



6th International conference on Intelligent Human Computer Interaction, IHCI 2014

Local Phase-Context for Face Recognition under Varying Conditions

Mohamed Dahmane^{a,*}, Langis Gagnon^a

^aComputer Research Institute of Montreal, 405, Ogilvy Ave. #101, Montreal (Quebec) H3N 1M3, Canada

Abstract

In this paper, we address the problem of face recognition of low-resolution images under varying light, illumination and blur using local texture based face representation. The main contribution is the texture representation using Phase-Context which is based on four-quadrant mask of the Fourier transform phase in local neighborhoods.

The contextual phase generates a more discriminative code filtering responses, and a more effective feature set than the Local Phase Quantization (LPQ) descriptor which is suffering from the influence of the noisy filter responses, the order relation breakdown of the generated codes, and the discretization effect of the quantization.

The experimental results on CMU-PIE, extended YALE-B and CAS-PEAL-R1 databases show that the Phase-Context methodology is more descriptive than LPQ, outperforming the widely used Local Binary Pattern (LBP), and Histogram of oriented Gradients (HOG).

© 2014 The Authors. Published by Elsevier B.V.

Peer-review under responsibility of the Scientific Committee of IHCI 2014.

Keywords: Face recognition ; Fourier transform ; Phase-Context ; Local Phase Quantization ; texture representation

1. Introduction

Biometric-based identification using face recognition have received very much attention for recognizing individuals. The computer vision researchers are more and more interested in automating physiological information capturing from human faces, a process that presents an important challenge arising from facial dynamics and other scene conditions. Several representations can be found in the literature¹ proposing different methodologies for visual face recognition. The aim is to obtain a descriptor of the face that is robust to variations in pose, lighting and facial expressions. Moreover, an efficient descriptor is that which captures facial attributes that account for the most significant variation between-person faces while minimizing within-person variation. Feature-based methods require precise and stable localization of facial landmarks which is not ensured with small face images. In this case, the appearance based-methods, that have attracted more and more attention are more suitable, however they are in the need to a suitable feature set reliably insensitive to facial dynamics but at the same time closely relevant to inter-person variations.

In this paper, we propose a new texture representation using local phase information named Phase-Context for efficient face recognition. Experimental results both on CMU-PIE, YALE-B and CAS-PEAL-R1 datasets show that

* Corresponding author. Tel.: +1-514-840-1234 (6976) ; fax: +1-514-840-1244.

E-mail address: dahmanmo@crim.ca

Phase-Context representation performs better compared to the state-of-the-art representations. The effectiveness is assessed against the widely used descriptors such as LPQ, LBP, and HOG.

We briefly discuss related work in Section 2, describe our method in Section 3. Experimental results are presented in Section 4 and conclusions are given in Section 5.

2. Related work

Many of previous study of face recognition have adopted the concept of holistic processing such as principal component analysis (PCA), independent component analysis (ICA) and Fisher's linear discriminant analysis (LDA). The initial success of eigenfaces strengthened the idea of matching images in compressed subspaces².

Recently, local methods have gained much interest following advances in texture description using local descriptors^{3,4,5,6,7}. Ahonen *et al.*⁸ proposed to use the histogram of LBP⁵ as a combination of micro-patterns for a global description of the face image. The LBP patterns encode the facial regions locally while the whole shape of the face is recovered by the global histogram. LBP codes are insensitive to monotonic gray scale transformation. In practice, this type of coding is sensitive to noisy images when the pixel intensity order is not preserved, due to the threshold function adopted by the operator. Many other LBP variants were proposed in the literature to improve the conventional LBP method⁹. Tan and Triggs²¹ proposed a generalization of LBP and Gabor feature cues combined to a Kernel Linear Discriminant Analysis.

Histogram of oriented Gradient (HOG) descriptors were proposed by Dalal and Triggs³ for person detection. Déniz *et al.*¹⁰ investigated HOG descriptors at different scales for face recognition. The main idea behind HOG is based on the local edge information. That is each window region can be described by the local distribution of the edge orientations and the corresponding gradient magnitude. The local distribution is described by the local histogram of oriented gradients which is formed by dividing the detection window into a set of small regions called cells, and within each cell integrate the magnitude of the edge gradient for each orientation bin. This is done for all of the pixels within the cell. To provide better invariance, the local HOG is normalized over the block of neighboring cells.

More recently, the Local Phase Quantization (LPQ) descriptor was used for face recognition¹¹. LPQ was originally designed by Ojansivu and Heikkila similarly to the LBP methodology as a texture descriptor⁶. LPQ is insensitive to image blurring, and it has proven to be a very efficient descriptor in face recognition from blurred as well as sharp images¹¹.

3. Fourier transform Phase

In the literature, one can find several attempts at designing feature sets using magnitude information. However, the phase can be a very important source of information^{12,13}.

The phase image is based on computing the local frequency using short-term Fourier transformation (equ.1) on local M by M neighborhood \mathcal{N}_x at each pixel position x of the image $f(x)$.

$$F_u(x) = \sum_{y \in \mathcal{N}_x} f(x-y) \exp(-j2\pi u^T y) \quad (1)$$

$F_u(x)$ is efficiently evaluated for all image positions $x \in \{x_1, \dots, x_N\}$ using precomputed 2D filters at four frequency points $u_1 = [a, 0]^T$, $u_2 = [0, a]^T$, $u_3 = [a, a]^T$, and $u_4 = [a, -a]^T$. The frequency parameter a is fixed to the lowest non-zero frequency ($1/M$) as in⁶. For each pixel position x , the Fourier coefficients are given by

$$G(x) = [F_{u_1}(x), F_{u_2}(x), F_{u_3}(x), F_{u_4}(x)]$$

where $F_{u_i}(x)$ is the complex term $\{F_{u_i}^{\Re}(x), F_{u_i}^{\Im}(x)\}$ in which \Re and \Im denote the real and the imaginary parts.

3.1. Local Phase Quantization

The local phase quantization (LPQ) descriptor⁶ is based on quantifying the Fourier transform phase by considering the sign of each component in $G(x)$. This is done by using a simple scalar quantizer

$$q_j(x) = \begin{cases} 1 & \text{if } g_j(x) \geq 0 \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

where $g_j(x)$ is the j^{th} component of the vector $G(x)$. The phase information is then represented as an integer codeword between 0 – 255 using the eight binary coefficients $q_j(x)$

$$f_{\text{LPQ}}(x) = \sum_{j=1}^8 q_j 2^{j-1} \quad (3)$$

In this paper, we refer to this type of coding as codeword-based coding.

Local phase quantization for face description¹¹ used block-based image description similarly to HOG³ and LBP⁸ methodologies. Face images are divided into small non overlapping regions to extract local 256–dimensional LPQ histograms which are then concatenated into a single code histogram. The final histogram is used as a feature vector to represent the face image. We used a block-grid size of 8×8 that was experimentally fixed.

3.2. Phase-Context

The atomic element in LPQ operator can be considered as the quantized code of the image phase. The encoding operators have to optimally balance between the accuracy and the robustness of the discretization against data variance.

The LPQ codeword quantizes the phase as an integer code denoting which quadrant the phase belongs to. Figure (1) shows some examples of frequency responses $F_{u_i}(x)$ at the frequency point u_i . As we can see, the distance between the generated codewords (numbers between parenthesis in the left subfigure) of a phase that belongs to quadrant Q-III and codewords of that which goes to Q-IV is $\text{dist}((0), (1)) = (1)$, whereas the distance between a codeword with phase in quadrant Q-III (i.e. (0)) and a codeword of that from Q-II (i.e. (2)) is (2). Quadrants Q-IV and Q-II are direct neighbors of Q-III and should give the same distances, that is why LPQ encoding is a histogramming-based approach. Another problem of the LPQ encoding is that it does not consider the noisy response cases.

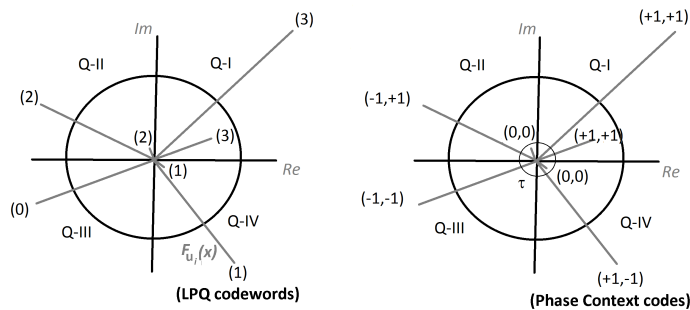


Fig. 1. Local Phase Quantization (left) vs. Phase-Context “no-quantization” (right)

Considering these issues, we propose to use an alternative approach that is based on a four-quadrant phase context assessment (Fig. 1). Rather to consider integer codewords encoding as in LPQ (equ. 3), in the Phase-Context, we keep the signs of the frequency responses and avoid any kind of quantization (equ. 4).

Furthermore, to reduce the influence of erroneous filter responses, we consider only phases with coefficient magnitude larger than some threshold value τ to avoid sign jittering effect of the noisy filter responses. For each frequency u_i we set a different τ that depends on the mean value of the magnitude $|F_{u_i}|$.

$$PC_{u_i}(x) = \begin{cases} (\text{sgn}(F_{u_i}^{\Re}(x)), \text{sgn}(F_{u_i}^{\Im}(x))) & \text{if } \|(F_{u_i}(x))\| \geq \tau \\ (0, 0) & \text{otherwise} \end{cases} \quad (4)$$




Fig. 2. Comparison of the distribution of code emergence frequency for Phase-Context pseudo-codes ($6561 = 3^8$ codes) on PIE vs. YALE-B database




Fig. 3. Comparison of the distribution of code emergence frequency for LBP codewords (256 codes) on PIE vs. YALE-B database




Fig. 4. Comparison of the distribution of code emergence frequency for LPQ codewords (256 codes) on PIE vs. YALE-B database

We computed the distribution of code emergence frequency for LBP and LPQ codes, and Phase-Context pseudo-codes in both CMU-PIE and YALE-B face databases (described below). To be able to represent the distribution of the code emergence frequency for the Phase-Context, we generate pseudo codes as $\sum_{j=1}^8 (q_j + 1) 3^{j-1}$, where $q_j \in \{-1, 0, 1\}$. The generated pseudo-codes are compatibles with LBP and LPQ codes since they are generated in the same manner as histograms.

Obviously, for Phase-Context the histogram distributions of the two databases are quite closer and similar patterns emerge (Fig. 2), however both LBP and LPQ codewords gave a much less similar characteristic signature of the faces between the two databases (Fig. 3) and (Fig. 4). Moreover, the y-axis ranges of the three sets of representation give an idea about the uniformity of the corresponding code distribution.

To measure the similarity between two face images, we calculated a scaled L_2 distance between the corresponding Phase-Context descriptors. First, we locally normalize the descriptor such as, for each frequency, it has zero-mean and unit L_2 -norm. Then, as shown by equ. (5), scaling factors were applied such as the sub-descriptor corresponding

to frequency u_i has a L2-norm equal to S_{u_i} which would put more emphasis on the frequencies that are relevant to the nature of data.

$$ScaledPC_{u_i}^{\mathcal{R},\mathcal{I}}(x) = S_{u_i} \times \left(\left(PC_{u_i}^{\mathcal{R}}(x) - \overline{PC_{u_i}^{\mathcal{R}}} \right) / \|PC_{u_i}^{\mathcal{R}}\|, \left(PC_{u_i}^{\mathcal{I}}(x) - \overline{PC_{u_i}^{\mathcal{I}}} \right) / \|PC_{u_i}^{\mathcal{I}}\| \right) \quad (5)$$

4. Experiments

Two face datasets with different lights and illuminations were used in the experiments (Fig. 5). The important statistics of these datasets are summarized below (see also table 1):

- The first dataset is the widely used CMU PIE face database¹⁴, which contains 32×32 gray scale face images of 68 persons. Each person has 42 facial images under different light and illumination conditions.
- The second dataset is the extended YALE-B face database^{15,16} which has 38 individuals with around 64 near frontal images under different illuminations per individual.

All the gray-level images were aligned by fixing the positions of the two eyes and normalized to a resolution of 32×32 pixels.



Fig. 5. Examples of aligned face images under different light and illumination conditions, CMU PIE (top) and Extended YALE-B (bottom)

Table 1. Statistics of the two databases

Dataset	Size	Dimensionality	#of persons
CMU PIE	2856	32×32	68
Extended YALE-B	2414	32×32	38

In both experiments, the Phase-Context was compared to LPQ, LBP, and HOG. In what follows, when we speak of Phase-Context we mean scaled Phase-Context. Using Principal Component Analysis (PCA), reducing the dimensionality of our Phase-Context descriptor by $4\times$ preserved over 99% of the total variance within the data for both databases.

The $LPB_{8,2}^{u_2}$ operator¹⁷ was used. For the LPQ, we use short-term Fourier transform with Gaussian window and without decorrelation. The frequency parameter of the operator is fixed to $a = 1/M$ whereas the size of the Gaussian window was set to $(M - 1)/4$ as in²⁰ with $M = 3$.

Similarly to LPQ, the LBP and HOG histograms were extracted independently from the non-overlapping 8×8 pixel blocks and then concatenated to build a global descriptor of the face. Recognition was performed using χ^2 histogram comparison. For the Phase-Context, the obtained feature vector is standardized such as, for each frequency, each vector in the data has zero-mean and a given L2-norm.

In the first experiment, we follow the same protocol as in¹¹. For each cycle, we selected randomly 23 images of each of the 68 persons in the database. One image per person was used as gallery image, and the remaining 22 images were used as probe images for testing. The mean recognition rates were then averaged over 10 000 randomized cycles. Note that the images we used were 32×32 pixels whereas authors in¹¹ cropped images to size 128×128 pixels with less extraneous background.

The average recognition rates with corresponding standard deviations (*s.d.*) for LBP, LPQ, HOG, and Phase-Context are given in table 2. Clearly, results show that the Phase-Context supports better the low resolution images with

higher accuracy (98.5 *s.d.* 0.5%) than LPQ which achieved poor performance (52.7 *s.d.* 2.4%). Whereas, LBP and HOG led to unsatisfactory recognition rates, respectively, 41.1 *s.d.* 2.2% and 37.3 *s.d.* 1.9%. Besides, a non-negative matrix factorization based approach¹⁸ that was tested exactly on the same images obtained only 75.4%.

Table 2. Recognition performance on CMU PIE

Method	GNMF ¹⁸	LBP	LPQ	HOG	Phase-Context
Rates (%)	75.4	41.1 <i>s.d.</i> 2.2	52.7 <i>s.d.</i> 2.4	37.3 <i>s.d.</i> 1.9	98.5 <i>s.d.</i> 0.5

To evaluate the effect of lighting and illumination changes on the recognition performance of each method, images were preprocessed using the Self-Quotient Image (SQI) normalization¹⁹. The accuracy of all of the methods improved with different proportions as the SQI was introduced (table 3). Phase-Context yielded the best recognition average rate (98.9 *s.d.* 0.4%) with a relative gain¹ as lower as 0.4%. The HOG feature set gave the worst performance and was more sensitive to illumination changes with a very high relative gain of 54%. This is quite natural since HOG is based on the edge information.

Table 3. Recognition performance on CMU PIE preprocessed face images

Method	LBP	LPQ	HOG	Phase-Context
Accuracy (%)	55.3 <i>s.d.</i> 2.7	67.4 <i>s.d.</i> 2.4	57.1 <i>s.d.</i> 3.4	98.9 <i>s.d.</i> 0.4
Gain (%)	34	28	54	0.4

The second experiment was conducted on the extended Yale-B dataset, with more difficult illumination conditions than PIE dataset. The recognition scores are reported in table 4. The Phase-Context still gave the best results (75.8 *s.d.* 2.1%) against (16.1 *s.d.* 1.6%) for the LPQ which gave comparable rate to HOG and LBP. We found that the scaling values $S_{u_1} = 1$, $S_{u_2} = 9$, $S_{u_3} = 25$, and $S_{u_4} = 9$ deliver the best average recognition rates over the two databases. Slightly lower results were achieved with unit scaling values. Generally, best results were obtained when less emphasis was put on frequency u_1 with small S_{u_1} .

Table 4. Recognition performance on YALE-B

Method	LBP	LPQ	HOG	Phase-Context
Rates (%)	13.8 <i>s.d.</i> 1.4	16.1 <i>s.d.</i> 1.6	11.7 <i>s.d.</i> 1.4	75.8 <i>s.d.</i> 2.1

Similarly to PIE database, the illumination normalization preprocessing on YALE-B face images improved the recognitions rate of all the methods. Obviously, from the scores of table 5, we can see that Phase-Context with (75.9 *s.d.* 2.0%) performs much better than Local Phase Quantization (21 *s.d.* 1.8%), Histogram of Oriented Gradients (21.2 *s.d.* 1.9%) or Local Binary pattern with only (14.1 *s.d.* 1.5%).

Globally, the performances achieved on YALE-B dataset explain the presence of more difficult conditions relative to the CMU-PIE database. Again, the Phase-Context achieved a quite lower gain against the other methods which means that it is less sensitive to illumination change.

Table 5. Recognition performance on YALE-B preprocessed face images

Method	LBP	LPQ	HOG	Phase-Context
Accuracy (%)	14.1 <i>s.d.</i> 1.5	21.0 <i>s.d.</i> 1.8	21.2 <i>s.d.</i> 1.9	75.9 <i>s.d.</i> 2.0
Gain (%)	2	30	81	0.13

¹ The relative gain is defined as $(\text{ACCURACY}_{\text{with pre-proc.}} - \text{ACCURACY}_{\text{without pre-proc.}}) / \text{ACCURACY}_{\text{without pre-proc.}}$.

To test the robustness of our descriptor against blur, a set of Gaussian blurring kernels $\sigma = [0 : 0.5 : 2]$ was used to generate blurred face images. We adopted the same protocol as in¹¹ for testing the methods. The probe images were blurred while the gallery images used one sharp (original) image per person.

Figures 6 and 7 show the mean recognition rates of the different methods relative to the standard deviations of the blurring kernel. $\sigma = 0$ means that sharp images were used as probes without blur.

The performance of Phase-Context is far better than HOG, LPQ, and LBP, despite the blurring kernel we used, with or without preprocessing. For preprocessed images, recognition rates drop more slightly for blurred images with $\sigma \leq 1$. For both databases, regardless preprocessing was done or not, Phase-Context produced better performance (for $\sigma \leq 1$) than all other methods even with no blur. HOG tolerated less blur than LBP and LPQ for $\sigma \leq 0.5$. All the control methods gave unsatisfactory recognition rates as blur increases from $\sigma = 1$.




Fig. 6. Average recognition rates on the CMU PIE with increasing blur, No preprocessing (left) vs. with preprocessing (right).

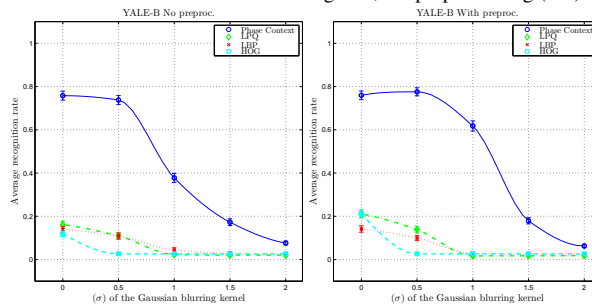


Fig. 7. Average recognition rates on the YALE-B with increasing blur, No preprocessing (left) vs. with preprocessing (right).

Furthermore, we include a comparison to the LBP+Gabor/Kernel Linear Discriminant Analysis (KLDA) approach²¹ proposed by Tan and Triggs. KLDA was used to extract discriminant information from the local appearance cues.

To be fair, we used the same standard experimental protocol from the CAS-PEAL-R1²² face database that contains 30 863 images of 1 040 individuals (595 males and 445 females). The gallery set contains 1 040 images, one image per person. As in²¹ we used the lighting probe set which contains 2 243 images. There is no overlap between the gallery and the probe sets, and all images were normalized in size to a resolution of 32×32 pixels.

Phase-Context (29.07%) was considerably more accurate than LBP+GABOR/KLDA (14.4%) which achieved more satisfying score than either feature set alone (Tab. 6). Note though that for training the KLDA on CAS-PEAL-R1, the authors used an extra training set, which contains 4 frontal images each of 300 subjects who were randomly selected from the full 1 040 subject data set.

Table 6. Comparison to a generalized LBP and Gabor feature sets.

Method	LBP/KLDA	GABOR/KLDA	LBP+GABOR/KLDA	Phase-Context
Accuracy (%)	8.5	10.87	14.4	29.07

5. Conclusion

In this paper, we addressed an important problem for face recognition which is feature extraction. We defined Phase-Context as a new feature set based on the phase of the Fourier transform. The idea is to use the four-quadrant mask of the Fourier transform phase in local neighborhoods and to reduce the influence of erroneous filter responses by privileging phases from filters with higher magnitude response since the phase of such filters is more stable. Also, the proposed Phase-Context can compensate for the order relation breakdown of the handcrafted codes generated by LPQ. Furthermore, it overcomes the discretization effect of the encoding process.

We showed that the proposed descriptor outperforms the current state-of-the-art methods in facial image analysis (e.g. LBP, LPQ, HOG). The series of tests conducted on low resolutions images from CMU-PIE, YALE-B and CAS-PEAL-R1 datasets showed that the Phase-Context is robust to lighting and illumination changes and has a high discriminative power with a significant recognition rate even with severely blurred images. It seems that these methods are more suited for higher resolutions, and that they are not adapted to be used on much smaller images.

Future work will consider the problem of face recognition in uncontrolled environments where face alignment is a crucial step.

Acknowledgments

This work has been made possible by investments through CANARIE's Research Software Program (www.canarie.ca), the AUTO21 Center of excellence (www.auto21.ca), the Ministère de l'Économie, Innovation et Exportation (MEIE) of Gouvernement du Québec, and the Natural Sciences and Engineering Research Council of Canada (www.nserc-crsng.gc.ca).

References

1. Li, S., Jain, A.: Handbook of Face Recognition. *Springer* (2005)
2. Draper, B., Baek, K., Bartlett, M., Beveridge, J.R.: Recognizing faces with PCA and ICA. *CVIU* **91**(1) (2003) 115–137
3. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: *Proc. CVPR*. (2005)
4. Lowe, D.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* **60**(2) (2004) 91–110
5. Ojala, T., Pietikainen, M., Maenpää, T.: Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on pattern analysis and machine intelligence* **24**(7) (2002) 971–987
6. Ojansivu, V., Heikkilä, J.: Blur insensitive texture classification using local phase quantization. In: *Proc. ICISP*. (2008)
7. Zhang, B., Gao, Y., Zhao, S., Liu, J.: Local derivative pattern versus local binary pattern: Face recognition with high-order local pattern descriptor. *Image Processing, IEEE Transactions on* **19**(2) (Feb 2010) 533–544
8. and A. Hadid, T.A., Pietikainen, M.: Face Recognition with Local Binary Patterns. In: *Proc. ECCV*. (2004) 469–481
9. Pietikäinen, M., Zhao, G., Hadid, A., Ahonen, T.: Computer Vision Using Local Binary Patterns. *CIV* **40**. Springer (2011)
10. Déniz, O., Bueno, G., Salido, J., De la Torre, F.: Face recognition using histograms of oriented gradients. *Pattern Recogn. Lett.* **32**(12) (September 2011) 1598–1603
11. Ahonen, T., Rahtu, E., Ojansivu, V., Heikkilä: Recognition of blurred faces using local phase quantization. In: *Proc. International Conference on Pattern Recognition*. (2008) 1–4
12. Oppenheim, A.V., Lim, J.S.: The Importance of Phase in Signals. *Proceedings of the IEEE* **69**(5) (May 1981) 529–541
13. Fleet, D., Jepson, A.: Stability of phase information. *IEEE Trans. Pattern Anal. Mach. Intelligence* **15**(12) (1993) 1253–1268
14. Sim, T., Baker, S., Bsat, M.: The CMU pose, illumination, and expression database. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **25**(12) (Dec 2003) 1615–1618
15. Georgiades, A., Belhumeur, P., Kriegman, D.: From few to many: illumination cone models for face recognition under variable lighting and pose. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **23**(6) (Jun 2001) 643–660
16. Lee, K., Ho, J., Kriegman, D.: Acquiring linear subspaces for face recognition under variable lighting. *IEEE Trans. Pattern Anal. Mach. Intelligence* **27**(5) (2005) 684–698
17. CVM : A general (LBP) implementation for Matlab, online. <http://www.cse.oulu.fi/CMV/Downloads/LBPmatlab> Accessed: March 2010.
18. Cai, D., He, X., Han, J., Huang, T.: Graph regularized nonnegative matrix factorization for data representation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **33**(8) (Aug 2011) 1548–1560
19. Wang, H., Li, S., Wang, Y., Zhang, J.: Local phase quantization for blur-insensitive image analysis. *Image and Vision Computing, Special Section: Opinion Papers*, **30**(8) (2012) 501–512
20. Rahtu, E., Heikkilä, J., Ojansivu, V., Ahonen, T.: Self quotient image for face recognition. In: *ICIP* (2004) 1397–1400
21. Tan, X., and Triggs, B.: Enhanced Local Texture Feature Sets for Face Recognition Under Difficult Lighting Conditions. In: *Image Processing, IEEE Transactions on*, **2** (6) (2010) 1635–1650
22. Cao, Bo and Shan, Shiguang and Zhang, Xiaohua and Gao, Wen : Baseline Evaluations on the CAS-PEAL-R1 Face Database. In: *Advances in Biometric Person Authentication* **3338** (2005) 370–378