

PERFORMANCE ASSESSMENT OF MORPHOLOGICAL DYNAMIC LINK ARCHITECTURE UNDER OPTIMAL AND REAL OPERATING CONDITIONS

C. Kotropoulos† A. Tefas† I. Pitas† C. Fernandez‡ F. Fernández‡

†Department of Informatics
Aristotle University of Thessaloniki
Box 451, Thessaloniki 540 06, GREECE
e-mail: costas@zeus.csd.auth.gr

‡Ibermática
Avenida del Partenon, 16-18
Campo de las Naciones, 28042 Madrid, SPAIN
e-mail: crfernan@ibermatica.es

ABSTRACT

In this paper, the performance of morphological dynamic link architecture (MDLA) is assessed under optimal and real operating conditions. It is shown that MDLA achieves a very low equal error rate on the extended M2VTS database which contains 295 persons' video data in 8 shots recorded under optimal conditions. However, its performance severely deteriorates, when it is applied to the IBERMATICA database that has been recorded under conditions simulating a real access-control to an automatic teller machine. The compensation for the image variations attributed to the variable recording conditions, i.e., changes in illumination, face size differences and varying face position, is addressed next. The use of simple and powerful pre-processing techniques aiming at compensating for the aforementioned variations prior to the application of MDLA is proposed. The first results obtained indicate that the proposed approach overcomes the image variations and stabilizes the performance of MDLA.

1. INTRODUCTION

The research activities in automatic face recognition have increased significantly over the past few years. This growth is mainly driven by application demands, such as identification for law enforcement and authentication for remote banking and access-control applications. A recent survey on face recognition can be found in [1].

Recently, a comparative study was performed for three well known face recognition techniques, namely, the eigenfaces, the auto-association and classification neural networks, and the elastic graph matching [2]. It was found that the eigenfaces worked well when the face images had relatively small lighting and moderate expression variations. Their performance deteriorated significantly, as lighting variation increased. On the contrary, the elastic graph matching was found relatively insensitive to variations in lighting, face position and expressions. Zhang et al. attributed the robustness of elastic graph matching to the use of Gabor filters and to the rigid and deformable matching stages of the recognition algorithm. Furthermore, they reported that the performance of the neural networks considered was upper bounded by that of the eigenface approach. In the closely related work [3], Adini et al. presented an empirical study that evaluated the sensitivity of several image representations (e.g., edge maps, directional and

non-directional derivatives of Gaussian filters, images convolved by 2-D Gabor-like filters) in changes of the illumination conditions. They found that all the aforementioned image representations were insufficient to overcome the variations due to changes in illumination direction, view-point and expressions.

Within the framework of EU-ACTS project M2VTS several frontal-face authentication algorithms have been developed. Among these, we mention the morphological dynamic link architecture (MDLA) and the morphological signal decomposition - dynamic link architecture [4], the elastic graph matching (residual matching and local discriminants) [5], the optimized robust correlation [6]. All algorithms were tested on the M2VTS database [7] which contains 37 persons' video data in 4 shots. Their performance has been assessed with respect to their receiver operating characteristic (ROC) curve on the M2VTS database. A scalar figure of merit that can be determined from a ROC is the Equal Error Rate (EER), i.e., the operating point having false acceptance rate equal to the false rejection rate. For example, MDLA achieved an EER of 3.7 % on that database¹, when discriminatory power coefficients were employed [4].

In this paper, we assess the performance of MDLA on two additional databases, namely, the extended M2VTS database (or XM2VTSDB shortly)[8] and the IBERMATICA database. The former database contains 295 persons' video data which include speech and image sequences of rotated heads. Eight recordings of the 295 persons have been collected under optimal recording conditions. Section 2 briefly describes the MDLA and evaluates its performance on XM2VTSDB under the two experimental configurations defined in [9]. However, both the M2VTS database and the XM2VTSDB were collected in a controlled environment, i.e., the size and the position of the face were constant, a uniform lighting was obtained by spotlight, the facial expression was neutral, etc. Motivated by [2, 3], we are interested in testing the performance of MDLA on a database recorded under conditions that simulate a real access-control in an automatic teller machine or the access to tele-services via INTERNET in a typical office environment. Towards this goal, a small database is collected under "real conditions" by Ibermática, the so-called IBERMATICA database. Several types of degradation are modeled in the latter database, such as, varying face size (i.e., scale differences), varying face position, changes in lighting and variable facial expressions. More details can be found in Section 3. Although, MDLA achieves a very low equal error rate (EER) on the XM2VTSDB, its performance severely deteriorates when MDLA is applied to the IBER-

This work was carried out within the framework of European ACTS project "Multi-modal Verification Techniques for Tele-services and Security Applications" (M2VTS).

¹5,328 client claims and 5,328 impostor claims.

MATICA database. Instead of searching for a robust image representation or a robust verification algorithm, the use of simple and powerful pre-processing techniques aiming at compensating for the aforementioned variations attributed to the variable recording conditions is proposed. Any authentication algorithm can be applied afterwards. It is shown that such an approach overcomes to an extent the image variations and stabilizes the performance of the MDLA. Similar results are expected for the other authentication algorithms developed within the M2VTS project as well. Section 4 presents the proposed face normalization techniques. Experimental results that quantify the success of the proposed techniques are also included in Section 4. The paper reports results on face verification and therefore, it complements the previously published works [2, 3].

2. PERFORMANCE OF MDLA ON THE XM2VTSDB

An alternative to linear techniques for generating an information pyramid are the scale-space morphological techniques. In the following, a brief description of MDLA and its performance evaluation on the XM2VTSDB is presented. In MDLA, we substitute the Gabor-based feature vectors used in elastic graph matching [2] by the *multiscale morphological dilation-erosion*. Let \mathcal{R} and \mathcal{Z} denote the set of real and integer numbers, respectively. Given an image $f(\mathbf{x}) : \mathcal{D} \subseteq \mathcal{Z}^2 \rightarrow \mathcal{R}$ and a structuring function $g(\mathbf{x}) : \mathcal{G} \subseteq \mathcal{Z}^2 \rightarrow \mathcal{R}$, the dilation of the image $f(\mathbf{x})$ by $g(\mathbf{x})$ is denoted by $(f \oplus g)(\mathbf{x})$. Its complementary operation, the erosion, is denoted by $(f \ominus g)(\mathbf{x})$. Throughout the paper, the *scaled hemisphere* is employed as a structuring function. The multiscale dilation-erosion of the image $f(\mathbf{x})$ by $g_\sigma(\mathbf{x})$ is defined by:

$$(f \star g_\sigma)(\mathbf{x}) = \begin{cases} (f \oplus g_\sigma)(\mathbf{x}) & \text{if } \sigma > 0 \\ f(\mathbf{x}) & \text{if } \sigma = 0 \\ (f \ominus g_{|\sigma|})(\mathbf{x}) & \text{if } \sigma < 0. \end{cases} \quad (1)$$

For $\sigma = -9, \dots, 9$, the outputs of multiscale dilation-erosion form the feature vector located at the grid node \mathbf{x} , i.e., $\mathbf{j}(\mathbf{x}) = ((f \star g_9)(\mathbf{x}), \dots, f(\mathbf{x}), \dots, (f \star g_{-9})(\mathbf{x}))$. An 8×8 sparse grid has been created by measuring the feature vectors $\mathbf{j}(\mathbf{x})$ at equally spaced nodes over the output of the face detection algorithm described in [10]. For a complete description of MDLA, the interested reader may refer to [4]. MDLA was tested on the XM2VTSDB according to the two configurations defined in [9]. A subset of 200 persons (i.e., training clients) is used for training. Let us denote by \mathcal{S}_1 the set of training clients. The design of a person authentication system based on MDLA is split in three procedures, namely, the training, the evaluation, and the test procedure. In the following, we describe the design of such a system in detail for the first configuration.

In the training procedure, three frontal images have been used for each client. The objective of the training procedure is to determine a threshold for each training client $T(X; Q)$ as follows. For $i, j = 1, 2, 3$ and $j \neq i$:

1. Compute the intra-class distances $d(X_i, X_j)$. Let $D(X, X) = \min_{i, j} \{d(X_i, X_j)\}$.
2. Compute the inter-class distances:

$$d(Y_i, X_j) \quad X \in \mathcal{S}_1, \quad Y \in (\mathcal{S}_1 - \{X\}) \quad (2)$$

$$\text{Let } D(Y, X) = \min_{i, j} \{d(Y_i, X_j)\}.$$

3. Let $OS_{(l)}\mathcal{A}$ denote the l -th order statistic of set \mathcal{A} . For $Q = 0, 1, \dots$ compute a threshold by:

$$T(X; Q) = OS_{(1+Q)}\{D(Y, X) \mid Y \neq X\} \quad (3)$$

It is evident that the training procedure results in 200 client claims and 39800 impostor claims. For several choices of Q , different pairs of false acceptance (FAR) and false rejection (FRR) rates are found. Accordingly, the receiver operating characteristic on the training set is obtained, which is plotted in Figure 1a. It is seen that an EER $\approx 3.0\%$ is achieved.

A second subset of 25 persons has been selected to evaluate the performance of any verification algorithm. Another three frontal images for each client have also been used. The evaluation client images and the evaluation impostor images build the evaluation set that is employed in the so-called evaluation procedure. By using the just-mentioned set, 600 client claims and 40,000 impostor ones are produced. Although there was a possibility to adapt the thresholds determined previously, we decided not to do so on purpose to keep the aforementioned claims disjoint to the claims that will be employed in the test procedure. The evaluation claims can be exploited then to train a fusion manager. The estimation of FAR and FRR parallels the steps described subsequently in the test procedure. The ROC on the evaluation set is shown in Figure 1a. It can easily be seen that an EER $\approx 8.0\%$ is obtained on the evaluation set.

A third subset of 70 persons is to constitute the set of test impostors. Two new frontal images for each client have been used as well. The test client images and the test impostor images form the test set. By using the test set, 400 client claims and 112,000 impostor ones can be produced. The objective of the test procedure is to provide estimates of FAR and FRR by using the test set. Let us denote the new frontal image of person X by X'_l , $l = 1, 2$. For each $X \in \mathcal{S}_1$, the number of false rejections is obtained as follows.

1. Compute the distances $d(X'_l, X_i)$ for $i = 1, 2, 3$.
2. Normalize the distances to the range $[0, 1]$: $z(X, X, i) = f(d(X'_l, X_i), T(X, Q))$.
3. Compute $D(X, X) = \min_i \{z(X, X, i)\}$.
4. Count a false rejection if $D(X, X) < 0.5$.

Let us denote the set of test impostor images by \mathcal{S}_3 . For each $X \in \mathcal{S}_1$ and $Y \in \mathcal{S}_3$ the number of false acceptances is obtained as follows:

1. Compute the distances $d(Y, X_i)$ for $i = 1, 2, 3$.
2. Normalize the distances to the range $[0, 1]$: $z(Y, X, i) = f(d(Y, X_i), T(X, Q))$.
3. Compute $D(Y, X) = \min_i \{z(Y, X, i)\}$.
4. Count a false acceptance if $D(Y, X) \geq 0.5$.

By varying the parameter Q , different thresholds are recalled from the training procedure and the ROC on the test set is created. The test ROC is shown overlaid in Figure 1a. The inspection of the figure reveals that an EER $\approx 6.57\%$ is achieved.

A second configuration is also tested that employs four training images for each client, two images for evaluation and another two images for testing. This configuration results in 400 evaluation client claims whereas the numbers of evaluation impostor claims, test client claims and test impostor claims are left intact. By adapting the just described training and evaluation procedures




Figure 1: Receiver Operating Characteristic (ROC) curves of MDLA on the training, evaluation and test sets of the extended M2VTS database (a) in configuration I, and (b) in configuration II. (c) ROC curves of MDLA applied to the IBERMATICA database with and without employing face normalization.

and by applying the previously presented test procedure, the ROCs plotted in Figure 1b are obtained. It is seen that the EER is approximately 1.5 % on the training set, 6.2 % on the evaluation set and 5% on the test set. For comparison purposes, we note that the EER on the M2VTS database (without any weighting) was 9.3% [4].

Table 1 summarizes the performance of MDLA with respect to the FAR and FRR achieved at three operating points defined on the evaluation set, namely, $FAR \approx 0$, $FAR \approx FRR$, and $FRR \approx 0$.

Table 1: Rates at $FAR \approx 0$, $FAR \approx FRR$, and $FRR \approx 0$ on XM2VTSDB in the two configurations of the experimental protocol. All rates are in %.

Conf.	Operating Point	Evaluation set		Test Set	
		FAR	FRR	FAR	FRR
I	$FAR \approx 0$	0.5	19	0.46	20
	$FAR \approx FRR$	8.11	8	8.23	6
	$FRR \approx 0$	48.42	1.66	46.63	0.75
II	$FAR \approx 0$	0.46	18.75	0.46	16.25
	$FAR \approx FRR$	6.04	6.5	6.17	3.5
	$FRR \approx 0$	36.85	1.25	34.67	0.75

3. PROBLEMS OCCURRED IN TESTS PERFORMED UNDER REAL CONDITIONS

The evaluation of any authentication algorithm on a database collected under well-controlled conditions is inadequate, when the algorithm is to be integrated into a platform for commercial exploitation. In such case, it is of utmost importance to evaluate the authentication performance of the algorithm on field tests. This is the case of the IBERMATICA database. Although the experiments reported in this paper have been conducted on 11 persons, the size of the database progressively increases (20 persons are currently included in the database). For each person two training sample images were stored in the database. A large number of test images was also recorded. In particular, more than twenty test images of each user were recorded, when he or she claims the correct identity (client images). In addition a couple of test images per user were

recorded, when he or she claims the identity of someone else in the database (impostor images). In total 514 client accesses were tested against 56 impostor ones. Several types of degradation are modeled in the database:

- (a) *Face size and position.* In practice it is very difficult to control the position of the subject with respect to the camera.
- (b) *Changes in illumination.* If a spotlight is not used, lighting variations occur. For example close to a window, the lighting depends strongly on the day-time and the weather.
- (c) *Facial expressions.* In practice it is almost impossible to control the mood of the subject. The smile causes probably the largest variation of facial expressions.

All the images in the database are recorded in 256 grey levels and they are of dimensions 320×240 pixels. Sample images from IBERMATICA database before normalization are shown in Figure 2(a). A set of field tests was implemented so that the performance of two authentication algorithms developed within the M2VTS research project, namely, the elastic graph matching and the morphological dynamic link architecture were evaluated under real conditions. The EERs obtained using the just-mentioned techniques are shown in Table 2. From the inspection of Table 2, it

Table 2: EERs achieved by the elastic graph matching and the morphological dynamic link architecture, when they are applied to the IBERMATICA database. The EERs are in %.

Method	EER
Elastic Graph Matching	25
Morphological Dynamic Link Architecture	35
Preprocessing + Morphological Dynamic Link Architecture	20

becomes evident that the performance of both algorithms depends strongly on the variable recording conditions. To alleviate such a dependency, a compensation for the aforementioned image variations is needed. To do so, the use of simple and powerful face normalization techniques prior to the application of any authentication algorithm is proposed in Section 4.

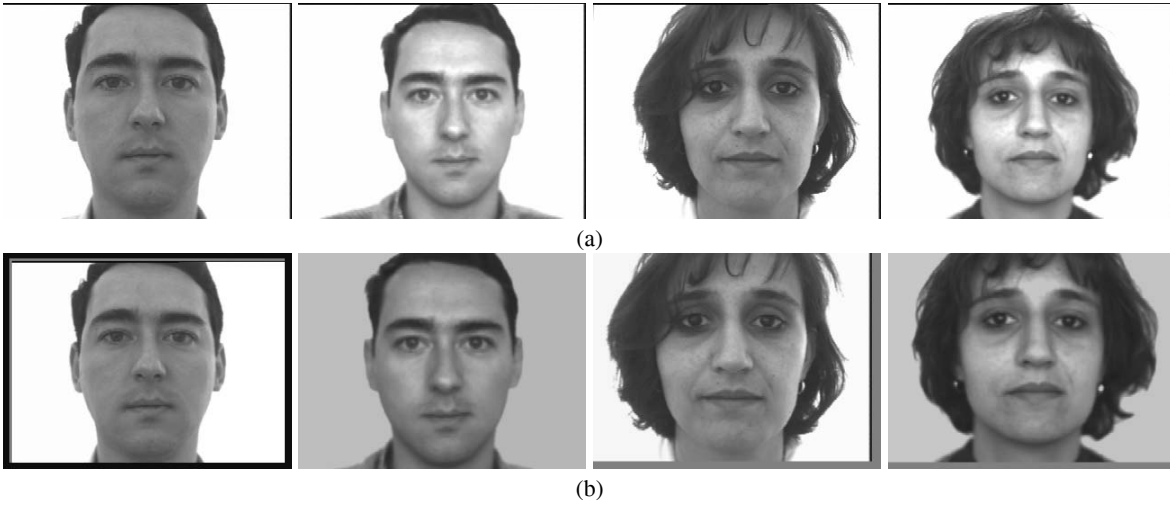


Figure 2: Sample images from IBERMATICA database (a) before (b) after normalization.

4. FACE NORMALIZATION TECHNIQUES

The proposed techniques are based on the detection of the facial region in the image and its splitting in two segments, the left segment and the right one. We assume that: (1) the background of the images is uniform, and (2) one only person appears in the scene. The steps of the algorithm are described in detail below.

Step 1: The oval shape of a face can be approximated by an ellipse. Therefore, the detection of the facial area in an image can be performed by detecting an object of elliptical shape. To do so, first we have to discard the image background. The starting point is an edge detection algorithm. By thresholding the resulting image after edge detection, a zero value is assigned to the background. Since the background is uniform and does not contain any edges, it can easily be discarded by employing a grass-fire algorithm. Accordingly, the image is segmented into two regions, one of which contains the facial region and the other contains the background.

Step 2: The next step is to model the face-like region by an ellipse using moment-based features [11]. Let us denote the face-like area by C and the best-fit ellipse by \mathcal{E} . An ellipse is defined by its center (x_0, y_0) , its orientation θ and the length a and b of its semi-major and semi-minor axes. The center of the ellipse is estimated by the center of mass of the region C . The orientation of the ellipse is computed by determining the angle between the axis of the least moment of inertia and the horizontal axis of the coordinating system, i.e.,

$$\theta = \frac{1}{2} \arctan \left(\frac{2\rho_{1,1}}{\rho_{2,0} - \rho_{0,2}} \right) \quad (4)$$

where $\rho_{i,j}$ denotes the (i, j) -central moment of the region C . The length of the semi-major axis a and the length of the semi-minor axis b can be computed by evaluating the least and the greatest moments of inertia. The least and the greatest moments of inertia of an ellipse with orientation θ , I_{min} and I_{max} , are given by,

$$I_{min} = \sum_{(x,y) \in C} [(y - y_0) \cos \theta - (x - x_0) \sin \theta]^2 \quad (5)$$

$$I_{max} = \sum_{(x,y) \in C} [(y - y_0) \sin \theta + (x - x_0) \cos \theta]^2 \quad (6)$$

where x, y denote the horizontal and vertical coordinate of a pixel. Accordingly, a and b are given by:

$$a = \left(\frac{4}{\pi} \right)^{\frac{1}{4}} \left[\frac{(I_{max})^3}{I_{min}} \right]^{\frac{1}{8}}, \quad b = \left(\frac{4}{\pi} \right)^{\frac{1}{4}} \left[\frac{(I_{min})^3}{I_{max}} \right]^{\frac{1}{8}} \quad (7)$$

respectively. To find the ellipse that models the given region best, we iteratively maximize the measure:

$$\mathcal{M} = \sum_{(x,y) \in \mathcal{E} \cap C} 1 - \sum_{(x,y) \in \mathcal{E} \cap C^c} 1 \quad (8)$$

where C^c denotes the complement of the region C (i.e., the background). The maximization of (8) corresponds to the maximization of the number of correctly modeled pixels (i.e., $(x, y) \in \mathcal{E} \cap C$) and the minimization of the number of incorrectly modeled pixels (i.e., $(x, y) \in \mathcal{E} \cap C^c$). In order to find a maximum of \mathcal{M} the following recursive procedure is proposed.

Step a. Calculate the parameters of the initial ellipse \mathcal{E}_0 and the measure \mathcal{M}_0 for the initial region C_0 .

At the i -th iteration of the algorithm:

Step b. Find the new region $C_i = C_{i-1} \cap \mathcal{E}_{i-1}$.

Step c. Calculate the parameters of the ellipse \mathcal{E}_i that fits C_i and the measure \mathcal{M}_i .

Step d. if $\mathcal{M}_i > \mathcal{M}_{i-1}$ go to **Step b.** Otherwise, the best ellipse is \mathcal{E}_{i-1} that has already been found.

By using this iterative algorithm the ellipse fitting becomes more robust to noise, i.e., to pixels that correspond to clothes, hair, etc.

Step 3: The first ellipse found is a coarse approximation of the facial area, because the hair, and in some cases, parts of the clothes are included in it. Thus, the skin region can not be segmented. To overcome this problem the ellipse is subdivided into its left and right segments with respect to the vertical axis. Moreover, the subdivision aids to compensate for the different lighting conditions which effect unevenly the two parts of the face.

Step 4: The next step is to apply a clustering algorithm to each segment of the ellipse, separately. By choosing $K=2$, a K -means algorithm hopefully succeeds to relate the skin-like areas with a

single cluster in each segment. Thus, the skin region is detected accurately. The mean intensity values of each face side are also estimated. They provide information for the lighting conditions during the recording procedure.

Step 5: The union of clusters in the left and right segment that correspond to skin-like areas is modeled by an ellipse using the algorithm described in **Step 2**. The quality of fit is measured again by (8). By applying Steps a–d, a finer approximation is obtained. The ellipse found at the last iteration is the best-fit ellipse we searched for.

In addition to the use of the just-described technique in facial area detection, the method provides: (i) an estimate of the face center that can be used in compensating for face translation, (ii) an estimate of the face width which is related to the length of the minor axis of the best-fit ellipse, and (iii) an estimate of the mean intensity value of each segment of the face (left and right) which is related to the lighting conditions of the recording procedure.

4.1. Illumination normalization

Lighting may cause uneven illuminations of the right and the left face segments. We assume that the illumination conditions in the left and right face segment are uniform. To compensate for the aforementioned effect, the mean intensities of both face segments should be equalized. Let m_L, m_R be the mean intensity values in the left and the right segment, respectively. The initial image $I(x, y)$ is transformed so that the left and right segments of the normalized image, $I_N(x, y)$, have the same (desired) mean intensity I_d :

$$I_N(x, y) = I_d \left(\frac{\frac{1}{m_R} - \frac{1}{m_L}}{1 + \exp\left(\frac{x_0 - x}{\lambda}\right)} + \frac{1}{m_L} \right) I(x, y) \quad (9)$$

where λ controls the slope of the sigmoidal function that appears in the denominator of the first fractional term inside parentheses. Let \mathcal{L} be the image region in which $\exp\left(\frac{x_0 - x}{\lambda}\right) \rightarrow \infty$ and \mathcal{R} be the image region in which $\exp\left(\frac{x_0 - x}{\lambda}\right) \approx 0$. These regions correspond to the left and the right segments respectively of the best-fit ellipse modeling the face. It can be easily proven that:

$$E[I_N(x, y), (x, y) \in \mathcal{L}] = E[I_N(x, y), (x, y) \in \mathcal{R}] = I_d, \quad (10)$$

which satisfies our objective.

4.2. Face position and size

Varying face position and size are easily compensated for, if the face is accurately approximated by an ellipse. The problem of varying face position can be solved by translating the initial image so that the center of the ellipse always coincides with the image plane center. The width of the face can be approximated by the length of the minor axis $2b$. Size (i.e., scale) normalization can be achieved by resizing the image with a horizontal scale factor $\frac{W_d}{2b}$ where W_d is the desired width of the normalized face. Image resizing is achieved by linear interpolation. Examples of normalized face images from IBERMATICA database are depicted in Figure 2(b).

4.3. Impact of face normalization on the authentication performance of MDLA

The ROC, when we compensate for the three types of degradation, i.e., scale, position and lighting variations is plotted in Figure 1c.

The ROC without applying the proposed face normalization technique is overlaid in the same plot as well. The EER drops to 20% when the face normalization technique is employed, as can be seen by inspecting Table 2. That is, a significant drop of 15% in EER (which amounts to a 43 % relative drop) is achieved.

5. REFERENCES

- [1] R. Chellapa, C.L. Wilson, and S. Sirohey, "Human and machine recognition of faces: A survey," *Proceedings of the IEEE*, vol. 83, no. 5, pp. 705-740, May 1995.
- [2] J. Zhang, Y. Yan, and M. Lades, "Face recognition: Eigenface, elastic matching and neural nets," *Proceedings of the IEEE*, vol. 85, no. 9, pp. 1423-1435, September 1997.
- [3] Y. Adini, Y. Moses, and S. Ullman, "Face recognition: The problem of compensating for changes in illumination direction," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 19, no. 7, pp. 721-732, July 1997.
- [4] C. Kotropoulos, A. Tefas, and I. Pitas, "Frontal face authentication using variants of Dynamic Link Matching based on mathematical morphology," in *Proc. of the IEEE Int. Conf. on Image Processing*, vol. I, pp. 122-126, Chicago, U.S.A., October 1998.
- [5] B. Duc, S. Fischer, and J. Bigün, "Face authentication with Gabor information on deformable graphs," *IEEE Trans. on Image Processing*, accepted for publication 1998.
- [6] J. Matas, K. Jonsson, and J. Kittler, "Fast face localisation and verification," in *Proc. of British Machine Vision Conference*, 1997.
- [7] S. Pigeon, and L. Vandendorpe, "The M2VTS multimodal face database," in *Lecture Notes in Computer Science: Audio- and Video- based Biometric Person Authentication* (J. Bigün, G. Chollet and G. Borgefors, Eds.), vol. 1206, pp. 403-409, 1997.
- [8] K. Messer, J. Matas, and J. Kittler, "Acquisition of a large database for biometric identity verification," in *Proc. of BioSig 98*, June 1998.
- [9] J. Luetin, and G. Maître, "Evaluation protocol for the extended M2VTS database (XM2VTSDB)," in *IDIAP Communication 98-05*, IDIAP, Martigny, Switzerland 1998.
- [10] C. Kotropoulos, and I. Pitas, "Rule-based face detection in frontal views," in *Proc. of IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP 97)*, vol. IV, pp. 2537-2540, Munich, Germany, April 21-24, 1997.
- [11] A.K. Jain, *Fundamentals of Digital Image Processing*. Prentice-Hall, Englewood Cliffs, NJ, 1989.