

Face Recognition in the Wild with the Probabilistic Gabor-Fisher Classifier

Simon Dobrišek, Vitomir Štruc, Janez Križaj, France Mihelič

Faculty of Electrical Engineering, University of Ljubljana, Tržaška cesta 25, SI-1000 Ljubljana, Slovenia

Abstract—The paper addresses the problem of face recognition in the wild. It introduces a novel approach to unconstrained face recognition that exploits Gabor magnitude features and a simplified version of the probabilistic linear discriminant analysis (PLDA). The novel approach, named Probabilistic Gabor-Fisher Classifier (PGFC), first extracts a vector of Gabor magnitude features from the given input image using a battery of Gabor filters, then reduces the dimensionality of the extracted feature vector by projecting it into a low-dimensional subspace and finally produces a representation suitable for identity inference by applying PLDA to the projected feature vector. The proposed approach extends the popular Gabor-Fisher Classifier (GFC) to a probabilistic setting and thus improves on the generalization capabilities of the GFC method.

The PGFC technique is assessed in face verification experiments on the Point and Shoot Face Recognition Challenge (PaSC) database, which features real-world videos of subjects performing everyday tasks. Experimental results on this challenging database show the feasibility of the proposed approach, which improves on the best results on this database reported in the literature by the time of writing.

I. INTRODUCTION

Face recognition in unconstrained settings, nowadays often referred to as face recognition in the wild, represents a challenging and still unsolved research problem. While early research on face recognition focused mainly on controlled conditions, recent organized efforts initiated by various institutions, research groups or agencies (such as [22], [21], [16], or [2]) with the goal of improving the state of the technology, have contributed significantly to developments in this field. However, open issues and challenges still remain.

The main reason for the difficulty of face recognition in the wild is the enormous facial-appearance variability encountered in unconstrained setting, such as variability caused by varying viewing angles, motion blur or out of focus imagery, self-occlusions and other similar factors. A couple of examples of the appearance variability encountered in unconstrained settings is presented in Fig. 1. Note how the appearance of the subjects depicted in the frames changes due to pose changes and different distances from the camera.

In this paper we try to address the problem of face recognition in the wild and build on one of the more popular approaches to face recognition, i.e., approaches exploiting Gabor filters. While a lot of advancements have been made in recent years in different sub-problems related to face recognition, such as face detection and registration [26], [28], facial-feature extraction [25], [27], [29], [12] or modeling-and-classification [9], [23], [20], various competitions and



Fig. 1. Sample frames from real-world videos (examples taken from the Point-and-Shoot Face Recognition Challenge (PaSC) database [2]). The sample frames show typical problems encountered with unconstrained face recognition, such as different head poses, non-frontal faces or low-resolution images.

group evaluations organized in the scope of top-tier biometric conferences (e.g., [8], [10]) show that Gabor-based methods regularly finish among the top performers, which suggests that Gabor features are highly discriminative and should be well suited for even the most difficult problems, such as face recognition in the wild.

The most relevant (Gabor-based recognition) technique for the work presented in this paper is the Gabor-Fisher Classifier proposed by Liu and Wechsler in [15]. The technique applies a battery of Gabor filters to the given facial image to extract oriented multi-scale features from the image and then constructs low-dimensional feature vectors by projecting the Gabor features into a feature space constructed with a modified version of Fisher’s discriminant analysis. The technique performed extremely well and was later also extended to other versions that included projections into other feature spaces, e.g. [13], [14], or inclusion of other types of Gabor features, e.g. [3].

In this paper we present a different approach and try to extend the Gabor-Fisher Classifier (GFC) using a probabilistic version of linear discriminant analysis (PLDA). With the extension, dubbed Probabilistic Gabor-Fisher Classifier (PGFC), we try to improve on the generalization capabilities of the GFC approach and make it suitable for face recognition in the wild. We keep the original processing pipeline and first extract facial features from a given input image using a battery of Gabor filters. We then construct vectors of Gabor features and subject them to PLDA to derive a low-dimensional representation of the input face image that can be used for recognition.

We assess our PGFC approach on the recently introduced Point and Shoot Face Recognition Challenge database and show that it is capable of delivering state-of-the-art performance even for the most challenging settings.

In summary, we present the following contributions in the paper:

- We introduce a novel Probabilistic Gabor-Fisher Classifier that extends the idea of [15] to a probabilistic setting;
- We evaluate the proposed approach in face verification experiments on the challenging Point-and-Shoot Face Recognition Challenge database; and
- We benchmark the PGFC approach against the state-of-the-art on the PaSC database with very encouraging results.

The rest of the paper is structured as follows: In Section II we present the novel Probabilistic Gabor-Fisher classifier, discuss the feature extraction and modeling procedure with the proposed approach and provide a summary of steps needed to train the PGFC technique. In Section III we present experiments with our approach and provide comparative results with state-of-the-art methods from the literature. We conclude the paper with some final thoughts and directions for future work in Section IV.

II. THE PROBABILISTIC GABOR-FISHER CLASSIFIER

In this section we present the proposed Probabilistic Gabor-Fisher classifier. We commence the section with a short overview of the proposed technique and then provide a detailed description of the feature extraction and modeling steps.

A. Overview

Motivated by the success of the Gabor-Fisher Classifier (GFC) [15], which has long presented one of the state-of-the-art approaches to face recognition, we propose a novel recognition technique for face recognition in the wild called Probabilistic Gabor-Fisher Classifier (PGFC). The proposed technique extends the idea of the GFC approach and applies a recent probabilistic extension of the Linear Discriminant Analysis (LDA) to the Gabor features extracted from facial images.

With the proposed approach a facial image is first filtered with a battery of Gabor filters. A compact feature representation is then derived from the outputs of the filtering operation through a projection into the Principal Component Analysis (PCA) and LDA subspaces. This projection reduces the dimensionality of the feature vectors to such an extent that a simplified version of the probabilistic Linear Discriminant Analysis (PLDA) can be employed, which ultimately produces estimates of the so-called identity variables that can be used for recognition.

B. Gabor Filters and Feature Extraction

Gabor filters represent one of the most popular tools for feature extraction from facial images. Their use in automatic face recognition system is motivated by two major




Fig. 2. Visual appearance of the real-part of the battery of 40 Gabor filters usually used for facial feature extraction.

factors: their computational properties and their biological relevance [3], [12].

Formally, two-dimensional Gabor filters are defined as follows [3], [14]:

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

where $x' = x \cos \theta_v + y \sin \theta_v$, $y' = -x \sin \theta_v + y \cos \theta_v$, and the parameters f_u and θ_v are computed as $f_u = f_{max}/2^{(u/2)}$ and $\theta_v = v\pi/8$.

As can be seen from the above equation, Gabor filters are comprised of a Gaussian kernel that is modulated by a complex plane wave with a center frequency of f_u and an orientation of θ_v . γ and η define the ratio between the center frequency and the size of the Gaussian envelope. When γ and η are set to fixed pre-defined values, the scale of the Gabor filters becomes dependent on the center frequency f_u . Commonly γ and η are set to the value of $\gamma = \eta = \sqrt{2}$. f_{max} denotes the maximum frequency of the filters and is usually set to $f_{max} = 0.25$.

When Gabor filters are used for facial feature extraction, a battery of filters is typically created first and is then used to extract multi-orientational and multi-scale features from a given input face image. For the task of face recognition the battery of filters commonly consists of Gabor filters with 5 different scales and 8 different orientations, i.e., $u = 0, 1, \dots, p-1$ and $v = 0, 1, \dots, r-1$, where $p = 5$ and $r = 8$ [13],[15], [3]. A visual example of the real part of the described battery of filters is shown in Fig. 2.

Since Gabor filters represent bandpass filters, features extracted with a given filter from the battery capture only information from a specific frequency band of the image. The total battery, however, covers a broader frequency band as shown in Fig. 3 where the power spectrum of the entire Gabor-filter battery is shown (together with the mirrored channels).

To extract features from facial images with the battery of filters, the given input image $I(x,y)$ is convolved with each of the filters from our battery $\psi_{u,v}(x,y)$. Here, the filter $\psi_{u,v}(x,y)$ is defined by the parameters f_u and θ_v :

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




Fig. 3. Coverage of the power spectrum for the entire battery of 40 Gabor filters typically used for facial feature extraction. The mirrored channels are shown as well.

The result of the filtering operation is a complex response, from which we retain only the magnitudes, i.e., $A_{u,v}(x, y) = \sqrt{\text{Re}(G_{u,v})^2(x, y) + \text{Im}(G_{u,v})^2(x, y)}$.

When deriving the Gabor magnitude feature vector from the given face image $I(x, y)$, we first compute the Gabor magnitude responses for the entire battery of 40 Gabor filters. This process results in the inflation of the dimensionality of the original pixel space by a factor of 40. This dimensionality problem is commonly approached by down-sampling of the magnitude responses and concatenation of the down-sampled magnitude responses into the final (still relatively high-dimensional) feature vector. To further reduce the dimensionality of the feature vector an additional subspace projection step or other kind of dimensionality reduction is commonly adopted.

In the experimental section we down-sample each of the Gabor magnitude responses using a rectangular uniformly-spaced sampling grid with 256 nodes to bring the combined dimensionality of the Gabor magnitude features to a manageable size and then subject the vectors of concatenated (and down-sampled) Gabor magnitude responses to probabilistic linear discriminant analysis for further processing as described in the next section.

C. Probabilistic modeling

Probabilistic Linear Discriminant Analysis (PLDA) represents a probabilistic extension of Linear Discriminant Analysis (LDA) [4] and was first introduced to the field of face recognition in [30], [9]. The technique decomposes a given feature vector into an identity-specific part that describes the class-membership of the feature vector and a so-called channel-specific part that accounts for the sample (or intra-class) variability of the feature vector [31]. Since the identity-specific and variability-specific part are separated by PLDA, the former can be used for identity inference.

PLDA has become relatively popular in recent years, mostly in the field of speaker recognition (see, e.g., [6], [5] [7], [23]). In this field, a specific variant of PLDA was developed that is simpler than the original formulation presented in [9] and was also used in our experiments. This simplified form, commonly referred to as *simplified PLDA*, is applicable only if the feature

vectors used for PLDA training are of a sufficiently low dimensionality. The technique can formally be described as follows: let $\{\eta_r : r = 1, \dots, R\}$ denote a collection of low-dimensional feature vectors extracted from a set of facial images of a specific individual. The simplified PLDA technique then decomposes each feature vector into the following form [31]:

$$\eta_r = m + \Phi\beta + \varepsilon_r, \quad (3)$$

where m stands for the average feature vector, Φ stands for the basis of the identity-specific subspace, β represents a latent identity vector drawn from the standard normal distribution, and ε_r denotes a sample-dependent residual term, which is again assumed to be normally distributed with zero mean and a full covariance matrix Σ [31].

Different from the common LDA technique, the parameters of the PLDA model $\{m, \Phi, \Sigma\}$ are not computed analytically, but are learnt from development data via an EM algorithm [6]. Once the PLDA model parameters are estimated, identity inference can be made using the hidden identity variable β that is estimated from the given input feature vector η_r .

In the case of our Gabor magnitude face representation, the feature vector η_r is produced from the concatenated (and down-sampled) magnitude responses by subjecting them to the whitened form of PCA, followed by the LDA projection step to produce feature vectors of sufficiently low-dimensions to be applicable for the simplified form of PLDA. In our experiments the dimensionality of the feature vectors η_r was set to 200, while the identity variable β exhibited the dimensionality of 150.

D. Summary of the PGFC method

The individual steps needed to estimate all open-parameters (i.e., the PCA and LDA transformation matrices and the PLDA model parameters) of the PGFC method are given below:

TRAINING

Data: Training or development set of N images

Result: PCA and LDA transformation matrices, PLDA model parameters $\{m, \Phi, \Sigma\}$

for *Each image in training set* **do**

→ compute the concatenated and down-sampled Gabor magnitude representation;

end

→ compute PCA and LDA transformation matrices;

for *Each image in training set* **do**

→ reduce the dimensionality of the Gabor magnitude representation using PCA and LDA;

end

→ estimate PLDA model parameters $\{m, \Phi, \Sigma\}$ from the computed vectors η_r .

Algorithm 1: Summary of the PGFC training procedure

Once the training is done, the estimated model-parameters and transformation matrices can be employed to compute the identity variable from the given input face image. This variable is later used for identity inference.

III. EXPERIMENTS AND RESULTS

In this section we describe the experimental evaluation of the PGFC technique. We discuss the results and provide comparisons with the state-of-the-art.

A. The database and experimental setup

To assess the performance of the PGFC technique on facial imagery captured in the wild, we select the PaSC database [2] for our experiments. The PaSC (Point and Shoot Face Recognition Challenge) database represents a recent database that contains video-data featuring 265 subjects. An important characteristic of the database, which sets it apart from other similar databases, is the fact that subjects featured in the video recordings do not face the camera directly, but instead perform various tasks while being recorded. Due to the outlined recording setup, the faces in the videos are often far from frontal and exhibit variability in terms of pose, self-occlusion, motion blur, poor focus and alike.

The PaSC database comprises a total of 2802 video recordings for experimentation and another 280 videos that can be utilized for training. 1401 of the videos designated for experimentation and recorded with a HD camera mounted on a tripod are labeled as “control” videos in the database, while the rest (i.e., another 1401 videos) represent videos acquired with a handheld camera. Using this data we conduct two types of experiments on the PaSC database, which are also defined by the experimental protocol associated with the data:

- *Experiments on the control data:* Here we compare all video recordings from the “control” group of recordings to each other and produce a similarity-score matrix of size 1401×1401 , and
- *Experiments on the handheld data:* Here we compare all video recordings from the “handheld” group of recordings to each other, again resulting in the similarity-score matrix of size 1401×1401 .

Among the similarity-scores from the computed matrices a little over 3000 scores belong to match pairs, while the rest (a little under 1 million scores) belong to non-match pairs. Note that a single similarity score is produced for a single video-to-video comparison with the PGFC method. To produce a single score for each comparison we first compute the average PLDA identity variable over all frames of the given (two) video sequences and then use the cosine similarity measure to produce a similarity score for a given pair of (mean) identity variables. The model parameters for PLDA (as well as PCA and LDA transformation matrices) are trained exclusively on the 280 PaSC training videos - no other data is used for this purpose.

We report our results in the form of the Receiver Operating Characteristic (ROC) curves and quantitatively in the form of the verification rate (VER) at the false accept rate




Fig. 4. Sample images from the PaSC database after cropping of the facial area and alignment based on the eye coordinates. Note that due to landmark localization errors the faces are very poorly aligned.

(FAR) of 1% (we will refer to this performance metric as $VER@0.01FAR$ in the remainder of the paper).

For our experiments we first processed all video recordings from the database with a Commercial Of The Shelf (COTS) face and landmark detector (we used the solution from PittPatt) to produce meta-data, such as facial locations and eye coordinates, based on which the faces were extracted from the video recordings. The faces were then aligned based on the eye-center locations, cropped to the size of 50×50 pixels and then photometrically normalized using the normalization technique presented in [11]. After the pre-processing the PGFC technique was applied to the extracted faces.

A few visual examples of the facial images from the PaSC database after cropping and alignment are shown in Fig. 4.


B. PGFC assessment and baseline comparisons

To provide an estimate of the performance of the proposed PGFC technique on the PaSC database, we produced the similarity-score matrices described in the previous section and then generated corresponding ROC curves for the controlled and handheld experiments. The results are shown in Fig. 5. Here, the performance of the baseline technique of the database - Local-Region PCA (LR-PCA) [21] - is also included in the comparison.

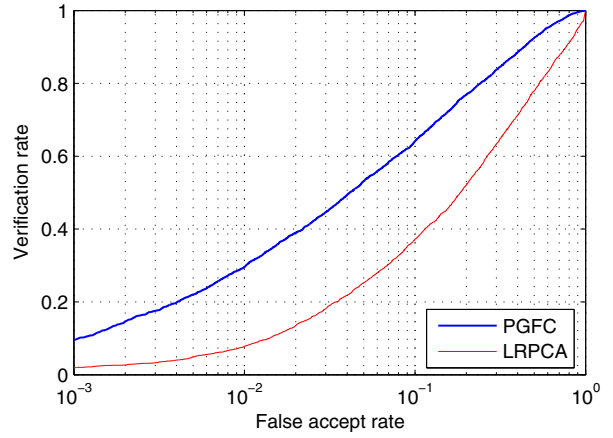
Note that the proposed PGFC technique produces significantly better verification performance than the baseline LR-PCA technique for both experiments. Our technique achieves the $VER@0.01FAR$ of 0.29 on the handheld videos and 0.35 on the control videos, while the baseline performance (in terms of $VER@0.01FAR$) on these videos is 0.08 and 0.10, respectively.

C. Comparisons with other descriptors

In our second experiment we benchmark the PGFC technique against combinations of different texture descriptors and PLDA. Specifically, we implement two additional techniques for our comparisons. The first technique (denoted as LBP-PLDA in Fig. 6) first extracts uniform local binary patterns - LBPs [27] - (from an 8-pixel neighborhood at the radius of 2) from image blocks of size 6 pixels and then represents the patterns of each block with 58-bin histograms.

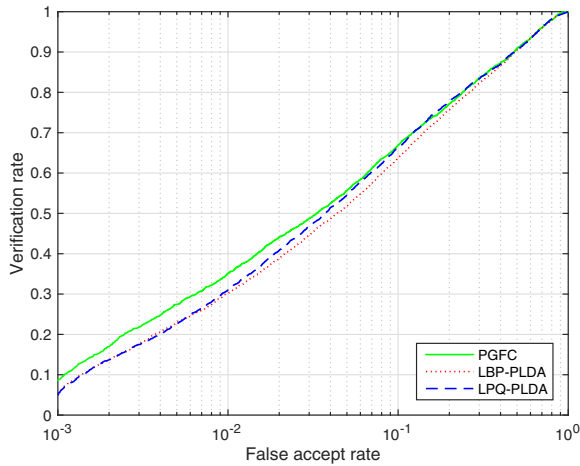


(a) Results for the controlled data

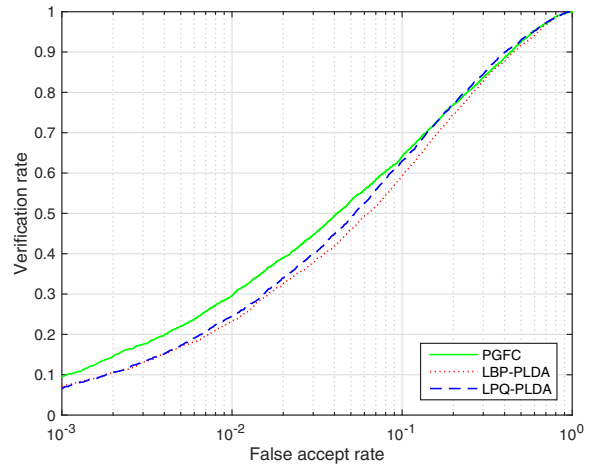


(b) Results for the handheld data

Fig. 5. ROC curves generated during the experiments. The blue curve shows the performance of the PGFC technique and the red one shows the performance of the baseline LR-PCA technique. The ROC curves are best viewed in color.



(a) Results for the controlled data



(b) Results for the handheld data

Fig. 6. ROC curves generated during the experiments on the controlled and handheld data. The green curve shows the performance of the PGFC technique, while the red and blue curves represent the performance of the LBP-PLDA and LPQ-PLDA techniques, respectively (best viewed in color).

The histograms are then concatenated into feature vectors that are subjected to PLDA for the final dimensionality reduction. Matching is again done using a simple cosine similarity. The second (denoted as LPQ-PLDA in Fig. 6) uses a similar approach as the first (with PLDA and cosine matching), but instead of LBP descriptors extracts Local Phase Quantization (LPQ [27]) patterns from image blocks of size 10×10 pixels and represents the extracted patterns of each block with 255-bin histograms. The concatenated histograms are then again used as feature vectors and input to PLDA. Note that the parameters for the LBP and LPQ descriptors were chosen in a way that resulted in the best performance of the LBP-PLDA and LPQ-PLDA on the PaSC database to ensure a fair comparison with our PGFC technique.

As shown in Fig. 6, the PGFC is again the top performer for the controlled and handheld data. It outperforms the LBP-PLDA technique, which achieves the VER@0.01FAR of 0.30

on the controlled data and the VER@0.01FAR of 0.24 on the handheld data, as well as the LPQ-PLDA approach, which achieves the VER@0.01FAR of 0.31 for the controlled and the VER@0.01FAR of 0.25 for the handheld data. The results suggest that among the tested descriptors Gabor features ensure the best overall performance.

D. Comparison with the state-of-the-art

While the proposed PGFC technique is clearly able to outperform the baseline LR-PCA approach and also compares favorably against combinations of different texture descriptors and PLDA on the PaSC database, the main question is how does it fare when compared against the state-of-the-art. To date not many results are available on this database in the literature. However, the best results (by the time of writing) on this dataset were reported during a face recognition challenge organized as part of IJCB 2014, where one of the experimental protocols (i.e., for the more

TABLE I

PERFORMANCE COMPARISON WITH STATE-OF-THE-ART METHODS FROM THE LITERATURE.

Method	VER@0.01FAR
LRPCA Baseline [21]	0.08
ISV-GMM [20]	0.05
LBP-SIFT-WPCA-SILD [17]	0.09
PLDA-WPCA-LLR [1]	0.19
Eigen-PEP [18], [19]	0.26
PGFC (ours)	0.29

challenging handheld experiment) used overlapped with one of the protocols adopted in our experiments. We, therefore, provide comparative results with this experimental protocol for different state-of-the-art techniques from the literature in Table I. The reader is referred to [10] for a detailed description of the techniques included in the table.

Note that our PGFC technique ensures the best performance for this protocol with the VER@0.01FAR of 0.29.

IV. CONCLUSION

We have presented a novel technique for face recognition called Probabilistic Gabor-Fisher Classifier that combines Gabor features and a simplified version of Probabilistic Linear Discriminant Analysis. We assessed the technique on videos from the challenging PaSC database and demonstrated state-of-the-art performance.

As part of our future work we plan to improve on the presented method by including more information from the Gabor filter outputs and seeking alternatives to the down-sampling procedure.

ACKNOWLEDGMENTS

The work presented in this paper was supported in parts by the national research program P2-0250(C) Metrology and Biometric Systems and the European Union's Seventh Framework Programme (FP7-SEC-2011.20.6) under grant agreement number 285582 (RESPECT). The support of COST Actions IC1106 and IC1206 is also appreciated.

REFERENCES

[1] V. Štruc, J. Križaj and S. Dobrišek. MODEST face recognition. In: *Int. Workshop on Biometrics and Forensics (IWBF'15)*, 2015.

[2] J. Beveridge, P. Phillips, D. Bolme, B. Draper, G. Givens, L. Yui Man, M. Teli, Z. Hao, W. Scruggs, K. Bowyer, P. Flynn, and C. Su. The challenge of face recognition from digital point-and-shoot cameras. In *Proc. of BTAS'13*, pages 1–8, 2013.

[3] V. Štruc and N. Pavešić. The Complete Gabor-Fisher Classifier for Robust Face Recognition. In *EURASIP Adv. in Sig. Proc.*, 2010.

[4] P. Belhumeur, J. Hespanha, and D. Kriegman, "Eigenfaces vs. Fisherfaces: Recognition Using Class Specific Linear Projection," *IEEE TPAMI*, vol. 19, no. 7, pp. 711–720, 1997.

[5] N. Dehak, P. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front-end factor analysis for speaker verification," *Comp. Vis. and Image Understanding*, vol. 19, no. 4, pp. 788–798, 2011.

[6] N. Brummer and E. de Villiers, "The speaker partitioning problem," in *Odyssey: The Speaker and Language Recognition Workshop*, 2010.

[7] D. Garcia-Romero and C. Espy-Wilson, "Analysis of i-vector length normalization in speaker recognition systems," in *Proc. of Interspeech 2011*, 2011.

[8] M. Günther, A. Costa-Pazo, C. Ding, E. Boutellaa, G. Chiachia, H. Zhang, M. de Assis Angeloni, V. Struc, E. Khoury, E. Vazquez-Fernandez, D. Tao, M. Bengherabi, D. Cox, S. Kiranyaz, T. de Freitas Pereira, J. Zganec-Gros, E. Argones-Rúa, N. Pinto, M. Gabbouj, F. Simões, S. Dobrišek, D. González-Jiménez, A. Rocha, M. Uliani Neto, N. Pavesic, A. Falcão, R. Violato, and S. Marcel. The 2013 face recognition evaluation in mobile environment. In *International Conference on Biometrics (ICB'13)*, June 2013.

[9] P. Li, Y. Fu, U. Mohammed, J. Elder, and S. J. Prince. Probabilistic models for inference about identity. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(1):144–157, 2012.

[10] J.R. Beveridge, H. Zhao, P.J. Flynn, Y. Lee, V.E. Liong, J. Lu, M. de Assis Angeloni, T. de Freitas Pereira, H. Li, G. Hua, V. Struc, J. Križaj, P.J. Phillips. The IJCB 2014 PaSC Video Face and Person Reconciliation Competition. *Int. Joint Conference on Biometrics (IJCB'14)*, 2014.

[11] N. Vu and A. Caplier. Illumination-robust face recognition using retina modeling. In *International Conference on Image Processing (ICIP'09)*, pp. 2335–2338, 2009.

[12] D. Cheng., X. Cao, F. Wen, and J. Sun. Computer Face Recognition Using Early Biologically Inspired Features. In *BTAS'13*, 1–6, 2013.

[13] C. Liu, "Capitalize on dimensionality increasing techniques for improving face recognition grand challenge performance," *IEEE Transactions on Image Processing*, vol. 11, no. 4, pp. 467–476, 2002.

[14] L. Shen, L. Bai, and M. Fairhurst, "Gabor wavelets and general discriminant analysis for face identification and verification," *Image and Vision Computing*, vol. 25, no. 5, pp. 553–563, 2007.

[15] C. Liu and H. Wechsler, "Gabor feature based classification using the enhanced fisher linear discriminant model for face recognition," *IEEE TPAMI*, vol. 28, no. 5, pp. 725–737, 2006.

[16] G. Huang, M. Ramesh, T. Berg, and E. Learned-Miller. Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments. In *Technical Report 07-49, University of Massachusetts*, 2007.

[17] M. Kan, S. Shan, D. Xu, and X. Chen. Side information based linear discriminant analysis for face recognition. In *Proc. of BMVC'11*, 2011.

[18] H. Li, G. Hua, Z. Lin, J. Brandt, and J. Yang. Probabilistic elastic matching for pose variant face verification. In *Proc. of CVPR'13*, pages 3499–3506, 2013.

[19] H. Li, G. Hua, X. Shen, Z. Lin, and J. Brandt. Eigen-PEP for Video Face Recognition. In *Proc. of ACCV'14*, 2014.

[20] C. McCool, R. Wallace, M. McLaren, L. El Shafey, and S. Marcel. Session variability modelling for face authentication. *IET Biometrics*, 2(3):117–129, 2013.

[21] P. Phillips, J. Beveridge, B. Draper, G. Givens, A. O'Toole, D. Bolme, D. J., Y. Man, S. H., and S. Weimer. An Introduction to the Good, the Bad & the Ugly Face Recognition Challenge Problem. In *Proc. of FG'11*, pages 346–353, 2011.

[22] P. Phillips, P. Flynn, T. Scruggs, K. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek. Overview of the Face Recognition Grand Challenge. In *Proc. of CVPR'05*, pages 947–954, 2005.

[23] B. Vesnicer, J. Zganec Gros, V. Štruc. Incorporating Duration Information into I-Vector-Based Speaker-Recognition Systems. In *Proc. of Odyssey'14*, 2014.

[24] C. Chan, M. Tahir, J. Kittler, and M. Pietikainen. Multiscale local phase quantization for robust component-based face recognition using kernel fusion of multiple descriptors. *IEEE TPAMI*, 35(7):1164–1177, 2013.

[25] D. Cheng., X. Cao, F. Wen, and J. Sun. Blessing of Dimensionality: High-dimensional Feature and Its Efficient Compression for Face Verification. In *Proc. of CVPR'13*, pages 3025–3032, 2013.

[26] G. Huang, M. Mattar, H. Lee, and E. Learned-Miller. Learning to Align from Scratch. In *Proc. of NIPS'12*, 2012.

[27] M. Pietikainen, A. Hadid, G. Zhao, and T. Ahonen. *Computer Vision using Local Binary Patterns*. Springer, 2011.

[28] X. Xiaoang and F. De la Torre. Supervised Descent Method and its Application to Face Alignment. In *Proc. of CVPR'13*, 2013.

[29] Y. Li, S. Shan, H. Zhang, S. Lao, and X. Chen. Fusing Magnitude and Phase Features for Robust Face Recognition. In *Proc. of ACCV'12*, pages 601–612, 2012.

[30] S. J. D. Prince and J. H. Elder. "Probabilistic linear discriminant analysis for inferences about identity," in *Proc. of ICCV*, 2007.

[31] V. Štruc, and J. Zganec-Gros and B. Vesnicer and N. Pavešić, "Patch-wise low-dimensional probabilistic linear discriminant analysis for face recognition" in *Proc. of ICASSP*, Vancouver, Canada, 2013.