

# A Wearable Face Recognition System for Individuals with Visual Impairments

Sreekar Krishna, Greg Little, John Black, and Sethuraman Panchanathan

Center for Cognitive Ubiquitous Computing (CUbiC)

Arizona State University

Tempe, AZ – 85281

{Sreekar.Krishna, Greg.Little, John.Black, and Panch}@asu.edu

## ABSTRACT

This paper describes the iCare Interaction Assistant, an assistive device for helping the individuals who are visually impaired during social interactions. The research presented here addresses the problems encountered in implementing real-time face recognition algorithms on a wearable device. Face recognition is the initial step towards building a comprehensive social interaction assistant that will identify and interpret facial expressions, emotions and gestures. Experiments conducted for selecting a face recognition algorithm that works despite changes in facial pose and illumination angle are reported. Performance details of the face recognition algorithms tested on the device are presented along with the overall performance of the system. The specifics of the hardware components used in the wearable device are mentioned and the block diagram of the wearable system is explained in detail.

## Categories and Subject Descriptors

J.0 General

## General Terms

Design, Experimentation, Performance.

## Keywords

Face Recognition, Wearable Computing, Assistive Device for Visually Impaired, Social Interaction Aide.

## INTRODUCTION

Humans, knowingly or unknowingly, participate in social interaction in their day-to-day life. Social interactions are the acts, actions, or practices of two or more people mutually oriented towards each other. Such interactions come in many forms - blinking, eating, reading, writing, dancing and walking. Vision plays such an important role in establishing and maintaining social interactions that it is sometimes challenging for individuals who are visually impaired to interact readily with their sighted counterparts.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ASSETS'05, October 9–12, 2005, Baltimore, Maryland, USA.

Copyright 2005 ACM 1-59593-159-7/05/0010...\$5.00.

Studies have shown that a significant portion of any information exchange between two humans is accomplished not with words, but with non-verbal-communication. Furthermore, most of these non-verbal-communications are facial gestures – though other bodily gestures also constitute a large portion.

As easy as it is for humans to understand body gestures, it has become the testing ground for intelligent machines. Assistive devices designed to facilitate social interactions are a good example of a type of machine intelligence that is still a long way from reality. In this paper, we discuss the smaller, but indispensable, problem of face recognition in the context of building a social interaction assistant to aide people who have visual impairments.

Face recognition has been an active area of research for the last decade, due to the availability of fast computing systems and increased security requirements in public places. This research has led to the development of improved algorithms, as well deployment of access control and identity verification systems, based on face recognition. Although there are numerous algorithms today that can achieve an acceptable level of recognition when face images are captured in a controlled environment, there are no algorithms capable of recognizing people reliably in real-world situations. Face recognition for any assistive device would require algorithms that are more robust than what is being achieved today by training algorithms on controlled face datasets.

Research focused on developing face recognition algorithms for security purposes tends to focus on finding methods that can achieve good recognition even when the person under surveillance wears disguises, such as facial hair, sun glasses and head gear. This requirement greatly limits the features that can be used to by face recognition algorithms, and tend to make them less than suitable for practical use in wearable devices. Contrast this with face recognition algorithms for assistive devices, which do not generally assume that the face being recognized is disguised. This allows any stable facial feature to be used for recognition – potentially providing a much more robust recognition. For example, the presence of a pair of eyeglasses could be regarded as an impediment to face recognition for security purposes, while the same eyeglasses could be used a cue for identifying an individual in a social

occasion. Recognizing this, our approach to face recognition is based on finding facial details that are unique to a particular face, even though they might be very vulnerable to deliberate disguise.

Irrespective of the differences between applications, one problem faced by all face recognition algorithms, results from changes in pose angle and changes in the illumination angle on the face. During group social interactions it is quite common to see frequent extreme changes in pose angle. (See the Theory section). The human brain deals with these problems by projecting a 2D retinal projection of a face into a pose angle and light invariant space, making it possible for us to recognize people despite such variations. Research [17] [18] [19] along such directions of thought have yielded promising results but not satisfactory.

This paper describes the research that we have conducted in pursuit of building a robust face recognition system for aiding people who are visually impaired or blind. The paper is divided into two major parts. The first part deals with the algorithmic side of the problem, describing the experiments conducted towards selecting a face recognition algorithm for a wearable face recognition system. The second part of the paper deals with the hardware aspects of the issue, detailing the choice of the different components that make up the wearable device, along with their performance details.

The rest of the paper is organized as follows. The Related Works section deals with the research that has gone into face recognition, and into building wearable devices for assisting the people who are visually impaired. The Theory section follows this, with a description of our approach to selecting the particular face recognition algorithm to be implemented on our wearable device, along with the hardware details. The Results section provides some insight into the performance of the device in a real world scenario, and the Future Work section offers a glimpse into possible extensions of the device, towards becoming a complete social interaction assistant.

## RELATED WORK

Face recognition has been an active area of research for the past three decades. Biometrics and law enforcement have been the most researched application areas for face recognition [1]. Researchers have used static images [2]-[5], video sequences [14] [15], infrared images [13] [16] and 3D range data [12] for achieving face recognition. While some researchers worked with the face image as a whole [2]-[5], many others have explored the possibility of analyzing face images by modeling the local characteristics of the face [7].

Among the most widely used and researched face recognition algorithms, five algorithms, namely Principle Component Analysis (PCA) [2], Linear Discriminant Analysis (LDA) [3], Bayesian Intrapersonal Classifier (BIC) [4], Hidden Markov Model (HMM) [5] and Elastic

Bunch Graph Matching (EBGM) [7], are probably the best known. In fact, these five algorithms have formed the basis for most of the research in the area of face recognition. The statistical approaches (including PCA, LDA and BIC) work on the face image as a whole, treating each face image as a point in a multidimensional space. The recognition rates of these algorithms depend heavily on the capture conditions, and slight changes in those conditions can result in a drastic reduction in the performance of the algorithms. HMM and EBGM are classified as network-based approaches, where the face image analysis is carried out by modeling the statistical and positional characteristics of the facial features into connected networks. The performance of such algorithms is dependent on the positional accuracy of feature extraction algorithms whose output can change drastically with slight changes in pose.

Parallel to the development of face recognition algorithms, the systematic empirical evaluation of these algorithms has resulted in FERET [8]-[10] and XM2VTS [11] protocols that have provided a basis for comparing and testing face recognition algorithms. Both of these protocols include a set of color or gray scale face images that are used to test algorithms. Detailed procedures are provided for analyzing the results of the experiments, in order to compare the performance of the algorithms. Though these protocols provide a basis for evaluating face recognition algorithms, no effort has been made to accurately record the two very important parameters of pose angle and illumination angle in the face images during the capture of the images. In this paper, we describe our work in establishing a new methodology for comparing the performance of face recognition algorithms, using a novel face database whose face images are very accurately calibrated with respect to pose and illumination angle.

Assistive devices for people who are blind have been of interest to both academia and industry. Most of the research has focused on developing navigational aides for the people who are blind, based on Global Positioning Systems (GPS) and infrared based proximity sensors. The decreasing size of navigational devices and computing elements has guided the technological advances in this area. Small form factor high definition cameras have also entered mass production recently, and this has motivated many developers to migrate towards the development of vision-based technologies for assisting people who are blind or visually impaired. Some of the noticeable ongoing projects include the *iCare project* [20] which is developing a *Reader*, an *Information Assistant*, an *Interaction Assistant*, and a *Haptic Interface* for the people who are blind, *vOICe* [21], which is a video-to-sound interface that translates video input into auditory excitations for the people who are blind, and *EyeTap* [22] which is a set of Personal Imaging Lab projects focused on personal imaging, mediated reality, and wearable computing. Researchers at Kyoto Institute of Technology

(<http://www.kit.ac.jp/english/index.html>) have created a wearable device to help people who are blind navigate along streets. Among all of these research projects, iCare Interaction Assistant is unique in being the only vision-based device specifically for helping people who are visually impaired involve more easily in social interactions.

## THEORY

### Choosing the Face Recognition Algorithm:

#### *The Database:*

As mentioned in the Related Works section, face recognition algorithms have always been tested on publicly available databases, such as the AT&T Database, The Oulu Physics Database, The XM2VTS Database, the Yale Face Database, the MIT Database, the CMU Pose, Illumination and Expression Database, The FERET Database, and the Purdue AR Database. In order to provide robust face recognition, an algorithm must be sensitive to subtle differences in image content that are useful for distinguishing between faces. However, equally important is its ability to disregard image content that is particular to the environment in which the image was captured, such as the illuminant. If the development of such an algorithm is based on a face database that was not captured with a range of pose angles and illumination angles, and if each image is not annotated with a precise set of values for those environmental variables, it is difficult to correlate face recognition failure (or success) with changes in these variables, and to refine the face recognition algorithm to be more tolerant of changes in these particular environmental variables. A face database that does not include a range of images to represent the values of each independent variable also complicates comparisons between different face recognition algorithms because two algorithms might have similar failure rates, even though they have failed for totally different reasons.

Some of the databases mentioned above have face images with a wide variety of pose angle and illumination angle variations. However, none of them use a precisely calibrated mechanism for acquiring these images. To address this issue (and to achieve a precise measurement of recognition robustness with respect to pose and illumination angle) we put together a database called FacePix [6], which contains face images with pose and illumination angles annotated in 1-degree increments. Figure 1 shows the apparatus that is used for capturing the face images. A video camera and a spot light are mounted on independent annular rings that can be rotated independently around a subject seated in the center. The angle markings on the platform and the face images are captured simultaneously into the frames of a video sequence, from which frames can be extracted as individual calibrated images.

The FacePix(30) database contains two sets of images for each of 30 different people. Each set contains (1) a set of

181 images with pose angles between  $-90$  and  $+90$  degrees, and (2) a set of 181 images with illumination angles between  $-90$  and  $+90$  degrees. The entire FacePix(30) database can be conceptualized as a 2D matrix of face images with 30 rows (representing the 30 different people), and 181 columns (representing all the angles from  $-90^\circ$  to  $+90^\circ$  at 1 degree increments).



**Fig. 1:** The face image capture setup.

All the face images (elements) in each matrix are 128 pixels wide and 128 pixels high. These face images are normalized, such that the eyes are centered on the 57<sup>th</sup> row of pixels from the top, and the mouth is centered on the 87<sup>th</sup> row of pixels. The pose angle images appear to rotate such that the eyes, nose, and mouth features remain centered in each image. Also, although the images are down sampled, they are scaled as much horizontally as vertically, thus maintaining their original aspect ratios. Figure 2 provides examples extracted from the database, showing pose angles and illumination angles ranging from  $-90^\circ$  to  $+90^\circ$  in steps of  $10^\circ$ .



**Fig. 2:** A subset of one face set taken from the FacePix(30) database, with Pose and Illumination angles ranging from  $+90$  degrees to  $-90$  degrees, in steps of 10 degrees.

#### *Comparative Study of Face Recognition Algorithms:*

Having built a database that captures the variations in pose and illumination, we selected four of the most widely used face recognition algorithms – PCA, LDA, BIC and HMM and plotted their recognition rate as the pose and illumination angles were varied over a range from  $-90$  to  $+90$  degrees, to produce a pair of “robustness curves”. (The robustness is the ability of the algorithm to learn a person’s face from a given set of pose or illumination images, and

then recognize that same person from a never-before-seen pose angle or illumination angle.

We ran several experiments on the FacePix(30) database, and combined the results of all these experiments to gauge the overall robustness of four different face recognition algorithms. Each experiment measured the degradation in recognition rate as an algorithm attempts to recognize probe (test) images that are farther and farther (in terms of pose or illumination angle) from the gallery (training) set. Each such experiment may be conceptualized as a function, with the following inputs:

1. Algorithm to test: PCA, LDA, BIC, or HMM
2. Database set: Pose angle, or Illumination angle
3. The Gallery (training) set list: One or more columns from a given database set, e.g., all the images at pose angles -90, 0, and +90 (NOTE: In this scheme, each gallery set contains only one image of each subject. However, some of the algorithms we tested needed multiple versions of each pose angle or illumination angle image. To satisfy these algorithms, we artificially manufactured 3 additional versions of each gallery image. One of these images was a low-pass filtered version of the original image, and two of these images were noisy versions of the original image.
4. The Probe set: The entire 2D matrix of the database set.

The output of this function is the “distance” of each probe image to the “nearest” image in the gallery set(s). Using these distances, we produced a rank ordering of the 30 people for each probe image. (The person with a rank of 0 was computed to be the “closest” to that probe image). These ranking numbers then provided a basis for computing the robustness ( $R$ ) for an algorithm trained with the chosen gallery sets.

The robustness at a particular angle  $\theta$  is given by

$$R(\theta) = 1 - \frac{2}{N-1} \left( \frac{1}{N} \sum_{i=1}^N r_{\theta_i} \right)$$

Where,  $N$  is the number of subjects in the database.

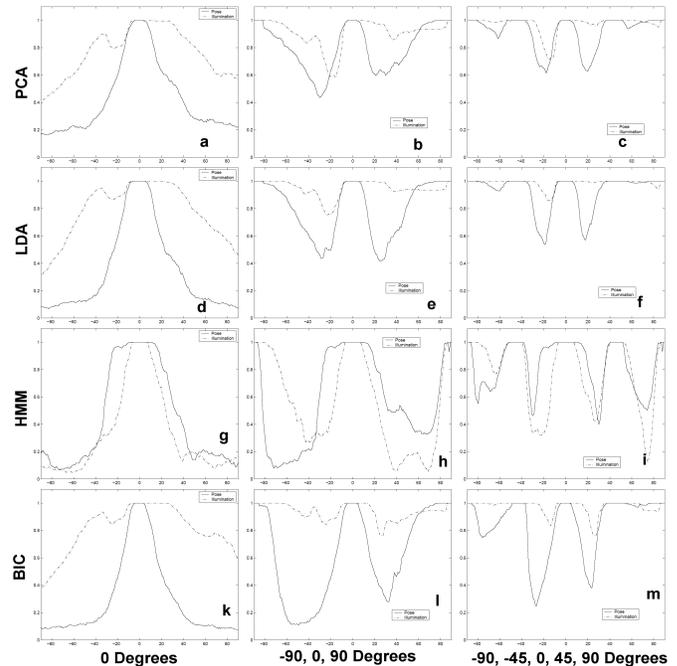
$r_{\theta_i}$  is the rank that is assigned for the  $i^{\text{th}}$  subject at the pose or illumination angle  $\theta$  (this rank value ranges from 0 to  $N-1$ ).

A Robustness value of 1 means that the recognition was accurate; while a value of 0 means that the recognition was no better than guessing randomly.

Figure 3 shows the robustness curves for all four face recognition algorithms, as a function of pose and illumination angles. The solid line shows the pose angle robustness, while the dotted line shows the illumination angle robustness. Each row in Figure 3 corresponds to one face recognition algorithm, and each column corresponds to a different training set. The first column shows the results

when the algorithms were trained with just the 0 degree (frontal) images, while the second column shows the results when trained with -90 degree (left profile), 0 degree (fontal), and +90 degree(right profile) images. The third column shows the results when trained with -90 degree, -45 degree, 0 degree, +45 degree, and +90 degree. Table 1 and Table 2 show the average robustness across all pose angles and all illumination angles respectively, while Table 3 and Table 4 show the average recognition rate across all pose angles and illumination angles.

It is clear that HMM is the poorest performing algorithm. From the roll off regions of the robustness curves, it is clear that the two subspace methods (PCA and LDA) have a more gradual roll off than the probabilistic methods (HMM and BIC). Accordingly, they have a better recognition rate across changes in both pose and illumination angles. The roll off rate is higher near 0 degrees (i.e. frontal views) than at the edges (i.e. profile views). This suggests that better overall robustness might be achieved by using a more densely spaced gallery set around the frontal region than towards the profile regions.



**Fig. 3:** Robustness curves for four widely used face recognition algorithms.

	0°	-90°, 0°, 90°	-90°, -45°, 0°, 45°, 90°
PCA	0.4306	0.8039	0.9178
LDA	0.3485	0.7985	0.9245
BIC	0.3003	0.6250	0.8443
HMM	0.4391	0.5764	0.8317

**Table 1:** Average Robustness for Pose

	0°	-90°, 0°, 90°	-90°, -45°, 0°, 45°, 90°
PCA	0.7715	0.9136	0.9735
LDA	0.7820	0.949	0.9888
BIC	0.3278	0.5483	0.7735
HMM	0.8108	0.9388	0.9740

**Table 2:** Average Robustness for Illumination

	0°	-90°, 0°, 90°	-90°, -45°, 0°, 45°, 90°
PCA	20.74%	50.53%	71.66%
LDA	20.70%	56.92%	78.67%
BIC	31.68%	41.27%	63.50%
HMM	18.42%	45.19%	69.47%

**Table 3:** Overall recognition rate for Pose changes

	0°	-90°, 0°, 90°	-90°, -45°, 0°, 45°, 90°
PCA	48.84%	71.71%	90.33%
LDA	53.04%	79.52%	94.92%
BIC	19.26%	37.38%	59.37%
HMM	49.80%	79.10%	93.54%

**Table 4:** Overall recognition rate for Illumination changes

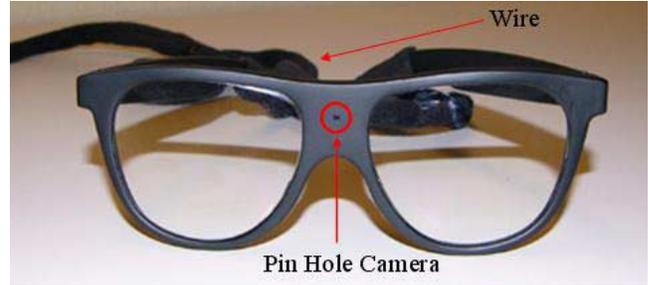
Comparing the results from each of the algorithms with respect to pose angle variance, LDA ranks first, followed by PC. Close behind is BIC, with HMM being the last. For illumination angle variance, LDA performs the best, followed by BIC, PCA and HMM, respectively.

These results were used as the basis for selecting algorithms that were tested on the wearable device. The performances of the tested algorithms are presented in the latter part of this paper.

### The Wearable Face Recognition Device:

The hardware used for building the assistive device essentially consists of three components.

1. An analog CCD camera used for acquiring the video. Fig 4 shows the camera glasses that are used for the assistive device. The camera has a 1/3" CCD with a light sensitivity of 0.2 Lux. The 92 degrees Field of View (FoV) provides a good coverage of the space in front of the user. The camera is powered using a 9V battery and the output is in NTSC video format.



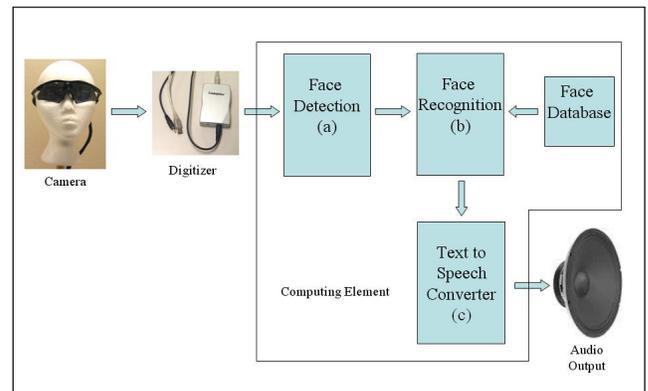
**Fig. 4:** Glasses used for the wearable face recognition system

2. Since the camera provides an analog video output, a digitizer is required to convert the composite video into a digital video format that can be used inside a computer for analysis. We used an Adaptec<sup>®</sup> video digitizer which converts the input signals into compressed AVI and transmits the AVI stream over a USB cable. The device driver is based on the standard Windows Driver Model (WDM) and appears to the programmer as a generic video capture device on the Windows<sup>®</sup> operating system.

3. A portable computing element (in our case a laptop) was used to execute the face recognition algorithm. We used a tablet PC with an Intel<sup>®</sup> Centrino 1.5 GHz processor and 512 MB of RAM. The choice of this particular laptop device was based on its small form factor. Fig 5 shows the frames acquired from the wearable device. The block diagram of the wearable assistive device is shown in the Fig 6.



**Fig. 5:** Frames from a video sequence obtained from the wearable device.



**Fig. 6:** Block Diagram of the wearable face recognition system.

*(a) The Face Detection Algorithm:*

The first step towards face recognition is to isolate the regions of the video frames where a human face exists. To this end, we used a face detection algorithm based on adaptive boosting [23]. A video frame acquired from the camera is divided into a number of overlapping regions of predetermined size. Each of these regions is analyzed for the presence of human face by using a bank of known filters (In this case, rectangular filters which together are representative of the intensity variations on a typical human face image. For example: the eye sockets on the face tend to be low intensity regions when compared with the forehead. A rectangular filter looking for such an intensity change in an image would have a width about equal to the width of the face, and a height divided into a white region (corresponding to the forehead), and a black region, (corresponding to the eye sockets). Analyzing every region extracted from the video frame is time consuming. To reduce the processing time, each region extracted from the video frame is passed through a cascade of filter banks. The filter bank in the beginning of the cascade has fewer filters, resulting in higher number of false positives, but a faster processing time. A filter bank at the end of the cascade has a large number of filters, and is capable of detecting a face with a better accuracy, but requires more processing time. The advantage is that the regions of the frame that have no resemblance to a face (a plain wall for example) is dropped at the beginning of the cascade as a non-face region, with very little lost processing time. On the other hand, a region with a face has to be accepted by all the filter banks to the end of the cascade. As expected, a region that somewhat resembles a face (but is not a face) will be dropped by a filter bank in the middle of the cascade. This results in very good processing time per frame, and makes face detection possible in real time. Fig 7 shows a video frame with a region marked as face. Fig 8 shows an example set of face images cropped out of a video sequence from the wearable device.



**Fig. 7:** Output from the face detection algorithm with the face region marked.



**Fig. 8:** Example set of face images from the wearable device

*(b) The Face Recognition Algorithm for the assistive device:*

Once a region in a video frame is identified as a face, it is analyzed in more detail, in an attempt to recognize the person. Inspecting the video frames in Fig 8, it can be seen that changes in pose angle are common, when compared to changes in the illumination. From Table 3, it is evident that LDA is the best performing algorithm with varying poses, followed by PCA. Since our Robustness results were obtained using a tightly controlled and calibrated face database, we tested both these algorithms on the face images that were extracted out of the video frames coming from the camera on the wearable device.

For testing the performance of our wearable device, 450 images of 10 different individuals were collected in an office environment (See Fig. 8). These images were then divided into two equal groups – one for training and the other for testing. Two experiments were carried out to compare the performance of PCA and LDA for face recognition on the wearable device.

*Experiment 1:*

1. The training set images were used for deriving the PCA and LDA subspaces, as described in [2] and [3] respectively.
2. The projection weights for all the training images were obtained by projecting them on to the subspaces that were derived in the Step 1.
3. For each subject in the training set, an average PCA or LDA projection weight vector was computed and stored as the individuals “identification vector”,  $\mathbf{i}$ .
4. When a face image had to be recognized, it was projected on to the PCA or LDA subspace that was derived in Step 1 and a weight vector,  $\mathbf{w}$ , was obtained.
5. The distance between the weight vector,  $\mathbf{w}$ , and all the individual identification vectors,  $\mathbf{i}$ , obtained in Step 3 were computed. The identification vector closest to  $\mathbf{w}$ , was chosen as the guess for the person in the test image.

*Experiment 2:*

1. The training set images were used to derive the PCA and LDA subspaces, as described in [2] and [3] respectively.

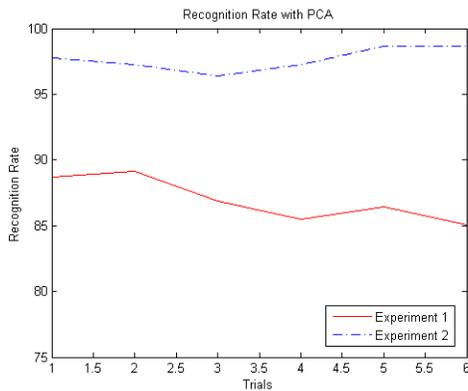
2. The projection weights for all the training images were obtained by projecting them on to the subspaces that were derived in the Step 1.
3. All the projection weights obtained in the previous step were labeled with the person to whom the image belonged, and were stored as the “identification vectors”.
4. When a face image had to be recognized, it was projected on the PCA or LDA subspace that was derived in the Step 1, and a weight vector,  $\mathbf{w}$ , was obtained.
5. The distance between the weight vector,  $\mathbf{w}$ , and all the “identification vectors” were computed. The label corresponding to the identification vector closet to  $\mathbf{w}$  was chosen as the guess for the person in the test image.

*(c)Text-to-speech Converter:*

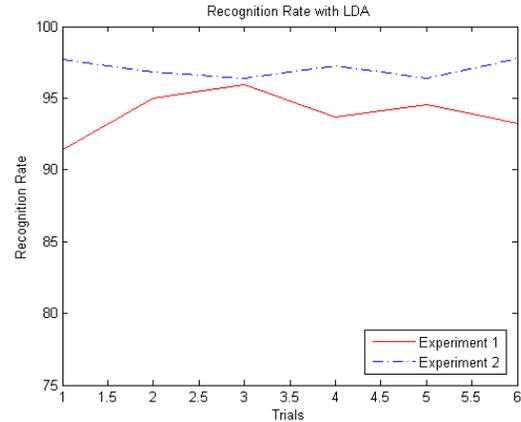
When the system produces a guess for the person in the video frame, the user is notified with an audio signal. Here, we used the Microsoft Speech Engine to convert the name of the identified individual from text to speech. This was fed to the headphones that the user wears. Upon further experimentation, we noticed that the face recognizer learnt the face images based on the environment where they were captured. Thus the face recognizer would sporadically recognize a certain person as some one else because the lighting conditions on the face changed momentarily. To accommodate for such a situation, the text-to-speech converter waits for the face recognizer to recognize the same individual in five consecutive frames before the name of the person is spoken out.

**RESULTS**

Fig. 9 and Fig. 10 show the performance of PCA and LDA face recognition algorithms on the images that were captured from the wearable device. Experiments 1 and 2 were conducted five times, with the training and testing images shuffled between trials. Fig. 11 shows the comparison of PCA and LDA for the same trails.



**Fig. 9** Recognition Performance using PCA

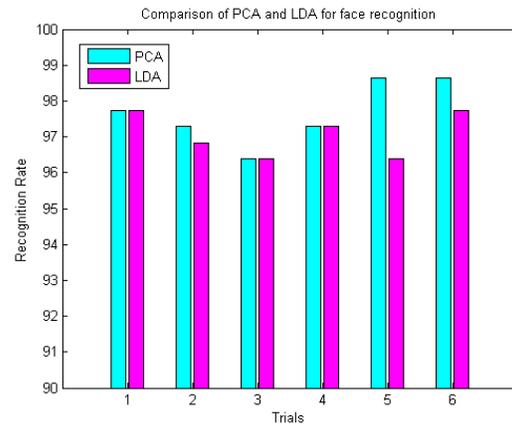


**Fig. 10** Recognition Performance using LDA

Table 5 shows the average time taken by PCA and LDA for recognizing a single face image, averaged over five trails.

	Experiment 1	Experiment 2
LDA	40.723 ms	12.656 ms
PCA	85.832 ms	40.583 ms

**Table 5:** Average time for recognizing single face image



**Fig. 11** Comparison of recognition performance between PCA and LDA

**DISCUSSION**

Looking at Fig. 9 and Fig. 10, it is evident that the performance in Experiment 2 is significantly higher than in Experiment 1, for both PCA and LDA. However, the average time for recognizing a single face is much higher in Experiment 2. This is due to the fact that  $N$  number of comparisons must be carried out, where  $N$  is the total number of face images in the training database. On the other hand, in case of Experiment 1, the total number of comparisons is equal to the number of subjects  $S$  in the database. For the experiments carried out here,  $S = 10$  and  $N = (450/2)$ . Though there is a significant difference in the execution times between the two, we chose Experiment 2 as

a model for our face recognition work, due to the higher accuracy of the face recognition.

Inspecting Fig. 11, it can be inferred that the performance of PCA is better than (or similar to) that of LDA. Further, the implementation complexity of PCA is lower than that of LDA. Though LDA is twice as fast as PCA, we opted for PCA to be the face recognition algorithm on the wearable device, due to its higher recognition rate.

## CONCLUSION AND FUTURE WORK

In this paper we have presented a wearable face recognition system and provided performance data for this device. Details of the method for selecting the most appropriate face recognition algorithm for this device were provided, along with a description of the hardware components that were used for the wearable system.

Having studied the performance of face recognition algorithms that treats a face image as a whole, experiments are being conducted to understand the performance of face recognition algorithms that model the local facial features of individuals. Simultaneous efforts are being made to acquire better performing cameras and low form factor computing elements such as handhelds and PDAs for the wearable device.

## ACKNOWLEDGMENTS

We want to thank Terri Hedgpeth, the project investigator for iCare. We also want to thank Ken Spector and David Paul for being part of the iCare focus group.

## REFERENCES

1. W. Zhao, R. Chellappa, and A. Rosenfeld. Face Recognition: A Literature Survey. *Technical Report CAR-TR948*, UMD CfAR, 2000.
2. M. Turk, and A. Pentland. Face recognition using Eigenfaces, *Proc. of IEEE Conference on Computer Vision and Pattern Recognition*, pp. 586–591, 1991.
3. K. Etemad, and R. Chellappa. Discriminant analysis for recognition of human face images. *Journal of Optical Society of America*, pp. 1724–1733, 1997.
4. B. Moghaddam, and A. Pentland. Probabilistic visual learning for object representation. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 19(7), pp. 696–710, July 1997.
5. A. Nefian, and M H Hayes III. Hidden markov models for face detection and recognition. *IEEE International Conference on Image Processing*, vol. 1, pp. 141–145, October 1998.
6. J. Black, M. Gargesha, K. Kahol, P. Kuchi, and S. Panchanathan. A framework for performance evaluation of face recognition algorithms. *ITCOM, Internet Multimedia Systems II, Boston*, July 2002.
7. L. Wiscott, J. M. Fellous, and C. von der Malsburg. Face Recognition by Elastic Buncg Graph Matching. *IEEE Trans. on Pattern Analysis and Machine Intelligence*. Vol. 19. pp. 775-779, 1997.
8. P. J. Phillips, P. Rauss, and S. Der. FERET (Face Recognition Technology) Recognition Algorithm Development and Test Report. *Technical Report ARL-TR 995*, U. S. Army Research Laboratory.
9. P. J. Phillips, H. Moon, P. Rauss, and S. A. Rizvi. The FERET Evaluation Methodology for Face-Recognition Algorithms. *Proc. IEEE conference on Computer Vision and Pattern Recognition*. pp 137-142, 1997.
10. P. J. Phillips, H. Moon, S. A. Rizvi, and P. Rauss. The FERET Testing Protocol. *Face Recognition: From Theory to Applications*. (H. Wechsler, P. J. Phillips, V. Bruce, F. F. Soulie, and T. S. Huang, eds.) Berlin: Springer-Verlag. Pp. 224-261, 1998.
11. K. Messer, J. Matas, J. Kittler, J. Luetin, and G. Maitre. XM2VTSDB: *The Extended M2VTS Database*. *Proc. International Conference on Audio- and Video-based Person Authentication*. pp. 72-77, 1999.
12. G. Gordon. Face Recognition Based on Depth Maps and Surface Curvature. *SPIE Proc. Vol. 1570: Geometric Methods in Computer Vision*. pp. 234-247, 1992.
13. J. Wilder, P. J. Phillips, C. H. Jiamg, and S. Wiener. Comaprison of Visible and Infra-Red Imagery for Face Recognition. *Proc. International Conference on Automatic Face and Gesture Recognition*. pp. 182-187, 1996.
14. K. G. Bahadir, U. B. Aziz, Y. Altunbasak, H. H. Monson III, and Russell M. Mersereau. Eigenface-Domain Super-Resolution for Face Recognition. *IEEE Trans. On Image Processing*, Vol. 12, No. 5, May 2003.
15. O. Yamaguchi, E. Fukui, and K. Maeda. Face recognition using temporal image sequence. *Proc. Third IEEE International Conference on Automatic Face and Gesture Recognition*. pp 318–323. April 1998.
16. X. Chen, P. J. Flynn, and K.W. Bowyer. PCA-based face recognition in infrared imagery: baseline and comparative studies. *IEEE International Workshop on Analysis and Modeling of Faces and Gestures*. Pp 127 – 134, Oct 2003.
17. F. J. Huang, Z. Zhou, H. Zhang, and T. Chen. Pose Invariant Face Recognition. *Proc. of the 4th IEEE International Conference on Automatic Face and Gesture Recognition*, Grenoble, France. pp.245-250, 2000.
18. B. Gokberk, L. Akarun, E. Alpaydin. Feature selection for pose invariant face recognition. *Proc. 16th International Conference on Pattern Recognition*. Volume 4. pp: 306 - 309. 2002.
19. S. Romdhani, V. Blanz, and T. Vetter. Face Identification by Fitting a 3D Morphable Model using Linear Shape and Texture Error Functions. *Computer Vision ECCV'02*, Copenhagen, Denmark. Vol. 4. pp: 3-19, 2002.
20. iCARE Projects. <http://cubic.asu.edu>.
21. The vOICE. <http://www.seeingwithsound.com>.
22. The EyeTap. <http://eyetap.org/>
23. P. Viola, and M. Jones. Robust Real-time Object Detection. *Second International Workshop on Statistical and Computational Theories of Vision – Modeling, Learning, Computing, and Sampling*, Vancouver, Canada, July 13, 2001.