

Face Recognition under Real-world Conditions

Ladislav Lenc¹, Pavel Král¹

¹*Department of Computer Science and Engineering, University of West Bohemia, Plzeň, Czech Republic
{llenc, pkral}@kiv.zcu.cz*

Keywords: Automatic Face Recognition, Czech News Agency, Confidence Measures, Gabor Wavelets, Scale Invariant Feature Transform

Abstract: This paper deals with Automatic Face Recognition (AFR). The main contribution of this work consists in the evaluation of our two previously proposed AFR methods in real conditions. At first, we compare and evaluate the recognition accuracy of two AFR methods on well-controlled face database. Then we compare these results with the recognition accuracy on a real-world database of comparable size. For such comparison, we use a sub-set of the newly created Czech News Agency (ČTK) database. This database is created from the real photos acquired by the ČTK and the creation of this corpus represents the second contribution of this work. The experiments show the significant differences in the results on the controlled and real-world data. 100% accuracy is achieved on the ORL database while only 72.7% is the best score for the ČTK database. Further experiments show, how the recognition rate is influenced by the number of training images for each person and by the size of the database. We also demonstrate, that the recognition rate decreases significantly with larger database. We propose a confidence measure technique as a solution to identify and to filter-out the incorrectly recognized faces. We further show that confidence measure is very beneficial for AFR under real conditions.

1 INTRODUCTION

Nowadays, Automatic Face Recognition (AFR) has been one of the most progressive biometric methods. The spectrum of the possible applications of the AFR is really broad. It reaches from the surveillance of people (police), access control to restricted areas to the labeling of the photos used in the recently very popular photo sharing applications or in the social networks.

Many algorithms for the face recognition were proposed. Most of them perform well under certain “good” conditions (face images are well aligned, the same face pose and lighting conditions, etc.). However, their performance is significantly degraded when these conditions are not accomplished. Many methods have been introduced to handle these limitations, but only few of them perform satisfactorily in a fully uncontrolled environment. The main goal of this paper is thus to evaluate two previously proposed methods (Lenc and Král, 2012a; Lenc and Král, 2012b) under real conditions. The evaluation will be performed on a huge database owned by the Czech News Agency (ČTK).

To improve the results of the AFR methods, using the confidence measures should be beneficial. There-

fore, we further use the previously proposed confidence measure approach (Lenc and Král, 2011) in the post-processing step. This way we can classify certain amount of faces with high recognition accuracy while the rest of faces remains unclassified.

The results of this work will be used by the ČTK in order to label the unlabeled photos in the large photo-database (about 2 millions pictures). Note that only few labeled images of every person are available. Unfortunately, the labeled images don’t contain only the head/face of the person, but also some other additional useless information (other people, background objects, etc.). Therefore, another goal of this work consists in the proposition of the algorithm for the corpus creation.

The following section gives an overview of some important AFR methods. Section 3 describes in more detail the two AFR methods we used in this work; namely the adapted Kepenekci method and the SIFT based Kepenekci approach. Confidence measures are also described in this section. Section 4 first introduces the corpora used for testing. Then, the algorithm for the corpus creation is presented. Further, there is the description of the experiments performed to evaluate the methods. In the last section we discuss the achieved results and give the direction of the

further research.

2 RELATED WORK

One of the first successful approaches is the Principal Component Analysis (PCA), so called Eigenfaces (Turk and Pentland, 1991). It is a statistical method that takes into account the whole image as a vector. Eigenvectors of the image matrix are used for face representation. The method is sensitive on variations in lighting conditions, pose and scale.

Another method, Fisherfaces, is derived from the Fisher's Linear Discriminant (FLD). A description of this method is presented in (Belhumeur et al., 1997). Similarly to the Eigenfaces, the Fisherfaces project an image into another, less dimensional, space. According to the authors, this approach should be insensitive to changing lighting conditions.

Another group of approaches use Neural Networks (NNs). Several NNs topologies were proposed. One of the best performing methods based on neural networks is presented in (Lawrence et al., 1997).

Hidden Markov Models (HMMs) are also successfully used for AFR (Samaria and Young, 1994). The method was tested on a dataset containing 5 images of each of the 24 individuals. Indicated recognition rate of this approach is 84%. For comparison, the Eigenfaces were tested using the same dataset and the recognition rate of 74% is reported.

In (Nefian and Hayes, 1998) another HMM-based approach is described. It is stated there, that the method significantly reduces the computational complexity in comparison with the older methods while the recognition rate remains the same.

In the last couple of years, several successful approaches based on Gabor wavelets were introduced (Shen and Bai, 2006). One of the first such methods was proposed by Lades (Lades et al., 1993). Some approaches (Shen, 2005) also combine the pre-processing with Gabor wavelets with well established methods such as Eigenfaces, Fisherfaces, etc. Another successful approach described in (Wiskott et al., 1999; Bolme, 2003) is Elastic Bunch Graph Matching (EBGM). Kepenekci proposes in (Kepenekci, 2001) an algorithm that outperforms the classical EBGM. Moreover, he addresses the main issue of the Elastic Bunch Graph Matching, manual labeling of the landmarks.

Recently, also the Scale Invariant Feature Transform (SIFT) (Lowe, 1999) is utilized for face recognition. The algorithm was originally developed for object recognition. The SIFT features has the ability to detect and describe local features in images. The

features are invariant to image scaling, translation and rotation. Moreover, they are also partly invariant to changes in illumination. When used in the AFR, the SIFT feature vectors of the reference and test images are compared using the Euclidean distance.

One of the first applications of this algorithm for the AFR is proposed in (Aly, 2006). It takes the original SIFT algorithm and creates the set of descriptors as described in Section 3. Another approach called Fixed-key-point-SIFT (FSIFT) is presented in (Krizaj et al., 2010).

3 METHOD DESCRIPTION

This section details the two AFR methods we used in this work: the adapted Kepenekci method (Lenc and Král, 2012a) and the SIFT based Kepenekci approach (Lenc and Král, 2012b).

3.1 Adapted Kepenekci method

3.1.1 Gabor Filter

Gabor filter is a sinusoid modulated with a Gaussian. A basic form of the two dimensional Gabor filter is shown in equation 1.

$$g(x, y; \lambda, \theta, \psi, \sigma, \gamma) = \exp\left(-\frac{\hat{x} + \gamma^2 \hat{y}^2}{2\sigma^2}\right) \cos\left(2\pi \frac{\hat{x}}{\lambda} + \psi\right) \quad (1)$$

where $\hat{x} = x \cos \theta + y \sin \theta$, $\hat{y} = -x \sin \theta + y \cos \theta$, λ is the wavelength of the cosine factor, θ represents the orientation of the filter and ψ is a phase offset, σ and γ are parameters of the Gaussian envelope, σ is the standard deviation of the Gaussian and γ defines the ellipticity (aspect ratio) of the function.

The set of 40 Gabor filters with different frequencies and orientations is used to extract local features. We consider only the real part of the wavelet response.

3.1.2 Face Representation

The image is convolved with all Gabor filters from the filter bank. 40 wavelet responses R_j , where $j = 1, \dots, 40$, are obtained. Each of these responses is scanned with a sliding window. Assume a square window W of size $w \times w$. All possible window positions within the response are evaluated. The center of the window, denoted (x_0, y_0) , is considered to be a fiducial point iff:

$$R_j(x_0, y_0) = \max_{(x, y) \in W} R_j(x, y) \quad (2)$$

$$R_j(x_0, y_0) > \frac{1}{w_i * h_i} \sum_{x=1}^{w_i} \sum_{y=1}^{h_i} R_j(x, y) \quad (3)$$

where $j = 1, \dots, 40$, w_i and h_i are image width and height respectively. The feature vector in point (x, y) is created as follows:

$$v(x, y) = \{x, y, R_1(x, y), \dots, R_{40}(x, y)\} \quad (4)$$

The resulting vector thus contains information about feature point coordinates and values of Gabor responses in this point.

3.1.3 Face Comparison

The cosine similarity (Tan et al., 2005) is employed for vector comparison. The similarity between two vectors thus takes the values in interval $[0, 1]$. Only the last 40 positions in the vector are considered.

Let us call T a test image and G a gallery image. For each feature vector t of the face T we determine a set of relevant vectors g of the face G . Vector g is relevant iff:

$$\sqrt{(x_t - x_g)^2 + (y_t - y_g)^2} < distanceThreshold \quad (5)$$

If no relevant vector to vector t is found, vector t is excluded from the comparison procedure. However, the most similar vector (from the relevant vector set) is used for the face similarity computation. The overall similarity of two faces OS is computed as an average of similarities between each pair of corresponding vectors as:

$$OS_{T,G} = mean\{S(t, g), t \in T, g \in G\} \quad (6)$$

Then, the face with the most similar vector to each of the test face vectors is determined. The variable C_i says how many times the gallery face G_i was the closest to some of the vectors of test face T . The similarity is computed as C_i/N_i where N_i is the total number of feature vectors in G_i . Weighted sum of these two similarities is used for similarity measure:

$$FS_{T,G} = \alpha OS_{T,G} + \beta \frac{C_G}{N_G} \quad (7)$$

The size of the sliding window is very important for the performance of this method. It determines the number of fiducial points detected and influences its accuracy. The higher the window size the less fiducial points are detected. On the other hand, searching larger window needs more computation time. In the comparison stage, the number of fiducial points determines the time needed. The above mentioned

threshold *distanceThreshold* also influences the accuracy and the run-time of this method. The smaller the value of this threshold is, the less comparisons are needed and the method works faster.

If more than one training example per person is used, we create a so called “composed face”. It means that all vectors extracted from the images of one person are put together and the resulting set of vectors is used as a gallery face.

3.2 SIFT based Kepenekci method

SIFT

The SIFT algorithm has basically four steps: extrema detection, removal of key-points with low contrast, orientation assignment and descriptor calculation (Krizaj et al., 2010).

To determine the key-point locations, an image pyramid with re-sampling between each level is created. It ensures the scale invariance. Each pixel is compared with its neighbors. Neighbors in its level as well as in the two neighboring (lower and higher) levels are examined. If the pixel is maximum or minimum of all the neighboring pixels, it is considered to be a potential key-point.

For the resulting set of key-points their stability is determined. Locations with low contrast and unstable locations along edges are discarded.

Further, the orientation of each key-point is computed. The computation is based upon gradient orientations in the neighborhood of the pixel. The values are weighted by the magnitudes of the gradient.

The final step is the creation of the descriptors. The computation involves the 16×16 neighborhood of the pixel. Gradient magnitudes and orientations are computed in each point of the neighborhood. Their values are weighted by a Gaussian. For each sub-region of size 4×4 (16 regions), the orientation histograms are created. Finally, a vector containing 128 (16×8) values is created.

The original SIFT algorithm is described in detail in (Lowe, 1999; Lowe, 2004b) and (Krizaj et al., 2010). An example of implementation of the SIFT algorithm can be found in (Lowe, 2004a).

Face comparison

To describe a face we create a set of the SIFT features. For the face comparison the same matching scheme is used as in the Kepenekci method 3.1.3. We use also a composed face (see above) if more than one training example per person is available.

3.3 Confidence Measure

As in many other works (Lleida and Rose, 1996; Jiang, 2005), our confidence measure for automatic face recognition is an estimate of the *a posteriori* class probability.

Let the output of our classifier be $P(F|C)$, where C is the recognized face class and F represents the face features. The likelihoods $P(F|C)$ are normalized to compute the *a posteriori* class probabilities as follows:

$$P(C|F) = \frac{P(F|C).P(C)}{\sum_{I \in \mathcal{FIM}} P(F|I).P(I)} \quad (8)$$

\mathcal{FIM} represents the set of all faces and $P(C)$ denotes the *prior* probability of the face class C .

The *relative confidence value* method as already presented in (Lenc and Král, 2011) is used. This approach computes the difference between the *best* hypothesis and the *second best* one:

$$P\Delta = P(\hat{C}|F) - \max_{C \neq \hat{C}} (P(C|F)) \quad (9)$$

Only the faces with $P\Delta > T$ are accepted. This approach aims at identifying the “dominant” faces among all the other candidates. T is an acceptance threshold and its optimal value is adjusted experimentally.

4 EXPERIMENTAL SETUP

4.1 Corpora

ORL Database

The ORL database was created at the AT & T Laboratories¹. It contains images of 40 individuals. 10 pictures for each person are available. The images contain just the face and have black homogeneous background. They may vary due to three following factors: 1) time of acquisition; 2) head size and pose; 3) lighting conditions. The size of the pictures is 92×112 pixels. The further description of this database is in (Li and Jain, 2005).

Czech News Agency (ČTK) Database

This face data-set has been created using the algorithm described in the next section. We created this corpus in order to evaluate the AFR methods on real-world data. All faces are extracted from the photographic data owned by the ČTK.



Figure 1: Examples of one face from the ČTK face corpus

The number of photographs containing human faces reaches 2 millions whereas the count of individuals is about 20 thousand. Out of this amount we created a sub-set containing faces of 1065 individuals. At least 4 images for each person are available. The pictures size is 128×128 pixels and have different background. The number of the face examples per person varies from 1 to 10 images.

Figure 1 shows an example of one person from this corpus.

This corpus has some important characteristics:

- large number of the different people,
- time span between images of one individual is more than 20 years,
- pose, tilt and lighting conditions vary,
- photos taken under real conditions,
- difficulties for the automatic face recognition.

4.2 Corpus Creation Algorithm

Unfortunately, the labeled part of the ČTK corpus contains not only the face in the photo (a whole person, more faces, some background objects, etc. available). Therefore, the pre-processing of the photos is necessary. In order to minimize the human costs, an semi-automatic corpus creation algorithm is proposed. It is composed of the following tasks:

1. face detection,
2. eyes detection,
3. head rotation,
4. image resizing and conversion to the grayscale,
5. identification and deletion of the incorrectly detected faces,
6. manual verification

To detect and to extract the faces we used the face detector implemented in the OpenCV library <http://opencv.willowgarage.com/wiki/>. It uses the well-known Viola-Jones (Viola and Jones, 2001) face detection algorithm. Once the faces in the image are detected, we try to correct the position and rotation of the head. Using Viola-Jones boosted cascade classifier trained for eye detection we determine the position of the eyes. If both eyes are successfully detected, we rotate and scale the face so that a horizontal position of the eyes is fixed. If the detection of the eyes is not successful the face is used as it is. Next, we resize the faces to the size of the 128×128 pixel and convert it to the grayscale.

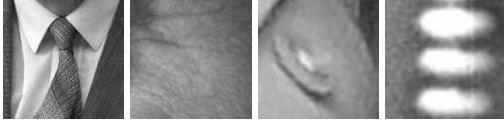


Figure 2: Examples of the incorrectly identified faces by the OpenCV face detector.

Unfortunately, the face detector identifies also a certain amount (up to 50 % depending on the quality of images) of non-face images (false positives errors, see Figure 2). These false detections must be removed. We used a meta-classifier for this task in order to distinguish two classes: face and non-face. We chose a neural network of the type Multi-Layer Perceptron (MLP) due to its simplicity and good classification results in our previous experiments (Král et al., 2006). The proposed MLP has three layers containing 1024 (size of the input face vector), 10 and 2 (face vs. non-face) nodes respectively. The input vector is created by scanning the image by window of the size 4×4 pixels. An average intensity value in the window is computed. To train the MLP, we used a set of manually selected face and non-face (250 faces and 250 non-faces) examples. The resulting corpus has been checked and corrected manually.

Figure 2 shows examples of false positives found by the face detector.

4.3 Experiments

The following experiments are proposed and evaluated in order to show:

1. the performance of the methods in the laboratory conditions,
2. a significant degradation of the recognition accuracy when the laboratory conditions are replaced with the real ones,
3. the further degradation when the number of recognized individuals is increasing and the number of the training examples is decreasing,
4. that the use of the confidence measure is beneficial in order to filter out the incorrectly recognized examples.

Both methods: adapted Kepenekci method (see Section 3.1) and the SIFT based Kepenekci approach (see Section 3.2) are evaluated on the two corpora described above. The cross-validation procedure is used to evaluate the approaches in all experiments.

Table 1: Recognition results on the ORL database

Method	Kepenekci	SIFT
Training Ex.	Recognition rate [%]	
9 of 10	99.50	100
8 of 10	99.58	99.86
7 of 10	98.75	100
6 of 10	98.39	99.82
5 of 10	97.42	99.50
4 of 10	96.00	99.17
3 of 10	93.48	97.41
2 of 10	90.63	91.88
1 of 10	78.89	81.52

Accuracy in the laboratory conditions

Table 1 shows the recognition rate of both methods on the ORL database when different number of training images (1-9) is used and the rest of images is used for testing. This table shows that the recognition scores of both approaches are comparable on the “small” and “artificial” corpus. Moreover, the obtained recognition rate is close to 100% when more examples is used. This experiment has been performed in order to show that both approaches perform well in the laboratory conditions on a small corpus.

Table 2: Recognition results on the ČTK database

Method	Kepenekci	SIFT
Training Ex.	Recognition rate [%]	
9 of 10	60.81	72.70
8 of 10	57.66	69.07
7 of 10	53.83	65.20
6 of 10	50.10	62.36
5 of 10	47.12	57.21
4 of 10	42.88	51.17
3 of 10	37.55	44.21
2 of 10	32.09	34.57
1 of 10	24.62	22.37

Accuracy in the real conditions with a small number of individuals

Table 2 displays the results of the methods on a subset of the real ČTK corpus of the comparable size as the ORL database when the number of training examples varies from 1 to 9. The sub-set contains 37 individuals, 10 images for each person. This experiment shows that the SIFT based Kepenekci approach outperforms significantly the adapted Kepenekci method (the difference about 11% when 9 training examples used). Therefore, we use only the SIFT method for further experiments.

Accuracy in the real conditions with a huge number of individuals

Table 3 details the relation among the recognition accuracy, the number of the individuals and the number of the training examples. This table shows that

¹<http://www.cl.cam.ac.uk/research/dtg/attarchive/facedatabase.html>

Table 3: Recognition results of the SIFT approach on the ČTK database using different number of training examples and the different amount of individuals

Training Ex.	Database size (individuals #)	Recognition rate [%]
9 of 10	37	72.70
8 of 9	88	56.70
7 of 8	194	49.94
6 of 7	367	41.92
5 of 6	595	33.51
4 of 5	841	27.90
3 of 4	1065	21.31

the recognition rate decreases significantly when the number of individuals increases and the number of training images decreases. Therefore, we use the confidence measure (see Section 3.3) to post-process the recognition results and to filter out the incorrectly recognized faces.

Confidence measure

Two experiments with the confidence measure are realized:

1. The number of individuals in the first experiment is close to 200 people. This number is chosen because it is sufficient in the current version of our AFR system for the needs of the ČTK. The number of training examples is 7, one example is used for testing.
2. In the second experiment, we evaluate the method on our whole corpus. The number of individuals is thus 1065 people (3 training examples).

Figure 3 depicts the recognition rate and the number of the classified faces of the first experiment (194 different individuals). The recognition accuracy without any confidence measure is only about 50%. This figure shows, that the recognition accuracy is close to 100% when more than 20% of images is classified. This result is interesting for our application because the main accent is to the correct classification. The amount of 20% of the correctly recognized faces represents a significant part of the corpus.

Figure 4 plots the recognition rate and the number of the classified faces of the second experiment (1065 different individuals). Without any confidence measure, the recognition rate is only 21%. This figure shows, that the recognition accuracy is close to 80% when more than 20% of images is classified. This result is not the perfect one. However, it is very promising and one of our future research directions focuses thus on the proposition of better confidence measure techniques.

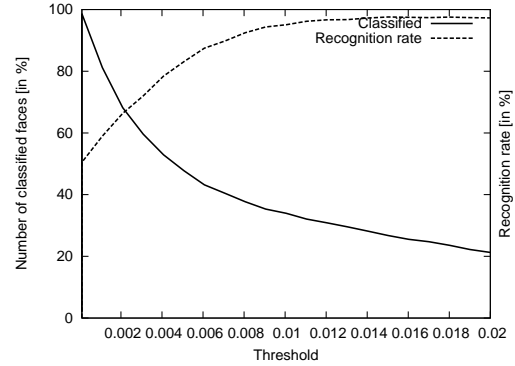


Figure 3: Face recognition rates and the numbers of the identified faces when 7 of 8 training examples used ($T \in [0; 1]$ and 194 individuals are recognized).

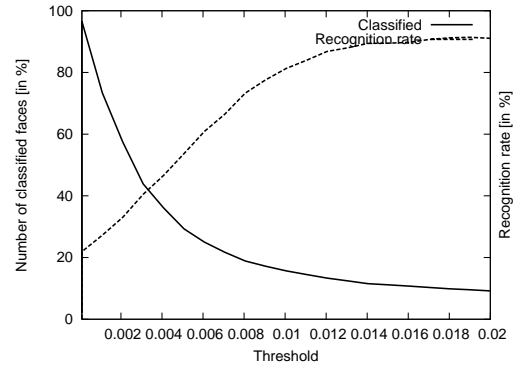


Figure 4: Face recognition rates and numbers of the identified faces when 3 of 4 training examples used ($T \in [0; 1]$ and 1065 individuals are recognized).

5 CONCLUSIONS AND PERSPECTIVES

In this paper, we have presented the difficulties of the AFR on the faces acquired in uncontrolled environment. Two previously proposed AFR methods (adapted Kepenekci method and the SIFT based Kepenekci approach) are described and evaluated on the ORL database (laboratory conditions) and on the sub-set of the ČTK database (real-world corpus). We have shown that the accuracy of both methods is close to 100% on the ORL database. The results on ČTK database show that the SIFT based Kepenekci method has significantly better recognition accuracy. It reaches 72.7% whereas the adapted Kepenekci method has recognition rate of 60.8% (9 training examples are used).

We further analysed the accuracy of the SIFT based Kepenekci approach if the corpus size increases and the number of training examples decreases. As supposed, the recognition rate decreases with increasing number of different recognized individuals. We

have also shown that the number of training examples influences the accuracy significantly. While 9 training examples is used, the recognition rate is 72.7%. If we use only 3 training examples the recognition rate is only 21.3%.

In the last two experiments we employed the confidence measure to post-process the recognition results. We compared the results when 7 training examples and 3 examples are used. The results show that using confidence measure is very beneficial for AFR under real-world conditions.

It is obvious that the AFR methods are nowadays capable to recognize faces perfectly under the condition: the acquisition of the face images must be controlled. If this condition is not accomplished, the task is much more difficult. Therefore the perspectives of the further work on recognition of real-world data lay more in the detection step than in the recognition itself. Further increase of image quality will ensure much better accuracy of the recognition. Using higher quality images and utilising the confidence measure will help to create a reliable recognition system.

ACKNOWLEDGEMENTS

This work has been partly supported by the UWB grant SGS-2010-028 Advanced Computer and Information Systems and by the European Regional Development Fund (ERDF), project NTIS - New Technologies for Information Society, European Centre of Excellence, CZ.1.05/1.1.00/02.0090. We also would like to thank Czech New Agency (ČTK) for support and for providing the photographic data.

REFERENCES

- Aly, M. (2006). Face recognition using sift features. CNS/Bi/EE report 186.
- Belhumeur, P. N., Hespanha, J. a. P., and Kriegman, D. J. (1997). Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Bolme, D. S. (2003). *Elastic Bunch Graph Matching*. PhD thesis, Colorado State University.
- Jiang, H. (2005). Confidence measures for speech recognition: a survey. *Speech Communication*, 45(4):455–470.
- Kepenekci, B. (2001). *Face Recognition Using Gabor Wavelet Transform*. PhD thesis, The Middle East Technical University.
- Král, P., Cerisara, C., and Klečková, J. (2006). Automatic Dialog Acts Recognition based on Sentence Structure. In *ICASSP'06*, pages 61–64, Toulouse, France.
- Krizaj, J., Struc, V., and Pavesic, N. (2010). Adaptation of sift features for robust face recognition.
- Lades, M., Vorbrüggen, J. C., Buhmann, J., Lange, J., and von der Malsburg, C. (1993). Distortion invariant object recognition in the dynamic link architecture. *IEEE Transactions On Computers*.
- Lawrence, S., Giles, S., Tsoi, A., and Back, A. (1997). Face recognition: A convolutional neural network approach. *IEEE Trans. on Neural Networks*.
- Lenc, L. and Král, P. (2011). Confidence measure for automatic face recognition. In *International Conference on Knowledge Discovery and Information Retrieval*, Paris, France.
- Lenc, L. and Král, P. (2012a). Gabor wavelets for automatic face recognition. In *38th International Conference on Current Trends in Theory and Practice of Computer Science*, Špindlerův Mlýn, Czech Republic.
- Lenc, L. and Král, P. (2012b). Novel matching methods for automatic face recognition using sift. In *8th AIAI (Artificial Intelligence Applications and Innovations) Conference*, Halkidiki, Greece.
- Li, S. and Jain, A. (2005). *Handbook of face recognition*. Springer-Verlag.
- Lleida, E. and Rose, R. C. (1996). Likelihood Ratio Decoding and Confidence Measures for Continuous Speech Recognition. In *ICSLP'96*, volume 1, pages 478–481, Philadelphia, USA.
- Lowe, D. (2004a). Software for sift.
- Lowe, D. G. (1999). Object recognition from local scale-invariant features. In *International Conference on Computer Vision*.
- Lowe, D. G. (2004b). Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 2.
- Nefian, A. V. and Hayes, M. H. (1998). Hidden markov models for face recognition. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*.
- Samaria, F. and Young, S. (1994). Hmm-based architecture for face identification. *Image and Vision Computing*.
- Shen, L. (2005). *Recognizing Faces - An Approach Based on Gabor Wavelets*. PhD thesis, University of Nottingham.
- Shen, L. and Bai, L. (2006). A review on gabor wavelets for face recognition. *Pattern Analysis & Applications*.
- Tan, P.-N., Steinbach, M., and Kumar, V. (2005). *Introduction to Data Mining*. Addison-Wesley.
- Turk, M. A. and Pentland, A. P. (1991). Face recognition using eigenfaces. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. Computer Vision and Pattern Recognition.
- Viola, P. and Jones, M. (2001). Rapid object detection using a boosted cascade of simple features. In *Conference on Computer Vision and Pattern Recognition*.
- Wiskott, L., Fellous, J.-M., Krüger, N., and von der Malsburg, C. (1999). Face recognition by elastic bunch graph matching. *Intelligent Biometric Techniques in Fingerprint and Face Recognition*.