An Analysis of Racial Bias in Facial Recognition Systems: Elastic Bunch Graph Matching with Principal Component Analysis

Patrick Wande
wpwande19@earlham.edu
Earlham College
Richmond, Indiana, USA

ABSTRACT

This paper investigates racial bias in the Principal Component Analysis (PCA) optimization of the Elastic Bunch Graph Matching (EBGM), a facial recognition technique. Utilizing the FERET database, this study assesses the algorithm's accuracy across diverse racial groups, focusing on non-Caucasian versus Caucasian faces. Through rigorous experimental design and statistical analysis, we evaluate performance disparities and aim to understand how PCA optimization impacts racial bias in facial recognition. The findings seek to inform strategies to enhance the fairness and inclusivity of such technologies.

Keywords: Facial recognition, EBGM, PCA, Optimization, Comparative Analysis, Racial Bias

1 INTRODUCTION

Facial recognition technology is a field that has been gaining prominence for its applications in various fields such as security, surveillance, and identity verification [4]. Despite its capabilities, EBGM is susceptible to the common biases found in machine learning algorithms, which can undermine its effectiveness across diverse populations. Addressing these biases is crucial for the practical and equitable application of facial recognition technologies.

Our research focuses on Principal Component Analysis (PCA) as a primary optimization technique for EBGM. We aim to determine whether optimizing EBGM with PCA can mitigate racial biases inherent in the algorithm. By employing PCA, we reduce the dimensionality of data in the FERET database—a dataset featuring individuals from various racial and ethnic backgrounds—and analyze how these changes affect the algorithm's performance in recognizing non-Caucasian faces compared to Caucasian ones. This study uses a detailed experimental framework and evaluates metrics such as accuracy and false positive/negative rates to discern performance variations across different racial groups.

Through this investigation, we endeavor to understand whether performance optimizations like PCA can effectively reduce racial biases within facial recognition algorithms, thereby improving their fairness and reliability in real-world applications

2 BACKGROUND AND RELATED WORKS

2.1 Elastic Bunch Graph Matching

The foundation of our research is the Elastic Bunch Graph Matching (EBGM). EBGM is quite proficient at recognizing individual faces from single images within large databases, where each person is represented by just one image. It has also made significant contributions to the field of facial recognition[4]

Its process of operation is described below:

(1) Jet Extraction: Jets refer to small texture descriptions at specific points in a facial image. These jets are essential for EBGM because they contain details about how a face looks. They are obtained by using Gabor wavelet transformations at key facial landmarks[7].
(2) Bunch Graph Construction: The jets are then grouped into bunch graphs, a unique data structure. In this structure, each facial landmark corresponds to a node, and each node contains jets representing the associated landmark. For example, if the eyes are chosen as landmarks, one node might contain jets for the right eyes of all model images, while another node holds the jets for the left eyes of all model images[6].
(3) Face Graph Representation: These bunch graphs together form a database of landmark descriptions. In EBGM, a face is represented as a graph with nodes and edges. Nodes represent facial features, like eyes, nose, and mouth, while edges show how these features relate to each other[6].
(4) Similarity Assessment: When landmarks in a new facial image are identified, its jets are extracted and compared with the corresponding jets in the database. A face graph for the new image is created, which contains its landmarks and jet values. This information is used to understand how similar the images are and complete the recognition process[6].

This research will aim to evaluate the effectiveness of Principal Component Analysis as an optimization of EBGM. This approach, introduced by Chen et al.[2], utilizes Principal Component Analysis (PCA) to process training and test images more efficiently. It picks the best candidates for EBGM, making it faster and better at recognizing faces, especially when dealing with lots of images[2].

The primary objective of this project is to empirically evaluate the performance of PCA optimization in terms of accuracy and its ability to recognize faces of different racial groups. Through systematic experimentation, we seek to gain insights into the effectiveness of this optimization technique.

Our research methodology involves utilizing facial images sourced from the FERET database. Renowned for its comprehensive collection of standardized frontal facial images, the FERET database offers a reliable and widely accepted dataset for conducting facial recognition experiments[2, 9].

2.2 Principal Component Analysis (PCA) in Elastic Bunch Graph Matching (EBGM)

Principal Component Analysis (PCA) is a dimensionality reduction and feature extraction technique widely used in various fields, including facial recognition[5]. In the context of EBGM, PCA is integrated into the algorithm to optimize its performance and enhance the efficiency of facial recognition.

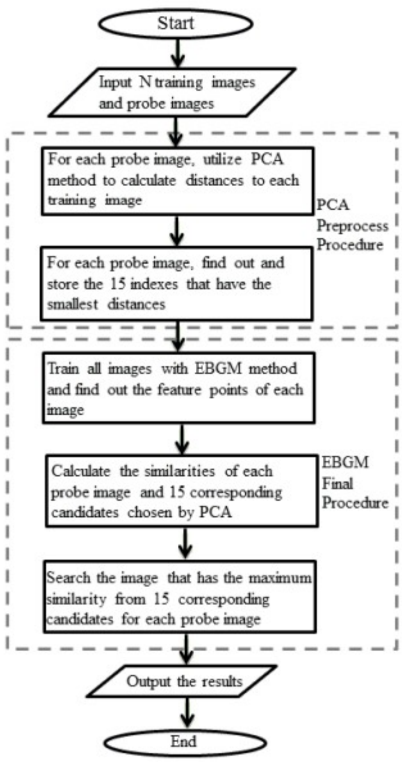

Figure 1: Work flow of optimizing EBGM with PCA. Source: Adapted from Xianming Chen et al., 2014 [13].

2.2.1 Core Concepts of PCA. PCA is a mathematical method that analyzes data by reducing the amount of information within the data and while retaining critical information. The core concepts of PCA include:

- **Dimensionality Reduction:** PCA aims to simplify data while retaining the essential information. It achieves this by projecting data onto a lower-dimensional subspace thereby changing how we see the data [5].
- **Eigenfaces:** In facial recognition, PCA often involves the concept of "eigenfaces," which are similar to special face patterns we get from a group of facial images. Eigenfaces represent the most significant variations in the dataset, making them valuable for recognition [8].
- **Feature Extraction:** PCA identifies the most relevant features in data by examining how variables within the data are related. These features are then used to describe and recognize faces [11].

2.2.2 Application of PCA in EBGM. In EBGM, PCA is applied to address challenges related to high-dimensional data and computational complexity. Chen et al. (2014) proposed a sequential hybrid approach in which PCA and EBGM are applied separately. In this model, PCA is used to process the training set and select a specific number of training images with the smallest distances between distinct features as candidates for EBGM. After this, for each test image, EBGM identifies the final match from the corresponding candidates. This approach yielded better recognition performance, particularly for larger galleries, and resulted in improved computational times [2]. The integration of PCA into the EBGM algorithm involves the following steps:

(1) **Preprocessing:** Before using EBGM, the pictures of faces are prepared. This includes standardization to a common size, noise reduction, and contrast enhancement using techniques like Median Filtering and Histogram Equalization [1, 5, 14].
(2) **PCA Feature Extraction:** PCA is used to extract the most important facial features from the processed facial images. These will be vital to the recognition process [5].
(3) **Hybridization with EBGM:** The extracted PCA features are then combined with the EBGM algorithm. This sequence optimizes EBGM's recognition performance and efficiency [2].
(4) **Recognition Process:** When a new facial image is then presented to the system, the PCA-enhanced EBGM compares it to the faces stored in the system and tries to find the best match for identification. [2].

By using PCA with EBGM, the work is made easier for the system, allowing it to work better with larger amounts of data and work against the problem of occlusion. This brings it a step closer to making it well-suited for real-world, immediate facial recognition applications.

For a deeper understanding of PCA and its application in facial recognition, readers are encouraged to explore relevant literature [5, 8, 11].

2.3 Racial Bias in Machine Learning

As explained previously, facial recognition technology has become a essential part of modern society, however, the rapid deployment and integration of this technology have outpaced considerations for its ethical implications, particularly concerning racial bias[12].

The algorithms powering facial recognition systems are not immune to the biases of the data on which they are trained. Studies have consistently found that these systems tend to perform worse for individuals from minority racial groups, especially African Americans[3]. This discrepancy arises partly because many datasets used to train these systems are disproportionately composed of images of Caucasian individuals, leading to less accuracy when recognizing faces from underrepresented groups[12]. Racial bias in machine learning, especially in facial recognition, stems from both the imbalance in training datasets and the methodologies used for developing algorithms. For example, algorithms developed in Western countries show a higher accuracy for Caucasian faces, whereas those developed in Asian countries tend to be more accurate for Asian faces. This suggests that the racial composition of the development team and the dataset can significantly influence the performance of facial recognition systems across different demographics[12]. Moreover, the challenge is not merely technical but also ethical. The over representation of minority groups in arrest databases, combined with the higher error rates of facial recognition systems for these groups, can lead to a higher incidence of mis-identification and wrongful arrests, perpetuating and exacerbating existing racial biases in law enforcement practices[3].

Recent advancements have been made in developing more equitable facial recognition technologies. One approach is adversarial training, where algorithms are trained to minimize the reliance on racial features that could lead to biased performance. This method involves modifying traditional training techniques to reduce the mutual information between facial features and sensitive attributes like race, thus aiming to make these attributes irrelevant for identity verification[12].

Another promising approach is the use of generative adversarial networks (GANs) to create balanced datasets. These networks can generate synthetic images of faces from underrepresented racial groups, thus enriching the training datasets and potentially reducing bias. By ensuring that facial recognition systems are exposed to a more diverse array of faces during training, the performance across different racial groups can be more uniform[12].

Despite these technological advances, significant challenges remain. The efficacy of adversarial and generative techniques in real-world applications is still under scrutiny, as these methods can sometimes fail to eliminate biases effectively. Moreover, the ethical implications of using synthetic images for training facial recognition systems are still being debated, particularly concerning privacy and the consent of individuals whose likenesses are used to generate these images[12].

As facial recognition technology continues to evolve, ongoing research and development efforts must prioritize not only the improvement of technical accuracy but also the fairness and equity of these systems. This will involve a concerted effort from researchers, developers, and policymakers to ensure that advancements in this technology do not come at the cost of perpetuating racial disparities.

3 EXPERIMENTATION PROCESS

3.1 Model

Before training, all images are collected into a single training folder after undergoing a standard preprocessing routine. The first step in this routine involves scanning the folder to identify the image with the smallest dimensions. This dimension serves as the baseline, and all other images in the training set are resized to match these dimensions. This, as well as the Median Filtering and Histogram Equalization done previously, ensures that all inputs to the model maintain uniformity, which is crucial for the subsequent steps of the feature extraction and matching processes.

Once preprocessing is completed, the base version of Elastic Bunch Graph Matching (EBGM) is applied. EBGM is a facial recognition technique that utilizes a graph-based representation of facial features. The initial phase of EBGM does not involve any dimensionality reduction; it focuses on identifying and using the raw pixel information and spatial relationships between facial landmarks directly extracted from the images.

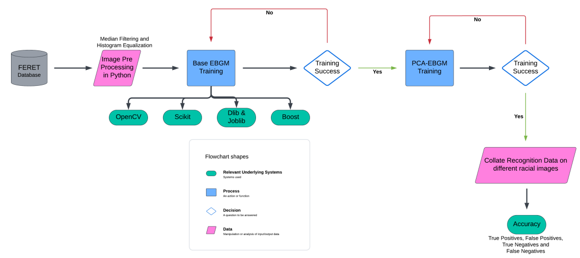

Figure 2: Flow of Software Architecture

After evaluating the base EBGM model, Principal Component Analysis (PCA) is integrated to enhance the system's performance. The integration begins with a dimension reduction procedure on the images. PCA simplifies the images by reducing the amount of information each one presents. This reduction is particularly beneficial for the landmark feature selection process, as it focuses the model's attention on the most significant features, thereby potentially improving recognition accuracy.

Following the dimension reduction, the model identifies facial landmarks in each image. These landmarks are used to create jet files, which represent local texture information around each landmark, and bunch graph files, which encapsulate the structural arrangement of these jets across the face. These files form the core of the model's facial representation.

In the final step, the model compares the jet and bunch graph files generated from the training images with those produced from the test images. For each test image, the model calculates similarity scores with all training images, which are presented as percentages. These scores indicate how closely features of the test image match those of each training image, thus facilitating the recognition process by quantifying similarity in a comprehensible manner.

3.2 Dataset

The photos used in this project will be obtained from the FERET database, a widely recognized source for facial image data with established evaluation methods[9]. The choice of the FERET database is motivated by several factors:

The FERET images are mostly frontal, meaning they're taken with the person's face aligned with the eyes[2, 9]. This consistency in image orientation is crucial for a comparative study of PCA and HSO optimizations.

To keep image processing consistent, the images will be resized to the dimensions of the smallest image. This way, there won't be a need to guess or interpolate new pixel information, which could blur the images.

Additionally, if necessary, the images will undergo standard image processing techniques, such as Median Filtering and Histogram Equalization[1, 14], to remove image noise and enhance contrast levels. Converting the images to grayscale simplifies their complexity and will make the process easier.

While the ultimate goal of facial recognition is to handle various image conditions, using these standardized techniques ensures a level playing field for our experiments.

The FERET database gives us a vast collection of facial images, making it an excellent choice for our comparative analysis of PCA and HSO optimizations within EBGM.

3.3 Racial Bias Experimentation

The primary goal of this experimentation process is to assess and identify any racial biases present in a facial recognition system. The experiment is designed to evaluate how well the system recognizes individuals from different racial groups when there are variations in the image. These variations include change in pose or expression. Due to the standardization and gray scaling of the images, the effect of change in lighting is no longer prevalent and will not be considered.

Specific images were selected from the dataset and organized into separate folders, each representing a different racial group. These are used to test the model against its much larger training data which contains other images of the same people. For each individual, several photographs were included to capture a range of expressions, angles. This variety ensures that the model learns to recognize the features of individuals from each racial group under varied circumstances.

Also, from the collection of images for each individual, one specific image was selected as the test image. This image was set aside and not included in the training dataset. The remaining images of the same individual were placed in the training folder. This approach ensures that while the model has prior knowledge of the individual's facial features from different images, it must rely on its learned predictive intelligence to recognize the test image, thereby simulating a more realistic and challenging scenario for the facial recognition system.

The Testing protocol for the model can be seen as this:

(1) **Model Training:** The facial recognition system was trained separately on the images from each racial group's folder. The training process involved adjusting the system to recognize and learn the distinct features and variations present in the images of individuals from each group.
(2) **Individual Recognition Test:** After training, the system was tested using the single image reserved for testing for each individual. This test image, as previously noted, was not included in the training set, requiring the system to generalize its recognition capabilities rather than simply matching the image to an existing one in the database.
(3) **Bias Assessment:** The accuracy of the facial recognition system in identifying the test image was recorded and compared across different racial groups. The metrics used for assessing the accuracy of results are explained below:
  - **True Positives (TP):** True Positives occur when the optimized EBGM techniques correctly identify a match. In facial recognition, TP represents successful identification when there's indeed a match [10].
  - **True Negatives (TN):** True Negatives occur when the optimized EBGM techniques correctly identify when a match is non-existent. TN represents successful identification when there is no match [10].
  - **False Positives (FP):** False Positives happen when the optimized EBGM technique claims to recognize a match that does not exist. These instances are crucial as they highlight recognition errors [10].
  - **False Negatives (FN):** False Negatives occur when the optimized EBGM technique fails to recognize a match that does exist [10].

Recognition accuracy in facial recognition is defined as the ratio of True predictions to the total number of predictions[10]. The calculation for recognition accuracy is as follows:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

4 RESULTS

4.1 Overview

The execution of the facial recognition model across different racial groups yielded insightful findings regarding its accuracy and potential biases. Here, we detail the performance of the model with respect to the different racial categories used in the study: Black, Caucasian, Asian, and Hispanic. The dataset predominantly comprised Caucasian faces, which influenced the training phase and, subsequently, the model's performance across groups.

The data presented in this section are averages derived from multiple test instances across different racial groups. For each racial group, the facial recognition system was tested using a set of images, and performance metrics were recorded. The metrics—True Positives (TP), False Positives (FP), True Negatives (TN), and False Negatives (FN)—are averaged from all tests conducted. Specifically, if the system is run on 20 images of a particular group, the total number of each metric (e.g., TP) is divided by the number of images (20 in this case) to obtain the average number of true positives per image for that group.

Similarly, the frequencies of similarity scores are averaged. For instance, the total frequency of scores falling within the 0-10% similarity range for all tests with all images is divided by the number of tests (20) to calculate the average frequency for that similarity range. This approach ensures that the results are normalized, providing a clear and consistent basis for comparison across different groups and tests.

4.2 Accuracy Measurement and Distribution of Similarity Scores

The distribution of similarity scores was analyzed for four racial groups: Black, Asian, Caucasian, and Hispanic individuals. The frequency of scores across various percentage bins is detailed below for each group.

4.2.1 *Black Individuals.*

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} = \frac{8 + 5936}{8 + 5936 + 999 + 36} \approx 87.45\% \tag{1}$$

4.2.2 *Asian Individuals.*

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} = \frac{5 + 5961}{5 + 5961 + 996 + 17} \approx 87.77\% \tag{2}$$

4.2.3 *Caucasian Individuals.*

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} = \frac{12 + 6122}{12 + 6122 + 653 + 10} \approx 90.25\% \tag{3}$$

4.2.4 Hispanic Individuals.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} = \frac{15 + 5932}{15 + 5932 + 821 + 29} \approx 85.21\% \tag{4}$$

Similarity Score Range (%)	Average Frequency
0-10	2204
10-20	2230
20-30	1129
30-40	241
40-50	161
50-60	186
60-70	197
70-80	139
80-90	113
90-100	202

Table 1: Average Frequency of similarity scores in facial recognition tests for Black individuals.

Similarity Score Range (%)	Average Frequency
0-10	0
10-20	2203
20-30	2047
30-40	1297
40-50	249
50-60	164
60-70	184
70-80	200
80-90	141
90-100	110

Table 2: Average Frequency of similarity scores in facial recognition tests for Asian individuals.

Similarity Score Range (%)	Average Frequency
0-10	0
10-20	2204
20-30	2230
30-40	1129
40-50	241
50-60	161
60-70	186
70-80	197
80-90	139
90-100	197

Table 3: Average Frequency of similarity scores in facial recognition tests for Caucasian individuals.

Similarity Score Range (%)	Average Frequency
0-10	0
10-20	2203
20-30	2110
30-40	1239
40-50	248
50-60	161
60-70	185
70-80	201
80-90	140
90-100	202

Table 4: Average Frequency of similarity scores in facial recognition tests for Hispanic individuals.

Each racial group exhibits unique characteristics in terms of similarity score distribution and accuracy, suggesting different levels of system performance and potential biases that may affect the reliability and fairness of the facial recognition technology.

4.3 Discussion and Future Recommendations

The analysis presented in the results section reveals notable variations in the performance of the facial recognition system across different racial groups. These disparities highlight critical challenges in ensuring fairness and equity in machine learning applications, especially in sensitive areas such as facial recognition.

4.3.1 Interpretation of Results.

Accuracy Discrepancies. The highest accuracy was observed for Caucasian individuals at approximately 90.25%, while the lowest was for Hispanic individuals at about 85.21%. Such discrepancies can stem from several factors including but not limited to the diversity and representation within the training datasets. The predominant presence of Caucasian faces in the training data likely contributed to the higher accuracy rates for this group, as the model was better attuned to their facial features compared to those of other racial groups.

Similarity Score Distribution. The distribution of similarity scores also underscores potential biases in the model's decision-making process. For instance, the higher frequencies of lower similarity scores among Black, Asian, and Hispanic groups suggest that the model is less confident or accurate in recognizing these faces compared to Caucasian faces. This could be due to a variety of factors such as under representation in training data, inadequate feature extraction techniques that fail to capture the diversity of facial features across races, or even algorithmic predispositions that favor certain facial metrics.

4.3.2 Implications for Bias in Machine Learning. These results clearly illustrate the issue of bias in machine learning algorithms, particularly in facial recognition technologies. Bias can manifest in various forms, often as a direct consequence of non-representative training data or flawed algorithm design. When a model is trained predominantly on data from one racial group, it becomes less effective at identifying and processing faces from other racial groups, leading to higher error rates and potentially discriminatory practices.

Ethical and Social Considerations. The ethical implications of these findings are significant. As facial recognition technologies become more integrated into various aspects of daily life, from security systems to hiring processes, ensuring that these systems do not propagate or exacerbate existing societal biases is crucial. The risk of misidentification or wrongful exclusion/inclusion poses serious concerns that could impact individuals' rights and freedoms.

4.3.3 Recommendations for Future Work. To mitigate these biases, several steps can be recommended:

- **Diverse Datasets:** Expand and diversify the training datasets to include equal representation of all racial groups. This approach helps in developing a more balanced model that can accurately identify and process faces from a wide range of racial backgrounds.
- **Algorithmic Audits:** Regular audits and updates of the algorithms to check for and correct biases that may occur over time or as a result of changes in data patterns.
- **Transparency and Accountability:** Maintain transparency about the data and methods used in training models, and implement accountability measures for failures that lead to discriminatory outcomes.

5 CONCLUSION

The primary objective of this project has been to explore racial bias in machine learning models, with a specific focus on facial recognition systems. This focus is particularly relevant given the level of integration of facial recognition technology into various aspects of daily life, from security to personal device access. As these systems become more widespread, it is important to ensure they operate fairly and accurately across diverse populations becomes increasingly critical.

To address this, we have adopted one of the more reputable facial recognition techniques, Elastic Bunch Graph Matching (EBGM), and enhanced it through hybridization with Principal Component Analysis (PCA). EBGM is known for its robust approach to facial recognition, leveraging a graph-based representation of facial features, while PCA is employed to reduce the dimensionality of the data, focusing the EBGM on the most salient features which are hoped to reduce bias by simplifying the model's inputs.

The experimentation process was designed to systematically evaluate the model across different racial groups, using a dataset that includes diverse racial representations to ascertain any biases in recognition accuracy. The results of these experiments have provided insightful revelations into the performance disparities across racial lines. The enhanced EBGM system demonstrated variable accuracy rates, with different racial groups experiencing different levels of recognition precision. These outcomes highlight the ongoing challenges and the necessity for continual improvement in the methodologies used in facial recognition technologies.

The results from our experiments reveal significant disparities in accuracy among the different racial groups tested. Specifically, the enhanced EBGM system equipped with PCA showed an accuracy of 87.45% for Black individuals, 87.77% for Asian individuals, 90.25% for Caucasian individuals, and 85.21% for Hispanic individuals. These variations in accuracy rates underscore the presence of inherent biases within the facial recognition system, with the highest accuracy observed for Caucasian faces—a likely consequence of the training dataset's composition which predominantly comprised Caucasian faces.

The distribution of similarity scores further illustrates these biases, with non-Caucasian groups generally receiving lower scores. This indicates a systematic issue where the model, despite the integration of PCA to prioritize salient features, still fails to accurately and consistently recognize individuals from racially diverse backgrounds.

These findings highlight a critical need for the development of facial recognition technologies that are both fair and effective across all demographic groups. It is clear that without intentional and rigorous efforts to correct these biases, facial recognition technologies could perpetuate existing racial prejudices and have a detrimental impact on individuals from underrepresented groups.

Recommendations for Improvement. To mitigate these issues, we recommend the following approaches:

- **Enhanced Dataset Diversity:** Augment training datasets with a more balanced representation of all racial groups to ensure the model learns to recognize a wider array of facial features.
- **Bias Mitigation Algorithms:** Develop and implement algorithms specifically designed to detect and reduce bias in the training data and model predictions.
- **Continuous Model Evaluation:** Regularly re-evaluate and update the model with new data, and continuously test the model's performance across different demographics to ensure improvements are sustained over time.
- **Regulatory Oversight:** Establish standards and guidelines for the ethical use of facial recognition technology, ensuring that these tools undergo rigorous bias assessment before deployment.

In conclusion, while the integration of PCA with EBGM has shown potential in enhancing the recognition accuracy of facial recognition systems, our study confirms that significant work remains to eliminate racial biases inherent in these technologies. Ensuring fairness in facial recognition is not only a technical challenge but also a moral imperative that requires concerted efforts from researchers, practitioners, and policymakers. As we advance, it is crucial that we continue to push for technologies that uphold the highest standards of equity and justice, fostering trust and inclusivity in our increasingly digitized world.

REFERENCES

[1] Sapana S. Bagade and Vijaya K. Shandilya. Use of histogram equalization in image processing for image enhancement. 2011.

[2] Xianming Chen, Chaoyang Zhang, and Zhaoxian Zhou. Improve recognition performance by hybridizing principal component analysis (pca) and elastic bunch graph matching (ebgm). 2014.

[3] Clare Garvie and Jonathan Frankle. Facial-recognition software might have a racial bias problem. 2021.

[4] Zhi-Qiang Liu Jialiang Liu. Ebgm with fuzzy fusion on face recognition. 2005.

[5] Ian T. Jolliffe. Principal component analysis. 2003.

[6] Norbert Krüger Christoph von der Malsburg Laurenz Wiskott, Jean-Marc Fellous. Face recognition by elastic bunch graph matching. 1997.

[7] Chao-Chun Liu and Dao-Qing Dai. Face recognition using dual tree complex wavelet features. 2009.

[8] Alex Pentland Matthew Turk. Eigenfaces for recognition. 1991.

[9] Rizvi S.A. Rauss P.J. Phillips P.J., Moon H. The feret evaluation methodology for face-recognition algorithms. ieee transactions on pattern analysis and machine intelligence. 2000.

[10] Inioluwa Deborah Raji and Genevieve Fried. About face: A survey of facial recognition evaluation. 2021.

[11] Mansoor Roomi and Dr.M.Parisa Beham. A review of face recognition methods. *International Journal of Pattern Recognition and Artificial Intelligence*, 27, 2013.

[12] Noura Al-Moubayed Toby P.Breckon Seyma Yucer, Samet Akcay. Exploring racial bias within face recognition via per-subject adversarially-enabled data augmentation. 2020.

[13] Phillips P.J. Zhao W., Chellappa R. and Rosenfeld A. Face recognition: A literature survey. 2003.

[14] Youlian Zhu and Cheng Huang. An improved median filtering algorithm for image noise reduction. 2012.