From Gabor Magnitude to Gabor Phase Features: Tackling the Problem of Face Recognition under Severe Illumination Changes

Vitomir Štruc and Nikola Pavešić
Faculty of Electrical Engineering, University of Ljubljana
Slovenia

1. Introduction

Among the numerous biometric systems presented in the literature, face recognition systems have received a great deal of attention in recent years. The main driving force in the development of these systems can be found in the enormous potential face recognition technology has in various application domains ranging from access control, human-machine interaction and entertainment to homeland security and surveillance (Štruc et al., 2008a).

While contemporary face recognition techniques have made quite a leap in terms of performance over the last two decades, they still struggle with their performance when deployed in unconstrained and uncontrolled environments (Gross et al., 2004; Phillips et al., 2007). In such environments the external conditions present during the image acquisition stage heavily influence the appearance of a face in the acquired image and consequently affect the performance of the recognition system. It is said that face recognition techniques suffer from the so-called PIE problem, which refers to the problem of handling Pose, Illumination and Expression variations that are typically encountered in real-life operating conditions. In fact, it was emphasized by numerous researchers that the appearance of the same face can vary significantly from image to image due to changes of the PIE factors and that the variability in the images induced by the these factors can easily surpass the variability induced by the subjects' identity (Gross et al., 2004; Short et al., 2005). To cope with image variability induced by the PIE factors, face recognition systems have to utilize feature extraction techniques capable of extracting stable and discriminative features from facial images regardless of the conditions governing the acquisition procedure. We will confine ourselves in this chapter to tackling the problem of illumination changes, as it represents the PIE factor which, in our opinion, is the hardest to control when deploying a face recognition system, e.g., in access control applications.

Many feature extraction techniques, among them particularly the appearance based methods, have difficulties extracting stable features from images captured under varying illumination conditions and, hence, perform poorly when deployed in unconstrained environments. Researchers have, therefore, proposed a number of alternatives that should compensate for the illumination changes and thus ensure stable face recognition performance.

Sanderson and Paliwal (Sanderson & Paliwal, 2003), for example, proposed a feature extraction technique called DCT-mod2. The DCT-mod2 technique first applies the Discrete Cosine Transform (DCT) to sub-regions (or blocks) of facial images to extract several feature sets of DCT coefficients, and then compensates for illumination induced appearance changes by replacing the coefficients most affected by illumination variations with the corresponding vertical and horizontal delta coefficients. The authors assessed the technique on images rendered with an artificial illumination model as well as on (real-life) images captured under varying illumination conditions. Encouraging results were achieved on both image types.

Another technique was proposed by Gao and Leung in (Gao & Lung, 2002). Here, the authors argue that the so-called Line Edge Maps (LEM) represent useful face representations for both image coding and recognition and that, moreover, they also exhibit insensitiveness to illumination variations. With this technique a given face image is first processed to extract edge pixels, which are then combined into line segments that constitute the LEMs. The authors showed that their technique successfully outperformed popular feature extraction approaches on various databases.

Liu and Wechsler (Liu & Wechsler, 2002) use the Gabor wavelet (or filter) representation of face images to achieve robustness to illumination changes. Their method - the Gabor Fisher Classifier (GFC), adopts a filter bank of forty Gabor filters (featuring filters of five scales and eight orientations) to derive an augmented feature vector of Gabor magnitude features and then applies a variant of the multi-class linear discriminant analysis called the Enhanced Fisher discriminant Model (EFM) to the constructed Gabor feature vector to improve the vector's compactness.

While all of the presented techniques ensure some level of illumination invariance, Gabor wavelet based methods received the most attention due to their effectiveness and simplicity.

The feasibility of Gabor wavelet based methods for robust face recognition is not evidenced solely by the large number of papers following up on the work of Liu and Wechsler (e.g., Shen et al., 2007; Štruc & Pavešić, 2009b), but also by the results of several independent evaluations (and competitions) of face recognition technology, where the techniques utilizing Gabor wavelets regularly resulted in the best performance (Messer et al., 2006; Poh et al. 2009).

It has to be noted that the Gabor face representation as proposed by Liu and Wechsler does not represent an illumination invariant face representation, but rather exhibits robustness to illumination changes due to the properties of the deployed Gabor filter bank. Since Gabor filters represent band limited filters, the filter bank is usually constructed in such a way that it excludes the frequency bands most affected by illumination variations. Furthermore, the Gabor magnitude features that constitute the augmented Gabor feature vector are local in nature, which again adds to the illumination insensitiveness of the computed Gabor face representation.

While the existing Gabor based methods are among the most successful techniques for face recognition, they still exhibit some shortcomings, which, when properly solved, could result in an improved recognition performance. These shortcomings can be summarized into the following main points:

- most of the existing techniques rely solely on Gabor magnitude information while discarding the potentially useful Gabor phase information (e.g., Liu & Wechsler, 2002; Liu, 2006; Shen et al., 2007; Štruc & Pavešić, 2009b),
- the deployment of a large filter bank (usually comprising 40 filters) results in an inflation of the data size by a factor equaling the number of filters in the filter bank triggering the need for down-sampling strategies, which often discard information that could prove useful for the recognition task and
- Gabor magnitude features ensure only partial insensitiveness to illumination changes resulting in the necessity for additional (robust) face descriptors.

To tackle the above issues, we propose in this chapter a novel face representation called the oriented Gabor phase congruency pattern (OGPCP), which, as the name suggests, is derived from the Gabor phase congruency model presented in (Kovesi, 1999). The proposed face representation is based on the phase responses of the Gabor filter bank rather than the magnitude responses and as such offers an alternative to the established Gabor magnitude based methods. As we will show in this chapter, the feature vector constructed from the OGPCPs is more compact (i.e., less dimensional) than the traditional Gabor magnitude representation of face images and also exhibits robustness to illumination changes. Thus, it represents a novel robust face representation capable of substituting or complementing the existing Gabor magnitude based recognition techniques.

The rest of the chapter is structured as follows: In Section 2, a brief review of the Gabor filter based methods is given. In Section 3, the novel face representation, i.e., the oriented Gabor phase congruency pattern is presented and the augmented Gabor phase congruency feature vector introduced. In Section 4, the classification rule for the experiments is highlighted, while the experimental databases are described in Section 5. The feasibility of the proposed features is assessed in Section 6. The chapter concludes with some final remarks and directions for future work in Section 7.

2. Review of Gabor filters for face recognition

2.1 Gabor filter construction

Gabor filters (sometimes also called Gabor wavelets or kernels) have proven to be a powerful tool for facial feature extraction. They represent band-limited filters with an optimal localization in the spatial and frequency domains. Hence, when used for facial feature extraction, they extract multi-resolutional, spatially local features of a confined frequency band. In the spatial domain, the family of 2D Gabor filters can be defined as follows (Lades et al., 1993; Liu & Wechsler, 2002; Liu, 2006; Shen & Bai, 2006; Štruc & Pavešić, 2009b):

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

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 $\theta_v = v\pi/8$. Each filter represents a Gaussian kernel function modulated by a complex plane wave whose centre frequency and orientation are defined by the parameters f_u and θ_v, respectively. The parameters γ and η determine the ratio between the centre frequency and

the size of the Gaussian envelope and, when set to a fixed value, these parameters ensure that Gabor filters of different scales and a given orientation behave as scaled versions of each other. It has to be noted at this point that with fixed values of the parameters γ and η the scale of the given Gabor filter is defined uniquely by its centre frequency f_u. Commonly the values of the parameters γ and η are set to $\gamma = \eta = \sqrt{2}$. The last parameter of the Gabor filters f_{max} denotes the maximum frequency of the filters and is commonly set to $f_{max} = 0.25$. When employed for facial feature extraction, researchers typically use Gabor filters with five scales and eight orientations, i.e., $u = 1, 2, ..., p - 1$ and $v = 1, 2, ..., r - 1$, where $p = 5$ and $r = 8$, which results in a filter bank of 40 Gabor filters (Liu, 2006; Shen et al. 2007; Štruc et al., 2008a).

It should be noted that Gabor filters represent complex filters which combine an even (cosine-type) and odd (sine-type) part (Lades et al., 1993). An example of both filter parts in 3D is shown on the left side of Fig. 1, while the real parts of the entire filter bank (commonly comprising 40 Gabor filters) are presented in 2D on the right hand side of Fig. 1.
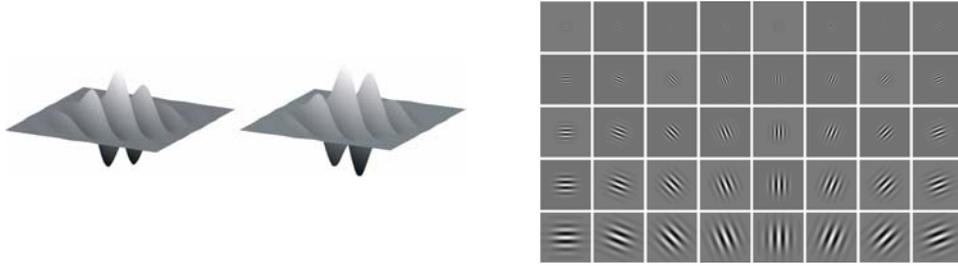

Fig. 1. Examples of Gabor filters: the real and imaginary part of a Gabor filter in 3D (left), the real part of the commonly employed Gabor filter bank of 40 Gabor filters in 2D (right)

2.2 Feature extraction with Gabor filters

Let $I(x, y)$ denote a grey-scale face image of size $a \times b$ pixels and let $\psi_{u,v}(x, y)$ represent a Gabor filter defined by its centre frequency f_u and orientation θ_v. The filtering operation or better said the feature extraction procedure can then be written as the convolution of the face image $I(x, y)$ with the Gabor wavelet (filter, kernel) $\psi_{u,v}(x, y)$, i.e. (Štruc & Pavešić, 2009b),

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

In the above expression, $G_{u,v}(x, y)$ represents the complex convolution output which can be decomposed into its real (or even) and imaginary (or odd) parts as follows:

$$E_{u,v}(x, y) = Re[G_{u,v}(x, y)] \text{ and } O_{u,v}(x, y) = Im[G_{u,v}(x, y)]. \tag{3}$$

Based on these results, both the phase ($\phi_{u,v}(x, y)$) as well as the magnitude responses ($A_{u,v}(x, y)$) of the filter can be computed, i.e.:

$$A_{u,v}(x,y) = \sqrt{E_{u,v}^2(x,y) + O_{u,v}^2(x,y)}$$
$$\phi_{u,v}(x,y) = \arctan(O_{u,v}(x,y)/E_{u,v}(x,y)) \tag{4}$$

As already stated in the previous section, most of the techniques found in the literature discard the phase information of the filtering output and retain only the magnitude information for the Gabor face representation. An example of the this information (in image form) derived from a sample face image is shown in Fig. 2.
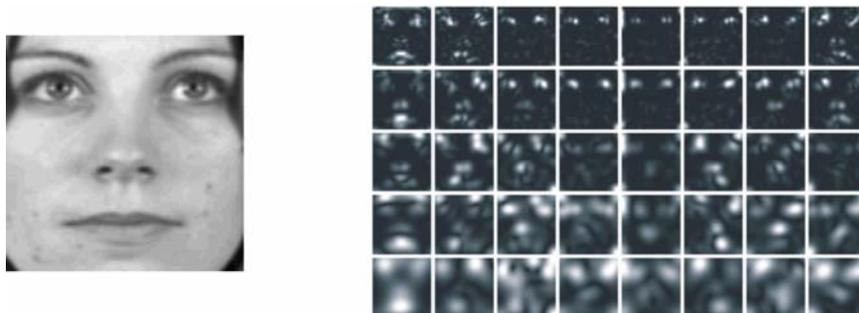

Fig. 2. An example of the Gabor magnitude output: a sample image (left), the magnitude output of the filtering operation with the entire Gabor filter bank of 40 Gabor filters (right)

2.3 The Gabor (magnitude) face representation

The first step when deriving the Gabor (magnitude) face representation of facial images is the construction of the Gabor filter bank. As we have pointed out several times in this chapter, most existing techniques adopt a filter bank comprised of 40 Gabor filters, i.e., Gabor filters with 5 scales ($u = 1, 2, \ldots, 5$) and 8 orientations ($v = 1, 2, \ldots, 8$).
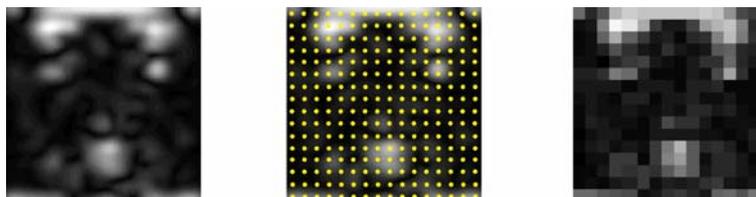

Fig. 3. Down-sampling of a magnitude filter response: an example of a magnitude response (left), an example of the magnitude response with a superimposed sampling grid (middle), a down-sampled magnitude response (right)

To obtain the Gabor (magnitude) face representation, a given face image is filtered with all 40 filters from the filter. However, even for a small image of 128×128 pixels, the magnitude responses of the filtering outputs comprise a pattern vector with 655360 elements, which is far too much for efficient processing and storage. To overcome this problem, down-sampling strategies are normally exploited. The down-sampling techniques reduce the dimensionality of the magnitude responses, unfortunately often at the expense of potentially useful discriminatory information. A popular down-sampling strategy is to

employ a rectangular sampling grid (as shown in Fig. 3) and retain only the values under the grid's nodes. This down-sampling procedure is applied to all magnitude responses, which are then normalized to zero mean and unit variance and ultimately concatenated into the final Gabor (magnitude) face representation or, as named by Liu and Wechsler (Liu & Wechsler, 2002), into the augmented Gabor feature vector.

If we denote the down-sampled magnitude response (in column vector form) of the Gabor filter of scale u and orientation v as $\boldsymbol{g}_{u,v}$, we can define the augmented Gabor (magnitude) feature vector $\boldsymbol{x}$ as follows:

$$\boldsymbol{x} = \left(\boldsymbol{g}_{0,0}^T, \boldsymbol{g}_{0,1}^T, \boldsymbol{g}_{0,2}^T, \dots, \boldsymbol{g}_{4,7}^T\right)^T. \tag{5}$$

It has to be noted that in the experiments presented in Section 6 of this chapter, we use images of 128×128 pixels and a rectangular down-sampling grid with 16 horizontal and 16 vertical lines, which corresponds to a down-sampling factor of 64. Nevertheless, even after the down-sampling, the augmented Gabor (magnitude) feature vector still resides in a very high-dimensional space (Shen et al., 2007) - in our case the dimensionality of the vectors still equals 10240. To make the processing more efficient, researchers commonly turn to so-called subspace projection techniques, e.g. (Liu, 2006; Shen et al. 2007; Štruc & Pavešić, 2009b). Two of these techniques, namely, the Principal Component Analysis (PCA) and the Linear Discriminant analysis (LDA), will also be adopted for our experiments. The description of these techniques is beyond the scope of this chapter, the reader is, however, referred to (Turk & Pentland, 1991) and (Belhumeur et al. 1997) for more information on PCA and LDA, respectively.

3. The Gabor (phase) face representation

This section introduces a novel face representation called oriented Gabor phase congruency pattern (OGPCP) and, consequently, the augmented Gabor phase congruency feature vector.

3.1 Background

Before we turn our attention to the novel representation of face images, i.e., to the oriented Gabor phase congruency pattern, let us take a closer look at why the Gabor phase information is commonly discarded when deriving the Gabor face representation.

Unlike the (Gabor) magnitude, which is known to vary slowly with the spatial position, the (Gabor) phase can take very different values even if it is sampled at image locations only a few pixels apart. This fact makes it difficult to extract stable and discriminative features from the phase responses of Eq. (2) and is the primary reason that most of the existing methods use only the (Gabor) magnitude to construct the Gabor feature vector (Zhang et al., 2007; Štruc et al., 2008a).

To the best of our knowledge, there are only a few studies in the literature that successfully derived useful features from Gabor phase responses for the task of face recognition, e.g., (Zhang et al., 2007; Bezalel & Efron, 2005; Gundimada & Asari, 2006; Gundimada et al., 2009). A common characteristic of these methods is the fact that they use features or face representations derived from the Gabor phase information rather than the "raw" phase

responses themselves or combine the phase information with other face descriptors to compensate for the variability of the Gabor phase.

Zhang et al. (Zhang et al., 2007), for example, adopt local histograms of the phase responses encoded via local binary patterns (LBPs) as face image descriptors and show that over small image regions the Gabor phase patterns exhibit some kind of regularity (in terms of histograms) and, hence, contain useful information for the task of face recognition. Other authors (e.g., Bezalel & Efron, 2005; Gundimada & Asari, 2006; Gundimada et al., 2009) incorporate the Gabor phase information by adopting the 2D phase congruency model of Kovesi (Kovesi, 1999) to detect edges in a given face image and deploy the resulting "edge" image for detection of interest points that are used with other image descriptors, such as Gabor magnitude features.

The face representation proposed in this chapter is related to the work of (e.g., Bezalel & Efron, 2005; Gundimada & Asari, 2006; Gundimada et al., 2009) 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 proposes a face representation that is only partially based on Kovesi's 2D phase congruency model and employs the proposed representation for recognition rather than solely for feature selection. As will be shown in the next section, the proposed face representation exhibits several desirable properties which overcome most of the shortcomings of the existing Gabor magnitude based methods.

3.2 The oriented Gabor phase congruency patterns

The original 2D phase congruency model as proposed by Kovesi in (Kovesi, 1999) was developed with the goal of robust edge and corner detection in digital images. However, as we will show, it can (though with a few modifications) also be used to encode phase information of the Gabor filter responses in a way that is useful for the task of face recognition.
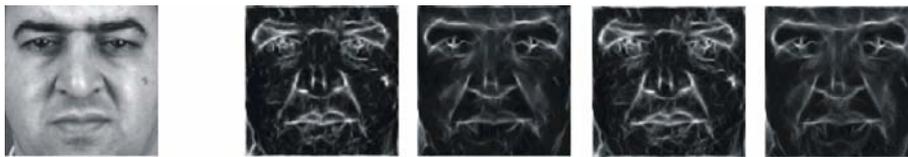

Fig. 4. Examples of phase congruency images (from left to right): the original image, the PCI for $p = 3$ and $r = 6$, the PCI for $p = 5$ and $r = 6$, the PCI for $p = 3$ and $r = 8$, the PCI for $p = 5$ and $r = 8$

Kovesi's original phase congruency model searches for points in an image where the log-Gabor filter responses over several scales and orientations are maximally in phase (Kovesi, 1999; Štruc & Pavešić, 2009a). Thus, a point in an image is of significance only if the phase responses of the log-Gabor filters over a range of scales (i.e., frequencies) display some kind of order. In the original approach, phase congruency is first computed for each of the employed filter orientations, while the results are then combined to form the final phase congruency image (PCI). Some examples of such images obtained with log-Gabor filters with p scales and r orientations are shown in Fig. 4. Note that the code used to produce the presented phase congruency images was provided by P. Kovesi and can be found at his homepage: http://www.csse.uwa.edu.au/~pk/Research/MatlabFns/index.html

While the presented approach is suitable for robust (in terms of noise, illumination variations and image contrast) edge and corner detection, its usefulness for facial feature extraction is questionable. As it was emphasized by Liu in (Liu, 2006), a desirable characteristic of any feature extraction procedure is the capability of extracting multi-orientational features. Rather than combining phase congruency information computed over several orientations, and using the result for construction of the facial feature vector, we therefore propose to compute an oriented Gabor phase congruency pattern (OGPCP) for each of the employed filter orientations and to construct an augmented (Gabor) phase congruency feature vector based on the results. Note that, differently from the original 2D phase congruency model proposed in (Kovesi, 1999), we use conventional Gabor filters as defined by Eq. (1) rather than log-Gabor filters (Štruc et al., 2008a).

Taking into account the original definition of the phase congruency, we can derive an oriented version of the phase congruency, which, when presented in image form, reveals the oriented Gabor phase congruency patterns (OGPCPs) for the v-th orientation:

$$OGPCP_v(x,y) = \frac{\sum_{u=1}^{p-1} A_{u,v}(x,y)\Delta\phi_{u,v}(x,y)}{\sum_u^{p-1}(A_{u,v}(x,y) + \epsilon)}, \tag{6}$$

where ϵ denotes a small constant that prevents division by zero and $\Delta\phi_{u,v}(x,y)$ stands for the phase deviation measure of the following form:

$$\Delta\phi_{u,v}(x,y) = \cos\big(\phi_{u,v}(x,y) - \bar{\phi}_v(x,y)\big) - \big|\sin(\phi_{u,v}(x,y) - \bar{\phi}_v(x,y))\big|. \tag{7}$$

In the above expression $\phi_{u,v}(x,y)$ denotes the phase angle of the Gabor filter (with a centre frequency f_u and orientation θ_v) at the spatial location (x,y), while $\bar{\phi}_v(x,y)$ represents the mean phase angle at the v-th orientation. Several examples of the OGPCPs for a sample image are shown in Fig. 5.

Fig. 5. Examples of OGPCPs (from left to right): the original image, the OGPCP for $\theta_v = 0°$ and $p = 2$, the OGPCP for $\theta_v = 0°$ and $p = 3$, the OGPCP for $\theta_v = 0°$ and $p = 4$, the OGPCP for $\theta_v = 0°$ and $p = 5$

Kovesi showed that the expression given in (6) can be computed directly from the filter outputs defined by (3); however, for details on computing the OGPCPs the reader should refer to the original paper (Kovesi, 1999).

It should be noted that the OGPCPs as defined by Eq. (6) represent illumination invariant (and contrast independent) face representations, since they do not depend on the overall magnitude of the filter responses. This property makes the OGPCPs a very useful image representation for face recognition.

3.3 The augmented Gabor phase congruency feature vector

The OGPCPs presented in the previous section form the foundation for the augmented Gabor phase congruency feature vector, which is computed from a given face image by the following procedure:

- for a given face image all r OGPCPs are computed for the chosen number of filter scales p (an example of all OGPCPs for a sample image with $r = 8$ and $p = 2$ is presented in Fig. 6.)
- the OGPCPs are down-sampled by a down-sampling factor ρ (in a similar manner as shown in Fig. 3),
- the down-sampled OGPCPs are normalized using the selected normalization procedure (zero mean and unit variance, histogram equalization, …), and
- the down-sampled and normalized OGPCPs in column vector form (denoted as $\boldsymbol{D}_v$) are concatenated to form the augmented Gabor phase congruency feature vector $\boldsymbol{x}$.

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

$$\boldsymbol{x} = \left(\boldsymbol{D}_0^T, \boldsymbol{D}_1^T, \boldsymbol{g}_2^T, \dots, \boldsymbol{g}_{r-1}^T\right)^T, \tag{8}$$

where T denotes the transform operator and $\boldsymbol{D}_v$, for $v = 1, 2, \dots, r-1$, stands for the vector derived from the OGPCP at the v-th orientation.
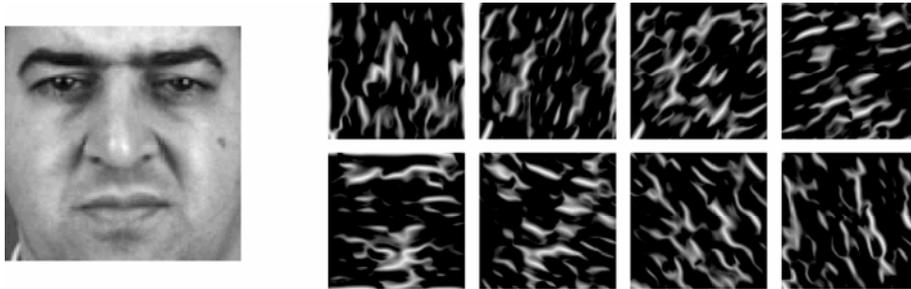

Fig. 6. An example of all OGPCPs: the original image (left), the OGPCPs (for $r = 8$) that form the foundation for construction of the augmented Gabor phase congruency feature vector

Note that in the experiments presented in Section 6 a down-sampling factor of $\rho = 16$ was used for the OGPCPs, as opposed to the Gabor magnitude responses, where a down-sampling factor of $\rho = 64$ was employed. This setup led to similar lengths of the constructed (Gabor) feature vectors of both methods and thus enabled a fair comparison of their face recognition performances. Furthermore, as the smaller down-sampling factor was used for the OGPCPs, less potentially useful information is discarded when oriented Gabor phase congruency patterns are employed for the face representation rather than the Gabor magnitude features.

As with the augmented Gabor magnitude feature vector, the Gabor phase congruency feature vector also resides in high-dimensional space and, hence, requires additional processing with, for example, subspace projection techniques to further reduce its dimensionality. In the experimental section, PCA and LDA are evaluated for this purpose.

4. Classification rule

In general, a face recognition system can operate in one of two modes: in the verification or identification mode (Štruc et al., 2008b).

When operating in the verification mode, the goal of the system is to determine the validity of the identity claim uttered by the user currently presented to the system. This is achieved by comparing the so-called "live" feature vector y extracted from the given face image of the user with the template corresponding to the claimed identity. Based on the outcome of the comparison, the identity claim is rejected or accepted. The verification procedure can formally be written as follows: given the "live" feature vector y and a claimed identity C_i associated with a user-template y_i, where $i \in \{1, 2, \ldots, N\}$ and N stands for the number of enrolled users, determine the validity of the identity claim by classifying the pair (y, C_i) into one of two classes w_1 or w_2 (Jain et al., 2004):

$$(y, C_i) \in \begin{cases} w_1, & if \ d(y, y_i) \geq \Delta \\ w_2, & otherwise \end{cases}, \tag{9}$$

where w_1 stands for the class of genuine identity claims, w_2 denotes the class of impostor identity claims, $d(\cdot, \cdot)$ denotes a function measuring the similarity of its arguments. In our case the similarity function takes the form of the cosine similarity measure, i.e.,

$$d(y, y_i) = \frac{y^T y_i}{\sqrt{y^T y y_i^T y_i}} \tag{10}$$

and Δ represents 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 which best matches the "live" feature vector. This can be formalized as follows: given a "live" feature vector y and a database containing N templates $y_1, y_2, \ldots, y_N$ of the enrolled users (or identities) $C_1, C_2, \ldots, C_N$, determine the most suitable identity, i.e., (Jain et al., 2004):

$$y \in \begin{cases} C_i, & if \ \max_i d(y, y_i) \geq \Delta, \ i = 1, 2, \ldots, N \\ C_{N+1}, & otherwise \end{cases}, \tag{11}$$

where $d(y, y_i)$ 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 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 **y.** It should be noted that, in the experiments presented in the experimental section, the user templates are constructed as the mean vectors of the feature vectors extracted from the enrollment face images of the users.

5. The databases and experimental configurations

This section presents the experimental databases used to determine the feasibility of the proposed augmented phase congruency feature vectors for face recognition. It commences by describing the two face databases, namely, the XM2VTS (Messer et al., 1999) and the Extended YaleB database (Georghiades et al., 2001; Lee et al., 2005), and proceeds by presenting the pre-processing procedure applied to the experimental images prior to the actual experiments.

5.1 The XM2VTS database
The XM2VTS database comprises a total of 2360 facial images that correspond to 295 distinct subjects (Messer et al., 1999). The images were recorded in controlled conditions (in front of a homogenous background, with artificial illumination, with frontal pose and a neutral facial expression, etc.), during four recording sessions and over a period of approximately five months. At each of the recording session, two recordings were made resulting in eight facial images per subject that are featured in the database. Since the time elapsed between two successive sessions was around one month, the variability in the images is mainly induced by the temporal factor. Thus, images of the same subject differ in terms of hairstyle, presence or absence of glasses, make-up and moustaches, etc. Some examples of the images from the XM2VTS database are shown in Fig. 7.

Fig. 7. Sample images from the XM2VTS database

The face verification experiments on the XM2VTS were conducted in accordance with the first configuration of experimental protocol associated with the database, known also as the Lausanne protocol (Messer et al., 1999). Following the protocol, the subjects of the database

were divided into groups of 200 clients and 95 impostors. Images corresponding to the subjects in these two groups were then partitioned into image sets used for:

- training and enrolment (3 images for each of the 200 clients) – this image set was used for training of the feature extraction techniques and for building client models/templates in the form of mean feature vectors,
- evaluation (3 images for each of the 200 clients and 8 images for each of the 25 evaluation impostors) – this image set was employed to determine the decision threshold for a given operating point of the face verification system and to estimate any potential parameters of the feature extraction techniques, and
- testing (2 images for each of the 200 clients and 8 images for each of the 70 test impostors) – this image set was used to determine the verification performance in real operating conditions (i.e., with predetermined parameters)

While the first image set featured only images belonging to the client group, the latter two image sets comprised images belonging to both the client and the impostor groups. The client images were employed to assess the first kind of error a face verification system can make, namely, the false rejection error, whereas the impostor images were used to evaluate the second type of possible verification error, namely, the false acceptance error. The two errors are quantified by two corresponding error rates: the false rejection and false acceptance error rates (FRR and FAR), which are defined as the relative frequency with which a face verification system falsely rejects a client- and falsely accepts an impostor-identity-claim, respectively. To estimate these error rates each feature vector extracted from an image of the client group was matched against the corresponding client template, while each of the feature vectors extracted from an impostor image was matched against all client templates in database. The described setup resulted in the following verification experiments: 600 client verification attempts in the evaluation stage, 40000 impostor verification attempts in the evaluation stage, 400 client verification attempts in the test stage and 112000 impostor verification attempts in the test stage (Messer et al., 1999; Štruc et al. 2008).

It has to be noted that there is a tradeoff between the FAR and FRR. We can select an operating point (determined by the value of the decision threshold) where the FAR is small and the FRR is large or vice versa, we can choose an operating point with a small FRR but at the expense of a large FAR. To effectively compare two face verification systems, an operating point that ensures a predefined ratio of the two error rates has to be selected on the evaluation image set or the values of the error rates must be plotted against various values of the decision threshold, resulting in the so-called performance curves. In this chapter we choose the latter approach and present our results in terms of two kinds of performance curves, namely, the Detection Error Tradeoff (DET) curves and the Expected Performance Curves (EPC), which plot the FAR against the FRR at different values of the decision threshold on the evaluation and test sets, respectively.

5.2 The Extended YaleB database

The Extended YaleB database was recorded at the Yale University and comprises 2432 frontal face images of 38 distinct subjects (Georghiades et al., 2001; Lee et al., 2005). It exhibits large variations in illumination, which is also the main source of variability in the

images of the Extended YaleB database employed in our experiments. Some examples of these images are shown in Fig. 8.

Fig. 8. Sample images from the Extended YaleB database

After removing a number of corrupt images from the database, a total of 2414 frontal face images with variable lighting were available for our experiments with each subject of the database being accounted for with a little more than 60 images. These images were then partitioned into 5 image subsets according to the extremity of illumination in the images, as proposed by Georghiades et al. in (Georghiades et al., 2001). The reader is referred to the original publication for more information on the partitioning.

The first image subset (denoted as S1 in the remainder) featured images captured in relatively good illumination conditions, while the conditions got more extreme for the image subsets two (S2) to five (S5). It should also be noted that the subsets did not contain the same number of images. The first subset, for example, contained 263 images, which corresponds to approximately 7 images per subject. The second subset contained 456 images, the third 455 images, the fourth 526 images and finally the fifth subset contained 714 facial images.

For our experiments we adopted the first subset for the training of the feature extraction techniques as well as for creating the user models/templates, and employed all remaining subsets for testing. Such an experimental setup resulted in highly miss-matched conditions for the recognition technique, since the test subsets featured images captured under varying illumination conditions, while the training images were acquired in controlled illumination conditions. Clearly, for a feature extraction technique to be successful, it has to extract stable features from the images regardless of the conditions present during the image acquisition stage. Furthermore, the experimental configuration is also in accordance with real life settings, as the training and enrollment stages are commonly supervised and, hence, the training and/or enrollment images are usually of good quality. The actual operational conditions, on the other hand, are typically unknown in advance and often induce severe illumination variations.

The results of our experiments on the Extended YaleB database are reported in terms of the rank one recognition rate, which corresponds to the relative frequency with which the test images from a given subset are recognized correctly.

5.3 Data pre-processing

Prior to the experiments, we subjected all images from both databases to a pre-processing procedure comprised of the following steps:

- a conversion of the (colour) face images to 8-bit monochrome (grey-scale) images – applicable only for the XM2VTS database,
- a geometric normalization procedure, which, based on the manually marked eye coordinates, rotated and scaled the images in such a way that the centres of the eyes were aligned and, thus, located at predefined positions,
- a cropping procedure that cropped the facial region of the images to a standard size of 128×128 pixels,
- a photometric normalization procedure that first equalized the histogram of the cropped facial images and then further normalized the results to zero-mean and unit-variance.

It should be noted that manual labelling of the facial landmarks is the only way to achieve a fair comparison of the recognition techniques, as it ensures that the differences in the observed recognition performances are only a consequence of the employed feature extraction techniques and not other influencing factors. Some examples of the pre-processed images (prior to photometric normalization) from the two databases are shown in Fig. 9.

Fig. 9. Examples of pre-processed images from the XM2VTS (left quadruple of images) and Extended YaleB (right quadruple of images) databases

6. Experiments, results and discussion

6.1 Baseline performance

In the first series of our recognition experiments, we aimed at determining the performance of some baseline face recognition techniques on the two test databases. To this end, we implement the popular Principal Component Analysis (PCA) (Turk & Pentland, 1991) and Linear Discriminant Analysis (LDA) (Belhumeur et al., 1997) techniques, also known as the Eigenface and Fisherface methods, and assess the techniques for different lengths (i.e., different Number Of Features - NOF) of the PCA and LDA feature vectors. The results of this assessment are presented in Fig. 10 for the XM2VTS database in the form of DET curves and in Table 1 for the Extended YaleB database (EYB) in the form of rank one recognition rates (in %). Considering the number of subjects and images in the databases, the maximum length of the feature vector for the PCA technique equals 599 for the XM2VTS database and 262 for the EYB database, while the maximum length for the LDA technique is 199 for the XM2VTS database and 37 for the EYB database.
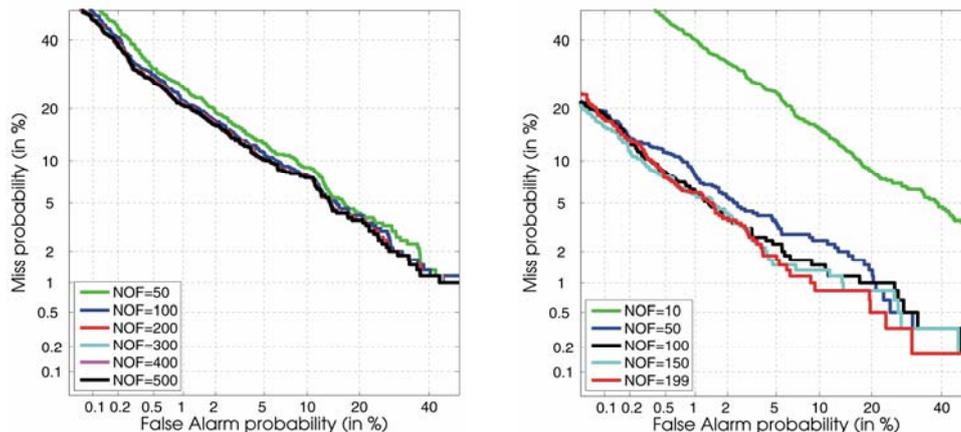
Fig. 10. DET curves of the baseline experiments on the evaluation image sets of the XM2VTS database: for the PCA technique (left), for the LDA technique (right)

NOF	PCA				NOF	LDA			
	S2	S3	S4	S5		S2	S3	S4	S5
10	56.6	29.5	11.2	15.6	5	98.3	56.9	9.9	13.6
50	93.4	54.7	16.7	21.9	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	55.0	16.7	22.0	30	100	99.3	53.6	47.6
200	93.6	55.0	16.7	22.0	37	100	99.8	56.3	51.0

Table 1. Rank one recognition rates (in %) for different lengths of the PCA and LDA feature vectors obtained on different subsets of the EYB database

Note that for the PCA technique the performance on the XM2VTS saturates when 200 features are used in the feature vectors. Similar results are also observed for the EYB database, where the performance on all subsets peaks with 150 dimensional feature vectors. For the LDA technique the best performance on both databases is achieved with the maximum number of features, i.e., 199 for the XM2VTS database and 37 for the EYB database. The presented experimental results provide a baseline face recognition performance on the two databases for the following comparative studies of the techniques using the augmented phase congruency feature vectors.

6.2 Baseline performance with the augmented phase congruency feature vector
In our second series of face recognition experiments we evaluate the performance of the PCA and LDA techniques in conjunction with the augmented phase congruency feature vectors and assess the relative usefulness of additional normalization techniques applied to the augmented feature vectors prior to the deployment of the subspace projection techniques PCA and LDA. We use five filter scales ($p = 5$) and eight orientations ($r = 8$) to construct the oriented phase congruency patterns, which in their down-sampled form

constitute the augmented phase congruency feature vectors, and apply the following normalization schemes to these vectors:

- after the down-sampling of the oriented Gabor phase congruency patterns, each down-sampled OGPCP is normalized to zero mean and unit variance prior to concatenation into the final augmented feature vector (denoted as ZMUV) – Fig. 11 (upper left corner),
- after the down-sampling of the oriented Gabor phase congruency patterns, each down-sampled OGPCP is first subjected to histogram equalization and then to zero mean and unit variance normalization prior to concatenation into the final augmented feature vector (denoted as oHQ) - Fig. 11 (upper right corner), and
- after the down-sampling of the oriented Gabor phase congruency patterns, the down-sampled OGPCPs are concatenated into a bigger image, which is subjected to the histogram equalization procedure and then to zero mean and unit variance normalization (denoted as HQ) - Fig. 11 (lower row).
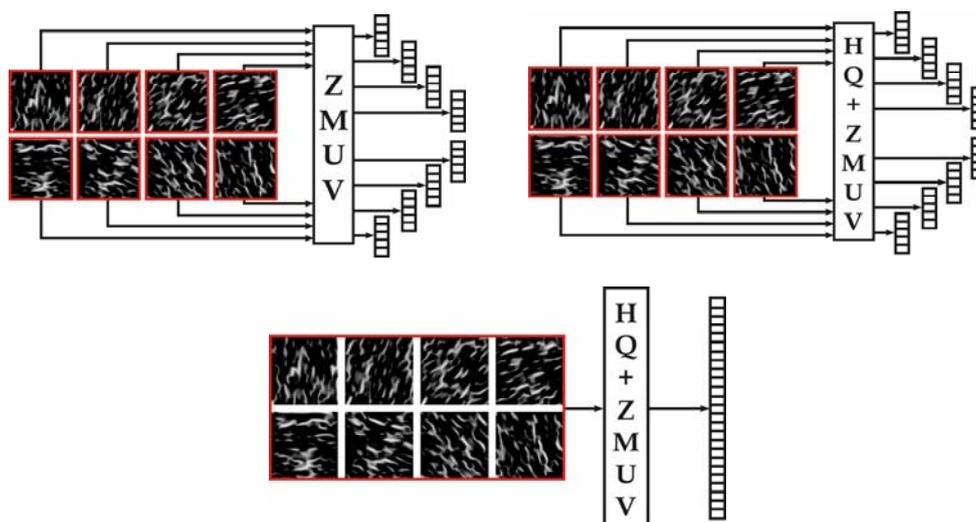

Fig. 11. Diagrams of the employed normalization procedures: ZMUV (upper left corner), oHQ (upper right corner), and HQ (lower row)

It should be noted that the number of scales and orientations that were used in our experiments, i.e., $p = 5$ and $r = 8$, was chosen based on other Gabor filter based methods presented in the literature – see, for example, (Liu & Wechsler, 2002; Shen et al., 2007; Štruc & Pavešić, 2009b). For the implementation of the subspace projection techniques the following feature vector lengths were chosen: 37 for LDA on EYB, 199 for LDA on XM2VTS, 150 for PCA on EYB and 200 for PCA on XM2VTS. These lengths were selected based on the baseline results from the previous series of experiments. However, since the number of features in the feature vectors is not the primary concern of this section, it could also be set differently.

The results of the experiments are again presented in the from of DET curves for the XM2VTS database in Fig. 12 and in the form of rank one recognition rates for the EYB database in Table 2.
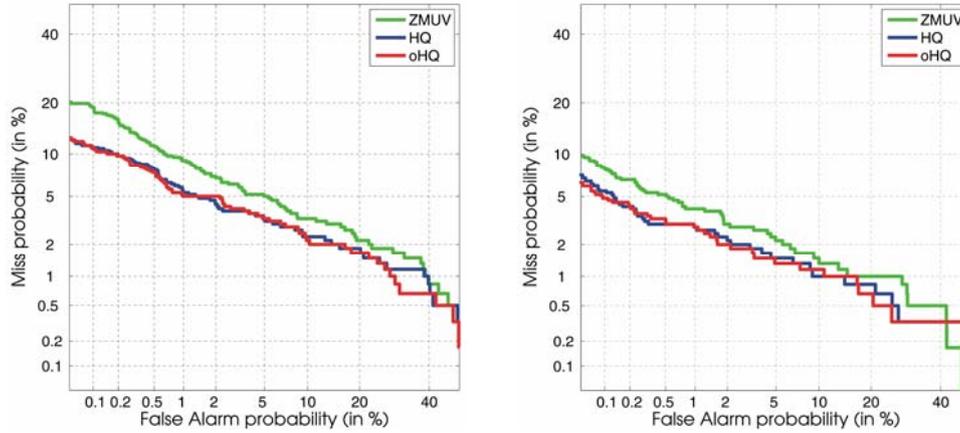

Fig. 12. DET curves of the comparative assessment of the normalization techniques on the evaluation image sets of the XM2VTS database: for the PCA technique (left), for the LDA technique (right)

Norm	PCA				Norm	LDA			
	S2	S3	S4	S5		S2	S3	S4	S5
ZMUV	100	99.1	83.4	92.7	ZMUV	100	99.8	88.8	93.8
HQ	100	99.1	81.6	89.8	HQ	100	100	86.1	94.8
oHQ	100	99.3	84.6	92.7	oHQ	100	100	87.1	94.8

Table 2. Rank one recognition rates (in %) for different normalization schemes of the augmented phase congruency vector prior to PCA and/or LDA deployment on different subsets of the EYB database

From the experimental results we can see that the traditional ZMUV technique resulted in the worst performance, while both the HQ and oHQ techniques achieved similar recognition rates on both databases. While the difference in their performance is statistically not significant, we nevertheless chose the oHQ technique for our following comparative assessments due to better results on the EYB database. Furthermore, if we compare the results obtained with the PCA and LDA techniques on the raw pixel data (Table 1) and the results obtained with the augmented feature vectors, we can see that the performance has improved significantly.

6.3 Impact of filter scales
In the third series of face recognition experiments, we assess the impact of the number of filter scales p in the Gabor filter bank on the performance of the PCA and LDA techniques applied to the augmented phase congruency feature vectors. We fix the angular resolution

of the filter bank to $r = 8$ and vary the value of the filter scales from $p = 2$ to $p = 5$. In all of the performed experiments we use the same dimensionality of the PCA and LDA feature vectors as in the previous section and adopt the oHQ technique for the normalization of the augmented feature vectors. We once more present the results of the described experiments in form of the DET curves for the XM2VTS database (Fig. 13) and in form of rank one recognition rates for the EYB database (Table 3).
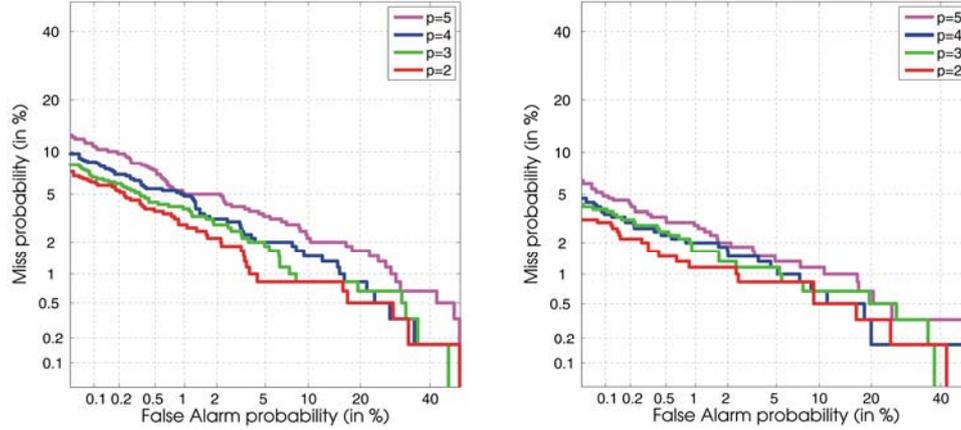

Fig. 13. DET curves generated for different numbers of filter scales employed during construction of the OGPCPs. The results were obtained on the evaluation image sets of the XM2VTS database: for the PCA technique (left), for the LDA technique (right)

p	PCA				p	LDA			
	S2	S3	S4	S5		S2	S3	S4	S5
5	100	99.3	84.6	92.7	5	100	100	87.1	94.8
4	100	100	91.8	94.8	4	100	100	94.5	94.4
3	100	100	93.4	95.2	3	100	100	96.4	96.4
2	100	100	93.0	92.3	2	100	100	94.7	96.6

Table 3. Rank one recognition rates (in %) on the EYB database for different numbers of filter scales employed during construction of the OGPCPs.

We can notice that on the XM2VTS database the verification performance steadily improves when the number of filter scales employed for the construction of the augmented phase congruency feature vector decreases. Thus, the best performance for the PCA (Fig. 13 left) as well as for the LDA (Fig. 13 right) techniques is observed with two filter scales, i.e., $p = 2$. Here, equal error rates of 2.15% and 1.16% are achieved for the PCA and LDA techniques, respectively.

Similar results are obtained on the EYB database. Here, the performance also increases with the decrease of used filter scales. However, the performance peaks with $p = 3$ filter scales. Since the improvements with the EYB database are not as pronounced as with the XM2VTS database, we chose to implement the construction procedure of the augmented phase congruency feature vector with 2 filter scales for the final comparative assessment.

6.4 Comparative assessment

In our last series of recognition experiments, we compare the performance of the PCA and LDA techniques on the proposed augmented phase congruency feature (PCF) vector with that of several established face recognition techniques from the literature. Specifically, we implement the following techniques for our comparative assessment: the Eigenface technique (PCA) (Turk & Pentland, 1991), The Fisherface technique (LDA) (Belhumeur et al., 1997), and the LDA and PCA techniques applied to the Gabor face representation (GF) proposed in (Liu & Wechsler, 2002).

All experiments on the XM2VTS database presented so far have been performed on the evaluation image sets, while the test image sets were not used. In this series of experiments we employ the test image sets for our assessment and implement all recognition techniques with all parameters (such as decision thresholds, feature vector lengths, number of employed filter scales, etc.) predefined on the evaluation image sets. Differently from the experiments presented in the previous sections, we do not present the results in the form of DET curves, but rather use the EPC curves. The choice of these performance curves is motivated by the work presented in (Bengio & Marithoz, 2004). Here, the authors argue that two recognition techniques cannot be compared fairly using DET curves, as in real life operating conditions a decision threshold has to be set in advance. In these situations the actual operating point may differ from the operating point the threshold was set on. To overcome this problem, the authors proposed the EPC curves, which plot the half total error rate (HTER=0.5(FAR+FRR)) against the parameter α, which controls the relative importance of the two error rates FAR and FRR in the expression: α FAR + $(1 - \alpha)$FRR. To produce the EPC curves, an evaluation image set and a test image set are required. For each α the decision threshold that minimizes the weighted sum of the FAR and FRR is computed on the evaluation image set. This threshold is then used on the test images to determine the value of the HTER used for the EPC curves.

Fig. 14. Examples of modified face images (from left to right): the original image, the modified image for $\tau = 40$, the modified image for $\tau = 80$, the modified image for $\tau = 120$, the modified image for $\tau = 160$

To make the final assessment more challenging, we introduce an artificial illumination change to the test sets of the XM2VTS database. To this end, we adopt the model previously employed in (Sanderson & Paliwal, 2003), which simulates different illumination conditions during the image acquisition stage by modifying the pre-processed face images $I(x, y)$, i.e.,

$$\tilde{I}(x, y) = I(x, y) + my + \tau, \tag{12}$$

where $x = 0, 1, \ldots, a - 1$; $y = 0, 1, \ldots, b - 1$; $m = -2\tau/(b - 1)$ and τ denotes the parameter that controls the "strength" of the introduced artificial illumination change. Sanderson and

Paliwal (Sanderson & Paliwal, 2003) emphasized that this model does not cover all illumination effects possible in real life settings, but is nevertheless useful for providing suggestive results. Some examples of the modified face images $\tilde{I}(x, y)$ obtained with the presented model for different values of the parameter τ are shown in Fig. 14.

The results of the final assessment are presented in Fig. 15 for the XM2VTS database and in Table 4 for the EYB database.
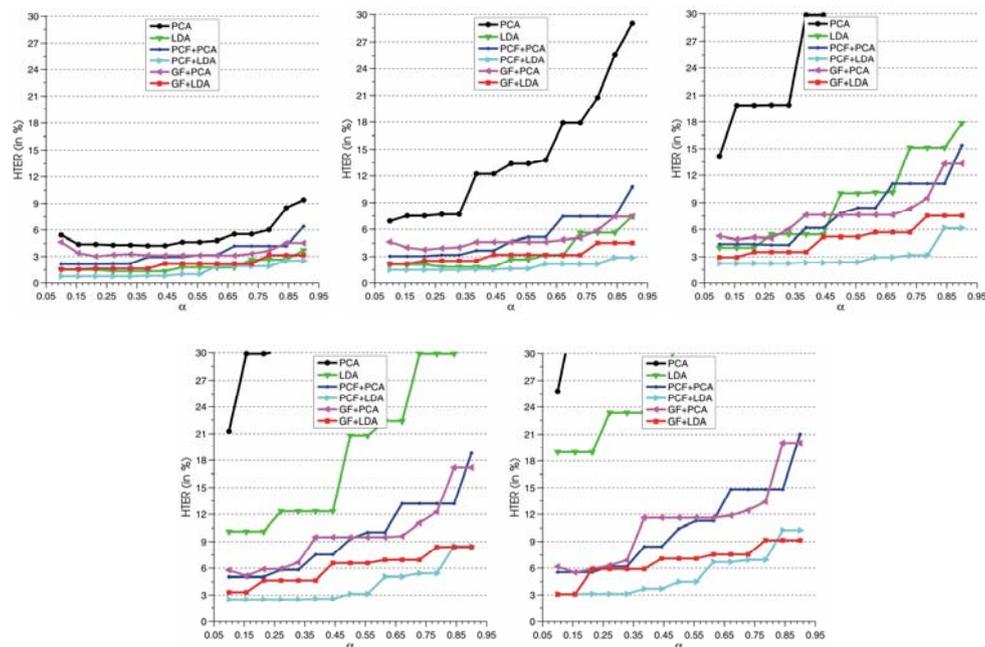

Fig. 15. EPC curves obtained on the test sets of the XM2VTS database for different values of the parameter τ (from left to right starting in the upper left row): with the original images, with the modified images with τ=40, with the modified images with τ=80, with the modified images with τ=120, with the modified images with τ=160,

Method	S2	S3	S4	S5
PCA	93.6	55.0	16.7	22.0
LDA	100	99.8	56.3	51.0
GF+PCA	100	97.8	77.9	85.2
GF+LDA	100	100	83.2	89.1
PCF+PCA	100	100	93.0	92.3
PCF+LDA	100	100	94.5	94.4

Table 4. Rank one recognition rates (in %) on the EYB for the comparative assessment

The first thing to notice from the presented results is that both the Gabor magnitude as well as the Gabor phase congruency features result in a significant improvement in the

recognition performance when compared to the raw pixel data and, furthermore, that both types of features result in a more robust performance in the presence of illumination changes. This fact is best exemplified by the recognition rates on subsets S4 and S5, where the increase in performance from the pixel data to the Gabor-based features is more than 60% (in absolute terms) for the PCA-based techniques and more than 25% (in absolute terms) for the LDA-based techniques.

In general, the augmented Gabor phase congruency feature vectors resulted in better performance in difficult illumination conditions than the Gabor magnitude features. While this improvement was only minimal on the XM2VTS database and its synthetically degraded versions, the results on the EYB database show improvements (on the image subset S4) of around 10% in absolute terms.

6.5 Discussion

From the experimental results presented in previous sections, we found that amongst the tested feature extraction techniques, the LDA technique combined with the Gabor magnitude and Gabor phase congruency features ensured the best recognition performance on all experimental databases (i.e., on the XM2VTS, the EYB and on the degraded versions of the XM2VTS databases). While both feature types significantly improved upon the techniques baseline performance with "raw" pixel intensity values, there are several differences in both feature types, which affect their usability in real-life face recognition systems.

First of all, as stated in a number of studies from the literature (e.g., Liu & Wechsler, 2002; Shen & Bai, 2006; Shen et al. 2007; Štruc & Pavešić, 2009b), the Gabor magnitude based methods require 40 Gabor filters, i.e., filters with five scales and eight orientations, to achieve their optimal performance. The same number of filters was also used in our experiments to obtain the performance presented in previous sections. The Gabor phase congruency features based methods presented in this chapter, on the other hand, require only 16 Gabor filters, filters with two scales and eight orientations, for an optimal performance. This fact makes the Gabor phase congruency methods significantly faster than the Gabor magnitude based methods.

Second of all, since there is only one output per employed filter orientation for the Gabor phase congruency based methods and not five, as it is the case with the Gabor magnitude based techniques, the increase in data is not that extreme for the proposed face representation.

Last but not least, we have to emphasize that in its optimized form (with two filter scales and eight orientations) the Gabor phase congruency techniques operate on a much narrower frequency band than the Gabor magnitude methods. Based on the experimental results presented in previous sections, we can in fact conclude that most of the discriminatory Gabor-phase information is contained in the OGPCPs obtained with Gabor filters of high frequencies ($u = 0,1$). In addition to the high frequency filters, the Gabor magnitude methods effectively also use the low frequency Gabor filters. This finding suggests that the Gabor phase congruency and Gabor magnitude features represent feature types with complementary information and could therefore be combined into a unified feature extraction technique which uses Gabor magnitude as well as Gabor phase information for face recognition.

7. Conclusion and future work

In this chapter we have proposed a novel face representation derived from the Gabor filter outputs. Unlike popular Gabor filter based methods, which mainly use only Gabor magnitude features for representing facial images, the proposed feature extraction technique exploits the Gabor phase information and derives the novel face representation named the Oriented Gabor Phase Congruency Pattern or OGPCP. This representation forms the foundation for the construction of the augmented Gabor phase congruency feature vector, which, similar to the established Gabor magnitude representations, can be combined with subspace projection techniques to form powerful and efficient feature extraction approaches. The feasibility of the proposed face representation (or features) was assessed on two publicly available datasets, namely, on the XM2VTS and on the Extended YaleB dataset. On both datasets, the proposed features resulted in a promising face recognition performance and outperformed popular face recognition techniques, such as PCA, LDA, the Gabor-Fisher classifier and others. The proposed features were shown to ensure robust recognition performance in the presence of severe illumination changes as well.

The future work with respect to the proposed Gabor phase congruency face representation, i.e., the OGPCP, will be focused on evaluating different strategies to combine the traditional Gabor magnitude face representation with the proposed Gabor phase congruency patterns of facial images.

8. References

Belhumeur, B.; Hespanha, J. & Kriegman, D. (1997). Eigenfaces vs. Fisherfaces: Recognition using Class specific Linear Projection. *IEEE Transactions on Pattern Analysis and Machine Intelligence,* Vol. 19, No. 7, 711-720, 0162-8828

Bengio, S. & Marithoz, J. (2004). The Expected Performance Curve: A New Assessment Measure for Person Authentication. *Proceedings of the Speaker and Language Recognition Workshop Oddyssey*, pp. 279-284

Bezalel, E. & Efron, U. (2005). Efficient Face Recognition Method Using a Combined Phase Congruency/Gabor Wavelet Technique. *Proceedings of the SPIE Conference on Optical Information Systems III*, doi:10.1117/12.618226, San Diego, CA, USA, August 2005, SPIE

Gao, Y. & Leung, M.K.H. (2002). Face Recognition Using Line Edge Map. *IEEE Transactions on Pattern Analysis and Machine Intelligence,* Vol. 24, No. 6, 764-779, 0162-8828

Georghiades, A.S.; Belhumeur, P.N. & Kriegman, D. (2001). 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, 643-660, 0162-8828

Gross, R.; Baker, S.; Metthews, I. & Kanade, T. (2004). Face Recognition Across Pose and Illumination, In: *Handbook of Face Recognition*, Li, S.Z. & Jain, A.K. (Ed.), 193-216, Springer, 0-387-40595-X, New York, USA

Gundimada, S. & Asari, V.K. (2006). 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, 25-31, 2070-3961

Gundimana, S.; Asari, V.K. & Gudur N. (2009). Face Recognition in Multi-sensor Images Based on a Novel Modular Feature Selection Technique. *Information Fusion*, article in press, 1566-2535

Jain, A.K.; Ross, A. & Prabhakar, S. (2004). An Introduction to Biometric Recognition. *IEEE Transactions on Circuits and Systems for Video Technology*, Vol. 14, No. 1, 4-20, 1051-8215

Kovesi, B. (1999). Image Features from Phase Congruency. *Videre: Journal of Computer Vision Research*, Vol. 1, No. 3, 1-26

Lades, M.; Vorbruggen, J.; Buhmann, J.; Lange, J.; Malsburg, C. von der; Wurtz, R. & Konen, W. (1993). Distortion Invariant Object Recognition in the Dynamic Link Architecture. *IEEE Transactions on Computers*, Vol. 42, No. 3, 300-311, 0018-9340

Lee, K.C.; Ho, J. & Kriegman, D. (2005). Acquiring Linear Subspaces for Face Recognition under Variable Lighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence,* Vol. 27, No. 5, 684-698, 0162-8828

Liu, C. & Wechsler, H. (2002). Gabor Feature Based Classification using the Enhanced Fisher Linear Discriminant Model for Face Recognition. *IEEE Transactions on Image Processing*, Vol. 11, No. 4, 467-476, 1057-7149

Liu, C. (2006). Capitalize on Dimensionality Increasing Techniques for Improving Face Recognition Grand Challenge Performance. *IEEE Transactions on Pattern Analysis and Machine Intelligence,* Vol. 28, No. 5, 725-737, 0162-8828

Messer, K.; Matas, J.; Kittler, J.; Luettin, J. & Maitre, G. (1999). XM2VTSDB: The Extended M2VTS Database. *Proceedings of the 2nd International Conference on Audio- and Video-Based Person Authentication*, pp. 72-77, Washington D.C., USA

Messer, K.; Kittler, J.; Short, J.; Heusch, G.; Cardinaux, F.; Marcel, S.; Rodriguez, Y.; Shan, S.; Su, Y.; Gao, W. & Chen, X. (2006). Performance Characterisation of Face Recognition Algorithms and Their Sensitivity to Severe Illumination Changes. *Proceedings of the IAPR International Conference on Biometrics*, pp. 1-11, 978-3-540-31111-9, Hong Kong, China, January 2006, Springer, Berlin/Heidelberg

Phillips, P.J.; Scruggs, W.T.; O'Toole, A.J.; Flynn, P.J.; Bowyer, K.W.; Schott, C.L. & Sharpe, M. (2007). FRVT 2006 and ICE 2006 Large-Scale Results, *NISTIR 7408*

Poh, N.; Chan, C.H.; Kittler, J.; Marcel, S.; McCool, C.; Argones-Rua, E.; Alba-Castro, J.L.; Villegas, M.; Paredes, R.; Štruc, V.; Pavešić, N.; Salah, A.A.; Fang, H. & Costen, N. (2009). Face Video Competition. *Proceedings of the IAPR International Conference on Biometrics*, pp. 715-724, 978-2-642-01792-6, Alghero, Italy, June 2009, Springer, Berlin/Heidelberg

Sanderson, C. & Paliwal, K. (2003). Fast Features for Face Authentication under Illumination Direction Changes. *Pattern Recognition Letters*, Vol. 24, No. 14, 2409-2419, 0167-8655

Shen, L. & Bai, L. (2006). A Review of Gabor Wavelets for Face Recognition. *Pattern analysis and applications*, Vol. 9, No. 2, 273-292, 1433-7541

Shen, L.; Bai, L. & Fairhurst, M. (2007). Gabor Wavelets and General Discriminant Analysis for Face Identification and Verification. *Image and Vision Computing*, Vol. 25, No. 5, 553-563, 0262-8856

Short, J.; Kittler, J. & Messer, K. (2005). Photometric Normalisation for Face Verification, *Proceedings of the 5th International Conference on Audio- and Video-Based Person Authentication*, pp. 617-626, 978-3-540-27887-0, New York, USA, July 2005, Springer, Berlin/Heidelberg

Štruc, V.; Vesnicer, B. & Pavešić, N. (2008a). The Phase-based Gabor Fisher Classifier and its Application to Face Recognition under Varying Illumination Conditions. *Proceedings of the 2nd International Conference on Signal Processing and Communication Systems*, pp. 1-6, 978-1-4244-4242-3, Gold Coast, Australia, IEEE, NJ

Štruc, V.; Mihelič, F. & Pavešić, N. (2008b). Face Recognition using a Hybrid Approach. *Journal of Electronic Imaging*, Vol. 17, No. 1, 1-11, 1017-9909

Štruc, V. & Pavešić, N. (2009a). Phase-congruency Features for Palm-print Verification. *IET Signal Processing*, Vol. 3, No. 4, 258-268, 1751-9675

Štruc, V. & Pavešić, N. (2009b). Gabor-based Kernel-partial-least-squares Discrimination Features for Face Recognition. *Informatica (Vilnius)*, Vol. 20, No. 1, 115-138, 0868-4952

Turk, M. & Pentland, A. (1991). Eigenfaces for Recognition. *Journal of Cognitive Neuroscience*, Vol. 3, No. 1, 71-86, 0898-929X

Zhang, B; Shan, S., Chen, X. & Gao, W. (2007). Histogram of Gabor Phase Patterns (HGPP): A Novel Object Representation approach for Face Recognition. *IEEE Transactions on Image Processing*, Vol. 16, No. 1, 57-68, 1057-7149