

A Control System for Assessing Commercial Face Recognition Software for Racial Bias

Orimolade Joseph Folorunso
Department of Electrical and
Electronics Engineering, Federal
University of Agriculture Abeokuta
PMB 2240, Alabata, Abeokuta,
Ogun State, Nigeria

Oronti Iyabosola Busola
School of Engineering,
University of Warwick
Coventry, CV4 7AL, United
Kingdom

Olopade Abdullah Oluwatosin
Department of Electrical and
Electronics Engineering
Federal University of Agriculture
Abeokuta,
PMB 2240, Alabata, Abeokuta,
Ogun State, Nigeria

ABSTRACT

Depending on how algorithms are trained, they could be significantly more accurate when identifying white faces than black ones. It has recently been shown that algorithms trained with biased data result in algorithmic discrimination. During training, an algorithm is given pairs of face images of the same person, and it learns to pay more attention to features that strongly indicate that the two images represent the same person. Recently, with facial recognition becoming more prevalent in law enforcement and consumer products, there is increasing concern that such systems are ominously less accurate for people with black skin. In this work, a database of still images is created with a total of 132 black-faces representing 22 individuals, having images cropped to pixel sizes of 100x100, 80x80, 60x60 and 40x40 respectively. Two separate face recognition algorithms are also developed using Principal Component Analysis (PCA) and Linear Discriminant Analysis (LDA). Performance indices of the PCA and LDA algorithms are assessed in terms of the recognition rate, error rate, false rejection rate (FRR), and false acceptance rate (FAR). The objective of this research is to provide a control measure for testing racially-biased error rates in commercially available face recognition software.

General Terms

Pattern Recognition and Artificial Intelligence.

Keywords

Face Library, Racially Biased Data, Principal Component Analysis, Linear Discriminant Analysis, Face Recognition Software, Performance Indices.

1. INTRODUCTION

The face is the only natural member of the body which serves as the primary focus of recognition and social association [1]. The human ability to recognize faces is quite remarkable. Since the onset of researches in artificial intelligence (AI), researchers have been working hard to transfer the same power of recognition exhibited in humans to machines, thus resulting in the technology called face recognition. Face recognition uses information harvested from an individual's facial features to achieve automated identification. It is a more complex problem than merely detecting the presence of a human face [2], [3]. Basically, the technology compares two face images in a fairly complex process to determine whether they represent the same person. Face recognition is intrinsically probabilistic in the sense that the results produced are not exact, but rather identifies possible matches [4]. Face recognition is thus very useful in finding a person within a

large database of faces, especially in law enforcement, access control, criminal investigations, image processing and bio-data database organization.

2. THE CURRENT CHALLENGE

It is common knowledge that companies such as HP, Apple, Microsoft, IBM and Google have released commercial software that perform automated facial analysis which is now inbuilt into most smart phones and laptops as a security mechanism. However, new research and real-life examples show that existing commercial facial recognition systems are more prone to either misidentify or fail to identify persons with darker skin tones, especially the black race. A few examples include earlier facial recognition systems from HP and Google which failed to recognize people with dark skin, a HP webcam failing to register black people in 2009, and Google Photos facial recognition software categorizing black people as gorillas in 2015 [5]. This in effect means that facial recognition technology is inherently subject to biases based on the data sets provided and the conditions under which algorithms are developed [4]. In September, 2017, when Apple debuted its new facial recognition system, Face ID, which securely unlocks iPhone X [6], the corporation immediately faced questions about how the product would overcome the security and bias problems exhibited in the past by similar facial recognition systems deployed by other companies.

In a recent research, Buolamwini and Gebu [7] built a dataset of 1,270 faces using the faces of politicians from countries with a high percentage of women in office, and tested the accuracy of three leading facial recognition systems, namely, Microsoft's, IBM's, and Megvii of China. These corporations were selected because gender classification features were offered in their facial analysis software, and their codes were publicly available for testing. The fallouts of the tests showed inaccuracies in gender identification based on an individual's skin color. Lighter-skinned subjects were better classified than darker individuals. Given the foregoing, it can be reasonably argued that facial recognition algorithms may be most accurate on the populations who developed them. This situation thus gives rise to the need for diverse data sets, as well as diversity among the people who create and deploy facial recognition technologies in order to avoid racial bias [4], [8], [9], [10].

Some existing studies related to performance evaluation of face recognition systems were outlined in [2]. A novel performance evaluation model (PEM) and tool were proposed to evaluate the performance of biometric recognition systems

and to enable comparison and evaluation of face recognition systems. Also, the National Institute of Standards and Technologies (NIST) conducts voluntary tests of facial-recognition vendors every four years. Although accuracy rates have improved drastically between each round of testing, many algorithms have been found to show marked differences in accuracy across race, gender, and other demographics [2], [11].

The objective of this study is to contribute to the diversity in the collections of images that accurately capture races of interest in face recognition algorithms, thereby increasing the consistency of testing regimes for existing systems, and making it easier to train new algorithms. The developed face recognition system is thus not designed as a commercial application, but as a control mechanism for testing mainstream face recognition software for racial bias. It is hoped that this work will aid in the development of a comprehensive set of bias tests for face recognition algorithms. The dataset is not yet publicly available, but can be obtained on request.

3. EXPERIMENTAL PROCEDURE

This work makes use of frontal shots of undergraduate students in the Federal University of Agriculture Abeokuta, Nigeria. The faces were captured with a camera to form a database of still black face images. Face recognition algorithms using PCA and adapted LDA were separately developed to carry out preprocessing, classification and recognition processes on the face database. Three major phases feature in the face recognition system design.

3.1 Face Library Formation

In this phase, face images of twenty-two different individuals were acquired using a W110 Sony Webshot 7.2 mega pixel camera. Six frontal shots of each individual were taken in a controlled environment with a pre-specified factor of illumination. Four images from each individual were used to represent a training dataset while the remaining two represented the test dataset. Twenty-two classes of training and test datasets for different individuals were thus employed to form the face database, which was then stored in a face library on a computer system. A cross-section of faces used in the database is illustrated in Figure 1(a).



Figure 1(a): Cross-section of cropped images



Figure 1(b): Resized face image

The captured images were cropped to accommodate only the frontal facial section, and then further resized to 100x100, 80x80, 60x60, and 40x40 pixel sizes respectively. Resizing of the images was done to further investigate the performance of the face recognition system at all specified image dimensions. These processed images formed the face database for the face

recognition system. Thereafter, the images were converted from color images into gray-scale, where each pixel point on the image is represented by an 8-bit value depending on its intensity. Hence, the face image was characterized by a two-dimensional N by N array of 8-bit intensity, and each image considered as a vector of dimension N^2 , such that a typical image size of 100×100 becomes equivalent to a point in a 10,000-dimensional space. These were the faces that were used eventually to train and test the system.

3.2 Training Phase

Given any dataset described by ‘ n ’ variables or features, the dataset can be described with a smaller set of new variables. This is known as reducing the dimensionality of a dataset. As part of the training process, the image vectors obtained in the preceding step were further subjected to the dimensionality reduction process, while still keeping their identity. This procedure was observed by both PCA and adapted LDA in the training phase.

3.2.1 PCA Training Algorithm

In PCA, a new set of feature images called principal components or Eigenfaces form a linear combination of the original. In this work, PCA for dimensionality reduction was applied to the image database as illustrated in [13], [14], [15] respectively. A summary of the procedure is shown as follows:

3.2.1.1 Average Face Vector Computation

For a set of training image set given by $\Gamma_1, \Gamma_2, \Gamma_3, \dots, \Gamma_M$, the average or mean face vector (Ψ) is obtained from the expression:

$$\Psi = \frac{1}{M} \sum_{n=1}^M \Gamma_n \quad (1)$$

3.2.1.2 Computation of Difference Vectors

Each face vector in the training image dataset differs from the average face vector by the difference vector, which is given by:

$$\Phi_i = \Gamma_i - \Psi \quad (2)$$

3.2.1.3 Covariance Matrix computation

The covariance matrix of the set of difference vectors is given by:

$$C = \frac{1}{M} \sum_{n=1}^M \Phi_n^T \Phi_n \quad (3)$$

$$C = A^T A, [N^2 \times N^2 \text{ matrix}] \quad (4)$$

Where, $A = [\Phi_1, \Phi_2, \Phi_3, \dots, \Phi_M]$, which is equivalent to a $[N^2 \times M]$ matrix.

3.2.1.4 Computation of Eigenvectors and Eigenvalues

In practice, the matrix in Equation 4 is very large and impractical, and we would need to obtain N^2 Eigenvectors (v_i) and Eigenvalues (λ_i) of the covariance matrix C . The dimensionality reduction process becomes relevant in this case. The Eigenvectors of a matrix AA^T can be obtained by first finding the solution to a $[M \times M]$ matrix given by $S = AA^T$. A common theorem in linear algebra enables the transformation of a $N^2 \times N^2$ matrix into a $M \times M$ matrix, thus reducing the order of the number of pixels in the images N^2 to the order of the number of images in the training set (M). A full analysis of this procedure is shown in [16].

Eigenvectors \mathbf{v}_i of S can then be computed using:

$$A A^T \mathbf{v}_i = u_i \mathbf{v}_i \quad (\text{Where, } u_i = A A^T) \quad (5)$$

By multiplying each of the Eigenvectors of $A A^T$ by A , we obtain;

$$A A^T A \mathbf{v}_i = u_i A \mathbf{v}_i$$

Such that:

$$C A \mathbf{v}_i = u_i A \mathbf{v}_i \quad (\text{Since } C = A A^T A) \quad (6)$$

Thus, $C = A^T A$ and $S = A A^T$ assume the same Eigenvalues, and their Eigenvectors are related as shown in Equation 6.

3.2.1.5 Eigenspace Computation by Ordering Eigenvectors

Eigenvectors \mathbf{v}_i are ordered according to their corresponding Eigenvalues from high to low, to allow for proper characterization of variations across images. Only Eigenvectors associated with non-zero Eigenvalues are kept, and the Eigenvector associated with the largest Eigenvalue is the one that reflects the greatest variance in the image. This matrix of Eigenvectors produce an orthonormal Eigenface basis V in a space called the Eigenspace, where each column of V is an Eigenface.

3.2.1.6 Projection of Difference Image Vector on Eigenspace

Each of the difference image vector Φ_i is projected into the Eigenspace. To do this, the dot product of the difference image with each of the ordered Eigenvectors is calculated as shown in Equation 7.

3.2.1.7 Covariance Matrix computation

The covariance matrix of the set of difference vectors is given by;

$$\bar{\Phi}_i = \sum_{i=1}^M V_i \cdot \Phi_i \quad (7)$$

Where, $\bar{\Phi}_i$ = Projected Images

V_i = Ordered Eigenvectors (Eigenfaces)

Φ_i = Difference Image Vectors

3.2.2 Adapted LDA Training Algorithm

LDA provides a class-specific linear method for dimensionality reduction. The traditional approach [13], [17] is adapted for this work using the Fisherface method. This method selects the optimal projection matrix W in such a way that the ratio of the between-class scatter to the within-class scatter is maximized. Further improvements on the LDA algorithm can be found in [18], [19], [20]. Generally, problems associated with face recognition algorithms include the small sample size (SSS) problem, where the number of images in the training set N is much smaller than the number of pixels in each image Γ_i , and the fact that LDA's optimizing criteria are not directly related to its classifying ability. To compensate for the above, the training algorithm is adapted as follows:

3.2.2.1 Incorporation of PCA Step

The PCA algorithm for dimensionality reduction is incorporated into the LDA structure. PCA is used as a preprocessing step so that eigenvectors associated with zero eigenvalues are discarded. This effectively removes the null space of the within-class scatter matrix of the training data set and reduces the dimension of the feature space to $N - X_i$, where X_i represents twenty-two classes of the training dataset. The steps are specified in the previous section.

3.2.2.2 Introduction of Optimizing Criterion

Here, the optimizing criterion is introduced and performed on the lower dimension PCA subspace. This criterion seeks to maximize the ratio of the between-class scatter to the within-class scatter. The between-class scatter matrix is defined as:

$$S_b = \sum_{i=1}^c N_i (\mu_i - \mu) (\mu_i - \mu)^T \quad (8)$$

While the within-class scatter matrix is given as follows;

$$S_w = \sum_{i=1}^c \sum_{X_k \in X_i} (X_k - \mu_i) (X_k - \mu_i)^T \quad (9)$$

Where,

μ_i = mean image of class X_i

N_i = number of samples in class X_i

3.2.2.3 Determining the Optimal Projection Matrix

The optimal projection matrix W is chosen such that the within-class scatter of the projected samples is made to be zero. Since S_w is a singular matrix, the image set is projected to a lower dimensional space ($X_i - 1$), such that the resulting S_w matrix is non-singular with orthonormal columns, and the image subspace contains Fisherfaces.

W is thus given by;

$$W^T = W_{pca}^T W_{fld}^T \quad (10)$$

Where,

$$W_{pca} = \arg \max_w |W^T S_w W| \quad (11)$$

$$W_{fld} = \arg \max_w \left| \frac{W^T W_{pca}^T S_b W_{pca} W}{W^T W_{pca}^T S_w W_{pca} W} \right| \quad (12)$$

The solution obtained defines the axes of the Fisher face space.

3.3 Testing Phase

3.3.1 Recognition Performance

The testing phase involves checking the recognition performance of the face recognition algorithm. Faces to be tested also pass through all image preprocessing stages as earlier outlined. Two scenarios play out. The two images per class set aside for the test dataset are verified for recognition purposes while some images not included in the training set are also tested for recognition. In the first scenario, the expectation is that the test dataset will be recognized as faces in the face database. On the other hand, the second scenario demands that the system will bring out a non-member faces outcome. The test image is projected on the computed Eigenspace and Fisher Face Space for the PCA and LDA algorithms respectively. Hence, the projection of the test image is compared to those of the training images.

3.3.2 Classification Tool

In this work, the Euclidean distance is the tool for which classification is established for both the PCA and adapted LDA algorithms. If two weight vectors are assumed to be points in an image space, a line segment between these two points give the Euclidean distance connecting them. Assuming a n -dimensional Euclidean space \mathbf{R}^n , vectors are identified with the scalar components;

$$\mathbf{a} = (a_1, a_2, a_3, \dots, a_{n-1}, a_n) \quad (13)$$

An arbitrary image point, Ω , can thus be described by the feature vector

$$\Omega = [a_1(\Omega), a_2(\Omega), a_3(\Omega), \dots, a_n(\Omega)] \quad (14)$$

Then, the distance between two points Ω_i and Ω_k is defined in Equation 14 as;

$$d(\Omega_i, \Omega_k) = \sqrt{\sum_{r=1}^n (a_r(\Omega_i) - a_r(\Omega_k))^2} \quad (15)$$

The Euclidean distance is given by $d(\Omega_i, \Omega_k)$, and it provides a measure of similarity between the two weight vectors Ω_i and Ω_k . If the distance is below a user defined threshold ε_k as given in Equation 16, the class to which the smallest Euclidian distance exists is the class to which the test image belongs, and the algorithm returns the image as “Known”. However, if an Euclidean distance threshold classification value for identification is exceeded, the test image is classified as “Unknown”;

$$\frac{|\Omega_i - \Omega_k|}{|\Omega_k|} \leq \varepsilon_k \quad (16)$$

where ε_k is the user defined threshold, Ω_i is the weight vector of the trained face image and Ω_k is the weight vector of the test image. The flow chart for the recognition algorithm is presented in Figure 2.

4. RESULTS AND DISCUSSION

PCA and LDA were both implemented with all acquired and preprocessed faces at 40x40, 60x60, 80x80 and 100x100 pixels (cropping level). A total of 22 classes (different

individual images) corresponding to 88 (4 faces per class) training faces and 44 (2 faces per class) test images were used to determine the performance of the two algorithms. Likewise, 29 impostors were used to calculate the False Acceptance Rate (FAR) of the Face Recognition System. MATLAB (R2014) was used to develop the program code which was initially run on an Intel(R) Pentium(R) computer with dual CPU T3400 at 2.16 GHz processing speed. Parameters considered in the assessment of the performance of the two algorithms include cropping dimensions, total number of test images, total number of identified face images, total number of unidentified faces, percentage recognition rate, percentage error rate, time taken to train face images, and time taken to identify face images.

4.1 Experimental Results

Both PCA and LDA implementations of the face algorithm were tested. For the PCA algorithm, recognition rate ranged between 81.8% and 90.9%, and the error rate between 9.1% and 18.2%; recognition rate for LDA alternated between 84.1% and 93.2%, and the error rate between 6.8% and 15.9%. Tables 1 and 2 show percentage recognition and error rate values for the PCA and LDA face recognition algorithms respectively at various cropping levels.

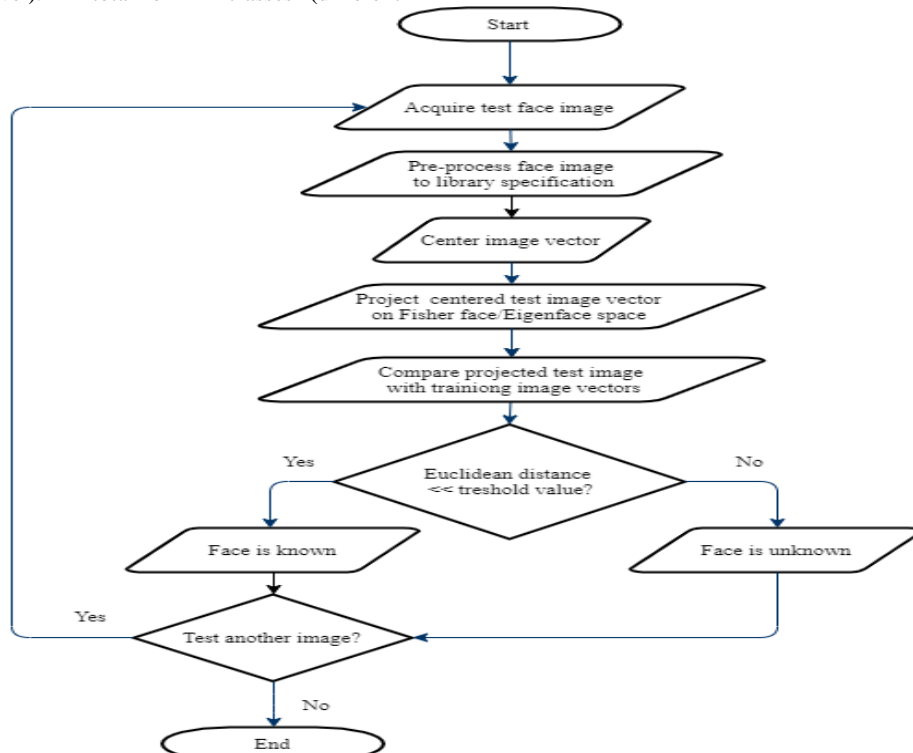


Figure 2: Flow chart of recognition phase of the adapted LDA/PCA face recognition system

Table 1. Percentage recognition and error rate of PCA algorithm at various cropping levels

| Resolution of Cropped Face Image | Total Number of Testing Images | Total Number of Identified Images | Total Number of Unidentified Images | Percentage Recognition Rate (%) | Percentage Recognition Error (%) | Time to Train face database (seconds) | Time to Identify an Image as Known or Unknown |
|----------------------------------|--------------------------------|-----------------------------------|-------------------------------------|---------------------------------|----------------------------------|---------------------------------------|---|
| 40 x 40 | 44 | 36 | 8 | 81.80 | 18.20 | 2.68 | 0.12 |
| 60 x 60 | 44 | 37 | 7 | 84.10 | 15.90 | 3.54 | 0.15 |
| 80 x 80 | 44 | 37 | 7 | 84.10 | 15.90 | 4.62 | 0.17 |
| 100 x 100 | 44 | 40 | 4 | 90.90 | 9.10 | 5.70 | 0.18 |

Table 2. Percentage recognition and error rate of adapted LDA algorithm at various cropping levels

| Resolution of Cropped Face Image | Total Number of Testing Images | Total Number of Identified Images | Total Number of Unidentified Images | Percentage Recognition Rate (%) | Percentage Recognition Error (%) | Time to Train face database (seconds) | Time to Identify an Image as Known or Unknown |
|----------------------------------|--------------------------------|-----------------------------------|-------------------------------------|---------------------------------|----------------------------------|---------------------------------------|---|
| 40 x 40 | 44 | 37 | 7 | 84.10 | 15.90 | 3.15 | 0.17 |
| 60 x 60 | 44 | 40 | 4 | 90.90 | 9.10 | 3.76 | 0.21 |
| 80 x 80 | 44 | 41 | 3 | 93.20 | 6.80 | 5.02 | 0.24 |
| 100 x 100 | 44 | 41 | 3 | 93.20 | 6.80 | 5.89 | 0.31 |

Again, from the test results obtained, we were able to compare the recognition rate and error rate of PCA and LDA individually in order to better observe general trends. Figures 3 and 5 illustrate this contrast at different cropping dimensions while Figures 4 and 6 show comparisons between training and test time of image sets.




Figure 3: Recognition rate of PCA and LDA algorithms

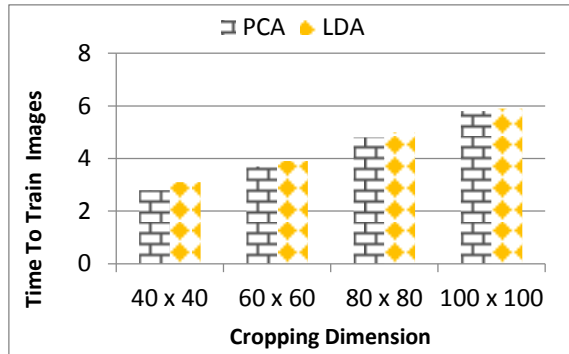


Figure 4: Training time for PCA and LDA algorithms

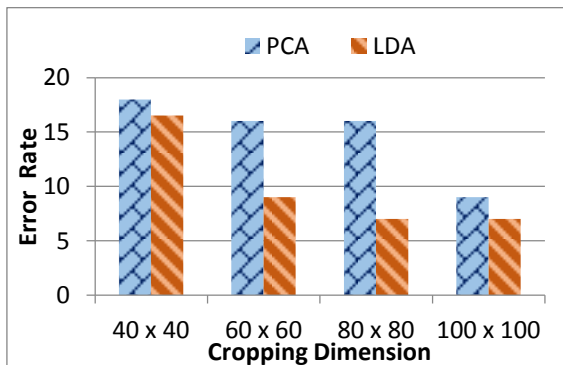


Figure 5: Error Rate of PCA and adapted LDA Algorithm




Figure 6: Test time for PCA and adapted LDA algorithms

Finally, the robustness and classification capabilities of the two algorithms is compared in terms of the false acceptance rate (FAR) using face images not admitted in the face database, as well as the false rejection rate (FRR) using the test dataset. Equations 17 and 18 show how to calculate for FAR and FRR values respectively while Tables 3 and 4 show the FAR and FRR values for the PCA and LDA algorithm at different crop dimensions.

$$\text{False Acceptance Rate (FAR)} = \frac{\text{Number of falsely identified face images}}{\text{Total number of face images used}} \times 100\% \quad (17)$$

$$\text{False Rejection Rate (FRR)} = \frac{\text{Number of rejected face images}}{\text{Total number of face images used}} \times 100\% \quad (18)$$

Table 3. False acceptance rate values

| Resolution of Cropped Face Images | Total Number of Images used in Testing | PCA | | LDA | |
|-----------------------------------|--|---------------------------------------|---------------------------------|---------------------------------------|---------------------------------|
| | | Total Number of False Accepted Images | False Acceptance Rate (FAR) (%) | Total Number of False Accepted Images | False Acceptance Rate (FAR) (%) |
| 80 x 80 | 29 | 5 | 15.9 | 3 | 10.3 |
| 100 x 100 | 29 | 1 | 3.4 | 0 | 0 |

Table 4. False recognition rate values

| Resolution of Cropped Face Images | Total Number of Images used in Testing | PCA | | LDA | |
|-----------------------------------|--|---------------------------------------|---------------------------------|---------------------------------------|--------------------------------|
| | | Total Number of False Accepted Images | False Acceptance Rate (FAR) (%) | Total Number of False Accepted Images | False Acceptance Rate(FAR) (%) |
| 80 x 80 | 29 | 7 | 15.9 | 3 | 6.8 |
| 100 x 100 | 29 | 4 | 9.1 | 3 | 6.8 |

4.2 Results Analysis

Visual outcomes of tests carried out on both PCA and adapted LDA face recognition algorithms represented in the preceding tables and graphs show that both algorithms record increases in recognition rates as crop dimensions increase. For PCA, recognition rate ranged between 81.8% and 90.9%, while varying between 84.1% and 93.2% for LDA. The LDA algorithm performs consistently better at recognizing faces, while recording less error rates at all crop dimensions. Conversely, the PCA algorithm spends less time in training and testing the face recognition system. This might be an advantage in situations where speed of program execution is required over accuracy.

Tables 3 and 4 show calculated values for both the FAR and FRR of the face recognition system using image dimensions of 100x100 and 80x80 pixel sizes respectively. It is evident from these tables that both FAR and FRR values decrease with increase in the crop dimensions of faces. Despite the lower values of these two test parameters at higher crop dimensions for the two algorithms, the PCA algorithm is still seen to be underperforming when compared with LDA. These results inevitably lead us to conclude that the face recognition system implemented with the LDA algorithm is a better one.

5. CONCLUSION AND FURURE WORK

A face database and recognition system has been developed using both PCA and adapted LDA algorithms, and the design procedures are presented in this work. The dimension of images in the face database was initially reduced using PCA. Training and testing of the developed system was conducted using both PCA and LDA algorithms. Analysis of the performance of the system at varying image crop dimensions is also presented. It was verified that for both algorithms, as cropping dimensions increase, recognition rate also increases, while the FRR, FAR and error rate reduce. This indicates that the more the facial features present in the face image, the more the percent recognition rate of the system. Training and test time for both the PCA and LDA is also shown to be accomplished in fractions of seconds. Overall, it was observed that the adapted LDA algorithm performed better than the PCA algorithm at recognizing faces.

This work can be put to two types of uses. First, the image database can be employed in the control testing of commercial systems oriented towards recognizing static images by obtaining the relative performance of particular implemented algorithms when fed with black faces compared to other demographic datasets. This procedure actually has the potential of facilitating more accurate, independent and standardized tests for racially biased error rates. Secondly, image datasets used to train and test specific commercial face recognition software can be fed into the control algorithm to verify claims of speed and accuracy. It is hoped that when these methods are employed, face recognition algorithms will

be better able to accurately recognize individuals irrespective of race.

Projections for future work include balancing gender presence in the compiled image database, and adopting a predictive model such as Deep Learning for improving the performance of the face recognition system, either independently or in combination with any of the statistical approaches used in this current effort. Similarly, future projections include adapting the system to operate in real-time mode as against the static mode in which it is currently implemented. Finally, possibilities for interfacing with external peripherals when deployed in access control systems will be explored.

6. REFERENCES

- [1] Omidiora, A. 2005. *Face recognition systems using OPCA, OLDA and ANN*. A PhD Thesis, Ladoke Akintola University of Science and Technology, Ogbomoso, Nigeria.
- [2] YongNyuo, S., Jason, K., YongJun, L., Woochang, S., Jin, Y. C. 2007. Performance Evaluation Model for the Face Recognition System. In: *Proceedings of the Frontiers in the Convergence of Biosciences and Information Technology, FBIT 2007*, pp. 704-708.
- [3] Huang, G. B., Ramesh, M., Berg, T., Learned-Miller, E. 2007. Labeled faces in the wild: a database for studying face recognition in unconstrained environments. cs.brown.edu.
- [4] Garvie, C., Bedoya, A. M., Frankle, J. 2016. The perpetual line-up; unregulated police face recognition in America. Georgetown Law, Center on Privacy and Technology, USA.
- [5] Conger, K. 2017. How apple says it prevented face ID from being racist. [Online]. Available at: <https://gizmodo.com/how-apple-says-it-prevented-face-id-from-being-racist-1819557448>
- [6] Apple, 2017. *Face ID security*. [Online]. Available at: https://www.apple.com/business/docs/site/FaceID_Security_Guide.pdf
- [7] Buolamwini, J., and Gebru, T. 2018. Gender shades: intersectional accuracy disparities in commercial gender classification. In: *Conference on Fairness, Accountability, and Transparency, Proceedings of Machine Learning Research*. 81:1–15.
- [8] Goode, L. 2018. *Biases are seeping into software*. [Online]. Available at: <https://www.theverge.com/2018/2/11/17001218/facial-recognition-software-accuracy-technology-mit-white-men-black-women-error> [Accessed 25 April 2018].
- [9] Lohr, S. 2018. Facial recognition is accurate, if you're a white guy. [Online]. Available at: <https://www.nytimes.com/2018/02/09/technology/facial-recognition-race-artificial-intelligence.html> [Accessed 25 April 2018].
- [10] Garvie, C., and Frankle, J. 2016. Facial-Recognition Software Might Have a Racial Bias Problem. [Online]. Available at: <https://www.theatlantic.com/technology/archive/2016/04/the-underlying-bias-of-facial-recognition-systems/476991/> [Accessed 28 April 2018].
- [11] Phillips, P. J., Jiang, F., Narvekar, A., Ayyad, J., O'Toole, A. J. 2011. An other-race effect for face

- recognition algorithms. *ACM Transactions on Applied Perception*, 8:2, pp. 14:1-14:5.
- [12] Klare, B. F., Burge, M. J., Klontz, J. C., Bruegge, R. W. V., Jain A. K. 2012. Face recognition performance: role of demographic information. *IEEE Transactions on Information Forensics and Security*, 7:6, pp. 1789-1797.
- [13] Belhumeur, P. N., Hespanha, J. P., Kriegman, D. J. 1997. Eigenfaces vs. fisherfaces: recognition using class specific linear projection. *IEEE Transactions On Pattern Analysis And Machine Intelligence*, 19:7, pp. 711-720.
- [14] Paul, L. T., Al Sumam, A. 2012. Face recognition using principal component analysis method. *International Journal of Advanced Research in Computer Engineering & Technology*, 1:9, pp. 135-139.
- [15] Kaur, R., Himanshi, E. 2015. Face recognition using principal component analysis. *IEEE International Advance Computing Conference*, pp. 585-589.
- [16] Turk, M., Pentland, A. 1991. Eigenfaces for recognition. *Journal of Cognitive Neuro-science*, 3:1, pp. 73-86.
- [17] Martinez, A. M., and Kak, A. C. 2001. PCA versus LDA. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23:2, pp. 228-233.
- [18] Dai, G., Qian, Y. 2004. Face Recognition Using Novel LDA-Based Algorithms. In: *Proceedings of the 16th European Conference on Artificial Intelligence, ECAI'2004*, including Prestigious Applicants of Intelligent Systems, PAIS 2004, Valencia, Spain
- [19] Shin, Y., Kim, J., Lee Y., Shin, W., Choi, J. 2007. Performance evaluation model for the face recognition system. *Frontiers in the Convergence of Bioscience and Information Technologies*, pp. 704-708.
- [20] Fischler, M. A., Elschlager, R. A. 1973. The representation and matching of pictorial structures. *IEEE*