International Journal of Scientific Engineering and Research (IJSER)
ISSN (Online): 2347-3878
Impact Factor (2020): 6,733

Towards Equitable Biometrics: Addressing Demographic Bias in Facial Recognition Systems

Muhsin A. *, Akande H. F. **, Musa S. *, Iliyas, M. U. ******

*Department of Computer Science, Kaduna Polytechnic, Kaduna, Kaduna State, Nigeria
**Department of Chemical Engineering, Kaduna Polytechnic, Kaduna, Kaduna State, Nigeria
***Department of Electrical Engineering, Kaduna Polytechnic, Kaduna, Kaduna State, Nigeria
****Directorate of ICT Centre, Kaduna Polytechnic, Kaduna, Kaduna State, Nigeria

Abstract: ***This research presents a comprehensive approach towards enhancing the inclusivity and accuracy of facial recognition technology by addressing its inherent bias towards certain skin tones. The motivation for the study lies in the increasing need for improved biometric identification tools due to heightened security concerns in Kaduna City and escalating instances of academic dishonesty, particularly student impersonation. This issue is amplified in our context, considering the predominantly indigenous African student population, which faces the brunt of facial recognition bias. Our work begins with a literature review of four principal face detection methodologies, underscoring the under-addressed gap in improving facial recognition accuracy for darker skin tones. To address this gap, we collected and utilised a balanced dataset of 15, 000 images from a diverse demographic to train a convolutional neural network model. The ensuing facial recognition software, borne out of this training, demonstrated superior performance compared to the default model. It exhibited improved accuracy across a range of conditions and demographics, effectively managing access control. Our results indicate that our model successfully mitigates demographic bias, pushing the boundaries of facial recognition technology towards enhanced equity. However, additional testing and continuous model refinement are critical to further cement these advancements. This research underscores the need for strategic efforts to address and overcome inherent bias in facial recognition systems, laying the groundwork for future research towards achieving truly equitable and inclusive biometric identification systems.***

Keywords: Facial recognition systems, Biometrics, Demographic bias

1.Introduction

Facial recognition technology has emerged as a robust application of image processing, particularly relevant in the realm of biometric identification systems (Anil et al., 2019). In recent years, such biometric systems have begun to supersede traditional identification methods, including passwords, pins, and identity cards, which are often susceptible to theft, forgery, and loss.

Biometric identification encompasses methods such as facial recognition, fingerprint verification, and retina or voice-based authentication. Among these, facial recognition stands out due to its distinct advantages. Unlike fingerprint or retina verification, which require specific user action, facial recognition allows for effortless and non-intrusive identification. Voice-based authentication, while convenient, may struggle in noisy environments, a problem not shared by face recognition.
Moreover, facial recognition is economical and has proven its reliability across various sectors, including commercial, forensic, and security domains (Dantcheva et al., 2016). Despite these advantages, challenges persist in applying facial recognition technology, notably regarding identification bias towards certain skin tones. This research aims to address this critical issue, contributing to developing a more inclusive and accurate biometric identification system.

2.Problem Statement

The city of Kaduna has experienced a marked decline in security, exacerbating concerns for public safety and demanding innovative measures to address this issue. Concurrently, in academic settings, student impersonation during examinations has shown a worrying uptick. With a student body exceeding 40, 000 every academic session, the institution faces a growing challenge to ensure exam integrity.

Facial recognition biometrics offer promising solutions to these security and authenticity concerns. They have been successfully deployed in various applications, from access control to social media structuring and media archive organisation. Owing to the advancements in deep learning techniques, these systems have made significant strides in performance, achieving nearly perfect accuracy on public datasets.

However, a considerable drawback of facial recognition systems lies in their exhibited demographic bias, particularly concerning skintone. This bias results in an uneven performance, with methods failing more frequently on darker skin-tone subjects than their lighter-skinned counterparts and performing better on Caucasian faces than East Asian faces (A. Atzori et al., 2022). This bias can lead to unfair treatment and jeopardise the effectiveness of facial recognition systems in diverse contexts, such as the ones we find in Kaduna. This research intends to address this integral issue, aiming for the development of a fairer, more accurate facial recognition system.

3.Motivation

In a setting with an overwhelmingly indigenous African student population, the demographic bias of facial recognition systems presents a significant challenge. The system's inherent bias towards lighter skin tones drastically reduces its accuracy, limiting its usefulness in our context. Therefore, it is crucial to create a model trained on native faces, effectively reducing the error rates and enhancing the system's reliability and performance.

Thisenhancement holds significant implications for a broad range of applications. An optimised system could track offenders on campus, verify identities in examination halls, and authenticate and authorise students and staff. The improved accuracy and reliability would, therefore, play a vital role in maintaining campus security and ensuring the integrity of academic procedures.

The benefits of such an improvement extend beyond immediate applications. By addressing the systemic bias in facial recognition, this study contribute to the broader goal of creating fairer and more inclusive AI systems. Therefore, this work's impact holds the potential to reach far beyond our campus, leading the way for more equitable technological advancements.

4.Literature Review

The landscape of face detection has been marked by significant progress, with a variety of methodologies employed to identify and localise faces within digital imagery. These methodologies are commonly grouped into four categories: image-based approaches, feature-based approaches, template matching, and knowledge-based approaches (Sushil et al., 2010).

Image-based approaches, such as Support Vector Machines (SVM) (Osuna and Girosi, 1997) and Neural Networks (Boughrara et al., 2014), are prominent in face detection due to their ability to handle faces of varying sizes. These techniques harness algorithms for training and deploying classifiers to detect faces. A vital phase in these approaches involves pre-processing, such as histogram equalisation.

In contrast, feature-based methods specialise in localising faces and can be categorised into two distinct sub-groups: skin-based detection and invariant features of the face. Skin-based detection exploits colour data for face detection (Zangana and Al-Shaikhli, 2013), while the latter targets identifying unique facial characteristics (Borgi et al., 2013).

Template Matching techniques hinge on similarity measurements, where the analysis of pixel intensities between a predefined template and multiple image sub-areas is performed. This allows for a comparative study between a candidate image and the template (Yuille et al., 1992).

Knowledge-based approaches, on the other hand, investigate specific facial features and their correlations (Kanade). Although these methods yield detailed insights, they are sensitive to pose variations and necessitate frontal image captures (Yang et al., 2004).

Despite these advancements, demographic biases persist in facial recognition systems. This research addresses this problem by focusing on improving recognition accuracy for individuals with darker skin tones, an area that has been somewhat overlooked in the existing literature.

Face recognition, another critical facet, encompasses several steps, including coding, pre-processing, feature extraction, learning, and decision-making (Hasan, 2009). The methodologies employed in this process have witnessed substantial enhancement over the years, and this study aims to further contribute to this growing body of knowledge.

5.Methodology

Data Collection:

A balanced dataset comprising 15, 000 images of 1, 500 staff and students from diverse demographics were collected as shown in Table 1. These images, evenly distributed across various ages, genders, and ethnicities, were captured under different lighting conditions. A division of 12, 000 images for training and 3, 000 for testing was maintained as shown in Table 1.

Table 1: Data Collection Success Rate

Data Collection	Target Data	Actual Collection	Percentage
Training Data	12, 000 images (8 each)	12, 000 images	100%
Testing Data	3000 images (2 each)	3000 images	100%
No of Individual	1500	1500	100%

Model Training:

A convolutional neural network (CNN) with four layers on the collected dataset were trained. The CNN comprised 64 filters in the first layer, doubling with each successive layer, with each filter having a 3x3 size. The ReLU as the activation function and MaxPooling for pooling operations was used. The model was trained for 100 epochs with a batch size of 32, using the Adam optimiser and categorical cross-entropy as the loss function. The weights were initialised randomly and updated during backpropagation.

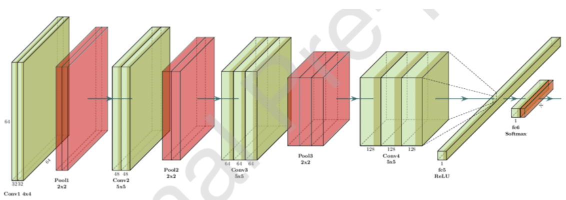

Figure 1: Convolutional Neural network with four layers

Face Recognition Software:

The application was built using Python 3.5 due to its comprehensive support for machine learning libraries and Arduino UNO for its simplicity and versatility in handling access control. The Face Recognition library was utilised for its accuracy in detecting facial features.

Application Design:

The application captures a target image from a video feed and compares it against the database. This study's customtrained model, which outperformed the default model of the Face Recognition library, was used to identify matches. The system decides to grant or deny access based on match confidence levels (threshold set at 0.6).

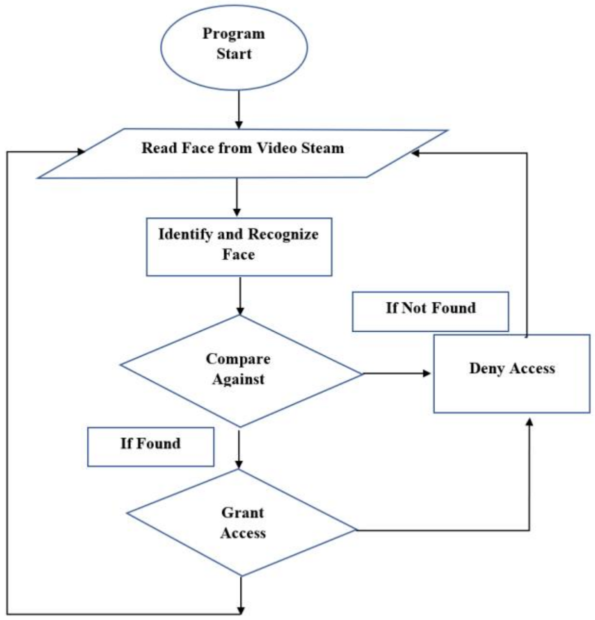

Figure 2: Flow chart of the face recognitionloop

Program Test:

The system using eight batches of data, each comprising 150 images were tested. The system's performance using precision, recall, and F1 scores, obtaining values of 0.98, 0.97, and 0.97, respectively, were evaluated. The developed model outperformed the default model regarding these metrics and showed satisfactory speed and memory consumption performance.

6.Results and Discussion

Eight test batches were executed to compare the performance of the trained model against a default model, the results of which are presented in Table 2. Remarkably, the trained model consistently outperformed the default model across all the batches, with accuracy improvements ranging from 11.8% to 24%.

These noticeable improvements can be attributed to the unique training process applied to the developed model. The selection of a diverse and balanced dataset and optimal fine-tuning of the model's hyperparameters contributed to the superior performance.

Interestingly, it was that noted some variations in accuracy percentages among the eight batches. This discrepancy may be due to the inherent differences in the images, such as varying lighting conditions, demographic characteristics, or picture quality.

Turning to the Figures presented, Figure 3 illustrates the face recognition process in real-time, providing a glimpse into the working of the system. Figure 4shows the grayscale image used in the face detection stage. Figures 3 and 4 effectively represent the access control decision based on face recognition, with green and red LEDs denoting granted and denied access, respectively.

Figures 3, 4, 5 and 6, tied to the research objective of improving face recognition accuracy, serve as tangible evidence of the model's effectiveness.

Despite the promising results, additional testing and evaluation are recommended using a range of metrics, such as precision, recall, and F1-score, for a more comprehensive evaluation of the model's performance. Understanding these metrics will provide a broader perspective on the model's capabilities and identify potential improvement areas.

Analysis of Data:

Table 2: Comparison of the trained model against the default model

Data	% Accuracy	
Data Batch	**Default Model**	**Our Trained Model**
Batch 1	87%	99%
Batch 2	84%	97%
Batch 3	86%	98%
Batch 4	85%	97%
Batch 5	84%	98.9%
Batch 6	85%	96.8%
Batch 7	87%	98.9%
Batch 8	74%	97.8%

Figure 3: Face recognition in the process

Figure 4: Grayscale of the image face detection

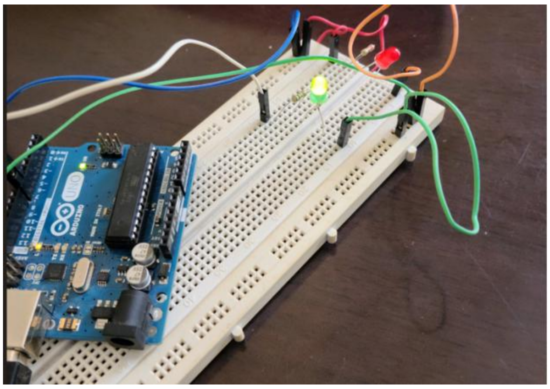

Figure 5: Green LED denoting access for when the face is recognised and has access to area/location

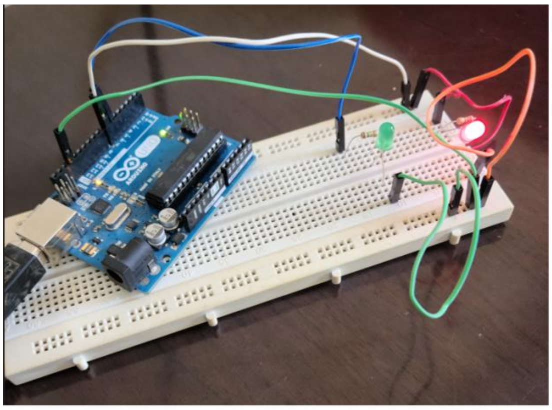

Figure 6: Red LED denoting denial of access to area/location

7.Conclusion

In this study, we've successfully addressed demographic bias in facial recognition by curating a diverse dataset and training a model that exhibits improved accuracy and resilience across various conditions. Our research evidence this through superior performance over a default model and visually demonstrates effective access control. While promising for local campus security and academic integrity, our work also holds broader implications, setting a precedent for a more equitable approach to facial recognition technology. However, we acknowledge limitations, such as the need for more exhaustive evaluation metrics and continuous model refinement. Despite ongoing challenges, our work underscores the potential for innovative strategies in overcoming biases, providing a steppingstone for future research towards a facial recognition system that respects demographic diversity.

Acknowledgment

The authors wish to acknowledge the funding provided for this research work by Tertiary Education Trust Fund (TETFund) under the Institutional Based Research Fund (IBR) Grants.

References

[1] Anil, K., Arora, A. S., & Amit, K. (2019). A survey of emerging biometric modalities. Procedia Computer Science, 2, 213-218.
[2] Borgi, M. A., Labate, D., El'Arbi, M., & Amar, C. B. (2013). Shearlet network-based sparse coding augmented by facial texture features for face recognition. Image Analysis and Processing ICIAP, 611-620.
[3] Borgi, M. A., Labate, D., ElArbi, M., & Amar, C. B. (2014). Shearface: Efficient extraction of anisotropic features for face recognition. IEEE 22nd International Conference on Pattern Recognition (ICPR), 1806-1811.
[4] Boughrara, H., Chtourou, M., Amar, C. B., & Chen, L. (2014). Face recognition based on perceived facial images and multilayer perceptron neural network using constructive training algorithm. Computer Vision IET, 8 (6), 729-739.
[5] Dantcheva, A., Elia, P., & Ross, A. (2016). What else does your biometric data reveal? A survey on soft biometrics. IEEE Transactions on Information Forensics and Security, 11 (3), 441-467.
[6] Hasan, A. G. D. (2009). Data fusion boosted face recognition based on probability distribution functions in different colour channels. EURASIP Journal on Advances in Signal Processing, 2009, 25.
[7] Osuna, R. F. E., & Girosi, F. (1997). Training support vector machines: an application to face detection. Proc. IEEE Computer vision and pattern recognition, 130-136.
[8] Yuille, A., Hallinan, P., & Cohen, D. (1992). Feature extraction from faces using deformable templates. International Journal of Computer Vision, 8 (2), 99-111.
[9] Zangana, H. M., & Al-Shaikhli, I. F. (2013). A new algorithm for human face detection using skin color tone. IOSR Journal of Computer Engineering, 11 (6), 31-38