AUDIO, VIDEO AND MULTIMODAL PERSON IDENTIFICATION IN A SMART ROOM

J. Luque, R. Morros, A. Garde, J. Anguita, M. Farrus,
D. Macho, F. Marqués, C. Martínez, V. Vilaplana, J. Hernando

Universitat Politècnica de Catalunya
Jordi Girona 1-3, Campus Nord Edifici D5
08034 Barcelona, Spain
{aluque, morros, agarde, jan, mfarrus, dusan, ferran, veronica, javier}@gps.tsc.upc.edu

Abstract. In this paper, we address the modality integration issue on the example of a smart room environment aiming at enabling person identification by combining speech and 2D face images. First we introduce the monomodal audio and video identification techniques and then we present the use of combined input speech and face images for person identification. The various sensory modalities, speech and faces, are processed both individually and jointly. It's shown that the multimodal approach results in improved performance in the identification of the participants.

1. Introduction

Person identification consists in determining the identity of a person from a data segment, such as a speech, video segment, etc. Currently, there is a high interest in developing person identification applications in the framework of smart room environments. In a smart room, the typical situation is to have one or more cameras and several microphones. Perceptually aware interfaces can gather relevant information to model and interpret human activity, behaviour and actions. Such applications face an assortment of problems such a mismatched training and testing conditions or the limited amount of training data.

In this work we present the audio, video and multimodal person identification techniques and the obtained results in the CLEAR'06 Evaluation Campaign inside the CHIL (Computers in the human interaction loop) project [1]. The CLEAR'06 Person Identification evaluation is a closed-set task, that is, all the possible speakers are known. Matched training and testing conditions and far-field data acquisition are assumed, as well as no *a priori* knowledge about room environment.

For acoustic speaker identification, the speech signals are parameterized using the Frequency Filtering (FF) [2] over the filter-bank energies, which is both computationally efficient and robust against noise. Next, in order to model the probability distribution of the parameters generated by each speaker, Gaussian Mixture Models (GMM) [3] with diagonal covariance are used.

In the case of visual identification, an appearance-based technique is used due to the low quality of the images. Face images of the same individual are gathered into groups. Frontal images within a group are jointly compared to the models for identification. These models are composed of several images representative of the individual. The joint recognition improves the performance of a face recognition algorithm applied on single images. Individual decisions are based on a PCA [4] approach given that the variability of the users' appearance is assumed to be low and so are the lighting variations.

Multimodal recognition involves the combination of two or more human traits like voice, face, fingerprints, iris, hand geometry, etc. to achieve better performance than using monomodal recognition [5]. In this work, a multimodal score fusion technique, Matcher Weighting with equalized scores, has been used. This technique has obtained an improvement for the correct identification rate on the closed-set 15/30 seconds training and 1/5/10/20 seconds testing conditions on the CLEAR'06 Evaluation task.

This paper is organized as follows: In Sections 2 and 3 an overview of the audio and video algorithms and techniques is given. Section 4 presents the technique for multimodal fusion. Section 5 describes the evaluation scenario and the experimental results. Finally, Section 6 is devoted to provide conclusions.

2. Audio Person Identification

The speaker identification (SI) task consists in determining the identity of the speaker of a speech segment. In this task it is usually assumed that all the possible speakers are known. For this evaluation, recordings using one microphone of an array from 26 speakers have been provided.

The first stage of current speaker recognition systems is a segmentation of the speech signal into regular segments. The speech signal is divided into frames of 30ms at a rate of 10ms. From each segment a vector of parameters that characterizes the segment is calculated. In this work we have used the Frequency Filtering (FF) parameters [2]. These parameters are calculated as the widely used Mel-Frequency Cepstral Coefficients (MFFC) [6] but replacing the final Discrete Cosine Transform of the logarithmic filter-bank energies of the speech frame with the following filter:

$$H(z) = z - z^{-1} \qquad\qquad (1)$$

Fig. 1 shows the calculation procedure of the FF parameters. These features have several interesting characteristics:

- They are uncorrelated.
- They are computationally simpler than MFCCs.
- They have frequency meaning.
- They have generally shown an equal or better performance than MFCCs in both speech and speaker recognition.

Fig. 1. Calculation procedure of the FF parameters

In order to capture the temporal evolution of the parameters the first and second time derivatives of the features are generally appended to the $\mathbf{o}_{FF}$ basic static feature vector. The so called delta coefficients [7] are computed using the following regression formula

$$\Delta\mathbf{o}_t(\mathrm{i}) = \frac{\sum_{\theta=1}^{\Theta} \theta\left(\mathbf{o}_{t+\theta}(\mathrm{i}) - \mathbf{o}_{t-\theta}(\mathrm{i})\right)}{2\sum_{\theta=1}^{\Theta}\theta^2} \qquad (2)$$

Where $\Delta\mathbf{o}_t(\mathrm{i})$ is the delta coefficient at time t computed in terms of the corresponding static coefficients $\mathbf{o}_{t-\Theta}$ to $\mathbf{o}_{t+\Theta}$. The same formula is applied to the delta coefficient with another window size to obtain acceleration coefficients.

For each speaker that the system has to recognize, a model of the probability density function of the parameter vectors is estimated. These models are known as Gaussian Mixture Models (GMM) [3], which is a weighted sum of Gaussian distributions:

$$\lambda_j = \sum_{m=1}^{M} w_m N\left(\mathbf{o}, \boldsymbol{\mu}_m, \boldsymbol{\Sigma}_m\right) \qquad j = 1,\ldots,J \qquad (3)$$

where λ_j is the model, $\mathbf{o}$ is the vector of parameters being modeled, J is the number of speakers, M is the number of Gaussian mixtures, w_m is the weight of the Gaussian m, and N is a Gaussian function of mean vector $\boldsymbol{\mu}_m$ and covariance matrix $\boldsymbol{\Sigma}_m$. The parameters of the models are estimated from speech samples of the speakers using the well-known Baum-Welch algorithm.

In the testing phase of a SI system a set of parameter vectors $\mathbf{O}=\{\mathbf{o}_i\}$ is calculated from the testing signal. After that, the likelihood that the vector $\mathbf{O}$ is produced by each speaker is calculated and it is chosen the speaker with the biggest likelihood, i.e.:

$$s = \arg\max_j\{L(\mathbf{O} \mid \lambda_j)\} \qquad (4)$$

where s is the recognized speaker and $L(\mathbf{O}|\lambda_j)$ is the likelihood that the vector $\mathbf{O}$ was generated by the speaker of the model λ_j.

3. Video Person Identification

In this section, the Visual Person ID task is presented. Recognition is stand-alone, taking detection and tracking for granted. That is, the system is semi-automatic.

We have developed for this task a technique for face recognition in smart environments. The technique takes advantage of the continuous monitoring of the scenario and combines the information of several images to perform the recognition. Appearance based face recognition techniques are used given that the scenario does not ensure high quality images. As the visual identification evaluation is a close-set identification task, models for all individuals in the database are created off-line using two sets of video segments: the first one consists on one segment of 15 s per each individual in the database, while the second one consists on one segment of 30 s per individual.

The proposed system works with groups of face images of the same individual. For each test segment, face images of the same individual are gathered into a group. Then, for each group, the system compares such images with the model of the person.

We first describe the procedure for combining the information provided by a face recognition algorithm when it is applied to a group of images of the same person in order to, globally, improve the recognition results. Note that the proposed system is independent of the distance measure adopted by the specific face recognition algorithm.

Combining groups of images. Let $\{x\}_i = \{\underline{x}_1, \underline{x}_2, ... , \underline{x}_M\}$ be a group of M probe images of the same person, and let $\{C\}_j = \{C_1, C_2, ... , C_S\}$ be the different models or classes stored in the (local or global) model database. S is the number of individual models. Each model C_j contains N_j images, $\{y\}_n^j = \{\underline{y}_1^j, \underline{y}_2^j, ..., \underline{y}_{Nj}^j\}$ where N_j may be different for every class. Moreover, let

$$d(\underline{x}_i, \underline{y}_n^j) : \Re^Q x \Re^Q \to \Re \tag{5}$$

be a certain **decision function** that applies to one element of $\{x\}_i$ and one element of $\{y\}_n^j$, where Q is the dimension of $\underline{x}_i$ and $\underline{y}_n^j$. It represents the decision function of any face recognition algorithm. It measures the similarity of a probe image $\underline{x}_i$ to a test image $\underline{y}_n^j$.

We fix a **decision threshold** R_d so that $\underline{x}_i$ and $\underline{y}_n^j$ represent the same person if $d(\underline{x}_i, \underline{y}_n^j) < R_d$. If, for a given $\underline{x}_i$ the decision function is applied to every $\underline{y}_n^j \in C_j$, we can define the **δ value** of $\underline{x}_i$ relative to a class C_j, δ_{ij} as

$$\delta_{ij} = \#\left\{ \underline{y}_n^j \in C_j \ / \ d(\underline{x}_i, \underline{y}_n^j) < R_d \right\} \tag{6}$$

That is, δ_{ij} counts the number of times that the face recognition algorithm matches $\underline{x}_i$ with an element of C_j. With this information, the δ-Table is built:

Table 1. δ-Table

	C_1	C_2	$\ldots$	C_S
+$\underline{X}_1$	δ_{11}	δ_{12}	$\ldots$	δ_{1S}
$\underline{X}_2$	δ_{21}	δ_{22}	$\ldots$	δ_{2S}
$\ldots$	$\ldots$	$\ldots$	$\ldots$	$\ldots$
$\underline{x}_M$	δ_{M1}	δ_{M2}	$\ldots$	δ_{MS}

Based on the δ-Table above, we define the following concepts:

- **Individual Representation of $\underline{x}_i$**: It measures the representation of sample $\underline{x}_i$ by class C_j:

$$R(\underline{x}_i, C_j) = \frac{\delta_{ij}}{N_j} \tag{7}$$

- **Total representation of $\underline{x}_i$:** It is the sum of the individual representations of $\underline{x}_i$ through all the classes:

$$\hat{R}(\underline{x}_i) = \sum_{j=1}^{S} R(\underline{x}_i, C_j) = \sum_{j=1}^{S} \frac{\delta_{ij}}{N_j} \tag{8}$$

- **Reliability of a sample $\underline{x}_i$ given a class C_j:** It measures the relative representation of sample $\underline{x}_i$ by class C_j considering that sample $\underline{x}_i$ could be represented by other classes:

$$\rho(\underline{x}_i, C_j) = \left\{ \begin{array}{l} \dfrac{R(\underline{x}_i, C_j)}{\hat{R}(\underline{x}_i)} = \dfrac{\delta_{ij}/N_j}{\sum_{k=1}^{S} \dfrac{\delta_{ik}}{N_k}} = \dfrac{\delta_{ij}}{N_j \cdot \sum_{k=1}^{S} \dfrac{\delta_{ik}}{N_k}}, \quad \text{if } \hat{R}(\underline{x}_i) > 0 \\ \\ \qquad\qquad 1, \qquad \text{if } \hat{R}(\underline{x}_i) = 0 \end{array} \right\} \leq 1 \tag{9}$$

The assignment $\rho_{ij} = \rho(\underline{x}_i, C_j) = 1$ when the total representation is zero will be commented when discussing the model updating.

- **Representation of C_j:** It estimates the relative representation of a group of samples $\{x\}_i$ by a class C_j. Weighting is performed to account for the contribution of the group $\{x\}_i$ to other classes:

$$R(C_j) = \frac{1}{M} \sum_{i=1}^{M} \rho_{ij} \cdot \delta_{ij} \tag{10}$$

- **Match Likelihood for class C_j:** It relates a class representation and its match probability. If $r = R(C_j)$, then:

$$ML(C_j) = \frac{1 - e^{\frac{-r^2}{\sigma^2}}}{1 - e^{\frac{-N_j^2}{\sigma^2}}} \tag{11}$$

where σ adjusts the range of $R(C_j)$'s values (see Fig. 2).

Fig. 2. Examples of ML for different σ values (N=5).

- **Relative Match Likelihood for a class C_j:** It relates the ML of a class Cj and the maximum ML of the other classes:

$$RML(C_j) = \begin{cases} \dfrac{ML(C_j)}{\max\limits_{k \neq j}(ML(C_k))} & if \ ML(C_j) \geq 0.5 \\ 0 & if \ ML(C_j) < 0.5 \end{cases} \tag{12}$$

This measure determines if the selected class (that with maximum ML) is widely separated from other classes. A minimum value of ML is required, to avoid analyzing cases with too low ML values.

Relying on the previous concepts, the recognition process is defined, in the identification mode, as follows:

1. Compute the δ-Table.
2. Compute the match likelihood (ML) for every model.
3. Compute the RML of the class with the highest $ML(C_j)$.

The group is assigned to the class resulting in a highest RML value. In this work, a PCA based approach [4] has been used. This way, the decision function is the Euclidean distance between the projections of $\underline{x}_i$ and $\underline{y}_n^j$ on the subspace spanned by the first eigenvectors of the training data covariance matrix:

$$d(\underline{x}_i, \underline{y}_n^j) = \left\| W^T \underline{x}_i - W^T \underline{y}_n^j \right\| \tag{13}$$

where W^T is the projection matrix.

The XM2VTS database [8] has been used as training data for estimating the projection matrix and the first 400 eigenvectors have been preserved. Due to the images being recorded continuously using the corner cameras, face images can not be ensured to be all frontal. Mixing frontal and non-frontal faces in the same models can be quite a problem for face recognition systems. To avoid this situation, eye coordinates are used to determine the face pose for each image. Only frontal faces are used for identification.

Note that, in our system, models per each person have been automatically generated, without human intervention. All images for a given individual in the training intervals are candidates to form part of the model. Candidate face bounding boxes are projected on the subspace spanned by the first eigenvectors of the training data covariance matrix W^T. The resulting vector is added to the model only if different enough from the vectors already present in the model.

4. Multimodal Person Identification

In a multimodal biometric system that uses several characteristics, fusion is possible at three different levels: feature extraction level, matching score level or decision level. Fusion at the feature extraction level combines different biometric features in the recognition process, while decision level fusion performs logical operations upon the monomodal system decisions to reach a final resolution. Score level fusion matches the individual scores of different recognition systems to obtain a multimodal score. Fusion at the matching score level is usually preferred by most of the systems.

Matching score level fusion is a two-step process: normalization and fusion itself [9], [10], [11], [12]. Since monomodal scores are usually non-homogeneous, the normalization process transforms the different scores of each monomodal system into a comparable range of values. One conventional affine normalization technique is z-score, that transforms the scores into a distribution with zero mean and unitary variance [10],[12].

After normalization, the converted scores are combined in the fusion process in order to obtain a single multimodal score. Product and sum are the most straightforward fusion methods. Other fusion methods are min-score and max-score that choose the minimum and the maximum of the monomodal scores as the multimodal score. Matcher Weighting is a fusion method where each biometric is weighted by a different factor proportional to the recognition result of the biometric, and in user weighting method different weighting methods are applied for every user.

Normalization and Fusion Techniques. Scores must be normalized before being fused. One of the most conventional normalization methods is z-score (ZS), which normalizes the global mean and variance of the scores of a monomodal biometric. Denoting a raw matching score as a from the set A of all the original monomodal biometric scores, the z-score normalized biometric x_{ZS} is calculated according to

$$x_{zs} = \frac{a - \text{mean}(A)}{\text{std}(A)} \tag{14}$$

where *mean(A)* is the statistical mean of *A* and *std(A)* is the standard deviation.

Histogram equalization (HE) is a non linear transformation whose purpose is to equalize the variances of two monomodal biometrics in order to reduce the non linear effects typically introduced by speech systems. The HE technique matches the histogram obtained from the speaker verification scores and the histogram obtained from the face identification scores, both evaluated over the training data. The designed equalization takes as a reference the histogram of the scores with the best accuracy, which can be expected to have lower separate variances, in order to obtain a bigger variance reduction.

In Matcher Weighting (MW) fusion each monomodal score is weighted by a factor proportional to the recognition rate, so that the weights for more accurate matchers are higher than those of less accurate matchers. When using the Identification Error Rates (IER) the weighting factor for every biometric is proportional to the inverse of its IER. Denoting w^m and e^m the weigthing factor and the IER for the *m*th biometric x^m and *M* the number of biometrics, the fused score *u* is expressed as

$$u = \sum_{m=1}^{M} w^m x^m \tag{15}$$

where

$$w^m = \frac{\dfrac{1}{e^m}}{\displaystyle\sum_{m=1}^{M} \dfrac{1}{e^m}} \tag{16}$$

Before carrying out the fusion process, histogram equalization is applied over all the previously obtained monomodal scores. Since the best recognition results have been achieved in the voice recognition experiments, the histogram of the voice scores has been taken as a reference in the histogram equalization.

After the equalization process, the weighting factors for both voice and face scores are calculated by using the corresponding Identification Error Rates, as in Eq. 16. Z-score normalization is also applied, and final fused scores are obtained by using Eq. 15.

5 Experiments and Discussion

5.1 Experimental set-up

A set of audiovisual recordings of seminars and of highly-interactive small working groups seminars have been used. These recordings were collected by the CHIL consortium for the CLEAR 06 Evaluation. The recordings were done according to the "CHIL Room Setup" specification [1]. A complete description of the different recordings can be found in [13]. Data segments are short video sequences and matching far-field audio recordings taken from the above seminars.

In order to evaluate how of the duration of the training signals affects the performance of the system two training durations have been considered: 15 and 30 seconds. Test segments of different durations (1, 2, 5, 10 and 20 seconds) have been used during the algorithm development and testing phases. A total of 26 personal identities have been used in the recognition experiments.

Each seminar has one audio signal from the microphone number 4 of the Mark III array. Each audio signal has been divided into segments which contain information of a unique speaker. These segments have been merged to form the final testing segments of 1, 5, 10 and 20 seconds (see Table 2) and training segments of 15 and 30 seconds. Video is recorded in compressed JPEG format, with different frame-rates and resolutions for the various recordings.

Far-field conditions have been used for both modalities, i.e. corner cameras for video and Mark III microphone array for audio. In the audio task only one array microphone has been considered for both development and testing phases. In the video task, we have four fixed position cameras that are continuously monitoring the scene. All frames in the 1/5/10/20 sec segments and all synchronous camera views can be used and the information can be fused to find the identity of the concerned person. To find the faces to be identified, a set of labels is available with the position of the bounding box for each person's face in the scene. These labels are provided each 1s. The face bounding boxes are linearly interpolated to estimate their position in intermediate frames. To help this process, an extra set of labels is provided, giving the position of both eyes of each individual each 200 ms.

Table 2. Number of segments for each test condition

Segment Duration	Number of segments	
	Development	Evaluation
1 sec	390	613
2 sec	182	0
5 sec	78	411
10 sec	26	289
20 sec	0	178

The metric used to benchmark the quality of the algorithms is the percentage of correctly recognized people from test segments.

5.2 Results

In this section we summarize the results for the evaluation of different modalities and the result improvement with the multimodal technique. Table 3 shows the correct identification rate for both audio and video modalities and the fusion identification rate obtained depending on the length of the used test files.

Related to acoustic identification task, it can be seen that the results are better as the segments length increases. Table 3 shows that for the different test segment lengths the recognition rate increases when more data is used to test the speaker models. It can also be seen in Table 3 that this rate is higher when the amount of training data increases. Overall, using the 30 seconds training segments, an improvement of up to 8% in the recognition rate is obtained with respect to the case where 15 seconds segments are used.

For the face identification evaluation, in general, these results show a low performance of the system. Results for the training set B (using a segment of 30s to generate the models) show only a slight increase of performance with respect to training set A. It can also be seen that the results improve slowly as the segments length increases.

Table 3. Percentage of correct identification for both audio and video unimodal modalities and multimodal fusion. First row shows the duration of test segments in seconds. Second row shows the number of tested segments. Train A and B are the training sets of 15 seconds and 30

Duration	Segments	Train A			Train B		
		Speech	Video	Fusion	Speech	Video	Fusion
1	613	75.0	20.2	76.8	84.0	19.6	86.2
5	411	89.3	21.4	92.0	97.1	22.9	97.1
10	289	89.3	22.5	94.1	96.2	25.6	98.0
20	178	88.2	23.6	96.0	97.2	27.0	98.9
All	1491	83.3	21.3	86.7	91.6	22.5	93.2

The reasons for this low performance are manifold: First of all, the system uses only frontal faces to generate the models and for recognition. However, most of the face views found in the recordings are non frontal. Another reason for the low percentage of correctly identified persons is the low quality of the images. The need to cover all the space in the room with four cameras results in small images, were the person's faces are tiny. In the worst cases, face sizes are only 13x13 pixels.

In addition, poor illumination conditions in some recordings causes cameras to work at large diaphragm apertures. As a result, the depth of field is very shallow and several images are out of focus. Other recordings present interlacing errors. Figure 3 shows several examples of all these problems.

Fig. 3.   (a)   (b)   (c)   (d)   (e)   (f)

Examples of face bounding boxes taken from several recordings shown at its relative sizes. Smallest image (c) is 13x13 pixels and larger image (a) is 29x47 pixels. Images are taken from the training segments of the AIT, IBM, ITC, UPC and UKA recordings. Images in the test sequences are similar.

Another problem is that, due to the fact that face bounding boxes are interpolated from the 1 s labels, our system is, in many cases, considering as 'frontal' faces that are not really frontal ones. Figures 1, (a), (b), (c) and (e) are examples of this situation.

This leads us to conclude that, under these conditions, a more elaborated technique should be used. For instance, non-frontal face views should be taken into account, as most of the views found in the recordings are non-frontal. Even in this case, person identification using face detection alone is probably not going to give good results in these conditions. Identification should be performed combining more features other than those obtained from face bounding-boxes.

The determination of the weighting factors for the multimodal fusion has been done by using the training signals of 30 seconds as a development set. The first 15 seconds have been used for training and the other 15 seconds for testing. The recognition results obtained in the evaluation for multimodal identification can also be seen in Table 3. Fusion results of both systems are also shown for the different lengths of the testing segments. Fusion correct identification rates are higher than the monomodal rates. When testing with segments of 1, 10 and 20 seconds the obtained fusion results outperform those obtained with both monomodal systems. When training with 30 seconds segments and testing with 5 seconds the fusion results are the same as those obtained in the monomodal voice system. The best fusion result is obtained when testing with files of 20 seconds.

6 Conclusions

In this paper we present two algorithms for visual and acoustic person identification in smart room environments. A Frequency Filtering technique has been used for audio recognition. For video, an approach based on joint identification over groups of images of a same individual using a PCA approach has been followed.

For the acoustic identification task, the results show that the presented approach is well adapted to the conditions of the experiments. For the visual identification task, the low quality of the images results in a low performance of the system. In this case,

results suggest that identification should be performed combining more features other than frontal face bounding-boxes.

To improve the obtained results, a multimodal score fusion technique has been used. Matcher Weighting with histogram equalized scores is applied to the scores of the two monomodal tasks. The results show that this technique can provide an improvement of the recognition rate in all train/test conditions.

7   Acknowledgements

References

1.   J. Casas, R. Stiefelhagen, et al, "Multi-camera/multi-microphone system design for continuous room monitoring," CHIL-WP4-D4.1-V2.1-2004-07-08-CO, CHIL Consortium Deliverable D4.1, July 2004

2.   Nadeu C., Macho D. and Hernando J., "Time and frequency filtering of filter-bank energies for robust HMM speech recognition", Speech Communication, Vol. 34, pp. 93-114, 2001.

3.   D. A. Reynolds, "Robust text-independent speaker identification using Gaussian mixture speaker models", IEEE Transactions on Speech and Audio Processing, Vol. 3, Nº 1, pp. 72-83, January 1995

4.   M. Kirby, L.Sirovich, "Application of the Karhunen-Loeve procedure for the characterization of human faces", IEEE Trans. PAMI, vol 12, no. 1, pp. 103-108, Jan. 1990.

5.   Bolle, R.M. Connell, J.H., Pankanti, S., Ratha, N.K., Senior, A.W., 2004. Guide to Biometrrics, Springer, New York.

6.   Davis S. B., Mermelstein P., "Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences", IEEE Transactions on Acoustic, Speech and Signal Processing, Vol. 28, pp. 357-366, 1980.

7.   Furui S., "Speaker independent isolated word recognition using dynamic features of speech spectrum", IEEE Transactions ASSP, No. 34, pp.52-59, 1986.

8.   K. Messer, J. Matas, J.V. Kittler, J. Luettin and G. Maitre, XM2VTSDB: The extended M2VTS Database, AVBPA99, 1999.

9.   Fox, N.A., Gross, R., Chazal, P., Cohn, J.F., Reilly, R.B., 2003. Person Identification Using Automatic Integration of Speech, Lip and Face Experts, *Proc. of the ACM SIGMM 2003 Multimedia Biometrics Methods and Applications Workshop (WBMA'03),* Berkeley, CA, pp. 25-32.

10.   Indovina, M., Uludag, U., Snelik, R., Mink, A., Jain, A., 2003. Multimodal Biometric Authentication Methods: A COTS Approach, *Proceedings of the MMUA 2003, Workshop on Multimodal User Authentication,* Santa Barbara, CA, pp. 99-106.

11.   Lucey, S., Chen, T., 2003. Improved Audio-visual Speaker Recognition via the Use of a Hybrid Combination Strategy, *The 4th International Conference on Audio- and Video-Based Biometric Person Authentication,* Guilford, U.K.

12. Wang, Yuan; Wangm Yunhong; Tan, T., 2004. Combining Fingerprint and Voiceprint Biometrics for Identity Verification: an Experimental Comparison, *Proceedings of the ICBA,* Hong Kong, China, pp. 663-670.
13. Djamel Mostefa et al. "CLEAR Evaluation Plan v1.1" http://isl.ira.uka.de/~nickel/clear/downloads/chil-clear-v1.1-2006-02-21.pdf