

Face Verification using Gabor Filtering and Adapted Gaussian Mixture Models

Laurent El Shafey ^{†*}, Roy Wallace [†] and Sébastien Marcel [†]

[†]Idiap Research Institute
CH-1920 Martigny, Switzerland

^{*}Ecole Polytechnique Fédérale de Lausanne
CH-1015 Lausanne, Switzerland

laurent.el-shafey,roy.wallace,sebastien.marcel@idiap.ch

Abstract: The search for robust features for face recognition in uncontrolled environments is an important topic of research. In particular, there is a high interest in Gabor-based features which have invariance properties to simple geometrical transformations. In this paper, we first reinterpret Gabor filtering as a frequency decomposition into bands, and analyze the influence of each band separately for face recognition. Then, a new face verification scheme is proposed, combining the strengths of Gabor filtering with Gaussian Mixture Model (GMM) modelling. Finally, this new system is evaluated on the BANCA and MOBIO databases with respect to well known face recognition algorithms. The proposed system demonstrates up to 52% relative improvement in verification error rate compared to a standard GMM approach, and outperforms the state-of-the-art Local Gabor Binary Pattern Histogram Sequence (LGBPHS) technique for several face verification protocols on two different databases.

1 Introduction

Face recognition has been an active research area in the computer vision community over the last three decades [ZCPR03, LJ05]. It is indeed a very challenging task. Faces highly vary in size, location, shape and colors. Their overall appearance is also influenced by lighting conditions, facial expression, occlusion or facial features, such as beards, mustaches and glasses.

One of the first popular approaches was introduced by Turk and Pentland [TP91] who proposed a method based on Principal Component Analysis (PCA). This technique, referred to as Eigenfaces, has contributed to the emergence of other holistic approaches such as Fisherfaces [BHK97], which is based on Linear Discriminant Analysis.

One of the most successful features used in face recognition are Gabor features. Their use is usually motivated by their interesting invariance properties with respect to translation, scale and rotation [KKK06] and by the fact that they are biologically inspired [Dau85]. Using a set of Gabor filters can also be considered as an effective way to perform a frequency decomposition of a signal. The well-known Elastic Bunch Graph Matching (EBGM) technique [WFKvdM97] indeed makes use of Gabor-based features. Fiducial points of the face such as eyes or mouth are then described by sets of Gabor wavelet components called jets. In [ZSG⁺05], an alternative approach is proposed that also makes use of a family of

Gabor filters. Referred to as Local Gabor Binary Pattern Histogram Sequence (LGBPHS), the feature extraction process consists of applying the Local Binary Pattern (LBP) operator [AHP06] to Gabor filtered images, computing histograms over the LBP coefficients and concatenating them into a feature vector. This system has been shown to achieve state-of-the-art on the well-known FERET database.

Another successful face recognition technique relies on feature distribution modelling with Gaussian Mixture Models (GMM) [CSM03, LC04]. The face is firstly divided into blocks (or parts). Features are then extracted from these blocks, before being modelled using a GMM. As this technique is non-holistic, it can deal well with inaccurate localization of the face [CSM03].

This paper proposes to combine the strengths of both Gabor-based features and GMM modelling in a novel face verification system. In particular, we propose to model the outputs of a set of Gabor filters with a corresponding set of GMMs. This combination of feature extraction and modelling techniques appears to be surprisingly simple, yet it provides substantial gains in accuracy compared to the baseline approach of using a single GMM and, to the best of our knowledge, it has not been proposed before. Furthermore, it presents a successful embodiment of a more general framework proposed here for face recognition, which consists of performing a frequency decomposition followed by the use of a set of highly specialized classifiers.

Our contribution is threefold. Firstly, we provide an analysis of what are the most discriminative Gabor filter frequencies for the face recognition task (section 5.3). Secondly, we propose a new scheme for face verification by combining Gabor filtering and GMM-based feature distribution modelling (section 4). Finally, we provide an evaluation of the proposed system with respect to the standard GMM baseline system and state-of-the-art LGBPHS technique (section 5). Experiments show that the proposed system outperforms the DCT GMM baseline in all cases and achieves better results than LGBPHS for four of the seven BANCA face verification protocols and for the MOBIO male protocol.

2 Gabor Filtering

Gabor-based features have been intensively used in vision since the work of Daugman [Dau85] who claimed that two-dimensional Gabor functions are good models of receptive fields of a large amount of cells located in the mammalian Primary Visual Cortex. In addition, their invariance properties to translation, scale and rotation also contributed to make them popular [KKK06]. In face recognition, it is indeed expected that this will provide robustness to pose, illumination and expression. Gabor-based features are generated by applying filters to an image. More precisely, Gabor filters belong to the well-known class of band-pass filters, meaning that they pass frequencies within a certain range and reject frequencies outside that range.

2.1 Gabor filters in the spatial domain

Most of the previous work applies Gabor filtering in the spatial domain [WFKvdM97, ZSG⁺05]. In this domain, a two-dimensional Gabor filter (1) is a Gaussian kernel modu-

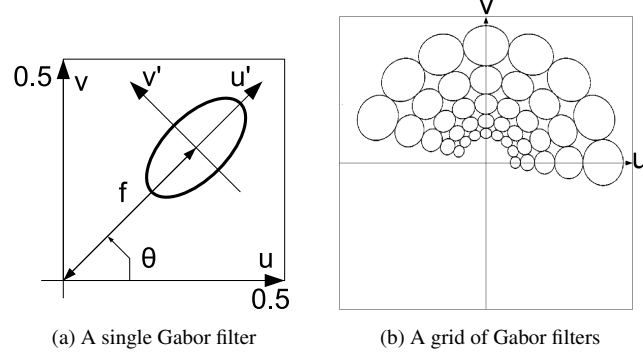


Figure 1: Representation of a single 2D Gabor filter (1a) and of a grid of Gabor filters (1b) in the frequency domain.

lated by a sinusoidal plane wave [KKK06]:

$$\begin{aligned} \psi_2(x, y) &= \frac{f^2}{\pi\gamma\eta} e^{-\left[\frac{x'^2 f^2}{\gamma^2} + \frac{y'^2 f^2}{\eta^2}\right]} e^{j2\pi f x'}, \\ x' &= x \cos(\theta) + y \sin(\theta), \\ y' &= -x \sin(\theta) + y \cos(\theta). \end{aligned} \quad (1)$$

The characteristics of this filter are thus specified by the following parameters: f , θ , γ and η which control the filter frequency, orientation and scale. They will be discussed in the following section.

2.2 Gabor filters in the Fourier domain

In this work, we perform Gabor filtering of images in the frequency domain, as this allows us to more naturally interpret the Gabor filtering process as a decomposition of the image information content in separate frequency bands.

In the frequency domain, the expression of a two-dimensional Gabor filter is given by the following equation:

$$\begin{aligned} \Psi_2(u, v) &= e^{-\frac{\pi^2}{f^2} [(u'-f)^2 \gamma^2 + v'^2 \eta^2]}, \\ u' &= u \cos(\theta) + v \sin(\theta), \\ v' &= -u \sin(\theta) + v \cos(\theta). \end{aligned} \quad (2)$$

This is a Gaussian whose center is at a distance f from the origin with an angle θ from the horizontal axis. The major axis is oriented along the θ direction. γ and η characterize the sharpness of the Gaussian along the major and the minor axes. u and v are the frequency coordinates in the reference system, whereas u' and v' are the coordinates in a θ rotated reference system (Figure 1a). The use of this filter retains frequencies around the center of the Gaussian, whereas other frequencies are rejected or strongly attenuated.

2.3 Gabor filters set

A family of such filters can then be obtained by considering different frequencies f and orientations θ . By processing face images using such a set of filters, information can be extracted from the face and modelled for face recognition. Previous work [WFKvdM97, ZSG⁺05] proposes to use a regular grid of filters with M frequencies and N orientations, where f_0 is the highest frequency, and k the ratio between consecutive frequencies (Figure 1b):

$$\begin{aligned} f_m &= \frac{f_0}{k^m}, m \in \{0, \dots, M-1\} \\ \theta_n &= \frac{n\pi}{N}, n \in \{0, \dots, N-1\} \end{aligned} \quad (3)$$

The design of this grid for face recognition has not been analyzed thoroughly in the literature. The values of $M = 5$ and $N = 8$ have become almost standard choices. However, the selection of the frequencies f_m is highly important as it determines what information is going to be extracted. This is controlled by the two parameters f_0 and k .

Finally, the variance (η and γ) of the Gaussians should also be specified. In order to evenly cover the frequency space while keeping redundant information at a low level, this is controlled by the percentage p of the peak magnitude at which adjacent filters intersect [MM96]. The same value of p is considered for both intersection between consecutive frequencies and consecutive orientations. In this case, we have:

$$\begin{aligned} \gamma &= \frac{(k+1) \sqrt{\ln\left(\frac{1}{p}\right)}}{(k-1)\pi}, \\ \eta &= \frac{1}{\tan\left(\frac{\pi}{2N}\right) \sqrt{\frac{\pi^2}{\ln\left(\frac{1}{p}\right)} - \frac{1}{\gamma^2}}}. \end{aligned} \quad (4)$$

Experiments with various settings of f_m and p (k being fixed) are presented in section 5.3 to analyze the modelling capability of various Gabor filters. Surprisingly, such an analysis has thus far not been sufficiently presented in the face recognition literature.

3 Modelling the Face with a GMM

An efficient and robust approach to face recognition relies on feature distribution modelling with GMMs. This approach was found to offer the best trade-off in terms of complexity, robustness and discrimination in [CSM03, LC04]. The main aspects of this system are presented in the remainder of this section.

3.1 Feature extraction

Face images are divided into blocks, blocks that overlap by a fixed percentage. Next, the Discrete Cosine Transform (DCT) is applied and a fixed number of the lowest frequency

coefficients are retained. These DCT coefficients are quite similar to Fourier coefficients as they characterize the frequencies contained in each block. The feature vector from each block is considered as a separate observation of a statistical distribution.

3.2 Adapted Gaussian Mixture Models

The statistical distribution of the features (DCT coefficients) for each client is modelled using a GMM with diagonal covariance matrices, as follows. Enrollment is achieved through adaptation of a Universal Background Model. Firstly, a world (background) model Ω_{world} is trained from a set of faces and is represented using a GMM. Finding the parameters of this GMM is performed by clustering and then using the Expectation-Maximization algorithm [DLR77]. Next, for each client, model adaptation is used to derive a client model Ω_{client}^i given one or more labelled enrollment images. Mean only adaptation was chosen for this work as it has been shown to be effective and require fewer observations [CSM03]. Finally, for classification of the features, \mathbf{x} , extracted from a test image, given a client model Ω_{client}^i , a log-likelihood score $h(\mathbf{x})$ can be computed, which corresponds to the log-likelihood ratio:

$$h(\mathbf{x}) = \ln(p(\mathbf{x}|\Omega_{client}^i)) - \ln(p(\mathbf{x}|\Omega_{world})) \quad (5)$$

As in [CSM03, LC04], each feature vector from an image is treated independently, meaning that log probabilities are simply summed. Previous work has used GMMs to model the distribution of DCT features extracted directly from facial images [CSM03, LC04]. This is referred to as the standard DCT-GMM approach. In contrast, this paper proposes to use GMMs to model features extracted from Gabor images, which will be shown to consistently improve face verification accuracy.

4 NOVEL GABOR DCT-GMM SYSTEM

The method proposed in this paper combines Gabor filtering with the DCT-GMM parts-based approach [LC04]. Gabor filters have interesting invariance properties to translation, scale and invariance [KKK06] and act as a band-pass filter. Therefore, processing each Gabor image separately allows us to train a highly specialized model for each particular frequency sub-band. These subsystems each perform enrollment and classification in the same way as the standard DCT-GMM approach described in the previous section, but are instead fed with Gabor images as inputs. This choice was motivated by the robustness of both of these techniques. An overview of the system is depicted in Figure 2.

4.1 Tan and Triggs preprocessing

For the proposed system, the illumination normalization algorithm introduced by Tan and Triggs [TT10] is used after cropping face images. Preliminary experimental results showed improved face recognition accuracy using the Tan and Triggs algorithm rather than using histogram equalization. An example of a cropped image processed by this algorithm is shown on Figure 3.

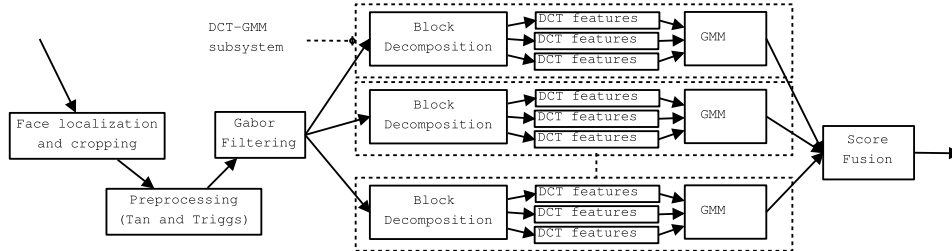


Figure 2: Overview of the proposed Gabor DCT-GMM system




Figure 3: A cropped image (3a), and the resulting image (3b) after applying the Tan and Triggs preprocessing and rescaling it in $[0, 255]$.

4.2 Gabor filtering

The Fast Fourier Transform (FFT) is applied to get a new representation of the image in the frequency domain. Next, each Gabor filter (given by Equation (2)) is used to process the image before a conversion back in the spatial domain through the inverse FFT. Regarding Gabor filters, their DC components are removed. The resulting image in the spatial domain has complex values. Like lots of previous approaches using Gabor [ZSG⁺05, SBF07], only the magnitude is kept. As a consequence, filters which are symmetric with respect to the origin in the frequency domain provide the same output image. This explains why orientations are only chosen in a half plane $[0, \pi]$.

4.3 DCT-GMM subsystems

Each of the Gabor filtered images is then considered separately and a dedicated DCT-GMM system (Section 3) is trained for each of these Gabor images.

Whereas the Gabor filtering stage (Section 3.2) extracted frequency bands from the image globally, in each subsystem, DCT coefficients are then computed locally over each block within a single Gabor image. The overall process can be seen as a frequency decomposition, followed by a spatial decomposition, and then a second frequency decomposition, but this time locally.

4.4 Fusion

In the standard DCT-GMM system, there is no need for late score fusion, because a single GMM is used to model all of the information contained in the face. In contrast, the

proposed Gabor DCT-GMM system needs to fuse the information provided by each of the specialized DCT-GMM subsystems. Linear fusion by sum rule is used here [Kit03].

5 Experiments and Results

5.1 Databases

There are several databases available to evaluate face recognition and verification systems. For these experiments, the BANCA English [PTBB⁺03] and the MOBIO databases are used, as they have the advantage of having separate training, development and test sets¹, as well as well-defined evaluation protocols for verification.

BANCA English is a bimodal database (face and speech) which contains images from 52 subjects divided into two groups g_1 and g_2 . Each subject recorded 12 sessions which were separated into three different scenarios: controlled (c), degraded (d) and adverse (a). A webcam was used in the degraded scenario, while an expensive camera was used in the controlled and adverse scenarios. From these scenarios, seven different protocols were defined, which specify the material to be used for enrollment and for testing. In three of the protocols, enrollment and test data is selected in matching scenarios, that is from controlled, degraded and adverse scenarios for the Mc, Md and Ma protocols, respectively. Two difficult protocols are also proposed, where enrollment occurs with controlled data and testing uses degraded or adverse data (Ud and Ua). Finally, two protocols combine data from various scenarios. For the Pooled protocol (P), enrollment uses data from the controlled scenario, and testing is performed on data from all the scenarios. For the Grand protocol (G), enrollment and testing use data from all the scenarios, that is, additional enrollment data is used compared to the P protocol.

Experiments were also conducted on the large and challenging MOBIO database. This is a bimodal face and speech database, which was captured to reflect potential real-world scenarios for person authentication on a mobile device (mobile phone or laptop). It contains more than 27,750 still images of 150 subjects extracted from videos, which were recorded during 12 different sessions. It provides separate evaluation protocols for males and females plus a common gender-independent world dataset.

5.2 Experimental setup and baseline systems

Hyperparameters (frequencies of the Gabor filter, number of Gaussians, number of DCT coefficients, etc.) of all the systems considered were tuned to minimize the equal error rate on group g_1 of the BANCA database, referred to as the tuning set. For each database, the global decision threshold is tuned on the corresponding development set. Results are presented as the half total error rate (HTER), that is, the mean of the false alarm rate and false rejection rate at the tuned threshold.

Three different baseline systems are used to evaluate the proposed Gabor DCT-GMM sys-

¹In the FRGC database, 153 subjects occur in both the training set as well as the test set, and there is no publicly-available development set. In the Labeled Faces in the Wild database, 758 image pairs in the training/development set (View 1) are exactly repeated in the test set (View 2).

tem: The baseline systems are referred to as Eigenfaces, DCT-GMM and LGBPHS.

Firstly, a traditional Eigenfaces baseline system [TP91] is evaluated, using the Euclidean distance for classification.

Secondly, the standard DCT-GMM approach is evaluated using a 16 by 16 block size with 75% overlap and 60 DCT coefficients, as suggested in [LC04]. For the GMM, 500 Gaussians were used to model the feature distributions. In the novel Gabor DCT-GMM approach, the same parameters are used for the DCT-GMM subsystems, except that the number of DCT coefficients was reduced to 30, which was found to be sufficient on the tuning set.

Finally, for the LGBPHS system [ZSG⁺05], Gabor filtering is first applied in the frequency domain, which results in a set of Gabor images for each face. The resulting Gabor images are also used for the proposed system. The specific Gabor filter parameters are discussed and given in the following section. Next, Local Binary Patterns (LBP) histograms are computed separately over each filtered image. Parameters used for the LBP are a radius $R = 2$, $P = 8$ neighbors and uniform patterns only, as it led to the optimal results on the tuning set.

For all the systems, face images are cropped and scaled to a size of 64 by 80 pixels, with a distance between the eyes of 33 pixels. For automatic face (eyes) detection and localization (on BANCA only), a Modified Census Transform-based detector [FE04] was used in combination with a method to reduce the number of false alarms [AMM10]. There were 5 images (out of 6540) where the system was not able to find the face. These images were excluded from the datasets for experiments using automatic face localization.

Tan and Triggs preprocessing [TT10] is used for all experiments with the default parameters proposed by the authors (gamma correction: $\gamma = 0.2$; DoG filtering $\sigma_0 = 1$ and $\sigma_0 = 2$; contrast equalization: $\alpha = 0.1$ and $\tau = 10$).

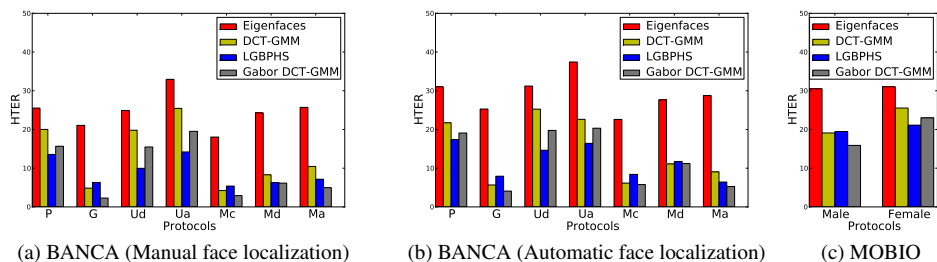


Figure 4: Half Total Error Rate (HTER) of the baseline systems (Eigenfaces, DCT-GMM, LGBPHS) and of the novel Gabor DCT-GMM system on the different protocols of the BANCA and MOBIO databases

5.3 Frequency bands analysis

For LGBPHS and the proposed Gabor DCT-GMM system, the Gabor filters used for feature extraction need to be chosen. Experiments were conducted on the tuning set to

determine filters which provide relevant information for the face recognition task. For this purpose, a grid of Gabor filters was defined, with five different central frequencies $f = [0.05, 0.15, 0.25, 0.35, 0.45]$, eight equally spaced orientations in $[0, \frac{7\pi}{8}]$. For each of these Gabor filters, a corresponding DCT-GMM subsystem was trained.

The influence of the choice of Gabor filter frequencies are reported in Table 1. Our experiments showed that this has a dominant impact on performance. Results depicted in this table are obtained by fusing scores of all considered orientations using the sum rule. Results are reported on the tuning set for the P and G protocols of BANCA. Trends were similar for the other protocols of the database. As shown in Table 1, medium fre-

Table 1: Half Total Error Rate (HTER) of Gabor-based systems by frequency, fusing over the 8 orientations on the tuning set (group g_1 of BANCA)

Frequency (f)	Gabor DCT-GMM		LGBPHS	
	P	G	P	G
0.05	36.34	16.67	28.72	15.64
0.15	17.33	3.66	14.80	8.54
0.25	21.79	4.95	10.84	5.90
0.35	29.49	11.28	12.13	7.87
0.45	31.54	14.80	18.13	16.67

quencies from 0.15 to 0.35 provide the best results for face verification, for both Gabor DCT-GMM and LGBPHS. Regarding orientations and scales of the filters, the results of the analysis showed less variations in performance. For subsequent experiments, as in [WFKvdM97, ZSG⁺05], $M = 5$ frequencies and $N = 8$ orientations are used to form the set of Gabor filters. Taking into account this analysis of Gabor filters showing that frequencies between 0.15 and 0.35 provide the most useful information, the parameters $f_{max} = 0.35$ and $k = 1.37$ were chosen. This results in a set of 40 filters at frequencies of $[0.35, 0.26, 0.19, 0.13, 0.10]$. In addition, the variance of the filters was set to $p = 0.7$, which was found to be a reasonable choice on the tuning set.

The Gabor images resulting from this filtering are used for both the LGBPHS system and the novel Gabor DCT-GMM system.

5.4 Comparison of approaches

Results from the baselines and the novel system are summarized in Figure 4. Overall, the proposed Gabor-DCT-GMM system significantly outperforms Eigenfaces and DCT-GMM for most of the protocols, and is more accurate than the LGBPHS system for some of them.

When compared to the DCT-GMM baseline, the proposed method is much better on both BANCA (between 22% and 52% relative improvements with manual annotations depending on the protocol) and MOBIO databases (10% and 17% relative improvements on the male and female protocols respectively). In fact, Gabor DCT-GMM combines several DCT-GMM subsystems, each of them specialized in modelling a particular frequency sub-

band. This can be considered as a divide and conquer method, where division occurs in the frequency space. The results suggest that such an approach can improve face recognition accuracy.

In addition, we evaluate the statistical significance of improvements in HTER for the previous comparisons. For this purpose, we use the methodology proposed by equation (15) and Figure 2 of [BM04], with a one-tailed test. For the BANCA Mc and P protocols, the improvement of the Gabor DCT-GMM compared to the DCT-GMM baseline is statistically significant at a level exceeding respectively 93.50% and 99.99%. On the MOBIO database, the number of probes is even larger, and the improvement is statistically significant at a level exceeding 99.99%.

In comparison to the state-of-the-art LGBPHS system, the relative performance of the proposed method is dependent on the protocol. Considering the BANCA database, in matched enrolling and testing conditions (protocols Mc, Md and Ma), the proposed system outperforms LGBPHS. Conversely, in unmatched conditions (protocols Ud, Ua and P), the Gabor DCT-GMM system is affected by the differences between the enrollment and the test set which additionally contains data from degraded or adverse conditions. The reason why LGBPHS performs better might be explained by the use of the LBP operator, which is robust to changes in absolute pixel intensities. Results suggest that this may make the LGBPHS system more robust to the differences in image capture conditions for these protocols.

On the G protocol, the Gabor DCT-GMM system outperforms the LGBPHS system with a relative improvement of 64%, when using manual annotations. This protocol has the particularity of containing three times more enrollment samples than the other ones. In addition, these samples come from the three conditions (controlled, degraded and adverse). The results suggest that when enrolled with additional samples that are well-matched to the testing conditions, the Gabor DCT-GMM approach performs particularly well, compared to LGBPHS.

On the MOBIO database, the results are also dependent on the protocol. For the male protocol, the proposed approach significantly outperforms LGBPHS with a relative improvement of 18%, whereas for female, there is a relative degradation of 9%. This might be explained by the use of a common gender-independent Universal Background Model for both protocols, which was trained with about three times as many male samples as female samples. This imbalance may have reduced the effectiveness of the GMM approaches for female clients in particular compared to LGBPHS.

Finally, when using automatic face localization on BANCA, the trends between systems are quite similar to the manual case. On the G protocol, the Gabor DCT-GMM system achieves 48% relative improvement compared to LGBPHS.

6 Conclusions and Future Work

In this paper, a novel extension to the DCT-GMM approach for face recognition was introduced, using Gabor filtering. In addition, an analysis of frequencies which provide relevant information for this particular task was conducted. The proposed system demon-

strates substantial improvements compared to the standard DCT-GMM approach (52% relative HTER improvement on the G protocol of BANCA). In addition, the proposed system outperforms a state-of-the-art LGBPHS technique on four of the seven BANCA face verification protocols, and on the male protocol of the MOBIO database.

Future work will further examine why LGBPHS outperforms the proposed system in unmatched training and testing conditions. In particular, this might be explained by the use of the LBP operator which is robust to changes in absolute pixel intensities, and therefore robust to differences in image capture conditions. Therefore, future work will aim at incorporating LBP processing within the Gabor DCT-GMM approach. Secondly, in the LGBPHS system, the location of the blocks is implicitly used during classification. Therefore, another possible future direction is to use information about the relative location of the blocks to improve the performance of the proposed approach. For example, this block location information might be embedded in the feature vector, as proposed for a DCT-GMM system in [CSB04]. Finally, future work should further investigate fusion schemes beyond the sum rule for combining the scores of individual Gabor filter GMMs, as well as cross-database evaluation.

Acknowledgments

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7) under grant agreements 238803 (BBfor2, www.bbfor2.net), 257289 (TABULA RASA, www.tabularasa-euproject.org) and 284989 (BEAT, www.beat-eu.org). The authors gratefully thank the EU for its financial support, and all project partners for a fruitful collaboration.

References

- [AHP06] T. Ahonen, A. Hadid, and M. Pietikäinen. Face description with Local Binary Patterns: Application to face recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 28:2037–2041, 2006.
- [AMM10] C. Atanasoaei, C. McCool, and S. Marcel. A principled approach to remove false alarms by modelling the context of a face detector. In *Proceedings of the British Machine Vision Conference*, pages 17.1–17.11, 2010.
- [BHK97] 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 (PAMI)*, 19:711–720, 1997.
- [BM04] S. Bengio and J. Mariéthoz. A Statistical Significance Test for Person Authentication. In *Proceedings of Odyssey: The Speaker and Language Recognition Workshop*, 2004.
- [CSB04] F. Cardinaux, C. Sanderson, and S. Bengio. Face Verification Using Adapted Generative Models. In *IEEE Conference on Automatic Face and Gesture Recognition (FG)*, pages 825–830, 2004.
- [CSM03] F. Cardinaux, C. Sanderson, and S. Marcel. Comparison of MLP and GMM Classifiers for Face Verification on XM2VTS. In *International Conference on Audio- and Video-based Biometric Person Authentication*, pages 911–920, 2003.

- [Dau85] J.G. Daugman. Uncertainty Relation for Resolution in Space, Spatial Frequency, and Orientation Optimized by Two-Dimensional Visual Cortical Filters. *Journal of the Optical Society of America A: Optics, Image Science, and Vision*, 2(7):1160–1169, 1985.
- [DLR77] A.P. Dempster, N.M. Laird, and D.B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, 39(1):1–38, 1977.
- [FE04] B. Froba and A. Ernst. Face detection with the Modified Census Transform. In *IEEE Conference on Automatic Face and Gesture Recognition (FG)*, pages 91–96, 2004.
- [Kit03] J. Kittler. Sum versus vote fusion in multiple classifier systems. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 25(1):110–115, 2003.
- [KKK06] J.K. Kamarainen, V. Kyrki, and H. Kälviäinen. Invariance Properties of Gabor Filter Based Features - Overview and Applications. *IEEE Transactions on Image Processing*, 15(5):1088–1099, 2006.
- [LC04] S. Lucey and T. Chen. A GMM Parts Based Face Representation for Improved Verification through Relevance Adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, pages 855–861, 2004.
- [LJ05] S.Z. Li and A.K. Jain. *Handbook of Face Recognition*. Springer, 2005.
- [MM96] B.S. Manjunath and W.Y. Ma. Texture Features for Browsing and Retrieval of Image Data. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 18(8):837–842, 1996.
- [PTBB⁺03] V. Popovici, J. Thiran, E. Bailly-Bailliere, S. Bengio, F. Bimbot, M. Hamouz, J. Kittler, J. Mariethoz, J. Matas, K. Messer, B. Ruiz, and F. Poiree. The BANCA Database and Evaluation Protocol. In *4th International Conference on Audio- and Video-Based Biometric Person Authentication, Guildford, UK*, volume 2688 of *Lecture Notes in Computer Science*, pages 625–638, Berlin, 2003.
- [SBF07] L. Shen, L. Bai, and M. Fairhurst. Gabor wavelets and general discriminant analysis for face identification and verification. *Image and Vision Computing*, 25:553–563, 2007.
- [TP91] M. Turk and A. Pentland. Face Recognition using Eigenfaces. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 586–591, 1991.
- [TT10] X.Y. Tan and B. Triggs. Enhanced Local Texture Feature Sets for Face Recognition Under Difficult Lighting Conditions. *IEEE Transactions on Image Processing*, 19(6):1635–1650, 2010.
- [WFKvdM97] L. Wiskott, J.-M. Fellous, N. Krüger, and C. von der Malsburg. Face Recognition by Elastic Bunch Graph Matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 19:775–779, 1997.
- [ZCPR03] W. Zhao, R. Chellappa, P.J. Phillips, and A. Rosenfeld. Face Recognition: A Literature Survey. *Surveys*, 35(4):399–458, 2003.
- [ZSG⁺05] 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 *IEEE International Conference on Computer Vision (ICCV)*, volume 1, pages 786–791, 2005.