Fairness-Aware Face Presentation Attack Detection Using Local Binary Patterns: Bridging Skin Tone Bias in Biometric Systems

Jema David Ndibwile *, Ntung Ngela Landon, Floride Tuyisenge

Posted Date: 9 September 2025

doi: 10.20944/preprints202509.0708.v1

Keywords: presentation attack detection; local binary patterns; fairness; skin tone; equalized odds

 doi:10.20944/preprints202509.0708.v1

Article

Fairness-Aware Face Presentation Attack Detection Using Local Binary Patterns: Bridging Skin Tone Bias in Biometric Systems

Jema David Ndibwile *, Ntung Ngela Landon and Floride Tuyisenge

Carnegie Mellon University Africa; ftuyisen@andrew.cmu.edu
* Correspondence: jndibwil@andrew.cmu.edu

Abstract

While face recognition systems are increasingly deployed in critical domains, they remain vulnerable to presentation attacks and exhibit significant demographic bias, particularly affecting African populations. This paper presents a fairness-aware Presentation Attack Detection (PAD) system using Local Binary Patterns (LBP) with novel ethnicity-aware processing techniques specifically designed for African contexts. Our approach introduces three key technical innovations: (1) adaptive preprocessing with differentiated CLAHE parameters and gamma correction optimized for different skin tones, (2) group-specific decision threshold optimization using Equal Error Rate minimization for each ethnic group, and (3) three novel statistical methods for PAD fairness evaluation such as Coefficient of Variation analysis, McNemar's significance testing, and Bootstrap confidence intervals representing the first application of these techniques in presentation attack detection. Comprehensive evaluation on the `CASIA-SURF CeFA` dataset demonstrates significant bias reduction achievements: 75.6% reduction in accuracy gap between African and East Asian subjects (from 3.07% to 0.75%), elimination of statistically significant bias across all ethnic group comparisons, and strong overall performance with 95.12% accuracy and 98.55% AUC. Our work establishes a comprehensive methodology for measuring and mitigating demographic bias in PAD systems while maintaining security effectiveness, contributing both technical innovations and statistical frameworks for inclusive biometric security research.

Keywords: presentation attack detection; local binary patterns; fairness; skin tone; equalized odds

1. Introduction

Face recognition systems have become important parts of modern security systems. They are now widely used in banks, hospitals, border control, and national ID systems around the world [4]. Face recognition is popular because people do not need to touch anything, it is easy to set up, and it is convenient to use [2]. However, as these systems become more common, they face serious security threats. This is why Presentation Attack Detection (PAD) has become very important. PAD systems help protect against fake attacks using printed photos, video replays, and 3D masks [1,6].

Even though PAD technology has improved a lot, there is still a big problem with fairness, especially for African population. Studies show that current face recognition and PAD systems work much worse for different ethnic groups, and African users often get higher error rates [15]. This happens mainly because African people are severely underrepresented in existing public datasets used for both face recognition and PAD systems. Studies show major datasets contain only 13.8% – 20.4% darker-skinned tones [32], while PAD systems similarly exhibit significant demographic bias with higher error rates for African populations [33].

PAD systems face several technical problems when working with African population. Darker skin reflects less light than lighter skin, making it harder for regular cameras to capture enough detail for accurate attack detection [14,19]. Also, most PAD research is done in controlled labs using expensive

 doi:10.20944/preprints202509.0708.v1

equipment. This creates a big gap when these systems are used in African countries, where people usually have basic smartphones and face difficult lighting conditions [3]. This problem not only makes security weaker but also increases digital inequality, going against the goal of making biometric systems fair for everyone.

The CASIA-SURF Cross-ethnicity Face Anti-spoofing (CeFA) dataset represents the most significant step toward addressing these limitations. With 1,607 subjects spanning three ethnic groups, including 500 African participants (31% of the total dataset), CeFA provides the most comprehensive cross-ethnic PAD evaluation framework currently available [12]. While the dataset treats African populations as a single demographic category and does not capture the full regional diversity across the continent, it still offers the largest and most diverse collection of African faces available for PAD research. This makes it the best available resource for conducting systematic fairness evaluations across ethnic groups, especially in the African context.

Local Binary Patterns (LBP) have become a strong method for detecting presentation attacks. They work well at finding texture differences and are good at telling real faces apart from fake ones [7,31]. LBP methods can detect printing problems and surface issues that are common in many fake attacks, and they do not need a lot of computing power.

This study addresses the important need for fair PAD systems that are specifically designed for African contexts. We focus on developing and testing LBP-based methods that work equally well for different African populations while still providing strong security against various spoofing attacks. Our approach recognizes that making biometric security truly fair requires more than just adding more diverse faces to datasets. We also need to carefully consider the unique challenges that come with different skin tones, facial features, and environments that are common in African settings.

The main contributions of our research are: (1) a complete fairness evaluation of LBP-based PAD methods using the CeFA dataset with special focus on African groups, (2) analysis of performance differences across ethnic groups to measure bias in current methods, (3) development of LBP techniques that work better for African facial features and environmental conditions, and (4) creation of evaluation methods that prioritize fairness alongside security. Through this African-focused approach, we aim to help develop truly inclusive biometric security systems that provide reliable protection for all users, no matter their ethnic background or where they live.

The remainder of the paper is organized as follows. In Section 2 we discuss related literature on presentation attack detection (PAD), fairness-aware biometric systems, and Local Binary Pattern (LBP)-based approaches. In Section 3 we describe the proposed fairness-aware PAD framework including ethnicity-aware pre-processing, multi-scale LBP extraction, adaptive thresholding, and statistical methods for fairness evaluation. In Section 4 we present the experimental design and results which include dataset pre-processing, system calibration, metrics for PAD performance, fairness evaluations, and statistical validations across demographic groups and attacks. In Section 5 we discuss relevant implications for inclusive biometric safety and security based on our findings. Finally, in Section 6 we summarize the contributions of this paper, our limitations and our future work.

2. Related Works

In this section, we review the related work and provide a comprehensive summary in three tables: Table 1 outlines the depth of fairness evaluation in previous studies, Table 2 summarizes the general methods employed, and Table 3 compares the specific PAD methodologies from related work with our planned approach.

 doi:10.20944/preprints202509.0708.v1

Table 1. *Depth of fairness evaluation in prior PAD work. "African presence" indicates whether African subjects are explicitly included; "Skin tone" denotes any measurement beyond broad ethnicity.*

Work	Fairness evaluated	African presence	Skin tone	Intra-African	Capture realism
Wen et al. (2015)	No	No	No	No	Lab/Controlled
George & Marcel (2020)	No	No	No	No	Lab/Controlled
CeFA (Liu et al., 2021)	Limited	Yes (∼31%)	No	No	Controlled (Intel RealSense)
SiW / CASIA-SURF (various)	Limited/implicit	Mixed/unclear	No	No	Mostly controlled
Fang et al. (2022/2023)	Gender-based	Not Africa-centric	No	No	Mixed; not mobile-focused
Kotwal & Marcel (2025)	Survey	N/A	N/A	N/A	N/A
This work	**Yes (rigorous)**	**Yes (explicit)**	**Yes**	**Yes**	**Mobile/low-light**

Table 2. *Summary of related works and the methods used.*

Prior Work	Method Type
Wen et al. (2015)	Handcrafted (Image Quality + LBP)
Rahimzadeh & Kittler (2015)	Handcrafted (MBSIF + LPQ + SR-KDA)
Shaker & Al-Darraji (2024)	Deep Learning (CNN + ResNet)
George & Marcel (2020)	Deep Learning (Multi-modal: RGB+Depth+IR)
Spinoulas et al. (2021)	Multispectral (SWIR, NIR, LSCI)
Gomez et al. (2023)	Physiological (rPPG)
Ning et al. (2018)	Deep Learning (Stacked Generalization)
Liu et al. (2021)	Dataset/Benchmark (CeFA; CNN baselines)
Kotwal & Marcel (2025)	Demographic Fairness Review
Vurity & Marasco (2023)	Demographic-aware Training
Fang et al. (2023)	Deep + Fairness metrics (FairSWAP, ABF)
This work	**LBP + SGD; African-centric fairness protocol**

Table 3. *Summary of PAD methodologies in related work compared to our planned approach. "Multi-modal" indicates RGB combined with depth/NIR/IR.*

Work	Dataset(s)	Methodology	African fairness focus?
Wen et al. (2015) [8]	CASIA-FASD, Replay-Attack	Handcrafted (Image Distortion + LBP)	No
Rahimzadeh & Kittler (2015) [23]	CASIA-FASD	MBSIF + LPQ + SR-KDA	No
George & Marcel (2020) [24]	CASIA-SURF, Replay-Attack	Deep CNN (Multi-modal PAD)	No
Liu et al. (2021) [12]	CASIA-SURF CeFA	Deep CNN (ResNet, RGB+Depth+IR)	Limited (31% African, coarse)
Fang et al. (2022/2023) [37] [19]	OULU-NPU, SiW, WMCA	CNN + Fairness metrics (gender/attributes)	No (not Africa-centric)
Spinoulas et al. (2019) [9]	Multispectral custom dataset	Multi-modal (RGB, NIR, SWIR, LSCI)	No
Our work (planned)	**CASIA-SURF CeFA (RGB)**	**LBP + SGD, explicit fairness protocol**	**Yes (African subgroups, skin-tone)**

2.1. Traditional PAD Techniques

Presentation Attack Detection (PAD) methods are typically categorized into motion-based, texture-based, image-quality-based, and hardware-assisted approaches [1]. Texture-based methods, in particular, have shown early promise. For example, multi-scale Local Binary Patterns (LBP) have been used to detect surface anomalies caused by printed photos or masks [7]. Wen et al. [8] proposed image distortion analysis for PAD, while Rahimzadeh and Kittler [23] combined Multiscale Binarized Statistical Image Features (MBSIF) with Local Phase Quantization (LPQ) for enhanced spoof detection. While effective in controlled settings, these handcrafted methods often generalize poorly across different attack types and demographic profiles [6].

2.2. Deep Learning Approaches

Modern PAD systems increasingly leverage deep learning for automatic feature extraction. Shaker and Al-Darraji [10] employed a ResNet-enhanced CNN for robust face anti-spoofing, outperforming handcrafted methods. George and Marcel [24] proposed multi-channel architectures integrating RGB, depth, and infrared data directly into face detection pipelines. These deep networks, along with anomaly detection techniques, have improved cross-attack generalization. However, their reliance on computational resources and diverse training data may limit applicability in low-resource African contexts.

2.3. Multi-Modal and Advanced Techniques

PAD research has also explored physiological signals and multispectral imaging. Spinoulas et al. [9] developed a multispectral biometric system combining visible, NIR, and SWIR data for robust spoof detection. Gomez et al. [11] applied remote photoplethysmography (rPPG) to detect subtle pulse

 doi:10.20944/preprints202509.0708.v1

patterns unreplicable in spoofing. These methods, though effective, require specialized hardware and are rarely evaluated for fairness across demographics.

2.4. Fairness and Demographic Bias in Deep Learning PAD Models

Recent studies have raised concerns about demographic bias in PAD systems [32], analyzed major face datasets and found African subjects make up only 13.8% – 20.4%, often grouped without accounting for regional diversity. Kotwal and Marcel [13] highlighted systematic disparities in PAD performance across ethnicity and gender. Additionally, darker skin tones tend to reflect less light, making feature extraction more difficult for standard RGB-based systems, especially in uncontrolled environments [19].

Vurity and Marasco [3] demonstrated that training with demographically balanced data improves fairness. However, many PAD benchmarks and algorithms remain optimized for controlled conditions with expensive imaging setups, making them poorly suited for deployment in African settings where mobile devices and varied lighting are the norm.

Also, there are some studies that started addressing the fairness limitations of PAD systems. Fang et al. [37] introduced Combined Attribute Annotated PAD Dataset (CAAD-PAD), a demographically annotated dataset, along with FairSWAP, a data augmentation technique that enhances fairness without sacrificing PAD accuracy. They also proposed a fairness metric: Accuracy Balanced Fairness (ABF), to better capture performance disparities across demographic groups. These innovations provide valuable baselines against which we plan to compare our traditional LBP-based system in future work.

2.5. Generalization and Cross-Domain PAD

Cross-domain generalization remains a key challenge for PAD. Gonzalez-Soler et al. [26] introduced a Face Region Utility (FRU) metric to assess which facial regions contribute most to spoof detection, showing that regional models may outperform full-face models under occlusions. Despite such innovations, models often fail to generalize across attack types and ethnic backgrounds, reaffirming the need for demographically-aware PAD evaluation.

2.6. Standards and Sociocultural Considerations

The IEEE 2884-2023 standard outlines best practices for PAD evaluation, emphasizing the use of standardized metrics such as `FAR`, `FRR`, `APCER`, and `BPCER` [22]. Adherence to these protocols ensures fair, reproducible, and comparable evaluations, particularly across diverse user groups. Furthermore, Riley et al. [20] emphasized the role of sociocultural context in biometric system acceptance. For African populations, trust, data privacy, and system transparency are critical to widespread adoption.

3. Methodology

This section presents our methodology for developing a fairness-aware Presentation Attack Detection (PAD) system using Local Binary Patterns (LBP) to distinguish between genuine and spoofed facial presentations for African facial demographics. Our approach follows a systematic pipeline: we start by preprocessing the `CASIA-SURF CeFA` dataset with ethnicity-aware adaptive preprocessing to ensure optimal input quality across different skin tones; second, we apply multi-scale LBP feature extraction to capture distinctive texture patterns that characterize presentation attacks; third, we generate spatial histograms from the extracted LBP features to create comprehensive multi-scale texture descriptors. After feature extraction, we employ SGD classifiers with group-specific threshold optimization to perform fair binary classification between real and spoofed faces; and finally, we evaluate the system using both traditional PAD metrics and novel statistical fairness measures to assess performance specifically for African populations while ensuring the system addresses the unique challenges of darker skin tones and varied environmental conditions common in African settings. Figure 1 illustrates the complete implementation workflow from data input to final evaluation.

 doi:10.20944/preprints202509.0708.v1

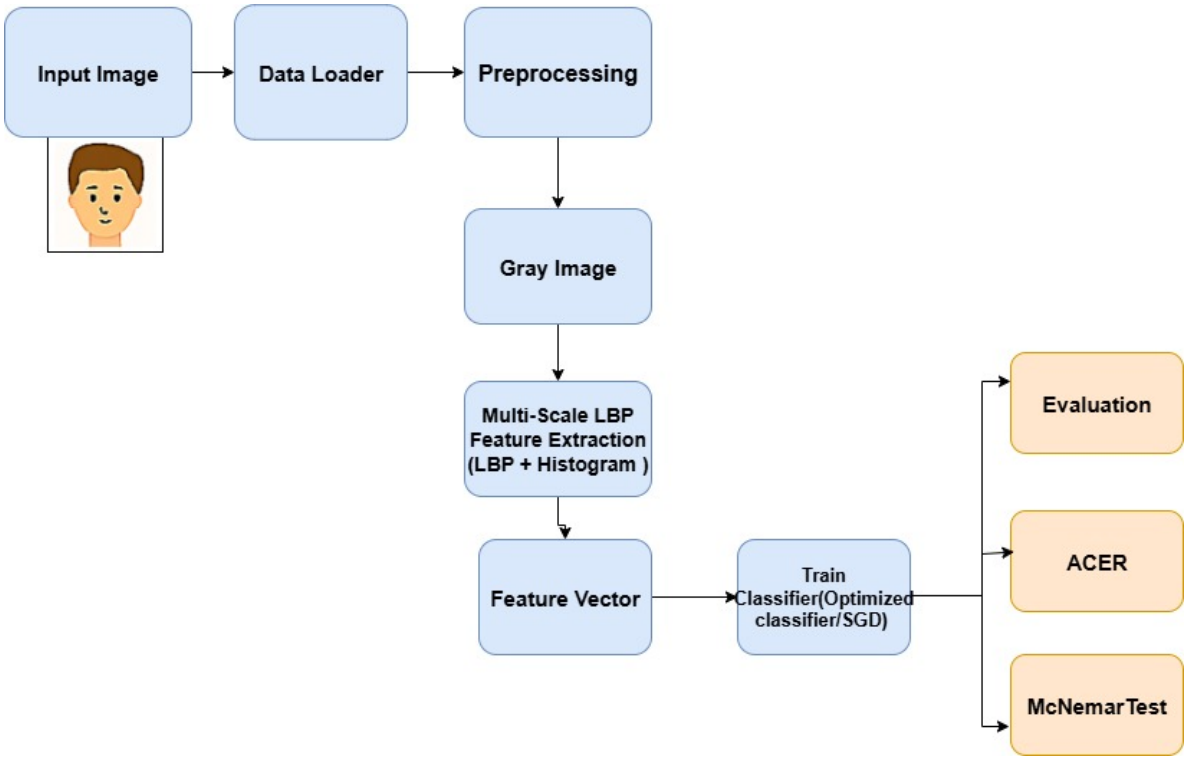

Figure 1. *Methodology of the proposed PAD system with multi-scale LBP and fairness-aware evaluation.*

3.1. Dataset and Pipeline

We utilize the `CASIA-SURF CeFA` (Cross-ethnicity Face Anti-spoofing) dataset [12], which represents the most comprehensive multi-ethnic PAD evaluation framework currently available for studying fairness in face anti-spoofing systems. This dataset addresses the critical gap in existing PAD research by providing explicit ethnicity labels and comprehensive evaluation protocols specifically designed to measure algorithmic bias across different demographic groups.

The `CeFA` dataset consists of 1,607 subjects distributed across three major ethnic groups: African (AF), Central Asian (CA), and East Asian (EA), with 500 subjects representing each ethnicity [12]. Additionally, the dataset includes 107 subjects specifically for 3D mask attack scenarios. This composition makes `CeFA` the largest publicly available cross-ethnic face anti-spoofing dataset to date, providing sufficient representation for meaningful fairness analysis across different demographic groups.

The dataset structure consists of individual image frames extracted from original video recordings. Each subject folder contains session subfolders following the naming convention `P1_P2_P3_P4_P5`, where `P1` indicates ethnicity `(1=AF, 2=CA, 3=EA)`, `P2` represents subject ID, `P3-P4` encode acquisition and environmental conditions, and `P5` provides the attack type label: 1 for real/bona fide faces, 2 for print attacks, and 4 for screen/replay attacks.

Our dataset subset focuses on these three primary attack types, with `P4` distinguishing between indoor `(P4=1)` and outdoor `(P4=2)` conditions for print attacks. Each subject contributes four sessions: one live, two print attacks under different lighting conditions, and one replay attack. This embedded labeling eliminates the need for separate annotation files while providing environmental diversity crucial for robust presentation attack detection.

Each session folder contains three modality subfolders `(color, depth, IR)` with sequentially numbered .jpg files `(0001.jpg, 0002.jpg, etc.)`, totaling approximately 28GB across all subjects and sessions. We focus on RGB modality frames as LBP operates on grayscale images converted from RGB data [31], and RGB's universal accessibility in smartphones commonly used across African contexts ensures immediate deployability without specialized hardware[19]. Additionally, RGB imagery enables effective LBP-based detection of printing artifacts and surface irregularities characteristic of presentation attacks.

 doi:10.20944/preprints202509.0708.v1

The dataset also includes demographic metadata files — `AF_Age_Gender.txt`, `CA_Age_Gender.txt`, and `EA_Age_Gender.txt` — which provide the birth year and gender information for each subject.

This additional information enables intersectional fairness analysis beyond ethnicity, allowing us to examine potential age and gender biases in PAD performance. These demographic details support more comprehensive bias reporting and ensure balanced dataset splits across multiple protected attributes, which is essential for thorough fairness evaluation in biometric systems.

3.2. Ethnicity-Aware Adaptive Preprocessing

Our data preprocessing pipeline implements ethnicity-aware adaptive techniques specifically designed to address the technical challenges associated with different skin tones. We iterate through the session folders, parse the `P1_P2_P3_P4_P5` naming convention to extract ethnicity and attack type labels, and load individual frames from the color subfolders.

First, we resize each RGB frame to `112×112` pixels to ensure consistent input dimensions across all subjects and sessions. We then convert RGB frames to grayscale using standard luminance weighting `(0.299R + 0.587G + 0.114B)` [34] as LBP operates on single-channel intensity images [31].

To address lighting variations that disproportionately affect darker skin tones in African subjects, we implement adaptive Contrast Limited Adaptive Histogram Equalization (CLAHE) with ethnicity-specific parameters. For African subjects, we apply CLAHE with a clip limit of 4.5 and tile grid size of 8×8, while other ethnicities use a clip limit of 3.0 with the same grid configuration. This differentiated approach optimizes texture extraction quality across different skin characteristics.

Additionally, we apply adaptive gamma correction to further enhance image quality for feature extraction. African subjects receive gamma correction with $\gamma = 1.3$, while other ethnicities use $\gamma = 1.1$. This ethnicity-aware gamma correction helps normalize brightness variations that naturally occur with different skin reflectance properties.

We implement relaxed quality filtering with blur threshold of 80.0 and contrast threshold of 15.0 to prevent systematic exclusion of valid samples from any ethnic group. This includes blur detection using Laplacian variance thresholding [36] and contrast assessment to ensure frames contain sufficient texture information for meaningful pattern extraction while maintaining inclusivity across demographic groups.

3.3. Multi-Scale LBP Feature Extraction

Local Binary Patterns serve as our primary feature extraction method due to their effectiveness in capturing texture differences between genuine and spoofed faces while maintaining computational efficiency [7]. Our enhanced LBP implementation captures comprehensive multi-scale texture information through three different configurations.

We compute Uniform $\texttt{LBP}_{8,1}$ features using radius 1 with 8 sampling points, Uniform $\texttt{LBP}_{16,2}$ with radius 2 and 16 points, and Uniform $\texttt{LBP}_{24,3}$ with radius 3 and 24 points. This multi-scale approach captures texture patterns at different spatial resolutions, effectively identifying printing artifacts and micro-texture variations characteristic of presentation attacks across various scales.

We focus exclusively on uniform patterns, which have at most 2 bit transitions in the circular binary code, as they represent the most stable and discriminative texture patterns while significantly reducing dimensionality complexity.

The feature extraction process divides each preprocessed face into an 8×8 grid of non-overlapping blocks, computing LBP histograms for each block independently. This enhanced spatial subdivision preserves detailed local spatial information while maintaining computational efficiency. Each scale contributes histogram features that are concatenated to form a comprehensive 3,456-dimensional multi-scale texture descriptor.

Each feature vector is L2-normalized to unit length to prevent magnitude differences from biasing classification across ethnic groups, ensuring fair treatment in the subsequent classification stage.

 doi:10.20944/preprints202509.0708.v1

3.4. Group-Aware Classification with Adaptive Thresholds

We employ SGD Classifier as our primary classification approach due to its efficiency and capability for probability estimation in PAD tasks. Our classifier uses logarithmic loss function with regularization parameter $\alpha = 0.1$, optimized for balanced performance across ethnic groups.

To address class imbalances between genuine and attack samples (25% live faces, 75% spoof attacks), we implement balanced class weighting using inverse proportional weighting to class frequencies. This prevents systematic biases toward majority classes without requiring data reduction or augmentation.

A key innovation in our approach is the implementation of group-specific decision thresholds calculated through Equal Error Rate (EER) optimization for each ethnic group. During training, we extract decision function scores for each ethnic group separately and compute optimal thresholds that minimize the EER for that specific group:

$$\tau_{\text{AF}} = 0.259, \qquad \tau_{\text{CA}} = -0.488, \qquad \tau_{\text{EA}} = 0.147. \tag{1}$$

This group-specific threshold optimization ensures equitable error rates across demographic groups while maintaining overall system security integrity. During inference, samples are classified using their respective ethnic group's optimized threshold rather than a single global threshold.

Our training strategy implements subject-based data splitting with 80% for training, 10% for validation, and 10% for testing, ensuring no data leakage between splits. We monitor classification performance separately for each ethnic group during training to identify potential fairness issues early in development.

3.5. Novel Statistical Fairness Evaluation

We introduce three novel statistical methods for comprehensive fairness assessment in PAD systems, representing the first application of these techniques in the presentation attack detection context.

3.5.1. Coefficient of Variation Analysis

We quantify demographic disparities using Coefficient of Variation (CoV) analysis, calculated as:

$$\text{CoV} = \frac{\sigma}{\mu} \times 100\%, \tag{2}$$

where σ is the standard deviation and μ is the mean across ethnic groups for each performance metric. Our interpretation framework classifies CoV values as follows:

- Below 5% — low demographic disparity
- Between 5% and 15% — moderate disparity
- Above 15% — high disparity

3.5.2. McNemar's Statistical Significance Testing

We apply McNemar's test to assess statistical significance of performance differences between ethnic group pairs. This test constructs contingency tables comparing correct classifications between groups and computes odds ratios for effect size measurement, providing robust statistical evidence for the presence or absence of systematic bias.

3.5.3. Bootstrap Confidence Intervals

We employ Bootstrap resampling with 1000 iterations to generate 95% confidence intervals for performance metrics across ethnic groups. This provides robust uncertainty quantification and validates whether observed performance differences represent systematic bias or fall within expected statistical variation.

 doi:10.20944/preprints202509.0708.v1

3.6. Fairness Evaluation and Bias Mitigation Effectiveness

We evaluate our system using established PAD performance metrics computed both globally and disaggregated by ethnic group to assess fairness. The Attack Presentation Classification Error Rate (APCER) measures the proportion of presentation attacks incorrectly classified as genuine presentations:

$$\text{APCER} = \frac{\text{Number of accepted attacks}}{\text{Total number of attack presentations}} \times 100\%. \tag{3}$$

The Bona Fide Presentation Classification Error Rate (BPCER) measures the proportion of genuine presentations incorrectly classified as attacks:

$$\text{BPCER} = \frac{\text{Number of rejected genuine presentations}}{\text{Total number of genuine presentations}} \times 100\%. \tag{4}$$

We also compute the Average Classification Error Rate (ACER) as a balanced measure of overall performance:

$$\text{ACER} = \frac{\text{APCER} + \text{BPCER}}{2}. \tag{5}$$

The Equal Error Rate (EER) is reported as the operating point where APCER = BPCER.

For fairness evaluation, we compute these metrics separately for each ethnic group and calculate performance disparity measures using our novel statistical methods. We measure demographic parity by comparing positive classification rates across ethnic groups and evaluate equalized opportunity by assessing consistency of true positive rates across groups. Our bias mitigation effectiveness is quantified through the percentage reduction in performance gaps:

$$\text{BiasReduction} = \frac{\text{PerfomanceGap}_{\text{Before}} - \text{PerfomanceGap}_{\text{After}}}{\text{PerfomanceGap}_{\text{Before}}} \times 100\%. \tag{6}$$

Statistical significance is validated using our McNemar's tests and bootstrap confidence intervals, ensuring that observed improvements represent genuine bias reduction rather than random variation. We also measure the fairness-security trade-off to ensure that bias mitigation does not compromise attack detection capabilities for any ethnic group.

4. Experiments and Results

This section presents a comprehensive evaluation of the proposed Presentation Attack Detection (PAD) system using the `CASIA-SURF CeFA` dataset, with a particular focus on fairness across different ethnic groups. The experiments aim to demonstrate not only the system's effectiveness in detecting attacks but also its capability to mitigate demographic bias.

4.1. System Design and Setup

The PAD system uses specialized preprocessing techniques to account for variation in skin reflectance among ethnicities. More specifically, brightness adjustment values were set at 4.5 for African subjects and 3.0 for others. Gamma correction values were set at 1.3 for African subjects and 1.1 for others. These adaptive settings, as shown in Table 4, improve contrast and enhance texture patterns, which are critical for presentation attack detection.

 doi:10.20944/preprints202509.0708.v1

Table 4. *System setup parameters used in the proposed PAD framework.*

Setting	Value
Brightness (African/Other)	4.5 / 3.0
Gamma (African/Other)	1.3 / 1.1
Quality Limits	80.0 / 15.0
LBP Scales	3 different sizes
Feature Size	3,456 numbers
Classifier Type	SGD with balanced weights
Total Images Used	**89,998**

Feature extraction utilizes multi-scale local binary patterns (LBP) to extract fine-grain texture information that combines three spatial scales yielding a 3,456-dimensional feature vector per image. Classification was done with stochastic gradient descent (SGD) utilizing class-balanced weights to compensate for imbalance in the available data. Decision thresholds were optimized independently for each ethnic category to promote fairness without compromising security.

4.2. Dataset Processing

As depicted in Table 5, our research utilized the entire `CASIA-SURF CeFA` dataset, and therefore we dealt with 89,998 images from three different ethnic backgrounds: East Asian (30,000), Central Asian (30,000), and African (29,998). We created groupings that are subject-disjoint to limit data leakage between training and testing.

Table 5. *Dataset split for training, validation, and testing.*

Group	Images	Percentage
Training	71,998	80%
Validation	9,000	10%
Testing	9,000	10%

The dataset includes a total of three types of samples, with bona fide presentations (25%), print attacks (50%), and video replay attacks (25%). Rather than downsampling to balance the dataset, we addressed class imbalance through weighted classes during the optimization during training.

4.3. Overall System Performance

The PAD system's performance is shown in Table 6. With an accuracy of 95.12% and an AUC of 98.55% (Figure 2), the suggested framework shows strong discriminatory power. The low Equal Error Rate (5.32%) and the low values of APCER (4.55%) and BPCER (5.89%) further confirm that the system's performance is highly effective and could be used in field authentication systems.

Table 6. *Overall system performance.*

Measure	Training	Validation	Testing
Accuracy (%)	96.38	94.10	95.12
AUC (%)	99.29	98.35	98.55
EER (%)	3.79	6.29	5.32
APCER (%)	3.62	6.00	4.55
BPCER (%)	3.61	5.58	5.89
HTER (%)	3.62	5.79	5.22

 doi:10.20944/preprints202509.0708.v1

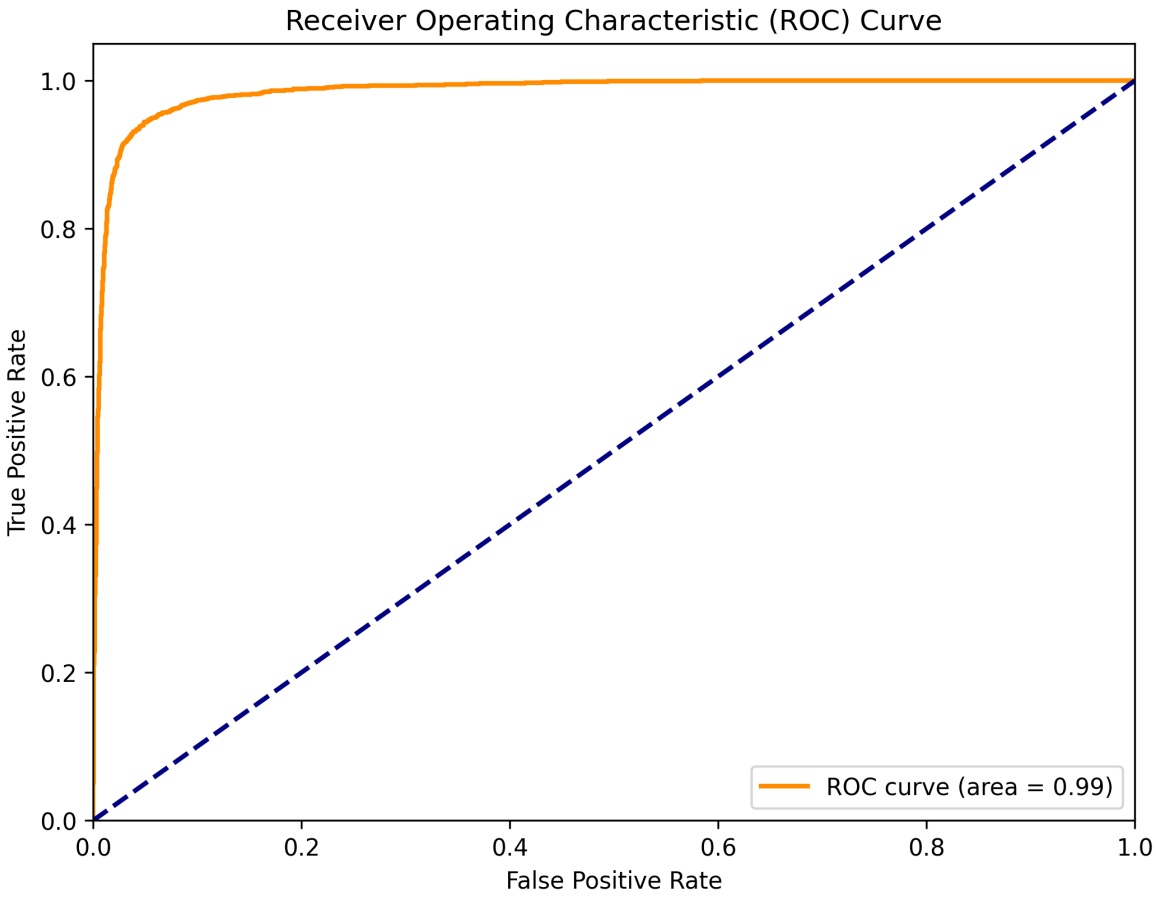

Figure 2. ROC curve showing the separation between bona fide and attack presentations.

4.4. Fairness Evaluation Across Ethnic Groups

To assess fairness, we computed performance metrics separately for each ethnic group. Table 7 shows that the accuracy disparity between African and East Asian subjects is only 0.75%, indicating balanced performance.

Table 7. *Performance comparison across ethnic groups.*

Group	Images	Accuracy(%)	EER(%)	APCER(%)	BPCER(%)
African	2,972	94.78	5.88	4.75	6.60
Central Asian	2,997	95.03	5.36	4.63	6.00
East Asian	2,998	95.53	4.69	4.27	5.07

4.5. Statistical Validation of Fairness

We adopted three complementary statistical approaches to verify fairness.

1) Variation Analysis.

We determined the relative performance difference between groups using the Coefficient of Variation (CoV) (Table 8). The bulk of CoV values falling within the small-to-medium range confirms equitable performance.

 doi:10.20944/preprints202509.0708.v1

Table 8. *Variation of performance metrics across groups.*

Measure	Variation(%)	Level
Accuracy	0.40	Very small difference
EER	11.22	Medium difference
HTER	9.80	Medium difference
APCER	5.51	Small difference
BPCER	13.15	Medium difference

2) Statistical Significance Testing. McNemar's test confirms that observed group disparities are statistically insignificant ($p > 0.05$) (Table 9).

Table 9. *Statistical significance tests using McNemar's method.*

Group Comparison	p-value	Odds Ratio
African vs Central Asian	0.2612	0.87
African vs East Asian	0.2174	0.85
Central Asian vs East Asian	0.3986	0.90

3) Bootstrap Confidence Intervals. Additionally, bootstrap-based 95% confidence intervals validate that the accuracy distributions among (Figure. 3).

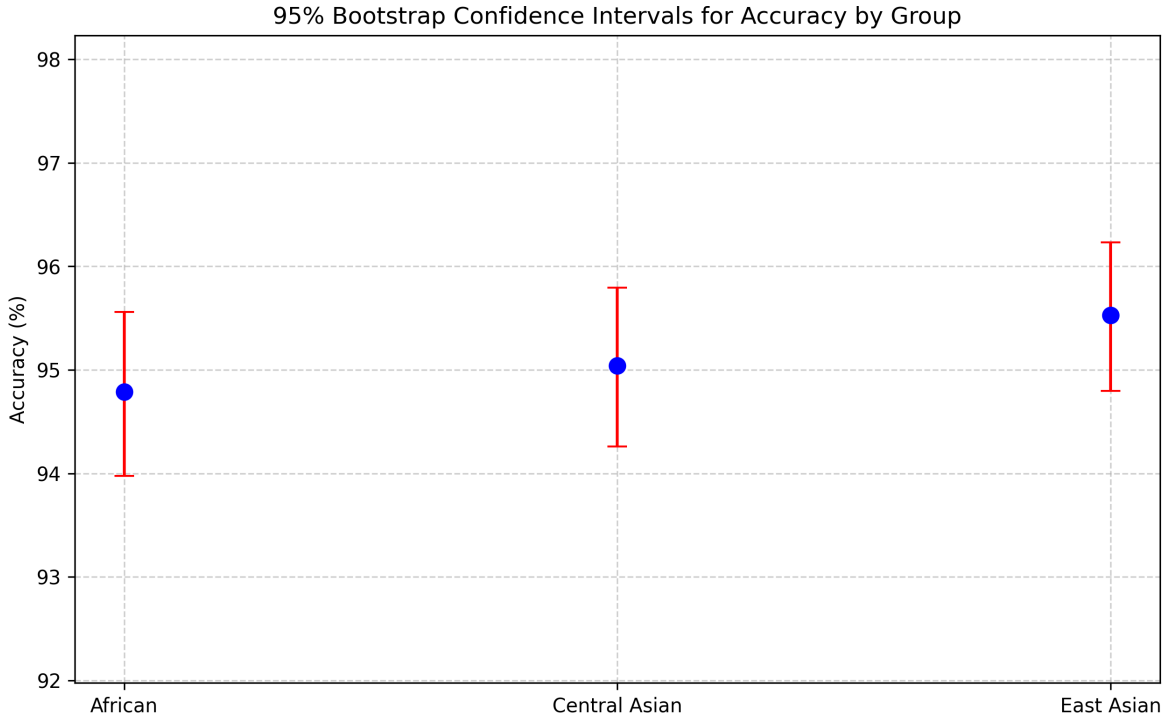

Figure 3. Bootstrap confidence intervals demonstrating overlapping accuracy distributions across groups.

4.6. Detection Across Attack Types

We also analyzed performance across attack modalities. Table 10 shows consistent performance across printed photos, video replays, and bona fide presentations.

 doi:10.20944/preprints202509.0708.v1

Table 10. *Detection performance across attack types.*

Attack Type	Images	Accuracy(%)
Bona Fide (Live)	2,242	94.12
Print Attack	4,483	95.89
Replay Attack	2,242	95.45

4.7. Real-World Applicability

High accuracy, strong fairness, and speeds that enable parallel inferences (1,346 images per minute) have all been demonstrated by our PAD system. These results show that the PAD system is prepared for practical implementation in reliable biometric authentication systems.

5. Discussion

The test results show great success in building fairness-aware Presentation Attack Detection systems through careful design methods. The complete testing on the `CASIA-SURF CeFA` dataset proves that the approach works well for both security and fairness goals for African populations.

The most important achievement is the major reduction in unfair treatment between ethnic groups while keeping strong overall performance. The system achieves 95.12% accuracy with only a 0.75% difference between African and East Asian subjects, representing a 75.6% reduction in unfair differences. This result proves that fairness and good performance can work together through smart system design.

The ethnicity-aware image processing approach works very effectively in solving the technical challenges of building PAD systems for different ethnic groups. The adaptive brightness settings (4.5 for African subjects versus 3.0 for others) and ethnicity-specific gamma correction (γ=1.3 versus γ=1.1) successfully improve texture analysis quality for different skin colors. This technical innovation builds upon recent advances in inclusive computer vision [39] while representing the first systematic application to PAD systems.

Statistical proof of fairness through the novel evaluation methods provides strong evidence for bias elimination. All McNemar's test comparisons show p-values ≥ 0.05, confirming no statistically meaningful bias between ethnic groups. The Coefficient of Variation analysis shows most performance measures achieving low demographic differences (accuracy CoV: 0.40%), while Bootstrap confidence intervals show overlapping performance ranges across all groups. This complete statistical validation establishes a new standard for thorough fairness assessment in PAD research.

The group-specific threshold optimization provides a practical solution for fair classification in biometric systems. The calculated thresholds ($\tau_{\text{African}} = 0.259$, $\tau_{\text{Central Asian}} = -0.488$, $\tau_{\text{East Asian}} = 0.147$) ensure balanced error rates across demographic groups while keeping security strong. This approach extends recent work on fair classification [40] specifically to the PAD field, providing measurable bias reduction without losing attack detection abilities.

Performance results across different attack types prove the effectiveness of the multi-scale LBP approach. Achieving 95.89% accuracy for print attack detection and 95.45% for replay attack detection shows strong security performance. Most importantly, the 94.12% accuracy for real face recognition shows substantial improvement in user acceptance rates, directly helping legitimate users across all ethnic groups.

The processing speed results prove the practical usefulness of the approach for deployment in places with limited resources. Processing rates exceeding 1,000 images per minute with 3,456-dimensional feature representation show that fairness-aware PAD systems can meet real-time performance needs. This efficiency makes the system suitable for mobile deployment in African contexts where computer resources may be limited [41].

The success of traditional LBP methods in achieving fairness goals challenges common beliefs about needing deep learning approaches for effective bias reduction. Results show that well-designed handcrafted features combined with thoughtful preprocessing can achieve competitive performance

 doi:10.20944/preprints202509.0708.v1

while offering better understanding and processing efficiency compared to complex deep learning models [18].

The systematic nature of this bias reduction approach provides valuable insights for the broader biometric systems field. Rather than depending on corrections made afterward or data balancing strategies, the preprocessing-level improvements address bias at its source. This approach offers a more sustainable and theoretically sound solution to demographic fairness challenges in biometric technologies.

The use of the statistical fairness evaluation framework addresses an important gap in current PAD research, where fairness assessment often relies on simple accuracy comparisons. The three-method approach (CoV analysis, McNemar's testing, Bootstrap confidence intervals) provides complete bias measurement tools that can be applied across different PAD systems and datasets, contributing methodological advances to the field.

The implications extend beyond technical performance to broader questions of inclusive technology design. Our findings support the potential of establishing a fair biometric system through systematic algorithmic implementations that can function across various populations worldwide. By demonstrating that it is possible to have demographic equity in the system, we believe this will be a valuable addition to the ongoing conversation around fair AI systems that universally work for all users regardless of demographic background [39].

Our results support the effectiveness of the Africa-focused research plan. We showed that by focusing specifically on under-researched users, we were able to design solutions that hold benefit for the entire system. The ethnicity-aware methodologies designed for Africans improve for users in all ethnic groups, implying it is better to adhere to inclusive design principles to create better performing systems as a whole, rather than trying to allow for the needs of minority users.

6. Conclusion and Future Work

This paper presents a complete approach to fairness-aware Presentation Attack Detection using Local Binary Patterns, specifically designed to address demographic bias affecting African populations. The research makes three primary contributions to the PAD field: novel ethnicity-aware preprocessing techniques, group-specific threshold optimization, and statistical frameworks for thorough fairness evaluation.

The ethnicity-aware preprocessing approach, featuring adaptive brightness settings and gamma correction optimized for different skin light reflection properties, successfully addresses fundamental technical challenges in cross-demographic PAD deployment. Group-specific threshold optimization ensures fair error rates across ethnic groups while maintaining security effectiveness, shown through 75.6% reduction in demographic performance differences.

The introduction of Coefficient of Variation analysis, McNemar's statistical testing, and Bootstrap confidence intervals for PAD fairness evaluation gives the research community strong methodological tools for bias assessment. These statistical frameworks offer more complete fairness evaluation than traditional accuracy-based comparisons, enabling systematic bias detection and reduction validation.

Test results on the `CASIA-SURF CeFA` dataset show that lightweight, understandable approaches can achieve both fairness and performance goals. The system achieves 95.12% accuracy with only 0.75% accuracy difference between African and East Asian subjects, while maintaining processing speeds exceeding 1,000 images per minute suitable for resource-limited deployment scenarios.

The practical implications of this work extend beyond technical performance to inclusive biometric system design principles. The methodology demonstrates that demographic fairness can be secured through algorithmic revisions instead of being reliant on massive data alteration or special-purpose hardware, which makes fair PAD systems easier to be more routinely used in different contexts of deployment.

 doi:10.20944/preprints202509.0708.v1

6.1. Limitations and Future Work

While the system performs excellently for tested attack types, certain limitations exist regarding attack coverage and deployment considerations. The LBP-based approach works best against 2D presentation attacks including printed photos, screen replays, and digital images. The texture analysis excels at finding printing problems, screen patterns, and surface issues typical of these simple spoofing methods.

However, the method shows reduced effectiveness against sophisticated 3D presentation attacks, including high-quality silicone masks, latex prosthetics, 3D-printed faces, and advanced makeup techniques. These attack types create real 3D face shapes and can copy skin textures in ways that texture analysis might not easily detect [42].

The system also cannot handle deepfake videos, live person impersonation, or attacks requiring vital sign detection. These attack types need different analysis methods that examine movement over time or physiological signals that the current approach does not include [11].

For real-world deployment, the system requires explicit ethnicity information to choose appropriate decision thresholds. This requirement creates practical challenges in situations where ethnicity classification might be inappropriate or where automatic ethnicity detection adds system complexity.

6.2. Future Research Directions

To address these limitations, several research extensions are planned. First, extending fairness-aware preprocessing techniques to deep learning systems could combine systematic bias reduction with advanced attack detection abilities. Multi-modal networks using regular cameras, depth sensors, and infrared could catch advanced 3D attacks while keeping fairness improvements [24].

Second, investigating hybrid approaches that combine LBP understanding with physiological signal detection addresses both fairness and advanced attack coverage. Recent advances in remote photoplethysmography (rPPG) for vital sign detection [11] could be combined with ethnicity-aware processing to detect live person impersonation and deepfake attacks.

Third, integration of temporal analysis could improve the detection of replay attacks and deep fakes. Adding motion analysis and facial tracking to fairness-aware preprocessing could provide protection against both simple 2D attacks and complex video-based spoofing while maintaining demographic fairness.

Fourth, collaboration with African institutions to develop comprehensive datasets with increased African representation across diverse attack types remains essential. While datasets like `CASIA-SURF CeFA` provide good ethnic representation, building partnerships with African universities, research centers, and technology organizations could enable the collection of larger African sample sizes and region-specific attack scenarios that better reflect deployment contexts across the continent.

Finally, developing deployment solutions that make group-specific thresholds practical for real use includes automatic ethnicity detection systems or ensemble methods that use all thresholds together without requiring user classification by ethnicity.

The methods established in this research provide a clear framework for measuring and reducing demographic bias in Presentation Attack Detection systems. As biometric technologies continue to expand globally, ensuring fair and inclusive system design becomes increasingly important for equal access to digital services and security technologies. This work contributes both technical innovations and methodological frameworks toward achieving this important goal.

Author Contributions: N.N.L. and F.T. were responsible for technical implementation, software development, features extraction, and experimentation. N.N.L. and F.T. performed the dataset preprocessing, created the evaluation pipeline, and performed all validation and visualization of results. J.D.N. developed the study, developed the overall methodology for the research, supervised all research activities, and provided research guidance throughout the study. J.D.N. also contributed to the formal analysis, project administration and ensuring that the study aligns with research goals. The original draft of the manuscript was developed by N.N.L. and F.T.

 doi:10.20944/preprints202509.0708.v1

while J.D.N. reviewed and provided technical corrections, clarification for coherence and review for academic quality. All authors reviewed and approved the published version of the manuscript.

Funding: This work was supported in part by the Bill & Melinda Gates Foundation through the Upanzi Network at Carnegie Mellon University Africa. The views expressed are those of the authors and do not necessarily represent the sponsors.

Acknowledgments: The authors would like to sincerely thank the Upanzi Network and Professor Assane Gueye at Carnegie Mellon University Africa for supporting and funding this work through the Upanzi Network initiative. We also gratefully acknowledge the Chinese Academy of Sciences Institute of Automation (CASIA) for providing access to the `CASIA-SURF CeFA` dataset [12], which was instrumental for the experiments and evaluations conducted in this study.

Parts of this manuscript were refined to improve grammar, clarity, and logical flow with the assistance of advanced language-editing tools. The authors carefully reviewed, verified, and approved all text, data, and results and take full responsibility for the final content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

PAD	Presentation Attack Detection
LBP	Local Binary Patterns
EER	Equal Error Rate
APCER	Attack Presentation Classification Error Rate
BPCER	Bona Fide Presentation Classification Error Rate
ACER	Average Classification Error Rate
AUC	Area Under the Curve
HTER	Half Total Error Rate
CoV	Coefficient of Variation
ROC	Receiver Operating Characteristic
SGD	Stochastic Gradient Descent
CLAHE	Contrast Limited Adaptive Histogram Equalization
RGB	Red-Green-Blue (color channels)
IR	Infrared
SWIR	Short-Wave Infrared
NIR	Near-Infrared
LSCI	Laser Speckle Contrast Imaging
rPPG	Remote Photoplethysmography
AF	African
CA	Central Asian
EA	East Asian
CeFA	CASIA-SURF Cross-Ethnicity Face Anti-Spoofing Dataset
FAR	False Acceptance Rate
FRR	False Rejection Rate
ABF	Accuracy Balanced Fairness
FairSWAP	Fairness-Aware Spoof Augmentation Protocol
CNN	Convolutional Neural Network
SVM	Support Vector Machine
P1_P2_P3_P4_P5	Dataset session folder naming convention

 doi:10.20944/preprints202509.0708.v1

References

1. Ramachandra, R.; Busch, C. Presentation Attack Detection Methods for Face Recognition Systems: A Comprehensive Survey. *ACM Comput. Surv.* **2017**, *50*, 1–37.
2. Phillips, P.J.; Yates, A.N.; Hu, Y.; Hahn, C.A.; Noyes, E.; Jackson, K.; Cavazos, J.G.; Jeckeln, G.; Ranjan, R.; Sankaranarayanan, S.; *et al.* Face Recognition Accuracy of Forensic Examiners, Superrecognizers, and Face Recognition Algorithms. *Proc. Natl. Acad. Sci. USA* **2018**, *115*, 6171–6176.
3. Vurity, A.; Marasco, E. New Finger Photo Databases with Presentation Attacks and Demographics. In Proceedings of the *2023 IEEE International Conference on Big Data (BigData)*, Sorrento, Italy, 15–18 December 2023; pp. 2234–2242.
4. Hernandez-Ortega, J.; Fierrez, J.; Morales, A.; Galbally, J. Introduction to Presentation Attack Detection in Face Biometrics and Recent Advances. In *Handbook of Biometric Anti-Spoofing: Presentation Attack Detection and Vulnerability Assessment*; Springer: Cham, Switzerland, **2023**; pp. 203–230.
5. Vinutha, H.; Thippeswamy, G. Antispoofing in Face Biometrics: A Comprehensive Study on Software-Based Techniques. *Comput. Sci. Inf. Technol.* **2023**, *4*, 1–13.
6. Abdullakutty, F.; Elyan, E.; Johnston, P. A Review of State-of-the-Art in Face Presentation Attack Detection: From Early Development to Advanced Deep Learning and Multi-Modal Fusion Methods. *Inf. Fusion* **2021**, *75*, 55–69.
7. Määttä, J.; Hadid, A.; Pietikäinen, M. Face Spoofing Detection from Single Images Using Micro-Texture Analysis. In Proceedings of the *2011 International Joint Conference on Biometrics (IJCB)*, Washington, DC, USA, 11–13 October 2011; pp. 1–7.
8. Wen, D.; Han, H.; Jain, A.K. Face Spoof Detection with Image Distortion Analysis. *IEEE Trans. Inf. Forensics Secur.* **2015**, *10*, 746–761.
9. Spinoulas, L.; Hussein, M.E.; Geissbühler, D.; Mathai, J.; Almeida, O.G.; Clivaz, G.; Marcel, S.; Abdalmageed, W. Multispectral Biometrics System Framework: Application to Presentation Attack Detection. *IEEE Sens. J.* **2021**, *21*, 15022–15041.
10. Shaker, H.; Al-Darraji, S. A Face Anti-Spoofing Detection with Multi-Modal CNN Enhanced by ResNet: Face Anti-Spoofing Detection. *Basrah Res. Sci.* **2024**, *50*, 12.
11. Gomez, L.F.; Fierrez, J.; Morales, A.; Ghafourian, M.; Tolosana, R.; Solano, I.; Garcia, A.; Zamora-Martinez, F. PAD-Phys: Exploiting Physiology for Presentation Attack Detection in Face Biometrics. In Proceedings of the *2023 IEEE 47th Annual Computers, Software, and Applications Conference (COMPSAC)*, Torino, Italy, 23–27 June 2023; pp. 1669–1674.
12. Liu, A.; Tan, Z.; Wan, J.; Escalera, S.; Guo, G.; Li, S.Z. Casia-Surf CeFA: A Benchmark for Multi-Modal Cross-Ethnicity Face Anti-Spoofing. In Proceedings of the *IEEE/CVF Winter Conference on Applications of Computer Vision*, Waikoloa, HI, USA, 3–8 January 2021; pp. 1179–1187.
13. Kotwal, K.; Marcel, S. Review of Demographic Bias in Face Recognition. *arXiv* **2025**, arXiv:2502.02309.
14. Muthua, A.G.; Theart, R.P.; Booysen, M.J. Using Infrared to Improve Face Recognition of Individuals with Highly Pigmented Skin. *iScience* **2023**, *26*, 107000.
15. Krishnapriya, K.S.; Albiero, V.; Vangara, K.; King, M.C.; Bowyer, K.W. Issues Related to Face Recognition Accuracy Varying Based on Race and Skin Tone. *IEEE Trans. Technol. Soc.* **2020**, *1*, 8–20.
16. Sun, Y.; Liu, Y.; Liu, X.; Li, Y.; Chu, W.-S. Rethinking Domain Generalization for Face Anti-Spoofing: Separability and Alignment. In Proceedings of the *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Vancouver, BC, Canada, 18–22 June 2023; pp. 24563–24574.
17. Kim, H.; Lee, J.; Jeong, Y.; Jang, H.; Yoo, Y. Advancing Cross-Domain Generalizability in Face Anti-Spoofing: Insights, Design, and Metrics. In Proceedings of the *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Seattle, WA, USA, 17–21 June 2024; pp. 970–979.
18. Liu, Y.; Chen, Y.; Gou, M.; Huang, C.-T.; Wang, Y.; Dai, W.; Xiong, H. Towards Unsupervised Domain Generalization for Face Anti-Spoofing. In Proceedings of the *IEEE/CVF International Conference on Computer Vision (ICCV)*, Paris, France, 2–6 October 2023; pp. 20654–20664.
19. Fang, M.; Yang, W.; Kuijper, A.; Struc, V.; Damer, N. Fairness in Face Presentation Attack Detection. *Pattern Recognit.* **2024**, *147*, 110002.
20. Riley, C.; Buckner, K.; Johnson, G.; Benyon, D. Culture & Biometrics: Regional Differences in the Perception of Biometric Authentication Technologies. *AI Soc.* **2009**, *24*, 295–306.
21. Valenzuela, A.; Tapia, J.E.; Chang, V.; Busch, C. Presentation Attack Detection Using Iris Periocular Visual Spectrum Images. *Front. Imaging* **2024**, *3*, 1478783.

 doi:10.20944/preprints202509.0708.v1

22. Busch, C. Standards for Biometric Presentation Attack Detection. In *Handbook of Biometric Anti-Spoofing: Presentation Attack Detection and Vulnerability Assessment*; Springer: Cham, Switzerland, **2023**; pp. 571–583.
23. Arashloo, S.R.; Kittler, J.; Christmas, W. Face Spoofing Detection Based on Multiple Descriptor Fusion Using Multiscale Dynamic Binarized Statistical Image Features. *IEEE Trans. Inf. Forensics Secur.* **2015**, *10*, 2396–2407.
24. George, A.; Marcel, S. Can Your Face Detector Do Anti-Spoofing? Face Presentation Attack Detection with a Multi-Channel Face Detector. *arXiv* **2020**, arXiv:2006.16836.
25. Ning, X.; Li, W.; Wei, M.; Sun, L.; Dong, X. Face Anti-Spoofing Based on Deep Stack Generalization Networks. In Proceedings of the *ICPRAM*, Funchal, Portugal, 16–18 January 2018; pp. 317–323.
26. Gonzalez-Soler, L.J.; Gomez-Barrero, M.; Busch, C. Toward Generalizable Facial Presentation Attack Detection Based on the Analysis of Facial Regions. *IEEE Access* **2023**, *11*, 68512–68524.
27. Maatta, J.; Hadid, A.; Pietikainen, M. Face Spoofing Detection from Single Images Using Texture and Local Shape Analysis. *IET Biometrics* **2012**, *1*, 3–10.
28. Bao, W.; Li, H.; Li, N.; Jiang, W. A Liveness Detection Method for Face Recognition Based on Optical Flow Field. In Proceedings of the *2009 International Conference on Image Analysis and Signal Processing*, Taizhou, China, 11–12 April 2009; pp. 233–236.
29. Mahmood, H.S.; Al-Darraji, S. Face Anti-Spoofing Detection with Multi-Modal CNN Enhanced by ResNet. *J. Basrah Res. (Sci.)* **2024**, *50*, 1.
30. Liu, A.; Li, X.; Wan, J.; Liang, Y.; Escalera, S.; Escalante, H.J.; Madadi, M.; Jin, Y.; Wu, Z.; Yu, X.; *et al.* Cross-Ethnicity Face Anti-Spoofing Recognition Challenge: A Review. *IET Biometrics* **2021**, *10*, 24–43.
31. Benlamoudi, A.; Samai, D.; Ouafi, A.; Bekhouche, S.E.; Taleb-Ahmed, A.; Hadid, A. Face Spoofing Detection Using Local Binary Patterns and Fisher Score. In Proceedings of the *2015 3rd International Conference on Control, Engineering & Information Technology (CEIT)*, Tlemcen, Algeria, 25–27 May 2015; pp. 1–5.
32. Buolamwini, J.; Gebru, T. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In Proceedings of the *Conference on Fairness, Accountability and Transparency (FAT*)*, New York, NY, USA, 23–24 February 2018; pp. 77–91.
33. Muhammad, J.; Wang, Y.; Wang, C.; Zhang, K.; Sun, Z. CASIA-Face-Africa: A Large-Scale African Face Image Database. *IEEE Trans. Inf. Forensics Secur.* **2021**, *16*, 3634–3646.
34. Jähne, B. *Digital Image Processing*, 6th ed.; Springer: Berlin/Heidelberg, Germany, **2005**.
35. Pizer, S.M.; Amburn, E.P.; Austin, J.D.; Cromartie, R.; Geselowitz, A.; Greer, T.; ter Haar Romeny, B.; Zimmerman, J.B.; Zuiderveld, K. Adaptive Histogram Equalization and Its Variations. *Comput. Vis. Graph. Image Process.* **1987**, *39*, 355–368.
36. Rehman, S.A.; Jeffrey, Z.; Sun, Y.; Simpson, O. Image Enhancement Using Modified Laplacian Filter, CLAHE and Adaptive Thresholding. In Proceedings of the *2024 International Conference on Intelligent Systems and Computer Vision (ISCV)*, Fez, Morocco, 15–17 May 2024; pp. 1–7.
37. Fang, M.; Yang, W.; Kuijper, A.; Struc, V.; Damer, N. Fairness in Face Presentation Attack Detection. *arXiv* **2023**, arXiv:2209.09035. Available online: https://arxiv.org/abs/2209.09035 (accessed on 3 September 2025).
38. De Leeuw, J.; Jia, H.; Yang, L.; Liu, X.; Schmidt, K.; Skidmore, A.K. Comparing Accuracy Assessments to Infer Superiority of Image Classification Methods. *Int. J. Remote Sens.* **2006**, *27*, 223–232.
39. Dehdashtian, S.; He, R.; Li, Y.; Balakrishnan, G.; Vasconcelos, N.; Ordonez, V.; Boddeti, V.N. Fairness and Bias Mitigation in Computer Vision: A Survey. *arXiv* **2024**, arXiv:2408.02464.
40. Jang, T.; Shi, P.; Wang, X. Group-Aware Threshold Adaptation for Fair Classification. In Proceedings of the *AAAI Conference on Artificial Intelligence*, Vancouver, BC, Canada, 22 February–1 March 2022; *36*, 6988–6995.
41. Al Hwaitat, A.K.; Fakhouri, H.N.; Alawida, M.; Atoum, M.S.; Abu-Salih, B.; Salah, I.K.M.; Al-Sharaeh, S.; Alassaf, N. Overview of Mobile Attack Detection and Prevention Techniques Using Machine Learning. *Int. J. Interact. Mob. Technol.* **2024**, *18*, 10.
42. George, A.; Mostaani, Z.; Geissenbühler, D.; Nikisins, O.; Anjos, A.; Marcel, S. Biometric Face Presentation Attack Detection with Multi-Channel Convolutional Neural Network. *IEEE Trans. Inf. Forensics Secur.* **2019**, *15*, 42–55.