Centrum voor Wiskunde en Informatica

REPORTRAPPORT

PNA

Probability, Networks and Algorithms

Probability, Networks and Algorithms

PNA

Survey of 2D face recognition methods for robust identification by smart devices

J.W.H. Tangelder

REPORT PNA-E0504 MAY 2005

CWI is the National Research Institute for Mathematics and Computer Science. It is sponsored by the Netherlands Organization for Scientific Research (NWO).
CWI is a founding member of ERCIM, the European Research Consortium for Informatics and Mathematics.

CWI's research has a theme-oriented structure and is grouped into four clusters. Listed below are the names of the clusters and in parentheses their acronyms.

Probability, Networks and Algorithms (PNA)

Software Engineering (SEN)

Modelling, Analysis and Simulation (MAS)

Information Systems (INS)

P.O. Box 94079, 1090 GB Amsterdam (NL)
Kruislaan 413, 1098 SJ Amsterdam (NL)
Telephone +31 20 592 9333
Telefax +31 20 592 4199

ISSN 1386-3711

Survey of 2D face recognition methods for robust identification by smart devices

ABSTRACT
The future generation of smart devices is envisioned to provide applications, services and smart environments that are sensitive to people's likings, personalized to their requirements, anticipatory of their behavior and responsive to their presence. A prerequisite for developing such attentive systems is the ability to recognize its user without requiring any user action, e.g. by face recognition from a stream of webcam images. Hence, recognition methods should be robust against changes in illumination, pose, facial expressions, etc. Therefore this report, surveys 2D face recognition methods with regard to applicability for robust identification using smart devices.

2000 Mathematics Subject Classification: 68T99
1998 ACM Computing Classification System: I.5.4
Keywords and Phrases: 2D face recognition; survey; robustness; smart devices
Note: Part of this research has been funded by the Dutch Bsik/BRICKS project.

Survey of 2D face recognition methods for robust identification by smart devices

Johan W.H. Tangelder
Centre for Mathematics and Computer Science (CWI)

May 4, 2005

Abstract

The future generation of smart devices is envisioned to provide applications, services and smart environments that are sensitive to people's likings, personalized to their requirements, anticipatory of their behavior and responsive to their presence. A prerequisite for developing such attentive systems is the ability to recognize its user without requiring any user action, e.g. by face recognition from a stream of webcam images. Hence, recognition methods should be robust against changes in illumination, pose, facial expressions, etc. Therefore this report, surveys 2D face recognition methods with regard to applicability for robust identification using smart devices.

1 Introduction

The future generation of smart devices is envisioned to provide applications, services and smart environments that are sensitive to people's likings, personalized to their requirements, anticipatory of their behavior and responsive to their presence. Examples of these technological developments are diverse: smart homes, access control for using services over the Internet, smart devices equipped with a web cam or camera (e.g. an intelligent mobile phone or PDA), gaming, and leisure (an intelligent family album able to browse photos of a particular person, given a few examples of mugshots containing that person [1]). User convenience, or ease of use, is an important issue in such an environment. Therefore, we require that recognition is transparent, i.e. it does not require any user specific action, like posing for a webcam. We assume that one or more webcams provide a continues flow of images containing occasionally some mugshots of a person to be recognized. Hence, a transparant face recognition system should be able to detect and recognize faces robustly from preferably one such a mugshot, or from a few of these mugshots. Therefore, recognition methods should be robust against changes in illumination, pose, facial expressions, etc.

Of course the application area of a face recognition methods, poses extract conditions, that may complicate or simplify face detection. E.g., for application in a home environment the following specific conditions have to be taken into account:

- **Frequently updates of face databases**:
  Since new faces will be frequently add to a home face database, and old ones will be deleted, limits the applicability of face methods requiring training on a full face database.
- **Uncontrolled environment**:
  In the home environment conditions like lighting conditions, facial orientation and expression cannot be controled. Hence, the quality of the obtained images is highly sensitive to the environmental factors.
- **Use of faces with approximate frontal pose**:
  People, when entering a door hall, usually look straight ahead. This allows to constrain face detection and recognition to faces with approximate frontal pose.

- **Low cost of incorrect identification**:
  Usually in a home environment security measures are not strict. Therefore, the risks of incorrect identification are not high.

- **Continuous flow of images**:
  A webcam can be applied to provide a continuous flow of images. This allows the selection of the best image(s) from a shot of a face consisting of a series of low quality images.

The problem of robust transparent face recognition in an uncontrolled environment can be formulated as follows: Given a sequence of mugshots of people using a service, maintain a list of people using the service, recognize if people belong to the list. Identify people from the list, and enroll unrecognized people.

Related problems described in the literature are

- **Watch List Face Surveillance**
  Given a watch list of faces find out if a face is on the watch list. If the face is on the watchlist, identify the face. Otherwise, reject the face [2]. Since this requires the availability of a reject option, watch list face surveillance is more difficult than recognition of a face, known to be present in the face database.

- **Face in a Crowd Recognition**
  Recognize faces from people in a crowd. Since most faces are partially visual due to occlusion [3], face in a crowd recognition is more difficult than face recognition from full frontal faces.

To improve face identification one or more of the following four preprocessing steps are recommended [4]: (1) translate, rotate and scale images to a fixed number of rows and columns of pixels, so that the centers of the eyes are placed on specific pixels; (2) apply a standard mask to remove background and hair; (3) perform histogram equalization in the masked facial pixels; (4) normalize face data to have zero mean and unit standard deviation.

Often, face identification is based on finding the nearest neighbour. In this approach, face similarity is defined directly by a distance between two images, or by a distance between the features of two images. The identity of a face image is determined as that of the most similar stored face image. The similarity of the nearest neighbour can be used as a measure of confidence in the identification. Furthermore, placing a threshold on this measure enables faces with no good matches to be categorised as unknow. For the watch list face surveillance problem, Li and Wechsler [2] propose an algorithm to set the threshold to optimize the recognition rate taking into account the reject option. An alternative to nearest neighbour similarity is using a Bayesian approach [5]. Practically, the major drawback of a Bayesian method is the need to estimate probability distributions in a high-dimensional space from very limited numbers of trainings samples per class. To avoid this problem, Moghaddam et al. [5] created a much simpler two-class problem from the multiclass problem by using a similarity based on a Bayesian analysis of image differences. Two mutually exclusive classes were defined: Ω_I, representing *i*ntrapersonal variations between multiple images of the same individual, and Ω_E, representing *e*xtrapersonal variations due to differences in identity. Assuming that both classes are Gaussian-distributed, likelihood functions $P(\Delta \mid \Omega_I)$ and $P(\Delta \mid \Omega_E)$ were estimated for a given intensity difference $\Delta = I_1 - I_2$. Given these likelihood functions and using the Maximum a Posteriori rule, two faces images are determined to belong to the same individual if $P(\Delta \mid \Omega_I) > P(\Delta \mid \Omega_E)$.

The goal of this article is to survey face recognition methods for application in such an uncontrolled environment. For a thorough introduction to face recognition we refer the reader to the book bij Gong et al. [6]. Zhao et al. [7] survey both still- and video-based face recognition methods. Kong et al. [8] survey visual and infrared face recognition methods. In contrast, we focus on transparant 2D face recognition methods. Since most face recognition systems work with grayscale images and face recognition using color images depend more on lighting conditions, we do not pay special attention to methods using color information.

For a discussion of 3D face recognition, we refer the reader to the BioSecure deliverable [9] describing the state-of-the art in face recognition.

In the next three section we discuss feature-based methods, appearance-based methods, and mixed methods for face recognition. The last section indicates directions for further research.

Face recognition methods can be classified into three groups [7].

- **Feature-based methods**
  Typically, in these methods, local features such as the eyes, nose, and mouth are first extracted and their locations (geometrical and/or appearance) are fed into a structural classifier.

- **Appearance-based methods**
  These methods, use the whole face region as the raw input to a recognition system.

- **Mixed methods**
  Combination of feature- and appearance-based methods may obtain the best results.

2 Feature-based methods

2.1 Structural matching

A number of earlier face recognition methods [10, 11] detect a set of geometrical features on the face such as the eyes, eye-brows, nose, and mouth. Properties and relations such as areas, distances, and angles between the feature points are used as descriptors for face recognition. However, the performance of face recognition depends on the accuracy of the feature location algorithm.

Another approach is to extract as features a line edge map (LEM) from a face edge map, based on a combination of template matching and geometrical feature matching [12]. The faces are encoded into binary edge maps using the Sobel edge detection algorithm. The similarity of face images is measured using a face feature scheme.

2.2 Elastic bunch graph matching

In Elastic Bunch Graph Matching (EBGM) [13] first a sparse grid is overlaid on the face image during and its nodes are "adjusted" to a set of fiducial points. The convolution of a set of 2D Gabor wavelets is computed at every grid-node and the output represents a local feature vector for that particular fiducial point of the face. Then, a flexible template comparison using a graph-matching algorithm is done. One of the most critical parts of EBMG is the accurate location of the grid-nodes at the fiducial points. To obtain a robust (to illulimition changes) and accurate detection of the fiducial points, Gonzalez-Jimenez and Albo-Castro [14] first extract ridges and valleys and then locates the fiducial points at the rigdes.

EBGM methods have good performance in general. However, they require a large-size image, e.g. 128x128. Hence, these methods can only be applied if a face is close enough to a webcam.

2.3 Hidden Markov Models

Early work of Hidden Markov Models (HMMs) [15] models human faces with a vertical top-to-bottom 1D HMM structure composed of superstates. Each superstate contains a horizontal left-to-right 1D Markov chain. A similar 1D HMM [16] uses the discrete cosine transform of the observation strip as a feature vector. In [17], the image is scanned in a zigzag fashion to form a 1D observation sequence. In embedded HMMs, an image is scanned in a 2D manner where each observation block retains vertical and horizontal indices for row and column, respectively [18, 19]. Embedded Bayesian networks, a generalized framework of embedded HMMs, show a significant complexity reduction [20]. Recently, a low-complexity 2D HMM structure [21] was derived based on the assumption of conditional independence among the neighboring observation blocks which enables the separation of the 3D state transition matrix into two 2D vertical and horizontal state transition matrices.

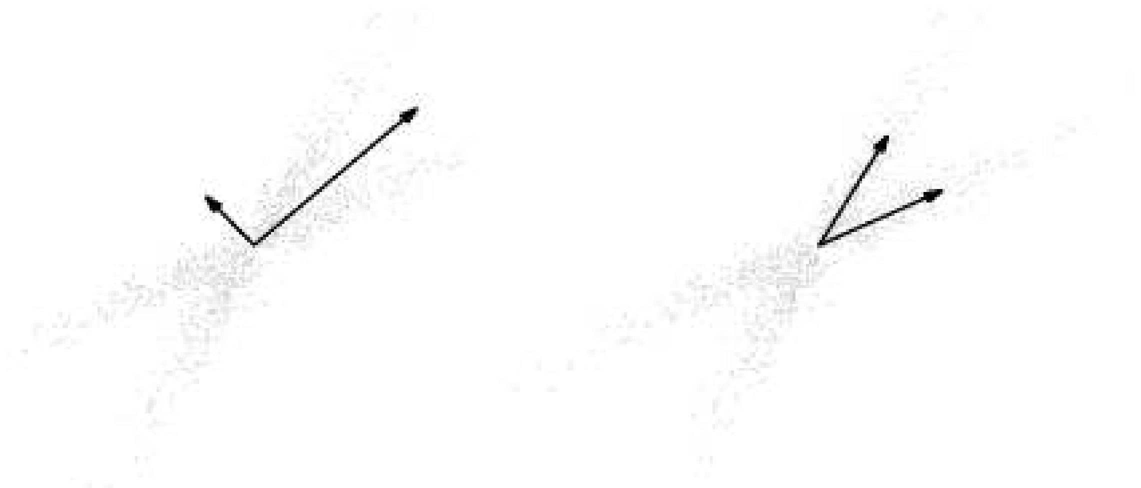

Figure 1: Example of a 2D data distribution and the corresponding principal components and independent components.

3 Appearance-based methods

3.1 Principal Components Analysis

Given images of a fixed size consisting of n columns and m rows an image can be represented by a vector of size $n \times m$. The face space is defined by the subset of images denoting faces. Principle Components Analysis (PCA) has been widely applied to capture the face space in a low-dimensional space spanned by orthogonal eigenvectors: An eigenface is computed by estimating the eigenvectors of the covariance matrix of a training set. The eigenvectors corresponding to the largest eigenvalues are taken as the principal components of an eigenface and capture the main modes of variations in the training data set. Provided that alignment and intensity normalizations are performed, PCA can also be used as a simple low-dimensional linear representation for face detection and identification [22] as follows. Given the eigenfaces, every face in the database can be represented as a vector of weights; the weights are obtained by projecting the image into the eigenface components by a simple inner product operation. When a new test image whose identification is required is given, the new image is also represented by its vector of weights. The identification of the test image can be done by locating the image in the database whose weights are the closest to the weights of the test image [22] or be using a probabilistic approach [5].

Advantages of using PCA are: (1) the reduced sensitivity to noise, (2) reduction of memory requirements, making (3) indexing in low-dimensional space much more efficient. However, (1) PCA depends on adequate position and view alignment, and (2) constant illumination. If these conditions are violated, the first eigenvectors will encode variations in pose and illumination. Therefore, PCA-based methods are not suitable for transparent face recognition under unconstrained conditions.

ICA (Independent component analysis) generates spatially localized features by producing bais vectors that are statiscally independent [23]. Figure 1 illustrates the difference between ICA and PCA. ICA is intimately related to the *b*lind source seperation problem, where the goal is to decompose an observed signal into a linear combination of unknown independent signals. All common ICA algorithms use an iterative approach to maximize independence [24].

Comparing recognition performance between PCA and ICA is complex, because differences in tasks, architecture, ICA algorithms, and distance metrics must be taken into account. The experiments by Draper et al. [24] show that for global matching tasks, requiring spatially overlaping feature vectors, like facial identify recognition, ICA performs better.

In order to have better identification, Linear Discriminant Analysis (LDA) can be applied to detect image variation due to external sources such as illumination and expression. Given a set of labelled face images, LDA aims to both reduce dimensionality and at the same time maximise the separability of different faces [25, 26]. LDA involves the eigenanalysis of a product of two matrices, where one of these matrices has to be inverted. The obtained eigenvectors, used as LDA representations bases, are called "fisherfaces". In contrast with ICA and PCA, LDA is a supervised learning technique that relies on class labels, whereas PCA is an unsupervised technique. One characteristic of both PCA and LDA is that they produce spatially global feature vectors. In other words, the basis vectors produced by PCA and LDA are non-zero for almost all dimensions, implying that a change to a single input pixel will alter every dimension of its subspace projection.

For the implementation of a face verification system on a smartcard Bourlai et. al. [27] developed a variant of the LDA method. In contrast with the conventional LDA method, their Client Specific Fisherface representation, requires only storing a single Fisher face per client on the smart card.

3.2 3D Morphable model

In this section, we briefly summarize the morphable model framework introduced by Blanz and Vetter [28, 29]. The 3D Morphable Face Model is a 3D model of faces that are learned from a set of $m = 200$ exemplar faces obtained with a laser scanner. Each exemplar face consists of n (approximately 70,000) vertices, for which complete sets of correspondences to all other faces are computed. The shape of a face is represented with a shape-vector $S = (X_1, Y_1, Z_1, X_2, \dots, Y_n, Z_n)^T \in \Re^{3n}$, which contains the X, Y, Z coordinates of its n vertices. Similarly, the texture of face is represented with a texture-vector $T = (R_1, G_1, B_1, R_2, \dots, G_n, B_n)^T \in \Re^{3n}$, where the R, G, B texture values are sampled at the same n points. A morphable model can be constructed using a data set of m exemplar faces; exemplar i is represented by the shape-vector S_i and the texture vector T_i. New shapes s and textures t can be generated by convex combinations of the shapes and textures of the m exemplar faces: $s = \sum_{i=1}^{m-1} a_i S_i$, $t = \sum_{i=1}^{m-1} b_i T_i$, $\sum_{i=1}^{m} a_i = \sum_{i=1}^{m} = 1$. To reduce the dimensionality of the shape and texture spaces, Principle Component Analysis (PCA) is applied separately on the shape and texture spaces:

$$s = \bar{s} + \sum_{i=1}^{m-1} \alpha_i \sigma_{s,i} s_i, t = \bar{t} + \sum_{i=1}^{m-1} \beta_i \sigma_{t,i} t_i, \tag{1}$$

Face recognition in the Morphable Face Model framework is based on fitting an image iteratively to a 3D Face by finding suitable model coefficients α_i and β_i. Moreover, the iterative process estimates the pose and illumination parameters. For initialization the system requires the user to select image coordinates of about seven feature points using an interactive tool.

Zhang and Samaras [30] enhanced the robustness to lighting conditions of the Morphable Face Model using a spherical harmonic representation for the illumination. They achieved high recognition rates for images under a wide range of illumination conditions, including multiple sources of illumination. Also, their method requires human intervention to select feature points.

Jiang et al. [31] propose an analysis-by-synthesis framework for face recognition with varying pose, illumination and expression. First, an efficient 2D-to-3D integrated face reconstruction approach is introduced to reconstruct a personalized 3D face model from a single frontal face image with neutral expression and normal illumination. Next, realistic virtual faces with different PIE are synthesized based on the 3D face to characterize the face subspace. Finally, face recognition is conducted based on these representative virtual faces.

To apply Morphable Face Models to transparant face identification, human intervention has to be avoided. Also, the high requirements on memory and processing time, make these methods less suitable to transparant face identification.

4 Mixed methods

Mixed methods try to improve recognition by combining feature-based methods and appearance-based methods.

4.1 Active Statistical Face Models

Active shape models introduced by Cootes et al. [32] are statistical shape models combining 1) a global Point Distribution Model (PDM) modeling the shape of an object and its variations using a set of landmark points, and 2) a set of Local Grey-Level (LGL) models capturing the local grey-level variations observed at each landmark point.

A PDM can be used to represent the shape of a face as a set of n labelled landmark points in a vector $x = (x_1, y_1, \ldots, x_n)$. The modes of variations of the face shape are captured by applying PCA on the deviation of the M example faces from the mean face shape. PCA is also used to model the Local Grey-Level (LGL) at the location of each landmark point of the PDM. An LGL model is learned for each landmark and together with the PDM these are used to construct an active shape model (ASM). An iteration search process is used to match an ASM to a novel face image based on the PDM model and the LGL models learned from the training examples.

Active appearance models (AAMs) PCA methods also represent the shape of a face as a set of n labelled landmark points in a vector $x = (x_1, y_1, \ldots, x_n)$. Like ASMs they use a global Point Distribution Model (PDM) modeling the shape of an object and its variations using a set of landmark points. In contrast to ASMs, the LGLs are replaced by 2) a texture variation model (TVM) modeling the textures observed at each landmark point. Both the PDM and TVM are represented by a PCA based represention. To some extent, AAM may give a quite good match to image texture, but when target image and the background vary significantly, it is still unable to locate feature points accureately. Meanwhile, both the training process and the searching procedure of AAM are quite complex and slow.

Chen et al. [33] present a method using a fast classifier to locate feature points candidates with a probabilistic output.

4.2 Other mixed methods

Chen et al. [34] use HMMs to model classes of face images and a set of Fisher scores is calculated through partial derivative analysis of the parameters estimated in each HMM. These Fisher scores are further combined with classic log-likelihood and appearance-based features to form feature vectors that exploit the advantages of both local and holistic features computed from a human face. Linear discriminant analysis (LDA) is then applied to analyze these feature vectors for face recognition. Performance improvements are observed over a stand-alone HMM method and a Fisher face method that uses only appearance-based feature vectors.

5 Conclusions and further research

First we summarize the discussed face recognition methods.

Appearance-based methods

Principle Components Analysis (PCA) has been widely and very succesfully applied to capture the face space in a low-dimensional space, spanned by orthogonal eigenvectors. ICA (Independent Component Analysis) methods generalize ICA methods by producing bais vectors that are statiscally independent, but not necessarily orthogonal. Draper et al. [24] show that for global matching tasks, requiring spatially overlaping feature vectors, like facial identify recognition, ICA performs better. But the applicability of PCA and ICA to robust face recognition is limited to environments with constant pose, illumination and expression, Therefore, both PCA and ICA are unsuitable for transparant face recognition under unlimited conditions. Linear Discriminant Analysis (LDA) can be applied to take into account image variation due to pose, illumination and expression.

Since the PCA, ICA, and LDA appearance based methods using 2D data are not applicable to robust 2D face identification, methods using a 3D morphoble face models have been introduced, which are robust to changes in pose and illumination. But to apply Morphable Face Models to transparant face identification, human intervation has to be avoided. Also, the high requirements on memory and processing time, make these methods less suitable to transparant face identification.

In general, appearance-based methods provide accurate recognition results with standard, well-illuminated frontal mug-shot images. However, performance often degrades rapidly with pose changes, non-uniform illumination, and background clutter.

Feature-based methods

Early face recognition methods apply structural matching to detect a set of geometrical features on the face such as the eyes, eye-brows, nose, and mouth. Properties and relations such as areas, distances, and angles between the feature points are used as descriptors for face recognition. However, the performance of these methods depends on the accuracy of the feature location algorithm. This methods evolved to elastic bunch graph matching methods, which have a good performance in general. However, they require a large-size image, e.g. 128x128. Hence, these methods can only be applied if a face is close enough to a webcam.

Mixed methods

For Active Shape Models (ASM) and Active Appearance Models (AAM), both the training process and the searching procedure of AAM are quite complex and slow.

From this summary we identify the following research issues for transparant face recognition under unlimited conditions.

- **Fusion with other biometrics** Fusion with other modalities enhances face recognition capabilities.

- **Robust feature point detection**
  Both EBGM and AAM/ASM methods require robust feature point extraction. Based on the first promising results [33, 14] more research is needed.

- **Robustness to expression variant faces**
  How can we robustly recognize a person's face for whom two face images differ in facial expression?

- **Robustness to illumination**
  How can we robustly recognize a person's face for whom two face images differ in illumination?

- **Robustness to pose**
  How can we robustly recognize a person's face for whom two face images differ in pose?

- **Robustness to occlusion**
  How can we robustly recognize a person's face for whom one or two face are partially occluded?

- **Face recognition with limited device capabilities**
  How can we robustly apply face recognition using smart devives with limited CPU power, storage capacity and bandwith?

Since the conditions for face recognitions depend on pose, facial expresssions, and the the environment, there will be no best face recognition method in general. Therefore, future research at CWI will focus on fusion of different modalities, fusion over time, and using 3D face models to enhance 2D face recognition capabilities.

References

[1] Zhang, L., Hu, Y., Li, M., Ma, W., Zhang, H.: Efficient propagation for face annotation in family albums. In: Proc. ACM Multimedia. (2004)

[2] Li, F., Wechsler, H.: Watch list face surveillance using transductive inference. In: ICBA. (2004) 23–29

[3] Gutta, S., Wechsler, H.: Face recognition using asymmetric faces. In: ICBA. (2004) 162–168

[4] Phillips, P., Moon, H., Rizvi, S., Rauss, P.: The FERET evaluation methodology for face-recognition algorithms. IEEE Transactions on Pattern Analysis and Machine Intelligence **22** (2000) 1090–1104

[5] B., M., A., P.: Probabilistic visual learning for object representation. IEEE Transactions on Pattern Analysis and Machine Intelligence **19** (1997) 696–710

[6] Gong, S., McKenna, S.J., Psarrou, A.: Dynamic Vision: From Images to Face Recognition. Imperial College Press, World Scientific Publishing (2000)

[7] Zhao, W., Chellappa, R., Phillips, P., Rosenfeld, A.: Face recognition: A literature survey. ACM Computing Surveys **35** (2003) 399–458

[8] Kong, S.G., Heo, J., Abidi, B.R., Paik, J., Abidi, M.A.: Recent advances in visual and infrared face recognition. Computer Vision and Image Understanding **97** (2005) 103–135

[9] Tistarelli, M.: State of the art in face recognition. Technical report, BioSecure Network of Excellence (2005) Deliverable D7.2.1.

[10] Cox, I., Ghosn, J., Yianilos, P.: Feature-based face recognition using mixture-distance. In: Proc. Int. Conf. Comput. Vis. Patt. Recogn. (1996) 209 – 216

[11] Manjunath, B., Chellappa, R., von der Marsburg, C.: A feature based approach to face recognition. In: Proc. IEEE Conf. Comput. Vis. Patt. Recogn. (1992) 373 – 378

[12] Gao, Y., Leung, M.: Face recognition using line edge map. IEEE Transactions on Pattern Analysis and Machine Intelligence **24** (2002) 764 – 779

[13] Wiskott, L., Fellous, J., Kruger, N., von der Malsburg, C.: Face recognition by elastic bunch graph matching. IEEE Transactions on Pattern Analysis and Machine Intelligence **19** (1997) 775 – 779

[14] Gonzalez-Jiminez, D., Alba-Casto, J.: Frontal face authentication through creaseness-driven gabor jets. In: ICIAR 2004, LNCS 3212. (2004) 660–667

[15] Samaria, F.: Face Recognition using Hidden Markov model. PhD thesis, Univ. of Cambridge (1994)

[16] Nefian, A., Hayes III, M.: Hidden markov models for face recognition. In: Proc. Int'l Conf. Acoustics, Speech, and Signal Processing, vol. 5. (1998) 2771–2774

[17] Kohir, V., Desai, E.: Face recognition. In: Proc. Int'l Conf. Image Processing. (2000) 309–312

[18] Nefian, A., Hayes III, M.: An embedded HMM for face detection and recognition. In: Proc. Int'l Conf. Acoustics, Speech, and Signal Processing, vol. 6. (1999)

[19] Nefian, A., M.H., H.: Maximum likelihood training of the embedded HMM for face detection and recognition. In: Proc. Int'l Conf. Image Processing. (2000)

[20] Nefian, A.: Embedded bayesian networks for face recognition. In: Proc. IEEE Int'l Conf. Multimedia and Expo. (2002) 133–136

[21] Othman, H., Aboulnasr, T.: A separable low complexity 2D HMM with application to face recognition. IEEE Trans. Patt. Anal. Mach. Intell. **25** (2003) 1229–1238

[22] Turk, M., Pentland, A.: Eigenfaces for recognition. Journal of Cognitive Neuroscience **3** (1999) 71 – 86

[23] Bartlett, M.: Face Image Analysis by Unsupervised Learning. Kluwer Academic publishers (2001)

[24] Draper, B.A., Vaek, K., Bartlett, M.S., Beveridge, J.R.: Recognizing faces with PCA and ICA. Computer Vision and Image Understanding **91** (2003) 115–137

[25] Swets, D., Weng, J.: Using discriminant eigenfeatures for image retrieval. IEEE Transactions on Pattern Analysis and Machine Intelligence **18** (1996) 831 – 836

[26] Bellhumeour, P., Hespanha, J., D.J., K.: Eigenfaces vs. Fisherfaces: recognition using class specific linear projection. IEEE Transactions on Pattern Analysis and Machine Intelligence **19** (1997) 711 – 720

[27] Bourlai, T., Messer, K., Kittler, J.: Performance versus computional complexity trade-off in face verification. In: ICBA. (2004) 169–177

[28] Blanz, V., Vetter, T.: A morphable model for the synthesis of 3D faces. In: SIGGRAPH. (1999) 187–194

[29] Blanz, V., Vetter, T.: Face recognition based on fitting a 3D morphable model. IEEE Transactions on Pattern Analysis and Machine Intelligence **25** (2003) 1063 – 1074

[30] Zhang, L., Samaras, D.: Pose invariant face recognition under arbitrary unknown lighting using spherical harmonics. In: Proc. Biometric Authentication. LNCS 3087 (2004) 10–23

[31] Jiang, D., Hu, Y., Yan, S., Zhang, L., Zhang, H., Gao, W.: Efficient 3D reconstruction for face recognition. Pattern Recognition **38** (2005) 787–798

[32] Cootes, T., Taylor, C., Cooper, D., Graham, J.: Active shape models - their training and application. Computer Vision and Image Understanding **61** (1995) 38 – 59

[33] Chen, L., Zhang, L., Zhang, H., Abdel-Mottaleb, M.: 3D shape constraint for facial feature localization using probabilistic-like output. In: FGR'04. (2004)

[34] Chen, L., Man, H., Nefian, A.V.: Face recognition based on multi-class mapping of Fisher scores. Pattern Recognition **38** (2005) 799–811