Acoustic-Labial Speaker Verification [1]

P. Jourlin [a,b], J. Luettin [a], D. Genoud [a], H. Wassner [a]

[a] *IDIAP, rue du Simplon 4, CP 592, CH-1920 Martigny, Switzerland*

[b] *LIA, 339 chemin des Meinajariès, BP 1228, 84911 Avignon Cedex 9, France*

Abstract

This paper describes a multimodal approach for speaker verification. The system consists of two classifiers, one using visual features and the other using acoustic features. A lip tracker is used to extract visual information from the speaking face which provides shape and intensity features. We describe an approach for normalizing and mapping different modalities onto a common confidence interval. We also describe a novel method for integrating the scores of multiple classifiers. Verification experiments are reported for the individual modalities and for the combined classifier. The performance of the integrated system outperformed each sub-system and reduced the false acceptance rate of the acoustic sub-system from 2.3% to 0.5%.

Key words: Person authentication; Bimodal speech; Decision fusion; Lip feature extraction; Hidden Markov Models

[1] This work has been performed within the framework of the M2VTS (*Multi Modal Verification for Teleservices and Security applications*) project, granted by the

1 Introduction

Automatic verification of a person's identity is a difficult problem and has received considerable attention over the last decade. The ability of such a system to reject impostors, who claim a false identity, becomes a critical issue in security applications. The use of multiple modalities like face, profile, motion or speech is likely to decrease the possibility of false acceptance and to lead to higher robustness and performance (e.g. Acheroy et al. (1996)). Brunelli and Falavigna (1995) have previously described a bimodal approach for person identification. The system was based on visual features of the static face image and on acoustic features of the speech signal. The performance of the integrated subsystem was shown to be superior to that of each subsystem.

The cognitive aspect of lip movements in speech perception has been studied extensively and the complementary nature of the visual signal has been successfully exploited in bimodal speech recognition systems. The fact that temporal lip information not only contains speech information but also characteristic information about a person's identity has largely been ignored, until recently, where Luettin et al. (1996a) have proposed a new modality for person recognition based on spatio-temporal lip features.

In this paper, we extend this approach and address the combination of the acoustic and visual speech modality for a speaker verification system. We

ACTS program.

describe the normalization and mapping of different modalities and the determination of a threshold for rejecting impostors. A scheme for combining the evidence of both modalities is described and we show that the performance of the multimodal system outperforms both unimodal subsystems.

2 The Database

The M2VTS audio-visual database has been collected at UCL (Catholic University of Louvain) and is described by Pigeon and Vandendorpe (1997). It contains 37 speakers (male and female) pronouncing in French the digits from zero to nine. One recording is a sequence of the ten digits pronounced continuously. Five recordings have been taken of each speaker, at one week intervals to account for minor face changes like beards and hairstyle. The images contain the whole head and are sampled at 25 Hz.

We have divided the database into 3 sets : the first three shots were used as training set, the 4^{th} shot as validation set and the 5^{th} shot as test set. The 5^{th} shot represents the most difficult recordings to recognize. This shot differs from the others in face variation (head tilted, unshaved), voice variation (poor voice SNR), or shot imperfections (poor focus, different zoom factor).

3 Lip Feature Extraction

We are interested in facial changes due to speech production and therefore analyse the mouth region only. Common approaches in face recognition are often based on geometric features or intensity features, either of the whole face or of parts of the face (e.g. Chellappa et al. (1995)). We combine both approaches, assuming that much information about the identity of a speaker is contained in the lip contours and the grey-level distribution around the mouth area. During speech production the lip contours deform and the intensities in the mouth area change due to lip deformation, protrusion and visibility of teeth and tongue.

These features contain information specific to the speech articulators of a person and to the way that person speaks. We aim to extract this information during speech production and to build spatio-temporal models of a speaking person.

3.1 Lip Model

Our lip model is based on active shape models (e.g. Cootes et al. (1994)) and has been described in detail by Luettin and Thacker (1997). It is used to locate, track and parameterize the lips over an image sequence of a speaking person.

Features are recovered from tracking results. They describe the shape of the inner and outer lip contours and the intensity at the mouth area. The shape features and the intensity features are both based on principal component analysis which was performed on a training set.

The intensity model deforms with the lip contours and therefore represents shape independent intensity information. This is an important property of the model. We obtain detailed shape information from the shape parameters and therefore would like the intensity model to describe intensity information which is independent of the lip shape and lip movements (e.g. Luettin et al. (1996b)).

3.2 Lip Tracking

Experiments were performed on all 5 shots of the M2VTS database. The database consists of colour images which were converted to grey-level images for our experiments. Several subjects have a beard or did not shave between different recordings. We used examples from the training set to build the lip model. The model was then used to track the lips over all image sequences of all three sets. This consisted of analysing over 27 000 images which we believe is the largest experiment reported so far for lip tracking. It is important to evaluate the performance of the tracking algorithm and we have previously attempted to do this by visually inspecting tracking results (e.g. Luettin and

Thacker (1997)). However this task is very laborious and subjective. Here we omit direct performance evaluation of the tracking algorithm. Instead we try to evaluate the combined performance of the feature extraction and the recognition process by evaluating the person recognition performance only. Person recognition errors might therefore be due to inaccurate tracking results or due to classification errors. Examples of lip tracking results are shown in Fig 1.

4 Speaker Verification

4.1 Test Protocol

We use the sequences of the training set (first 3 shots) of the 36 customers for training the speaker models. The validation set serves for computing the normalization and mapping function for the rejection threshold and the test set is used for the verification tests. Subject 37 is only used as impostor, claiming the identity of all 36 customers. Each customer is also used as an impostor of the 35 other customers. The verification mode is text-dependent and based on the whole sequence of ten digits.

For the verification task, we make use of a *world* model, which represents the average model of a large number of subjects (500 speakers for the acoustic model and 36 for the labial one). For each digit we compute the corresponding customer likelihood and the world likelihood. We can so obtain a customer likelihood $L_c(O)$ and a world likelihood $L_w(O)$ for all speech data. The difference between the ratio of the two scores and the threshold t is then mapped to the interval $[0, 1]$ using a sigmoïd function (e.g. Genoud et al. (1996)) :

$$S(c) = \frac{1}{1 + \exp\left(-\left(\frac{L_c(O)}{L_w(O)} - t\right)\right)} \tag{1}$$

If a final score $S(c)$ is equal to 0.5, no decision is made; if it is below 0.5, the speaker is rejected else he is accepted. Several methods have been proposed in order to find an a priori decision threshold, according to various criteria, e.g. Equal Error Rate and Furui's (1994) method. Due to the small amount of speech data for each speaker we calculated a customer independent threshold, based on a dichotomic method. The use of this method implies that the function of verification errors is convex. This function is computed on the *validation* set and the value, for which the number of false acceptance and false rejection errors is minimum, is used as threshold value.

4.2 Acoustic Speaker Verification

Since the word sequences are known in our experiments, we use a HMM based speech recognition system to segment the sentences into digits. The recognizer uses the known sequence of digit word models, which were trained on the *Polyphone* database of IDIAP (e.g. Chollet et al. (1995)), to find the word boundaries.

Each digit HMM has been trained with 110 to 200 examples of 835 speakers. The segmentation is performed on all three sets. The segmented training set is used to train one model for each digit and speaker. These models are called *customer* models. The acoustic parameters are Linear Prediction Cepstral Coefficients with first and second order derivatives. Each vector has 39 components.

We used left-right HMMs with between 2 and 7 emitting states, depending on the digit length. Each state is modelled with one single Gaussian mixture with diagonal covariance matrix. The same configuration is used for the world model. The world model is trained on the *Polyphone* database using 300 examples from 500 speakers for each digit.

When an access test is performed, the speech is first segmented into digits. The test protocol described above is applied, where the customer and world likelihoods are obtained by the product of all digit likelihoods, using the customer and world models, respectively.

The mapping function, obtained from the validation set, is used in the test set to map the score into the confidence interval. On the test set, we obtain a false acceptance rate of 2.3% and a false rejection rate of 2.8%. The identification rate of the 36 speakers was 97.2% (see Fig 2 and Table 1). However, it is well worth noting that only 36 tests were conduced for identification and false rejection but 1332 (36×37) tests for false acceptance.

4.3 Labial Speaker Verification

For segmenting the labial data we use the previous acoustic segmentation. Lip features can improve speech recognition results (e.g. Jourlin (1996)), but they do not provide enough information to segment speech into phonemic units. Lip movements may be useful in addition to the acoustic signal for segmenting speech, especially in a noisy acoustic environment (e.g. Mak and Allen (1994) ; Jourlin et al. (1995)). We did not use visual information for segmentation since our acoustic models were trained on a very large database and are therefore more reliable than our labial models. We used the same scoring method for labial verification as for acoustic verification, except the world model, which was trained on the 36 customers from the M2VTS database. Labial data has a four times lower sampling frequency than acoustic data. The number of emitting states was therefore chosen to be 1 or 2, depending on the digit length. The parameter vectors consisted of 25 components: 14 shape and 10 intensity parameters and the scale. The same test protocol, which was used for acoustic experiments, was now used for labial verification. On the test data set, we obtained a false acceptance rate of 3.0% and a false rejection rate of 27.8%. The identification rate was 72.2% (see Fig 3 and Table 1).

4.4 Acoustic-Labial Verification

The acoustic-labial score is computed as the weighted sum of the acoustic and the labial scores. Both scores have been normalized as described in the previous sections. The process uses individual threshold values for each modality and maps the scores into a common confidence interval. The normalization process is a critical point in the design of an integration scheme and is necessary to ensure that different modalities are mapped into the same interval and share a common threshold value. The different modalities are now normalized but they provide different levels of confidence. We therefore need to weight the contribution of each modality according to their confidence level. The weight is α for the acoustic score and $1 - \alpha$ for the labial one. The same dichotomic algorithm, used to compute the thresholds, is now used to find the optimal weight α. The function of verification results on the validation data is used for the dichotomic search, for the same reasons as described for threshold search.

The following results were obtained on the test set : using a weight of 0.86, we obtain a false acceptance rate of 0.5%, a false rejection rate of 2.8% and a correct identification rate of 100.0%. The absolute gain over the acoustic system is a 1.8% reduction in cumulated verification errors (FA+FR) and an increase of 2.8% in the identification rate (ID). Fig 4 shows the effects of weighting on acoustic-labial results, when the acceptance threshold is optimaly fixed for each modality. Table 1 sums up the results.

We have followed a data-driven approach for fusion, where data fusion is present at different levels. At the first stage, learning and decoding of the labial models use the segmentation obtained from the acoustic models. The first score normalization is performed by normalizing the scores with respect to a world model for each modality. The final normalization is obtained by finding an optimal mapping in the interval $[0, 1]$ for each modality. At this stage, the two scores are normalized, but we know that each modality has different levels of reliability. So, the last level of the fusion process is to find the optimal weight for the two sources of information (see Figure 5).

5 Conclusion

We have described a novel approach for multi-modal person authentication based on acoustic and visual HMM speaker models. Experiments were performed for speaker identification and verification on one of the largest audio-visual speech databases.

Results have shown that the performance of the visual system was considerably lower than the performance of the acoustic system. This could be due to the lower frame rate of the visual features, lip tracking errors, inappropriate features, or an inappropriate modelling method. More research is necessary to investigate these issues. The integrated acoustic-labial system increased the identification rate of the acoustic system from 97.2% to 100% and reduced the false acceptance rate from 2.3% to 0.5%. These results show that acoustic and labial features contain complementary information and that the performance of an acoustic speaker recognition system can be improved by the use of visual information.

References

Acheroy, M., C. Beumier, J. Bigün, G. Chollet, B. Duc, S. Fischer, D. Genoud, P. Lockwood, G. Maitre, S. Pigeon, I. Pitas, K. Sobottka and L. Vandendorpe (1996). Multi-modal person verification tools using speech and images. In: *European Conf. on Multimedia Applications, Services and Techniques,* Louvain-la-neuve, Belgium, 747-761.

Brunelli, R. and D. Falavigna (1995). Person identification using multiple cues *IEEE Trans. Pattern Anal. Mach. Intell.* 17 (10), 955-966.

Chellappa, R., C.L. Wilson and S. Sirohey (1995). Human and machine recognition of faces: A survey. *Proceedings IEEE* 83 (5), 705-740.

Chollet, G., J-L. Cochard, A. Constantinescu and P. Langlais (1995). Swiss french polyphone and polyvar : telephone speech databases to study intra and inter speaker variability. Technical Report, IDIAP, Martigny, Switzerland.

Cootes, T.F., A. Hill, C.J. Talor and J. Haslam (1994). Use of active shape models for locating structures in medical images. *Image and Vision Computing* 12 (6), 355-365.

Furui, S. (1994). An overview of speaker recognition technology. In: *ESCA Workshop on Automatic Speaker Recognition Identification Verification*, Martigny, Switzerland, 1-9.

Genoud, D., F. Bimbot, G. Gravier and G. Chollet (1996). Combining methods to improve speaker verification decision. In: *Internat. Conf. Speech and Language Processing*, Philadelphia, PA, 1756-1759.

Jourlin, P., M. El-Bèze and H. Méloni (1995). Bimodal speech recognition. In: *Internat. Workshop on Automatic Face and Gesture Recognition*, Zurich, 320-325.

Jourlin, P. (1996). Handling disynchronization phenomena with HMM in connected speech. In: *European Signal Processing Conference*, Trieste, Italy, 133-136.

Luettin, J., N.A. Thacker and S.W. Beet (1996a). Speaker identification by lipreading. In : *Internat. Conf. Spoken Language Processing*, Philadelphia, PA, 62-65.

Luettin, J., N.A. Thacker and S.W. Beet (1996b). Speechreading using shape and intensity information. In: *Internat. Conf. Spoken Language Processing*, Philadelphia, PA, 58-61.

Luettin, J. and N.A. Thacker (1997). Speechreading using Probabilistic Models. *Computer Vision and Image Understanding* 65(2), 163-178.

Mak, M.W. and W.G. Allen (1994). Lip-Motion analysis for speech segmentation in noise. *Speech Communication* 14 (3), 279-296.

Pigeon, S. and L. Vandendorpe (1997). The M2VTS Multimodal Face Database

(Release 1.00). In: J. Bigün, G. Chollet and G. Borgefors, Ed.,*Audio- and Video-based Biometric Person Authentication*, Springer-Verlag, Berlin, 403-409.

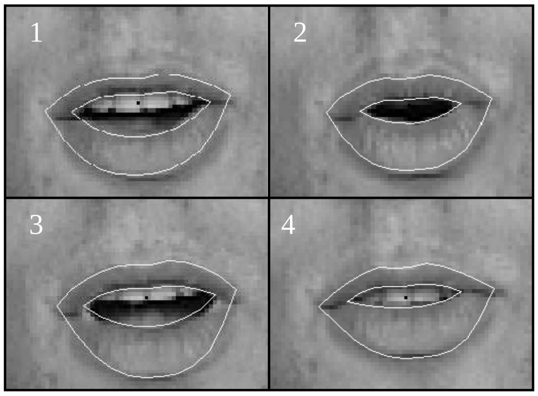

Fig. 1. Examples of lip tracking results

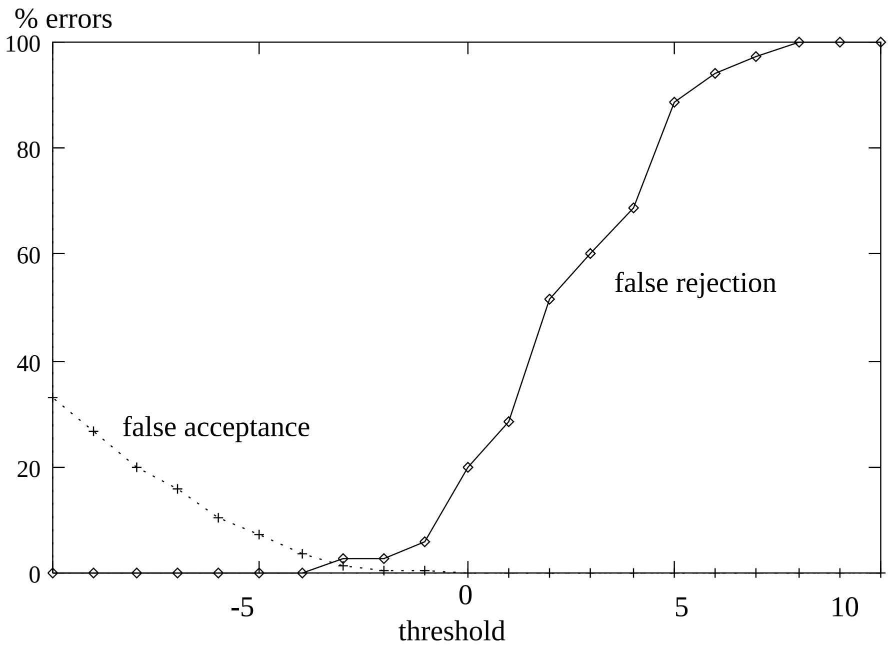

Fig. 2. Acoustic verification results (validation set)

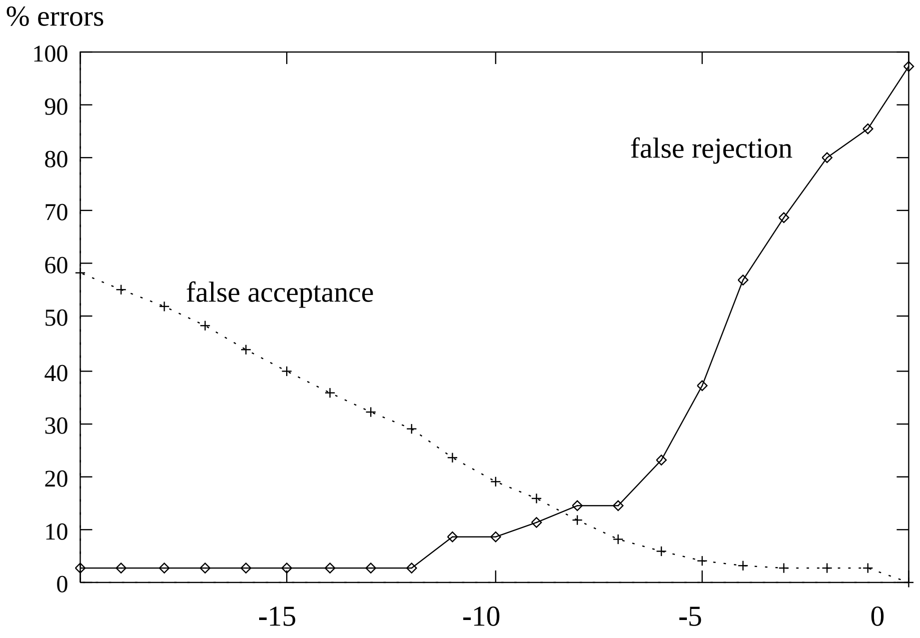

Fig. 3. Labial verification results (validation set)

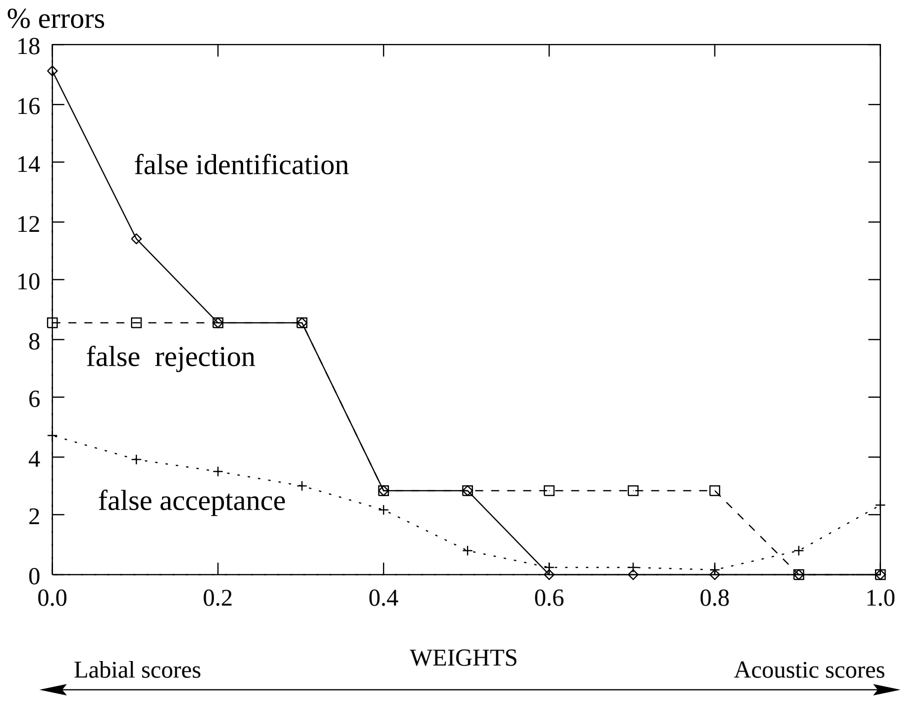

Fig. 4. Results with different weights (validation set)

Table 1

Results on validation and test set : ID is the correct identification, FA the false acceptance and FR the false rejection rate.

Type of score	Validation			Test		
	ID	FA	FR	ID	FA	FR
Acoustic	100.0	2.5	0.0	97.2	2.3	2.8
Labial	82.3	4.9	8.8	72.2	3.0	27.8
Bimodal	100.0	0.6	0.0	100.0	0.5	2.8
Number of tests	36	1332	36	36	1332	36

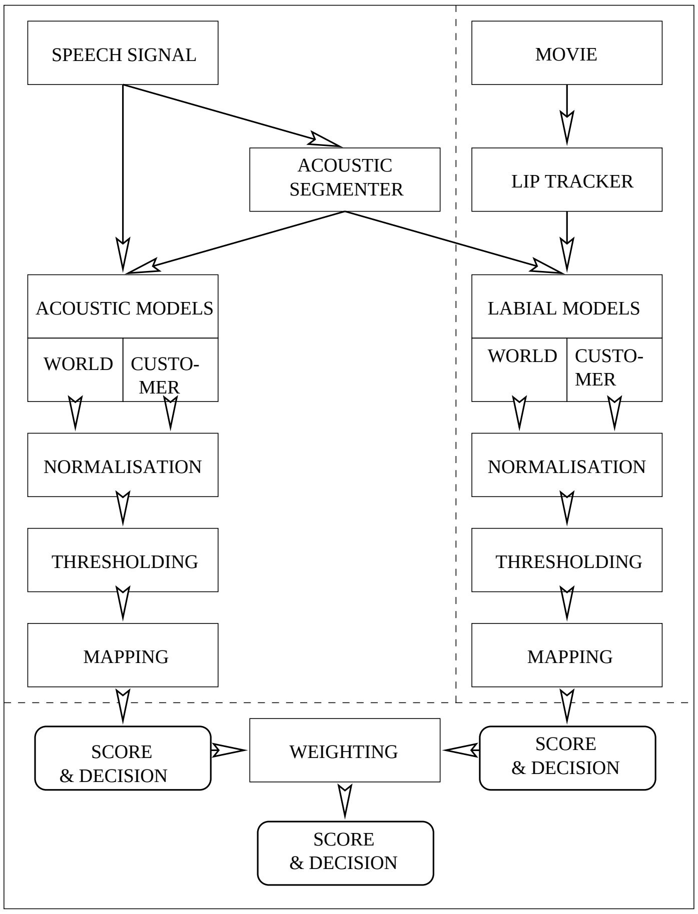

Fig. 5. Overview of the acoustic-labial system.