the computed value of the decision threshold Δ . When plotting the HTER (obtained on the test image sets) against different values of the weighting factor α , an example of the EPC is generated.

For the identification experiments we provide results not in the form of error rates, but rather in form of recognition



FIGURE 12: Sample images of one subject from the AR database taken at the first photo session. The images exhibit the following characteristics: (upper row—from left to right) neutral face (N), happy face (H), angry face (A), screaming face (SC), left light (L), right light (R), and frontal light (F); (lower row—from left to right) occluded by glasses (G), occluded by glasses and lit from left (GL), occluded by glasses and lit from right (GR), occluded by a scarf (S), occluded by a scarf and lit from left (SL), and occluded by a scarf and lit from right (SR).

TABLE 1: Results of the face identification experiments on the EYB database for varying lengths of the PCA and LDA feature vectors.

NOF	PCA				NOF	LDA			
	subset 2	subset 3	subset 4	subset 5		subset 2	subset 3	subset 4	subset 5
10	56.6	29.5	11.2	15.6	5	98.3	56.9	9.9	13.6
50	93.4	54.9	16.7	22.0	10	100	85.3	27.2	29.7
100	93.6	54.9	16.7	22.0	20	100	97.8	47.0	43.7
150	93.6	54.9	16.7	22.0	30	100	99.3	53.6	47.6
200	93.6	54.9	16.7	22.0	37	100	99.8	56.3	51.0

rates. To this end, we compute the so-called rank one recognition rate (ROR) for each of the probe (test) sets of the given database. Here, the ROR is defined as follows:

$$\text{ROR} = \frac{n_{si}}{n_s} 100\%, \quad (22)$$

where n_{si} denotes the number of images successfully assigned to the right identity and n_s stands for the overall number of images trying to assign an identity to.

In addition to the ROR, we also make use of cumulative match characteristic (CMC) curves, which represent performance curves for biometric recognition systems operating in identification mode. While the ROR carries information about the percentage of images where the closest match in the database corresponds to the correct identity, it is sometimes also of interest whether the correct identity is among the top r ranked matches (where $r = 1, 2, \dots, N$ and N denotes the number of subjects in the database of the biometric recognition system). This is especially important for law enforcement applications, where the top r matches can additionally be inspected by a human operator. When computing the recognition rate for the r th rank, the identification procedure is considered successful if the correct identity is among the top r ranked matches and is considered unsuccessful otherwise. Plotting the computed recognition rates against the corresponding rank results in an example of the CMC curve.

8. Experiments and Results

This section presents the experiments with the CGFC technique. It commences by describing the basic preprocessing

preceding the assessment of the proposed face recognition approach and continues by analyzing the results of the assessment.

8.1. Image Preprocessing. Before we turn our attention to the experiments, let us say a few words on the preprocessing preceding our experiments. Since we are concerned with the recognition of faces from digital images and not the performance of facial detectors, we assume that all facial images are properly localized and aligned. In any case, the results presented in the remainder of the paper can be considered as an upper bound on the performance with a properly working face detector. The reader is referred elsewhere for details on how to obtain properly localized face images, for example, [45, 46].

To localize the facial region in the experimental images, we use the eye coordinates provided with the four databases. Based on these coordinates, we first rotate and scale the images in such a way that the centers of the eyes are located at predefined pixel positions. This procedure ensures that the images are aligned with each other. Next, we crop the facial region to a standard size of 128×128 pixels and finally normalize the cropped region using histogram equalization followed by the zero-mean and unit-variance normalization. Some examples of the facial images from the XM2VTS, EYB, AR and FERET databases processed with the described procedure are presented in Figure 13.

8.2. The Baseline Performance. The first series of recognition experiments assesses the performance of some baseline face recognition techniques and adjusts the parameters of the techniques for the best possible performance. It should be




FIGURE 13: Examples of preprocessed images of (from top to bottom): the XM2VTS database, the EYB database, the AR database, and FERET database.

noted that, at this point, only two out of four experimental databases (i.e., the XM2VTS and EYB databases) are used in the experiments. The findings from this series of experiments are ultimately employed with the remaining two databases in the remainder of this section. Such an experimental configuration reflects real-life settings, where the parameters of the adopted recognition technique have to be set in advance on an independent (generic) database due to the fact that the actual test images are not available in the optimization stage.

We select the Principal Component Analysis (PCA) [47] and the Linear Discriminant Analysis (LDA) [33, 48] as our baseline techniques and assess their performance with different numbers of features (NOFs) in the PCA and LDA feature vectors. Considering the number of subjects and available training images in each of the two databases, the maximum length of the feature vector for the PCA method is 599 for the XM2VTS database and 264 for the EYB database, while the maximum number of features constituting the LDA feature vectors equals 199 for the XM2VTS database and 37 for the EYB database.

With the presented experimental setup in mind, let us first turn to the results of the assessment on the XM2VTS database presented in Figure 14. Here, the graphs depicted in Figures 14(a) and 14(c) represent DET curves generated during the experiments with the PCA and LDA techniques, respectively. The plots in Figures 14(b) and 14(d), on the other hand, were created by computing the HTER at three characteristic operating points, that is, at $\text{FAR} = \text{FRR}$, at $\text{FAR} = 0.1\text{FRR}$, and at $\text{FAR} = 10\text{FRR}$, on the DET curves and thus in a sense represent cross sections of the DET curves at

the three operating points. Note that for the PCA technique the DET curves are generated only for feature vectors with 300 or less features, while the cross sections (in Figure 14(b)) contain the entire span of feature vector lengths. The reason for such a presentation of the results lies in the fact that the performance of the PCA technique saturates at feature vector lengths of around 200, and, thus, the DET curves, if shown up to the maximum number of features possible, would be illegible.

Different from the PCA case, where the improvements in the face verification performance become marginal with the increase of the number of features once a certain feature vector length is reached, the performance of LDA technique steadily increases with the increase of the feature vector length. This setting is evidenced by the DET curves as well as their cross sections presented in Figures 14(c) and 14(d), respectively.

Similar observations as with the XM2VTS database can also be made with the EYB database, only this time for the task of face identification. From the results of the experiments on the EYB database presented in Table 1 in form of rank one recognition rates, we can again see that the PCA technique saturates in performance with less than the maximum possible feature vector length and reaches its top performance with 150 features. The LDA technique, on the other hand, once more requires the entire set of (in this case 37) features to obtain the best possible performance achievable with this technique. It is evident that, for the optimal performance, the LDA technique requires 100% of the features comprising its feature vectors, while the PCA approach saturates in its performance with a feature vector length which on both databases ensures that approximately 98% of the (training set) variance is retained. (Note that the 200 features of the XM2VTS correspond to 97.67% of the variance, while the 150 features of the EYB database account for 98.44% of the variance in the training data.) The role of the presented experimental results is twofold: (i) they provide a baseline performance on the two databases for the following comparative studies with the phase-based and complete Gabor-Fisher classifiers, and (ii) they serve as a guideline for selecting the feature vector lengths on the remaining two databases.

8.3. The Baseline Performance of the Gabor-Based Classifiers.

The second series of face recognition experiments assesses the performance of the classical Gabor-Fisher Classifier and the novel Phase-based Gabor-Fisher Classifier and, furthermore, evaluates the relative usefulness of additional normalization techniques applied to the augmented (Gabor magnitude and Gabor phase congruency) feature vectors prior to the deployment of LDA for dimensionality reduction. It has to be noted that commonly only a zero-mean and unit-variance normalization is applied to the Gabor magnitude features, usually with the justification of adjusting the dynamic range of the responses to a common scale. However, as will be shown in this section, the same result can also be achieved with other normalization techniques, which can also have a positive effect on the final face recognition

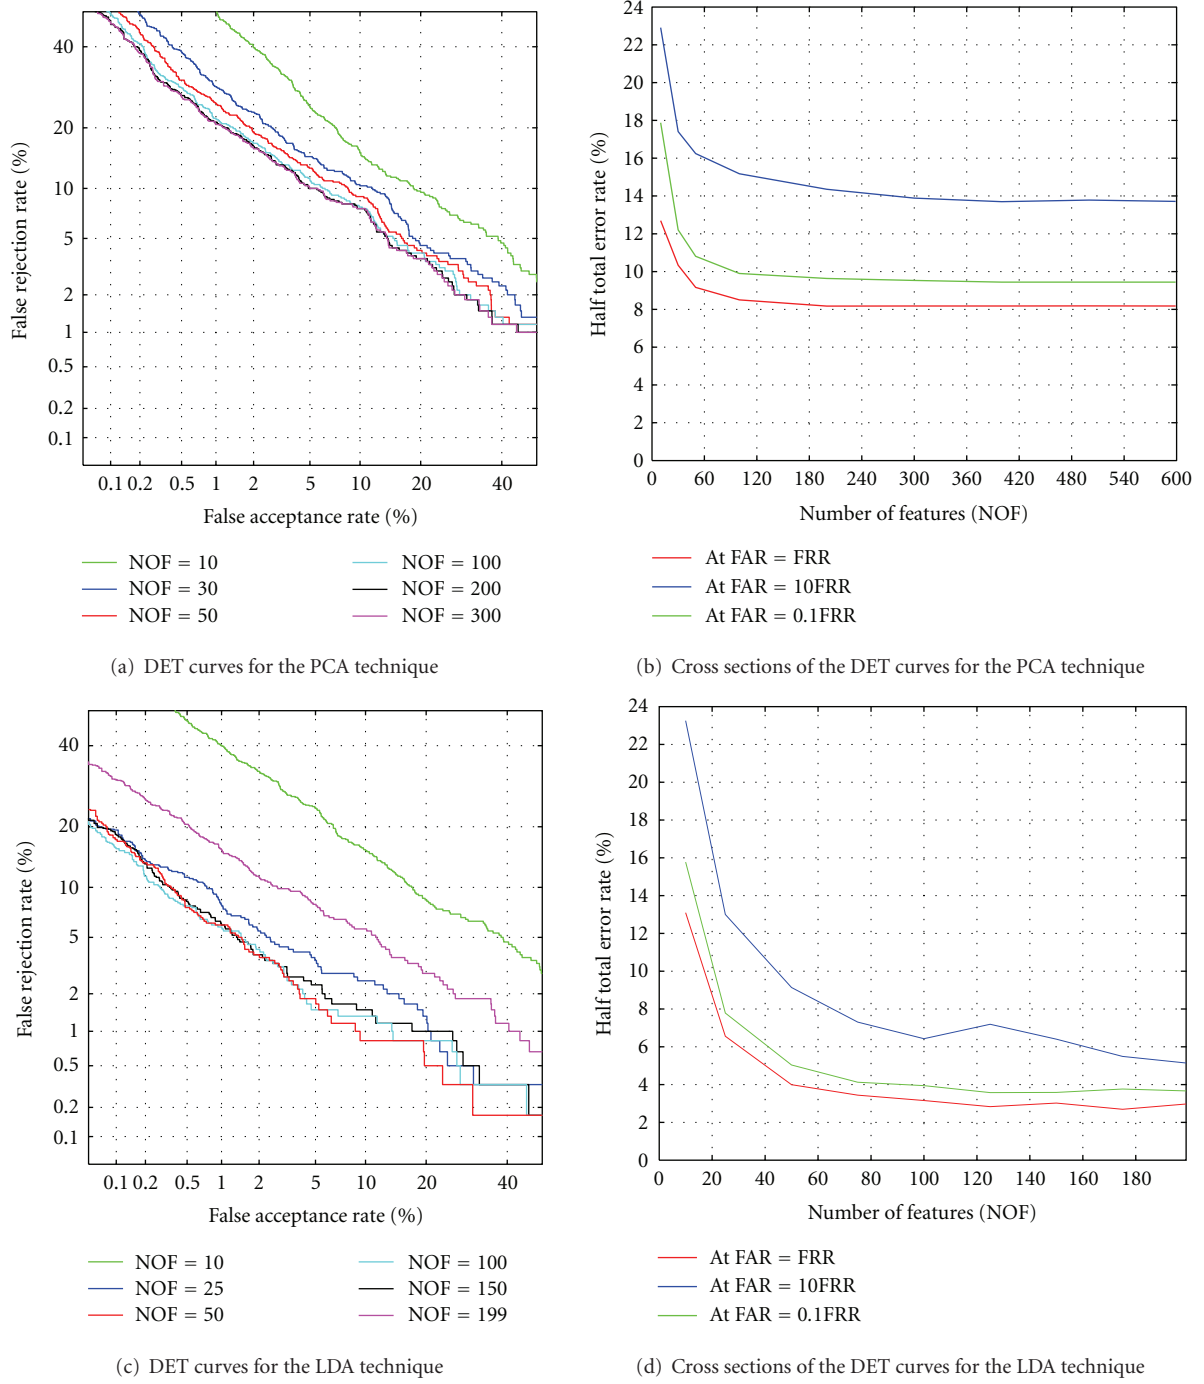


FIGURE 14: Results of the face verification experiments on the XM2VTS database for varying lengths of the PCA and LDA feature vectors.

performance. Note that again only the XM2VTS and EYB databases are employed in the experiments.

We implement the traditional Gabor Fisher Classifier as well as the Phase-based Gabor-Fisher Classifier with a Gabor filter bank containing filters of five scales ($p = 5$) and eight orientation ($r = 8$). Such a filter bank is the most common composition of Gabor filters used for deriving the Gabor face representation [6, 10, 11, 13, 14, 19, 20, 22], and is therefore

also chosen for this series of experiments. We follow three different strategies to normalize the downsampled Gabor magnitude responses (GMRs) and the oriented Gabor phase congruency images (OGPCIs):

- (i) after the downsampling of the GMRs (or OGPCIs), each downsampled GMR (or OGPCI) is normalized to zero-mean and unit-variance before concatenation

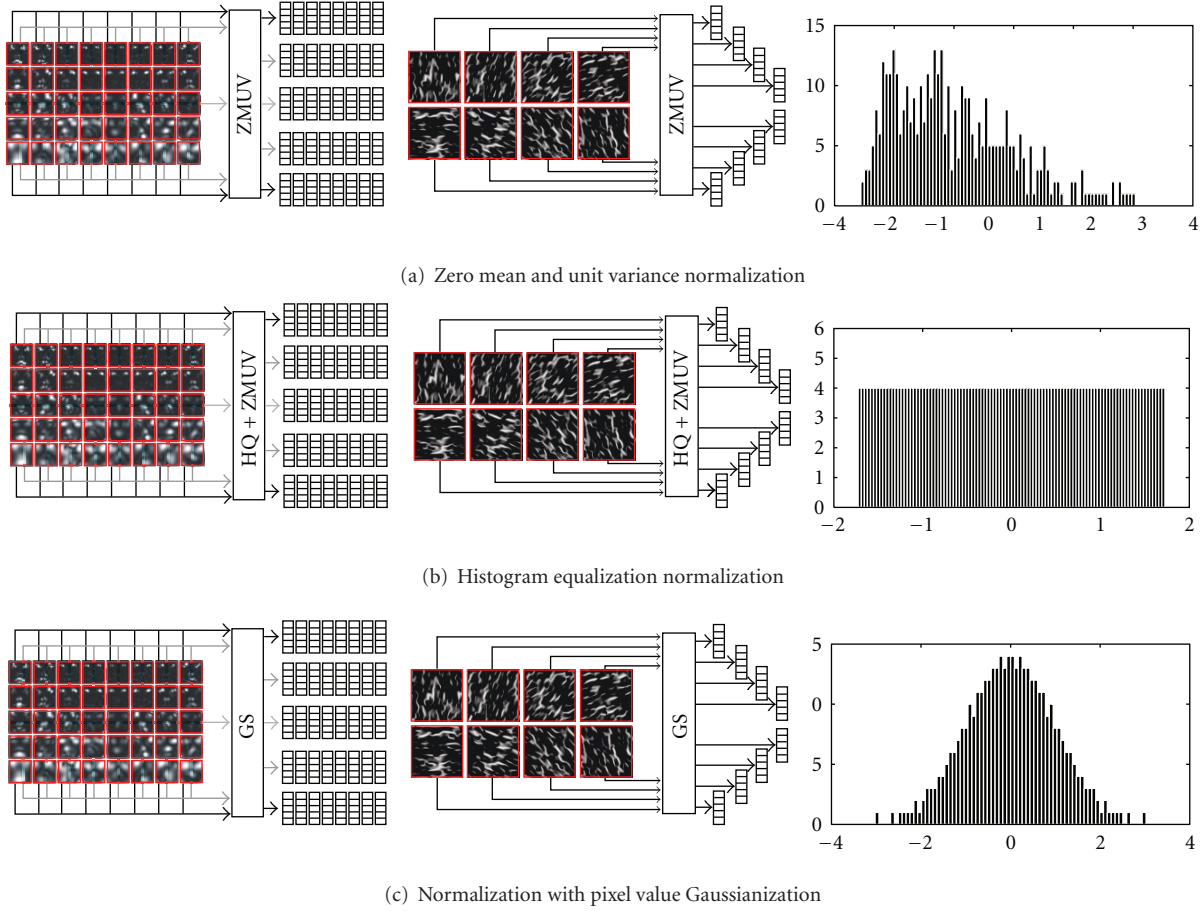


FIGURE 15: Visualization of the applied normalization schemes.

into the final augmented Gabor magnitude (or phase) feature vector (denoted by ZMUV),

- (ii) after the downsampling of the GMRs (or OGPSIs), each downsampled GMR (or OGPCI) is first subjected to histogram equalization followed by zero-mean and unit-variance normalization before concatenation into the final augmented Gabor magnitude (or phase) feature vector (denoted by HQ),
- (iii) after the downsampling of the GMRs (or OGPSIs), each downsampled GMR (or OGPCI) is first subjected to gaussianization [49] before concatenation into the final augmented Gabor magnitude (or phase) feature vector (denoted by GS). (It should be noted that the term gaussianization refers to the remapping of the histogram of an image or pattern vector to a normal distribution with predefined parameters. In our case the target distribution is $\mathcal{N}(0, 1)$.)

The described strategies are also shown in Figure 15, where the most left of each image-triplet depicts the normalization procedure applied on the downsampled GMRs, the center image depicts the normalization procedure applied to the

OGPCIs, and the most right image shows the impact of the applied normalization procedure on the histogram of the downsampled GMR (or OGPCI).

For the implementation of the subspace projection technique (LDA) the following feature vector lengths were chosen: 199 for the XM2VTS database and 37 for the EYB database. These lengths were selected based on preliminary experiments, which suggested the same result as the baseline experiments from the previous section, that is, that the best performance with LDA applied on augmented Gabor (magnitude and phase congruency) feature vectors is obtained with the maximum feature vector length.

The results of this series of experiments are presented in Figure 16 for the XM2VTS database and in Table 2 for the EYB database. If we first focus on the PBGFC technique, we can notice that overall the best performance was achieved with the help of the HQ technique. Figure 16(a) clearly shows that the DET curve generated during the experiments with the HQ techniques outperforms the remaining two normalization techniques at almost all operating points. Similarly, the technique also results in the best identification performance on three out of four test (probe) subsets of the EYB database when compared to any of the remaining two normalization techniques.

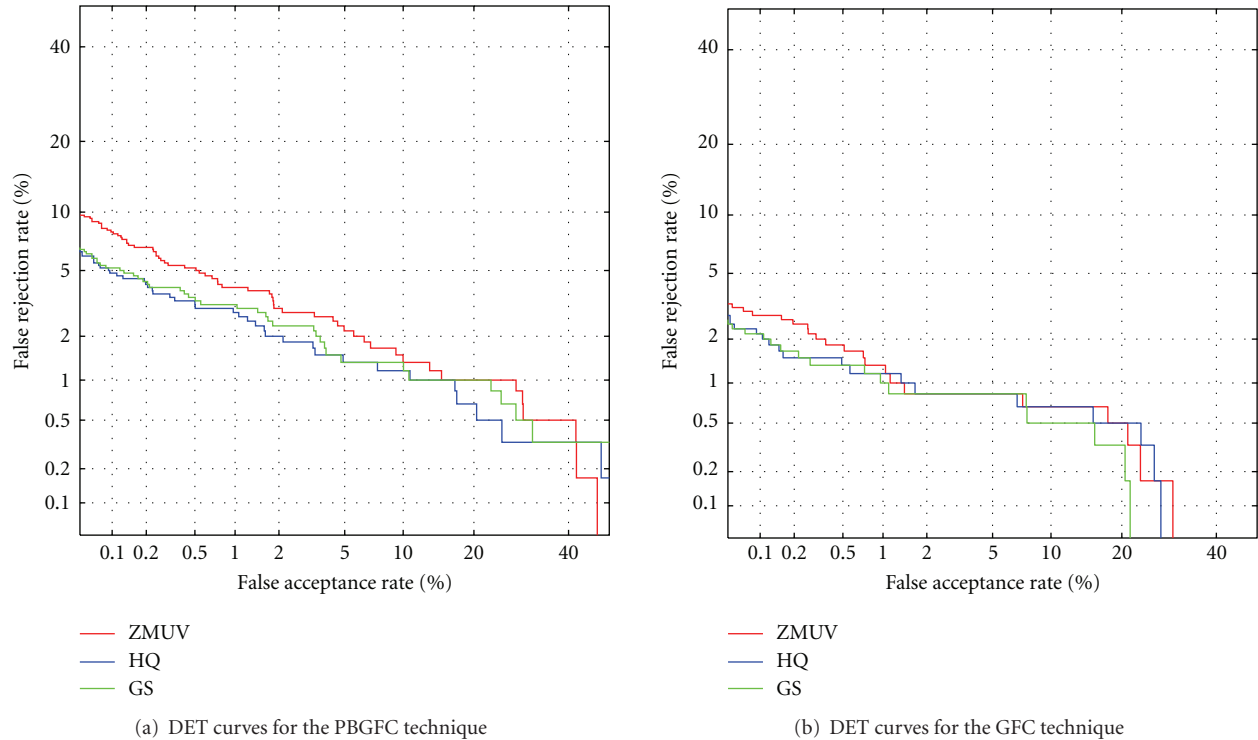


FIGURE 16: Comparison of the impact of different normalization techniques on the face verification performance on the XM2VTS database.

Differently from the PBGFC technique, we observe the best performance for the GFC method with the GS normalization technique both on the XM2VTS database as well as on the EYB database. The result is rather expected since the Gaussian distribution most appropriately reflects the sparse nature of the Gabor wavelet face representation. When compared to the baseline results obtained with the PCA and LDA techniques, both Gabor-based classifiers significantly improve upon the baseline performance on both experimental databases. However, putting this issue aside, we can conclude that this series of recognition experiments suggests that the PBGFC technique should be implemented with the HQ normalization technique, while the GFC method should be combined with the GS normalization procedure and that these combinations should be used in the following comparative assessments.

8.4. Impact of Filter Scales. The third series of experiments evaluates the impact of the number of the filter scales p in the Gabor filter bank on the performance of the PBGFC technique. We fix the angular resolution of the filter bank to $r = 8$ and gradually change the value of the employed filter scales for phase congruency computation from $p = 2$ to $p = 5$. In all of the experiments we set the feature vector lengths to their maximum value and adopt the HQ technique for normalization of the augmented phase congruency feature vectors. It should be noted that we do not assess the impact of the filter scales on the performance of the GFC techniques since various studies on this topic can be found in the

literature [10, 11], all with the same result—that 5 scales and 8 orientations result in the best performance.

From the results presented in Figure 17 for the XM2VTS database and Table 3 for the EYB database we can notice that differently from the GFC technique, the PBGFC does not perform at its optimum with 5 filter scales. Rather, the best performance for the XM2VTS database is observed with only 2 filter scales, that is, $p = 2$. Here, an equal error rate of 1.16% is achieved with the PBGFC approach using 2 filter scales only. Similar results are also observed on the EYB database, where again the recognition performance increases with the decrease of used filter scales. However, the performance peaks with $p = 3$ filter scales.

Based on this series of experiments, we chose to implement the construction procedure of the augmented phase congruency feature vector with 2 scales for the XM2VTS database and 3 scales for the EYB database for the inclusion into the Complete Gabor-Fisher Classifier that will be assessed in the next section.

8.5. Optimization of the Complete Gabor-Fisher Classifier. Up until now, the experiments focused on a separate assessment of the Gabor-Fisher and Phase-based Gabor-Fisher Classifiers for face recognition. However, the results obtained so far suggest that the combination of both techniques into a unified framework, that is, into the Complete Gabor-Fisher Classifier, could improve upon the recognition performance achievable with either of the two Gabor-based classifiers alone.

TABLE 2: Comparison of the impact of different normalization techniques on the rank one recognition rates (in %) on the EYB database.

Norm. tech.	PBGFC				GFC			
	subset 2	subset 3	subset 4	subset 5	subset 2	subset 3	subset 4	subset 5
ZMUV	100	99.8	88.8	93.8	100	100	83.2	89.1
HQ	100	100	86.1	94.8	100	100	82.3	89.1
GS	100	100	84.8	93.0	100	100	84.6	92.2

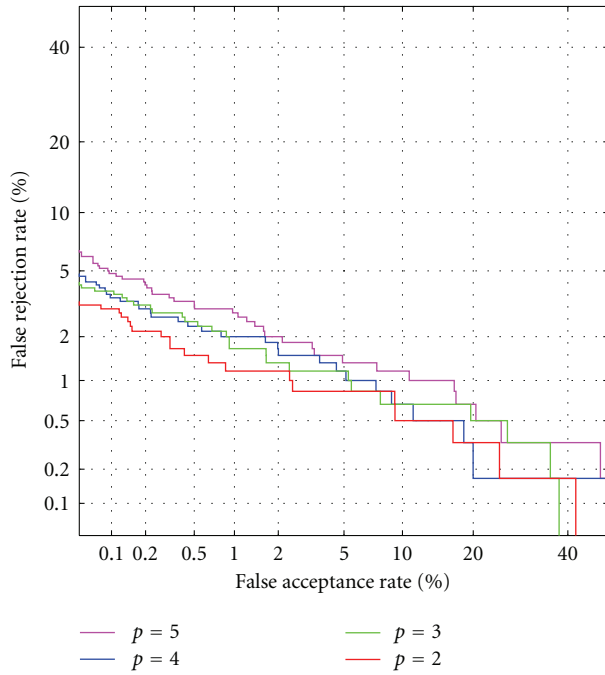


FIGURE 17: DET curves generated for different numbers of filter scales used for computing the OGPCIs.

TABLE 3: Rank one recognition rates (in %) obtained with PBGFC technique on the EYB database for different numbers of filter scales employed during construction of the OGPCIs.

No. of scales	subset 2	subset 3	subset 4	subset 5
$p = 5$	100	100	86.1	94.8
$p = 4$	100	100	94.5	94.4
$p = 3$	100	100	96.4	96.4
$p = 2$	100	100	94.7	96.6

The experimental results presented in Section 8.4 showed that the PBGFC requires less than the 40 filters needed by the GFC technique to achieve optimal face recognition performance. Thus, the PBGFC technique operates on a much narrower frequency band than the GFC approach with most of the discriminatory Gabor phase congruency information being contained in the OGPCIs obtained with Gabor filters of high frequencies. (Note that the number of filter scales is directly proportional to the filters banks coverage of the frequency domain.) In addition to the high frequency filters, the GFC method effectively uses also the

TABLE 4: Rank one recognition rates (in %) obtained with the CGFC approach on the EYB database for four different values of the fusion parameter γ .

γ	subset 2	subset 3	subset 4	subset 5
$\gamma = 0.1$	100	100	84.8	95.5
$\gamma = 0.3$	100	100	91.8	98.2
$\gamma = 0.5$	100	100	95.4	98.6
$\gamma = 0.7$	100	100	98.5	98.2

low frequency Gabor filters, which suggests that the Gabor phase congruency and Gabor magnitude features represent feature types with complementary information and can therefore be combined into a highly efficient unified Gabor-based face recognition approach [6].

As suggested in Section 5, we build the Complete Gabor-Fisher Classifier by combining the GFC and PBGFC techniques at the matching score level [34]. Recall that in this setting the final CGFC similarity score depends on the proper selection of the fusion parameter γ . To assess the robustness of the fusion scheme, the fourth series of face recognition experiments on the XM2VTS and EYB databases evaluates the performance of the CGFC technique with respect to different values of the fusion parameter γ , where $\gamma \in [0, 1]$.

The results obtained with the XM2VTS database are presented in Figure 18. Here, Figure 18(a) shows DET curves obtained at three different values of the fusion parameter γ , while Figure 18(b) depicts the HTER at the same three characteristic operating points as in the case of Figures 14(b) and 14(d). From the two graphs we can see that the fusion of two Gabor-based classifiers is quite robust, as the recognition performance for a wide range of parameter values of γ improves upon the performance of the individual classifiers or, as exemplified by the cross section of the DET curves at the equal error operating point (i.e., FAR = FRR) in Figure 18(b), performs at least as well as the better of the two Gabor-based classifiers.

Similar findings can also be made with the EYB database. Here, the recognition rates on the test subsets 2 and 3 are 100% regardless of the value of the fusion parameter γ . On the test subsets 4 and 5, however, the performance peaks at parameter values in the range from $\gamma = 0.4$ to $\gamma = 0.8$ as shown in Figure 19. The actual rank one recognition rates for four different values of the fusion parameter on all four test images subsets are presented in Table 4. Again, we can see that among the listed values of the fusion parameter, the values of $\gamma = 0.5$ and $\gamma = 0.7$ result in the best performance.

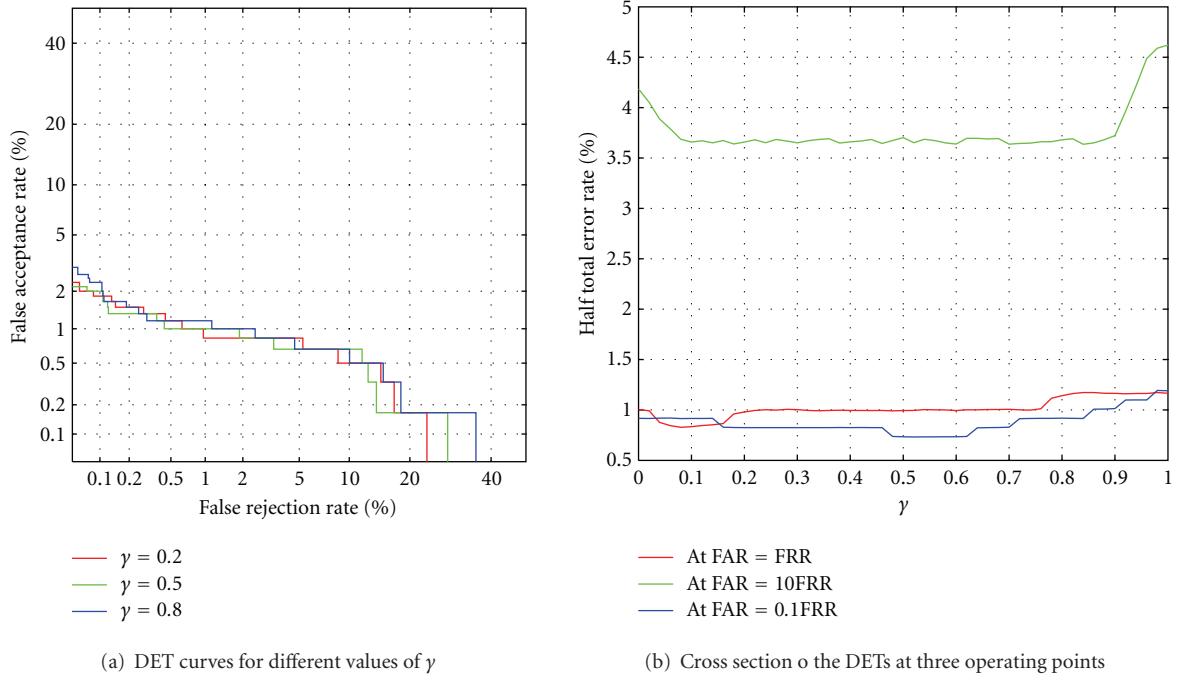


FIGURE 18: Assessment of the CGFC technique on the XM2VTS database.




FIGURE 19: Rank one recognition rates (in %) obtained during the optimization of the CGFC method on the EYB database.

Considering the results of this series of experiments, we select a value of $\gamma = 0.7$ for the calculation of the CGFC matching scores and use this value for the implementation of the CGFC approach on all four experimental databases.

8.6. Recognition in the Presence of Illumination Variations, Partial Occlusions, and Facial Expression Changes. Up to this point we have assessed only the impact of different parameter values and normalization techniques on the performance of the CGFC technique using only two out of four experimental

databases. In the last series of recognition experiments, however, we finally make use of all four databases and evaluate the robustness of the proposed CGFC approach to various sources of image variability commonly encountered in the field of face recognition.

First, we compute the feature vector lengths employed with the PCA and LDA techniques for the AR and FERET databases. As suggested by the experimental results obtained in Section 8.2, we fix the LDA feature vector length to its maximum possible value, which for the AR database equals

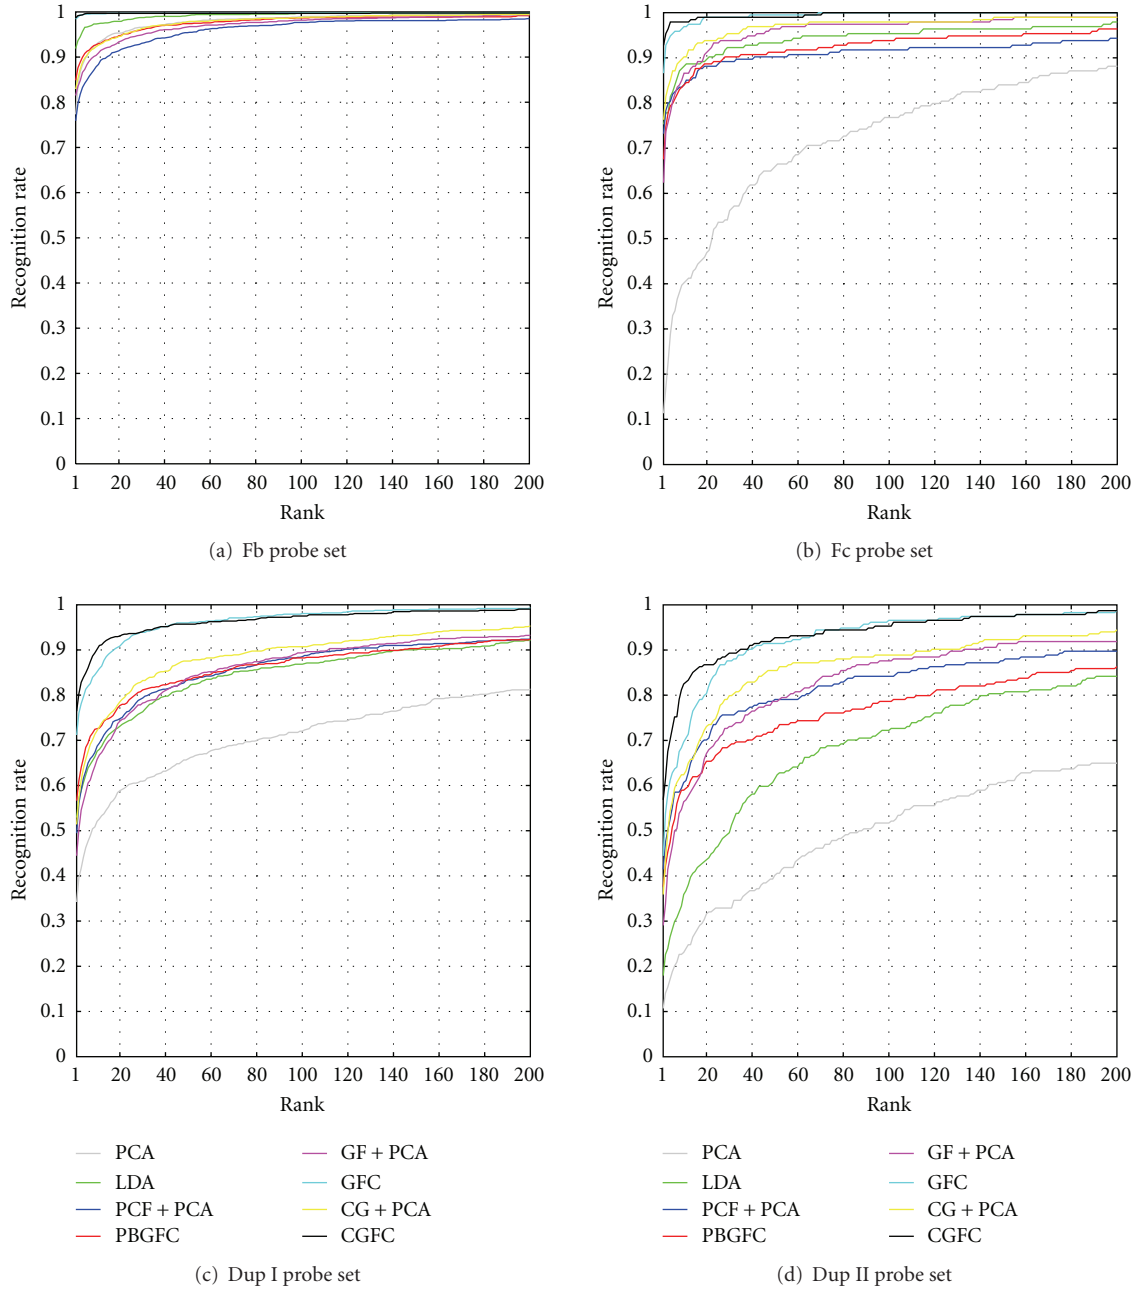


FIGURE 20: CMC curves obtained with different probe sets of the FERET database.

FIGURE 21: Examples of rendered face images (from left to right): the original face image, and the rendered image for $\tau = 40$, the rendered image for $\tau = 80$, the rendered image for $\tau = 120$, the rendered image for $\tau = 160$.

to $d' = 99$, and for the FERET database takes the value of $d' = 427$. For the PCA approach we compute the feature

vector length in such a way that approximately 98% of the training data variance is preserved, resulting in the feature vector length of $d' = 448$ for the FERET database and in the feature vector length of $d' = 358$ for the AR database. Next, we implement the PBGFC technique for the AR and FERET databases using three filter scales ($p = 3$), and use the value of $\gamma = 0.7$ for construction of the Complete Gabor-Fisher Classifier.

Using the presented parameters, we first evaluate the performance of the CGFC technique on the standard probe sets, that is, the Fb, Fc, Dup I, and Dup II probe sets, of the FERET database. The comparative results of the CGFC

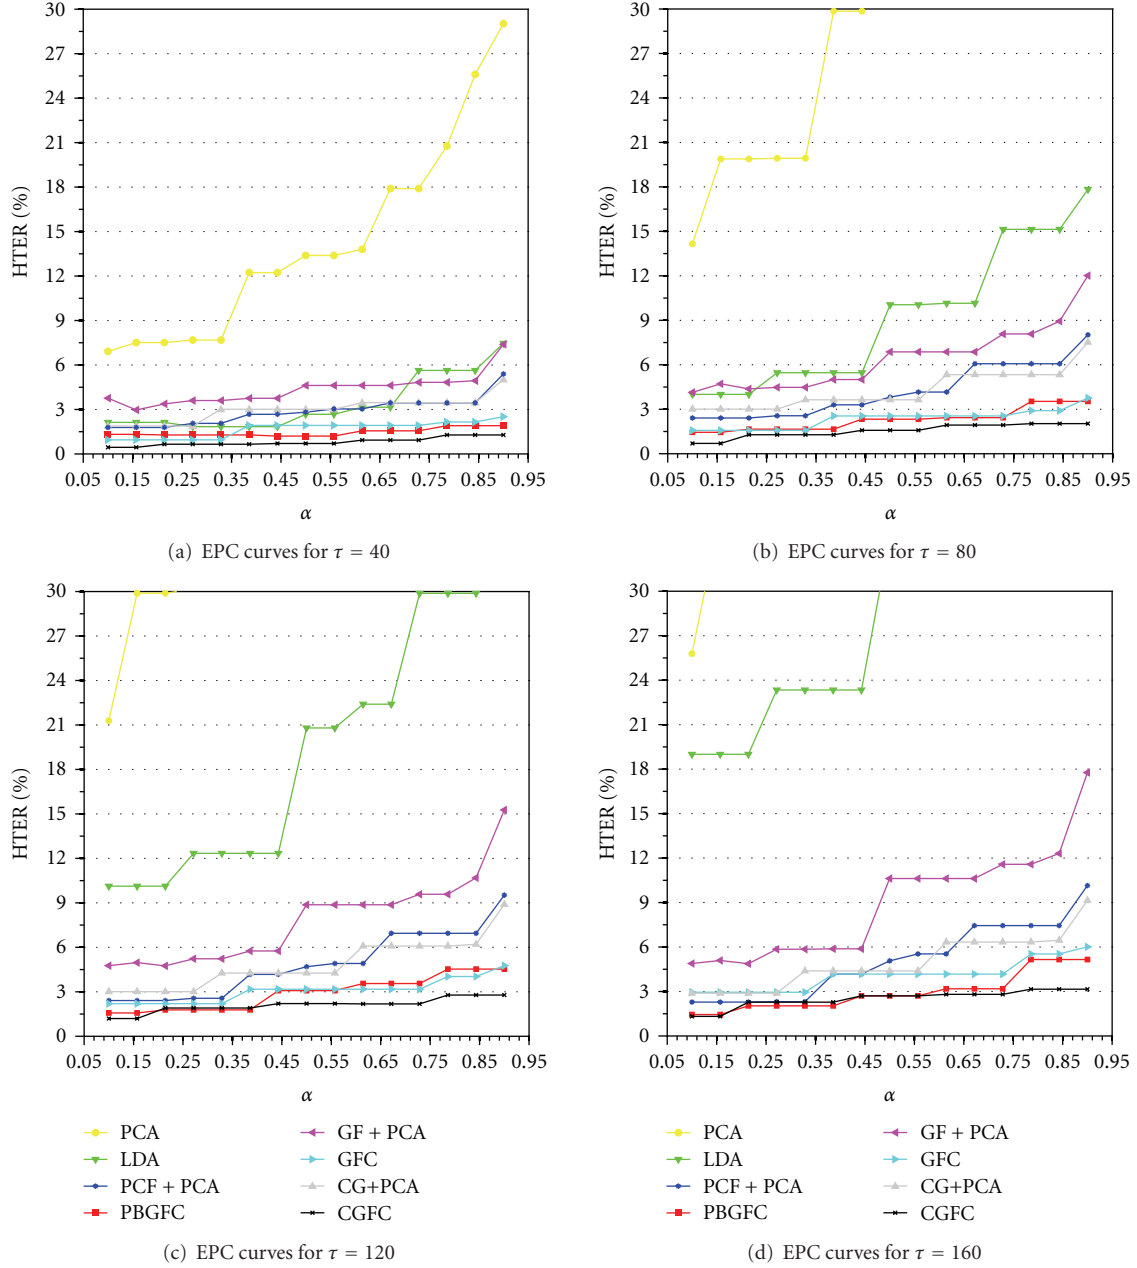


FIGURE 22: EPC curves obtained on the test sets of the XM2VTS database for different values of the parameter τ .

technique and our own implementations of seven popular face recognition techniques are presented in form of CMC curves in Figure 20. Here, the Eigenface technique (PCA) [47], the Fisherface technique (LDA) [33, 48], the Phase-based Gabor-Fisher Classifier (PBGFC) [6], the Gabor-Fisher Classifier (GFC) [10], the PCA technique applied to the augmented Gabor magnitude and Gabor phase congruency feature vectors (GF+PCA and PCF+PCA, resp.) as well as the PCA technique in combination with the complete Gabor face representation (CG+PCA) were adopted for the comparison. In addition to the graphical results, we also present the rank one recognition rates (RORs) for the baseline methods PCA and LDA, for the complete Gabor representations combined

with the PCA and LDA techniques (i.e., CG+PCA and CGFC, resp.) and three state-of-the-art methods, namely, the Local Gabor Binary Pattern Histogram Sequence (LGBPHS) [50] approach, a Local Binary Pattern-(LBP) [51] based technique and the best performing method from the original Sep96 FERET evaluation (BYS) [40]—see Table 5. (Note that n/a stands for the fact that the results for the specific probe set were not provided in the original publication.)

Note that the CGFC technique results in competitive recognition rates on all probe sets. It achieves the ROR of 98.7% on the Fb probe set, which comprises images with different facial expressions as in the gallery set, the ROR of 92.8% on the Fc probe set, where the probe (test) images

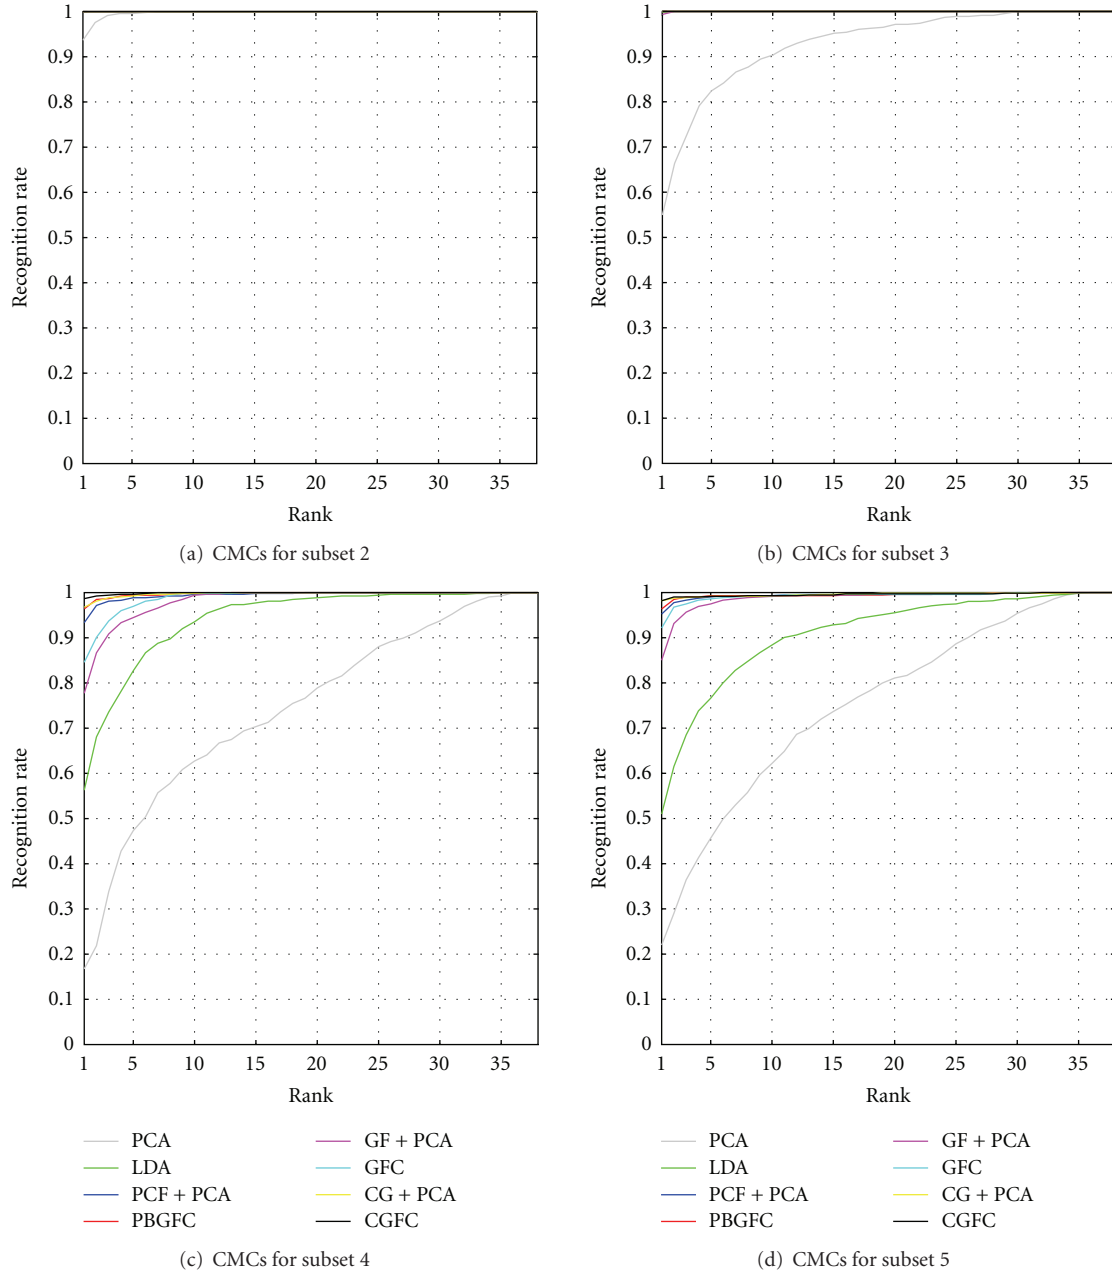


FIGURE 23: CMC curves obtained on the four probe subsets of the EYB database.

feature illumination variations, and RORs of 77.2% and 57.3% on the Dup I and Dup II sets, respectively, where the goal is to assess aging effects on the performance of the given face recognition technique. As already indicated in the introduction, the rather good performance of the CGFC method on the Fb probe set can be linked to the local nature of the Gabor features, which ensures robustness to changes in facial expression, while the robustness to illumination changes evidenced by the recognition rates on the Fc probe set can be related to frequency band coverage of the Gabor filter bank. Despite the competitiveness of the proposed approach in our experiments, it should, nonetheless, be noted that the FERET database does not define a standard

training set for the construction of the PCA/LDA subspace and, thus, the results of the comparison with other methods from the literature should be interpreted with this fact in mind.

Let us now turn our attention to the XM2VTS database. All experiments on this database conducted so far have been performed on the evaluation image sets, while the test image sets were not used. In this series of experiments, however, we employ the test image sets for our comparative study and implement all recognition techniques with their hyperparameters (such as decision thresholds, feature vector lengths, number of employed filter scales, etc.) predefined using the evaluation image sets. As suggested in Section 7.5, we report

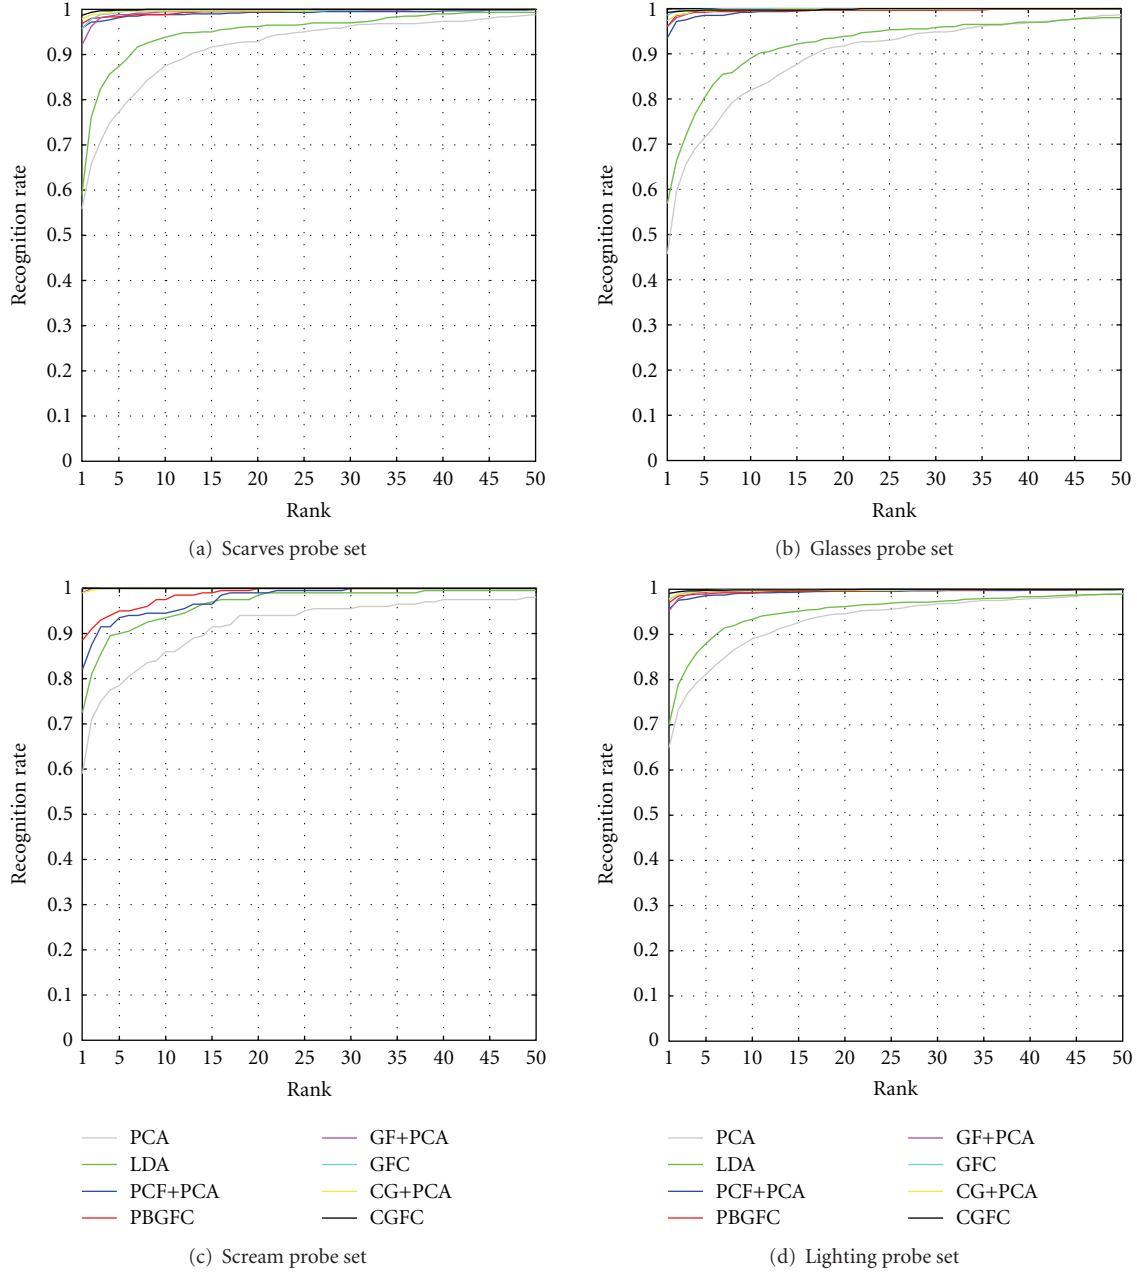


FIGURE 24: CMC curves obtained on the four probe sets (scarves, glasses, scream, and lighting) of the AR database.

the results in form of EPC curves. To make the assessment more challenging, we introduce an artificial illumination change to the test image sets from the XM2VTS database and adopt the artificial illumination model of Sanderson and Paliwal [52] for rendition of the facial images. The model simulates different illumination conditions during the image capturing process by modifying the preprocessed facial images $I(x, y)$, that is,

$$\tilde{I}(x, y) = I(x, y) + mx + \tau, \quad (23)$$

where $x = 0, 1, \dots, a-1$; $y = 0, 1, \dots, b-1$; $m = -2\tau/(b-1)$; τ denotes the parameter that controls the “strength” of the

introduced artificial illumination change [6]. The authors of [52] stressed that their model (see (23)) does not cover all possible illumination effects in real life settings, but is nevertheless useful for providing suggestive results regarding the robustness of the assessed face recognition techniques to illumination changes. Some examples of facial images rendered with the presented model for different values of the parameter τ are shown in Figure 21. Due to the extensive experimental section, we do not report the performance on the original test sets, but provide only comparative results on the harder (degraded) test image sets. Similar as with the FERET database, we implement seven face recognition techniques from the literature for our comparison.




FIGURE 25: CMC curves for the challenging *All* probe set of the AR database.

TABLE 5: Rank one recognition rates (in %) obtained on different probe sets of the FERET database.

Method	Fb	Fc	Dup I	Dup II
PCA	77.3	11.4	34.2	10.7
LDA	91.9	75.3	52.9	18.0
CG+PCA	82.9	62.4	51.4	35.9
CGFC	98.7	92.8	77.2	57.3
LGBPHS [50]	94.0	97.0	68.0	53.0
LBP [51]	93.0	51.0	61.0	50.0
BYS [40]	82.0	37.0	52.0	32.0

TABLE 6: Rank one recognition rates (in %) obtained on the EYB database during the comparative assessment.

Method	subset 2	subset 3	subset 4	subset 5
PCA	93.6	55.0	16.7	22.0
LDA	100	99.8	56.3	51.0
CG+PCA	100	100	96.8	98.2
CGFC	100	100	98.5	98.2
LBP [53]	100	99.8	97.3	87.5
LS [39]	100	100	85.0	n/a
HI [54]	100	100	97.3	n/a
GA [55]	100	100	98.6	n/a

The results of this series of face recognition (verification) experiments are presented in Figure 22 in the form of EPC curves. The first thing to notice from the presented results is that the CGFC method systematically outperforms all other techniques, significantly improving upon the baseline performance of the PCA and LDA methods. Moreover, it also results in the most robust performance in the presence of (artificial) lighting variations, again due to the properties of the Gabor filter bank. To evaluate the performance of

the CGFC technique on real illumination changes and, consequently, to further demonstrate the robustness of the proposed technique to illumination changes, we perform the next round of experiments on the EYB database

The final recognition experiments on the EYB database are again conducted in accordance with the experimental setup presented in Section 7.2. From the results presented in the form of CMC curves in Figure 23, we can see that, while on the probe subsets 2 and 3 almost all evaluated methods resulted in the recognition rate of 100% for all ranks, only the CGFC method is able to retain an almost perfect recognition rate for the more challenging probe subsets 4 and 5 as well. In the comparison with state-of-the-art techniques from the literature, that is, a LBP-based method [53], the Linear Subspace (LS) approach [39], the Harmonic Image (HI) technique [54], and the Gradient Angle (GA) procedure [55], presented in Table 6, the CGFC technique again resulted in competitive RORs making it a suitable choice for various (security, defense, etc.) applications requiring robust recognition in difficult illumination conditions.

Last but not least, we evaluated the CGFC technique on the AR database, where the goal was to assess the robustness of the proposed method against extreme facial expression changes, disguise (partial occlusion of the facial area), illumination changes and different combinations of the listed image characteristics. When looking at the CMC curves generated during the recognition experiments on the scarves, glasses, scream, and lighting probe sets (Figure 24), we can see that the CGFC method achieved the ROR of more than 98% on all probe sets. These results suggest that our method is highly robust to facial occlusions (which usually occur when the subject to be identified tries to conceal his identity through disguise), to extreme facial expression changes and, as already shown in the previous experiments, to illumination variations. When looking at the results obtained with the most challenging probe set, that is, the *All* probe set, shown in Figure 25, we can notice that the overall ROR on this set again reaches a value of more than 99%. To the best of our knowledge the recognition rates obtained on the different probe sets of the AR database represent the best results achieved on this database and published in the literature. This fact is also evidenced in the comparison with state-of-the-art techniques from the literature presented in Table 7. Here, the RORs are shown for the SubSampling (SS) method from [8], the recently proposed Sparse Representation Classifier (SRC) [9], and the Attributed Relational Graph-(ARG) based method [56] as well as for four techniques implemented to produce the CMC curves in Figure 24. The reader should note that, while the results taken from the literature were obtained using similar training images, the probe (test) sets for the SS, SRC and ARG techniques were either smaller in size or featured only one degradation at a time. For example, while we adopted all images featuring a scarf for our *scarves* probe set regardless of the illumination conditions, the SS, SRC and ARG methods produced their results based on only scarf images taken in the same lighting conditions as the training (target) images. Similar observations could also be made for

TABLE 7: Rank one recognition rates (in %) obtained on the five probe sets of the AR database.

Method	scarves	glasses	scream	lighting	all
PCA	55.8	45.7	59.0	64.9	64.7
LDA	58.8	56.8	72.5	70.0	71.6
CG+PCA	97.8	97.3	99.0	97.9	98.3
CGFC	98.7	99.2	100	99.1	99.2
SS [8]	93.0	84.0	87.0	n/a	n/a
SRC [9]	93.5	97.5	n/a	n/a	n/a
ARG [56]	85.2	80.7	66.7	98.5	n/a

the remaining probe sets. The performance of the CGFC method on the challengingly designed probe sets of the AR database offers a final demonstration of the competitiveness of the proposed approach.

9. Conclusion

In this paper we have proposed a novel face classifier for face recognition called the Complete Gabor-Fisher Classifier. Unlike the majority of Gabor filter-based methods from the literature, which mainly rely only on the Gabor magnitude features for representing facial images, the proposed classifier exploits both Gabor magnitude features as well as features derived from Gabor phase information. The feasibility of the proposed technique was assessed on four publicly available databases, namely, on the XM2VTS, FERET, AR and Extended YaleB databases. On all datasets, the proposed technique resulted in a promising face recognition performance and outperformed several popular face recognition techniques, such as PCA, LDA, the Gabor-Fisher Classifier and others. The proposed method was also shown to ensure robust recognition performance in the presence of extreme facial changes, severe lighting variations and partial occlusions of the facial area.

The source code used in all of our experiments will be made freely available. The interested reader is referred to <http://luks.fe.uni-lj.si/en/staff/vitomir/index.html> for more information.

Acknowledgments

The presented research has been supported in parts by the national research program P2-0250(C) Metrology and Biometric Systems, 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] R. W. Ives, Y. Du, D. M. Etter, and T. B. Welch, "A multidisciplinary approach to biometrics," *IEEE Transactions on Education*, vol. 48, no. 3, pp. 462–471, 2005.
- [2] A. K. Jain, A. Ross, and S. Prabhakar, "An introduction to biometric recognition," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 14, no. 1, pp. 4–20, 2004.
- [3] J.D. Woodward, N. M. Orlans, and P. T. Higgins, *Biometrics*, McGraw–Hill, New York, NY, USA, 2002.
- [4] V. Štruc and N. Pavešić, "Phase congruency features for palm-print verification," *IET Signal Processing*, vol. 3, no. 4, pp. 258–268, 2009.
- [5] A. Franco and L. Nanni, "Fusion of classifiers for illumination robust face recognition," *Expert Systems with Applications*, vol. 36, no. 5, pp. 8946–8954, 2009.
- [6] V. Štruc, B. Vesnicher, and N. Pavešić, "The phase-based Gabor Fisher classifier and its application to face recognition under varying illumination conditions," in *Proceedings of the 2nd International Conference on Signal Processing and Communication Systems (ICSPCS '08)*, pp. 1–6, Gold Coast, Australia, December 2008.
- [7] Y. Gao and M. K. H. Leung, "Face recognition using line edge map," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 6, pp. 764–779, 2002.
- [8] S. Fidler, D. Skočaj, and A. Leonardis, "Combining reconstructive and discriminative subspace methods for robust classification and regression by subsampling," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 3, pp. 337–350, 2006.
- [9] J. Wright, A. Y. Yang, A. Ganesh, S. S. Sastry, and Y. Ma, "Robust face recognition via sparse representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 2, pp. 210–227, 2009.
- [10] C. Liu and H. Wechsler, "Gabor feature based classification using the enhanced Fisher linear discriminant model for face recognition," *IEEE Transactions on Image Processing*, vol. 11, no. 4, pp. 467–476, 2002.
- [11] L. Shen and L. Bai, "A review of Gabor wavelets for face recognition," *Pattern Analysis and Applications*, vol. 9, no. 2, pp. 273–292, 2006.
- [12] L. Nanni and D. Maio, "Weighted sub-Gabor for face recognition," *Pattern Recognition Letters*, vol. 28, no. 4, pp. 487–492, 2007.
- [13] L. Shen, L. Bai, and M. Fairhurst, "Gabor wavelets and general discriminant analysis for face identification and verification," *Image and Vision Computing*, vol. 25, no. 5, pp. 553–563, 2007.
- [14] V. Štruc and N. Pavešić, "Gabor-based kernel partial-least-squares discrimination features for face recognition," *Informatica*, vol. 20, no. 1, pp. 115–138, 2009.
- [15] P. Kovesi, "Image features from phase congruency," *Videre: Journal of Computer Vision Research*, vol. 1, no. 3, pp. 1–26, 1999.
- [16] A. Eleyan, H. Özkaramanli, and H. Demirel, "Complex wavelet transform-Based face recognition," *EURASIP Journal on Advances in Signal Processing*, vol. 2008, Article ID 185281, 13 pages, 2008.

- [17] V. Kyrki, J.-K. Kamarainen, and H. Kälviäinen, "Simple Gabor feature space for invariant object recognition," *Pattern Recognition Letters*, vol. 25, no. 3, pp. 311–318, 2004.
- [18] M. Lades, Jan C. Vorbrüggen, J. Buhmann, et al., "Distortion invariant object recognition in the dynamic link architecture," *IEEE Transactions on Computers*, vol. 42, no. 3, pp. 300–311, 1993.
- [19] C. Liu, "Capitalize on dimensionality increasing techniques for improving face recognition grand challenge performance," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 5, pp. 725–737, 2006.
- [20] L. Shen and L. Bai, "Information theory for Gabor feature selection for face recognition," *EURASIP Journal on Applied Signal Processing*, vol. 2006, Article ID 30274, 11 pages, 2006.
- [21] B. Zhang, S. Shan, X. Chen, and W. Gao, "Histogram of Gabor phase patterns (HGPP): a novel object representation approach for face recognition," *IEEE Transactions on Image Processing*, vol. 16, no. 1, pp. 57–68, 2007.
- [22] B. Zhang, Z. Wang, and B. Zhong, "Kernel learning of histogram of local Gabor phase patterns for face recognition," *EURASIP Journal on Advances in Signal Processing*, vol. 2008, Article ID 469109, 8 pages, 2008.
- [23] E. Bezael and U. Efron, "Efficient face recognition method using a combined phase congruency/Gabor wavelet technique," in *Optical Information Systems III*, vol. 5908 of *Proceedings of SPIE*, pp. 1–8, San Diego, Calif, USA, September 2005.
- [24] S. Gundimada and V. K. Asari, "A novel neighborhood defined feature selection on phase congruency images for recognition of faces with extreme variations," *International Journal of Information Technology*, vol. 3, no. 1, pp. 25–31, 2006.
- [25] S. Gundimada, V. K. Asari, and N. Gudur, "Face recognition in multi-sensor images based on a novel modular feature selection technique," *Information Fusion*, vol. 11, no. 2, pp. 124–132, 2010.
- [26] Y. Guo, G. Zhao, J. Chen, M. Pietikäinen, and Z. Xu, "A new Gabor phase difference pattern for face and ear recognition," in *Proceedings of the International Conference on Computer Analysis of Images and Patterns*, vol. 5702 of *Lecture Notes in Computer Science*, pp. 41–49, Münster, Germany, September 2009.
- [27] Y. Guo and Z. Xu, "Local Gabor phase difference pattern for face recognition," in *Proceedings of the 19th International Conference on Pattern Recognition (ICPR '08)*, pp. 1–4, Tampa, Fla, USA, December 2008.
- [28] L. Qing, S. Shan, X. Chen, and W. Gao, "Face recognition under varying lighting based on the probabilistic model of Gabor phase," in *Proceedings of the International Conference on Pattern Recognition*, vol. 3, pp. 1139–1142, Hong Kong, August 2006.
- [29] T. Ojala, M. Pietikäinen, and T. Mäenpää, "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 7, pp. 971–987, 2002.
- [30] T. Ojala, M. Pietikäinen, and T. Mäenpää, "A generalized local binary pattern operator for multiresolution gray scale and rotation invariant texture classification," in *Proceedings of the 2nd International Conference on Advances in Pattern Recognition*, pp. 397–406, Rio de Janeiro, Brazil, March 2001.
- [31] S. Venkatesh and R. Owens, "An energy feature detection scheme," in *Proceedings of the International Conference on Image Processing*, pp. 553–557, Singapore, 1989.
- [32] T. Savič and N. Pavešić, "Personal recognition based on an image of the palmar surface of the hand," *Pattern Recognition*, vol. 40, no. 11, pp. 3152–3163, 2007.
- [33] P. N. Belhumeur, J. P. Hespanha, and D. J. Kriegman, "Eigenfaces vs. fisherfaces: recognition using class specific linear projection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, no. 7, pp. 711–720, 1997.
- [34] A. Ross, K. Nandakumar, and A. K. Jain, "Introduction to multibiometrics," in *Handbook of Biometrics*, A. K. Jain, P. Flynn, and A. Ross, Eds., pp. 271–292, Springer, New York, NY, USA, 2008.
- [35] V. Štruc, F. Mihelič, and N. Pavešić, "Face authentication using a hybrid approach," *Journal of Electronic Imaging*, vol. 17, no. 1, pp. 1–11, 2008.
- [36] B. Vesnicher and F. Mihelič, "The likelihood ratio decision criterion for nuisance attribute projection in GMM speaker verification," *EURASIP Journal on Advances in Signal Processing*, vol. 2008, Article ID 786431, 11 pages, 2008.
- [37] K. Messer, J. Matas, J. Kittler, J. Luetttin, and G. Maitre, "XM2VTSDB: the extended M2VTS database," in *Proceedings of the 2nd International Conference on Audio- and Video-Based Person Authentication*, pp. 72–77, Washington, DC, USA, 1999.
- [38] K.-C. Lee, J. Ho, and D. J. Kriegman, "Acquiring linear subspaces for face recognition under variable lighting," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 5, pp. 684–698, 2005.
- [39] A. S. Georgiades, 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*, vol. 23, no. 6, pp. 643–660, 2001.
- [40] P. J. Phillips, H. Moon, S. A. Rizvi, and P. J. Rauss, "The FERET evaluation methodology for face-recognition algorithms," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 10, pp. 1090–1104, 2000.
- [41] A. M. Martinez and R. Benavente, "The AR face database," Tech. Rep. 24, CVC, New York, NY, USA, June 1998.
- [42] A. M. Martinez, "Recognizing imprecisely localized, partially occluded, and expression variant faces from a single sample per class," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 6, pp. 748–763, 2002.
- [43] K. Delac, M. Grgic, and S. Grgic, "Independent comparative study of PCA, ICA, and LDA on the FERET data set," *International Journal of Imaging Systems and Technology*, vol. 15, no. 5, pp. 252–260, 2005.
- [44] S. Bengio and J. Marithoz, "The expected performance curve: a new assessment measure for person authentication," in *Proceedings of the the Speaker and Language Recognition Workshop Odyssey*, pp. 279–284, Toledo, Spain, 2004.
- [45] P. Viola and M. J. Jones, "Robust real-time face detection," *International Journal of Computer Vision*, vol. 57, no. 2, pp. 137–154, 2004.
- [46] A. Wagner, J. Wright, A. Ganesh, Z. Zhou, and Y. Ma, "Towards a practical face recognition system: robust registration and illumination by sparse representation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR '09)*, pp. 597–604, June 2009.
- [47] M. Turk and A. Pentland, "Eigenfaces for recognition," *Journal of Cognitive Neuroscience*, vol. 3, no. 1, pp. 71–86, 1991.
- [48] W. Zhao, A. Krishnaswamy, R. Chellappa, D. L. Swets, and J. Weng, "Discriminant analysis of principal components

- for face recognition,” in *Face Recognition: From Theory to Applications*, H. Wechsler, P. J. Phillips, V. Bruce, F. F. Soulie, and T. S. Huang, Eds., pp. 73–85, Springer, Berlin, Germany, 1998.
- [49] V. Štruc, J. Žibert, and N. Pavešić, “Histogram remapping as a preprocessing step for robust face recognition,” *WSEAS Transactions on Information Science and Applications*, vol. 6, no. 3, pp. 520–529, 2009.
 - [50] W. Zhang, S. Shan, W. Gao, X. Chen, and H. Zhang, “Local Gabor binary pattern histogram sequence (LGBPHS): a novel non-statistical model for face representation and recognition,” in *Proceedings of the International Conference on Computer Vision*, pp. 786–791, Beijing, China, October 2005.
 - [51] T. Ahonen, A. Hadid, and M. Pietikäinen, “Face recognition with local binary patterns,” in *Proceedings of the European Conference on Computer Vision*, vol. 3021, pp. 469–481, Prague, Czech Republic, May 2004.
 - [52] C. Sanderson and K. K. Paliwal, “Fast features for face authentication under illumination direction changes,” *Pattern Recognition Letters*, vol. 24, no. 14, pp. 2409–2419, 2003.
 - [53] X. Tan and B. Triggs, “Enhanced local texture feature sets for face recognition under difficult lighting conditions,” in *Proceedings of the Analysis and Modelling of Faces and Gestures*, vol. 4778 of *Lecture Notes in Computer Science*, pp. 168–182, Springer, 2007.
 - [54] R. Basri and D. Jacobs, “Lambertian reflectance and linear subspaces,” in *Proceedings of the 8th International Conference on Computer Vision*, pp. 383–390, Vancouver, Canada, July 2001.
 - [55] H. F. Chen, P. N. Belhumeur, and D. W. Jacobs, “In search of illumination invariants,” in *Proceedings of the International Conference on Computer Vision and Pattern Recognition (CVPR ’00)*, pp. 254–261, South Carolina, SC, USA, June 2000.
 - [56] B.-G. Park, K.-M. Lee, and S.-U. Lee, “Face recognition using face-ARG matching,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 12, pp. 1982–1988, 2005.