




# Data-Driven Fairness Generalization for Deepfake Detection

Uzoamaka Ezeakunne<sup>1</sup><sup>a</sup>, Chrisantus Eze<sup>2</sup><sup>b</sup> and Xiuwen Liu<sup>1</sup><sup>c</sup>

<sup>1</sup>Department of Computer Science, Florida State University, Tallahassee FL, U.S.A.

<sup>2</sup>Department of Computer Science, Oklahoma State University, Stillwater OK, U.S.A.

**Keywords:** Deepfake Detection, Image Manipulation Detection, Fairness, Generalization.

**Abstract:** Despite the progress made in deepfake detection research, recent studies have shown that biases in the training data for these detectors can result in varying levels of performance across different demographic groups, such as race and gender. These disparities can lead to certain groups being unfairly targeted or excluded. Traditional methods often rely on fair loss functions to address these issues, but they under-perform when applied to unseen datasets, hence, fairness generalization remains a challenge. In this work, we propose a data-driven framework for tackling the fairness generalization problem in deepfake detection by leveraging synthetic datasets and model optimization. Our approach focuses on generating and utilizing synthetic data to enhance fairness across diverse demographic groups. By creating a diverse set of synthetic samples that represent various demographic groups, we ensure that our model is trained on a balanced and representative dataset. This approach allows us to generalize fairness more effectively across different domains. We employ a comprehensive strategy that leverages synthetic data, a loss sharpness-aware optimization pipeline, and a multi-task learning framework to create a more equitable training environment, which helps maintain fairness across both intra-dataset and cross-dataset evaluations. Extensive experiments on benchmark deepfake detection datasets demonstrate the efficacy of our approach, surpassing state-of-the-art approaches in preserving fairness during cross-dataset evaluation. Our results highlight the potential of synthetic datasets in achieving fairness generalization, providing a robust solution for the challenges faced in deepfake detection.


## 1 INTRODUCTION


Deepfake technology, which combines "deep learning" and "fake," represents a significant advancement in media manipulation capabilities. Leveraging deep learning techniques, it enables highly convincing facial manipulations and replacements in digital media. While technologically impressive, this capability poses serious societal risks, particularly in spreading misinformation and eroding public trust. In response, researchers have developed various deepfake detection techniques that have shown promising accuracy rates in identifying manipulated content (Xu et al., 2023; Zhao et al., 2021a; Ezeakunne and Liu, 2023; Guo et al., 2022; Pu et al., 2022b; Wang et al., 2022).


However, a critical challenge has emerged in the form of fairness disparities across demographic groups as described in (Trinh and Liu, 2021; Xu

et al., 2022; Nadimpalli and Rattani, 2022; Masood et al., 2023). Studies have revealed that current detection methods perform inconsistently across different demographics, particularly showing higher accuracy rates for individuals with lighter skin tones compared to those with darker skin tones (Trinh and Liu, 2021; Hazirbas et al., 2021). This bias creates a concerning vulnerability where malicious actors could potentially target specific demographic groups with deepfakes that are more likely to evade detection.

While recent algorithmic approaches have shown promise in improving detection fairness when training and testing data use similar forgery techniques (Ju et al., 2024), the real-world application presents a more complex challenge. Deepfake detectors, typically developed and trained on standard research datasets, must ultimately operate in diverse real-world environments where they encounter deepfakes created using various, potentially unknown forgery techniques. This generalization capability is crucial for practical deployment, particularly on social media platforms where manipulated content proliferates.

<sup>a</sup> <https://orcid.org/0009-0006-5524-5138>

<sup>b</sup> <https://orcid.org/0000-0001-5440-4316>

<sup>c</sup> <https://orcid.org/0000-0002-9320-3872>




Figure 1: Fairness Generalization Comparison: A detector’s accuracy can generalize well to unseen test sets A and B, maintaining consistent detection performance. However, while fairness metrics are preserved on test set B, they fail to generalize to test set A, highlighting the challenge of achieving consistent fairness across different datasets.

While current methods of deepfake detection focus almost entirely on detection accuracy, it is very critical to also consider how well these approaches generalize to fairness. As shown in Figure 1, if a detector was trained on a dataset and its performance for fairness and detection on the training set are shown on the left bars, and the same detector is tested on an unseen dataset (test set A) at the middle, the detection performance, mostly in terms of model accuracy may generalize but fairness may fail to generalize. On the other hand, when tested on another unseen dataset (test set B), it can retain its fairness and detection performance. As can be seen in Figure 1, the detector’s detection performance can be generalized to both datasets but could fail to ensure fairness generalization on test set A.

Our research addresses these challenges through a comprehensive approach that combines data-centric and algorithmic solutions. We propose a novel framework that utilizes self-blended images (SBI) for synthetic data generation to balance demographic representation in the training data, addressing a key source of fairness disparities. This is coupled with a multi-task learning architecture that simultaneously optimizes for both detection accuracy and demographic fairness. The architecture employs an EfficientNet backbone for feature extraction, with separate heads for deepfake detection and demographic classification.

To enhance generalization capabilities, we implement Sharpness-Aware Minimization (SAM) optimization (Lin et al., 2024), which helps find more robust solutions by flattening the loss landscape. Our framework also incorporates a carefully designed loss function that balances classification accuracy, demographic prediction, and fairness constraints. This approach represents a significant advancement in creating deepfake detection systems that are not only accurate but also fair and generalizable across different

demographic groups and forgery techniques.

As demonstrated in our experimental results, while traditional detectors may maintain detection accuracy across different test sets, they often fail to preserve fairness when encountering new data. Our method addresses this limitation by specifically optimizing for both performance metrics, ensuring consistent fairness across different demographic groups even when tested on previously unseen datasets.

Therefore, our contributions are as follows:

- We propose synthetic data balancing, to balance training datasets across demographics to aid fair learning in deepfake detection.
- Our framework introduces a multi-task architecture, leveraging two classification heads; one for deepfake detection and another for learning the demographic features of the datasets. This enables the system to not only detect forgery but also ensure that it is aware of demographic biases, making it more equitable across different groups.
- The multi-task learning approach, combined with sharpness-aware loss optimization and robust feature extraction, makes the system more effective at generalizing the fairness and detection performances to unseen datasets.
- We performed intra-dataset and cross-dataset evaluations to show the improved performance of this method.

## 2 RELATED WORK

### 2.1 Deepfake Detection

To train deepfake detection models that perform well during training and in practice; on unseen datasets, several studies (Li et al., 2020a; Zhao et al., 2021b; Guan et al., 2022; Li, 2018) have introduced novel methods for manually synthesizing a variety of face forgeries similar to deepfakes. These methods help deepfake detection models learn more generalized representations of artifacts, improving their ability to identify deepfakes across different scenarios.

The authors in (Chen et al., 2022) proposed enhancing the diversity of forgeries through adversarial training to improve the robustness of deepfake detectors in recognizing various forgeries. FaceCutout (Das et al., 2021) uses facial landmarks and randomly cuts out different parts of the face (such as the mouth, eye, etc.) to improve the robustness of deepfake detection. Face-Xray (Li et al., 2020a) on the other hand, involves generating blended images

(BI) by combining two different faces using a global transformation and then training a model to distinguish between real and blended faces. The authors in (Shiohara and Yamasaki, 2022) expanded on this concept by creating a synthetic training dataset with self-blended images (SBI), which are created by applying a series of data augmentation techniques to real images. SBI has demonstrated even better generalization to previously unseen deepfakes compared to Face-Xray (Li et al., 2020a). After synthesizing forged images (i.e., pseudo deepfakes) using specialized augmentation techniques, they trained a binary classification model to perform the detection task.

## 2.2 Fairness in Deepfake Detection

Despite efforts to enhance the generalization of deepfake detection to unseen data, limited progress has been made in mitigating biased performance when testing on both the same (intra-domain) and different (cross-domain) dataset distributions.

Recent studies have highlighted significant fairness issues in deepfake detection, revealing biases in both datasets and detection models (Masood et al., 2023). The work done by (Trinh and Liu, 2021) and (Hazirbas et al., 2021) discovered substantial error rate differences across demographic groups, while (Pu et al., 2022a) found that the MesoInception-4 model by (Afchar et al., 2018) was biased against both genders. The authors in (Xu et al., 2022) advocated for diverse dataset annotations to address these biases, and (Nadimpalli and Rattani, 2022) introduced a gender-balanced dataset to reduce gender-based performance bias. However, these techniques led to only modest improvements and required extensive data annotation. Furthermore, (Ju et al., 2024) worked to improve fairness within the same dataset (intra-domain), but did not address fairness generalization between different datasets (cross-domain).

## 2.3 Fairness Generalization in Deepfake Detection

Very little work has been done to achieve fairness generation in deepfake detection. The recent work done in (Lin et al., 2024) proposed a novel feature-based technique to achieve fairness generalization and preservation. They combine disentanglement learning, fairness learning, and loss optimization to preserve fairness in deepfake detection. In the disentanglement learning module, they utilized a disentanglement loss to expose demographic and domain-agnostic forgery features. While the fairness learning module fused the two features from the disentan-

glement learning module to obtain predictions for the samples. Their framework, however, was trained on a dataset with imbalanced demographic groups which introduced some bias to their model.

Our framework approaches and solves fairness generalization using a data-driven approach. Our research prioritizes fairness and eliminates bias due to dataset imbalance by using synthetic datasets to balance the dataset distribution across the different groups. We also used a multi-task learning approach, combined with fairness and loss sharpness-aware optimization and robust feature extraction to achieve significant performance improvements on both intra-dataset evaluation and cross-dataset evaluation.

## 3 PROPOSED METHOD

Our method aims to enhance both the generalization and fairness of deepfake detection models through a data-centric strategy. This section provides a detailed overview of the problem formulation, followed by the process of generating our synthetic dataset. We then outline our approach for addressing dataset imbalance and conclude with a discussion of our multi-task learning framework, designed to simultaneously optimize for accuracy and fairness across demographic groups.

### 3.1 Problem Setup

Given an input face image  $I$ , the goal is to train a deepfake detection model  $g$  that classifies whether  $I$  is a deepfake (fake) or a real face. The model  $g$  is defined as:

$$g : I \rightarrow \{0, 1\} \quad (1)$$

where  $g(I)$  indicates the predicted label, with  $g(I) = 1$  denoting a deepfake and  $g(I) = 0$  denoting a real face.

*Detection Accuracy:* The primary goal is to maximize the detection accuracy. Given a set of face images  $\{I_i\}_{i=1}^N$  with true labels  $\{y_i\}_{i=1}^N$ , the overall accuracy of the model  $g$  is:

$$\text{Accuracy} = \frac{1}{N} \sum_{i=1}^N \mathbf{1}(g(I_i) = y_i) \quad (2)$$

where  $\mathbf{1}(\cdot)$  is the indicator function that equals 1 if the prediction matches the true label.

*Fairness Constraint:* The detector should perform fairly across different demographic groups. Let  $D$  represent the demographic group associated with

an image, and  $\mathcal{D}_k$  be the set of images from the demographic group  $k$ . Define  $\text{Accuracy}_k$  as the accuracy of the detector on images from demographic group  $k$ :

$$\text{Accuracy}_k = \frac{1}{|\mathcal{D}_k|} \sum_{i \in \mathcal{D}_k} \mathbf{1}(g(I_i) = y_i) \quad (3)$$

The fairness of the detector is assessed by minimizing the maximum disparity in accuracy across different demographic groups:

$$\text{Fairness} = \max_{k,l} |\text{Accuracy}_k - \text{Accuracy}_l| \quad (4)$$

where  $k$  and  $l$  denote different demographic groups. The objective is to minimize this fairness disparity.

### 3.2 Synthetic Data Balancing

To enhance fairness in deepfake detection, we utilize a technique based on self-blended images (SBI), as proposed in (Shiohara and Yamasaki, 2022). This technique involves generating synthetic images to balance the demographic distribution of the dataset, thereby improving fairness across various demographic groups. The dataset samples are shown in Figure 2. Each real sample (upper row) has a corresponding fake sample (bottom row), so the dataset is also balanced across real and fake classes. The process of data balancing using a synthetic dataset is detailed in the sections below.

#### Generation of Synthetic Images

Let  $I = \{I_1, I_2, \dots, I_n\}$  represent the set of real face images, where each  $I_i$  belongs to a specific demographic group. The goal is to generate a set of self-blended images  $\mathcal{S} = \{S_1, S_2, \dots, S_m\}$  using the following steps:

**Base Image Selection:** Select a diverse subset of real face images,  $I_{base} \subset I$ , ensuring coverage across demographic categories.

**Augmentation Process:** Apply a series of transformations to generate self-blended images. For a base image  $I_i$ , the synthetic image  $S_j$  is generated using a blend function  $B$  as follows:

$$S_j = B(I_i, T_k(I_i)) \quad (5)$$

where  $T_k$  represents a transformation operation such as scaling, rotation, or color adjustment, and  $B$  is the blending function that combines  $I_i$  with  $T_k(I_i)$ .

### Dataset Balancing

To balance the dataset, we increase the representation of underrepresented demographic groups:

**Demographic Representation:** Define the demographic groups  $\mathcal{D} = \{D_1, D_2, \dots, D_r\}$ . Let  $I_{D_k} \subset I$  be the set of images belonging to demographic group  $D_k$ . To balance the dataset, we create synthetic images  $S_{D_k}$  for each group  $D_k$ :

$$\mathcal{B} = I \cup \mathcal{S} \quad (6)$$

where  $\mathcal{B}$  is the balanced dataset and  $\mathcal{S} = \bigcup_{k=1}^r S_{D_k}$ .

This dataset  $\mathcal{B}$ , is balanced across all the demographics, and also equal in terms of the number of real and fake samples.

### 3.3 Multi-Task Learning

We propose a dual-task learning framework to enhance the fairness of deepfake detection models by leveraging demographic information. The architecture shown in Figure 3, is designed to improve the generalization of deepfake detection across various demographic groups, reducing bias in prediction accuracy between these groups.

The input images are passed through the EfficientNet encoder (Tan, 2019), which extracts high-dimensional feature representations. These features capture both low-level information (like textures and edges) as well as high-level semantic features (such as facial structure and potential tampering artifacts). The extracted features serve as the input for the real/fake classification head and the demographic head.

**Real/Fake Classification Head:** This head consists of a Multi-Layer Perceptron (MLP) that processes the extracted features and predicts whether the input image is real or fake. The output of this MLP is a probability score indicating the likelihood that the image is a deepfake or a genuine one. This branch focuses on learning the subtle cues and artifacts associated with the deepfake generation, such as inconsistencies in lighting, facial details, or blurring artifacts around the face.

**Demographic Head:** In parallel to the real/fake classification, the features are passed through a second MLP for demographic group classification. This MLP is designed to categorize the input into one of eight demographic groups, such as Black-Male (B-M), White-Female (W-F), Asian-Male (A-M), etc. The purpose of this head is to ensure that the model learns demographic-specific features that are independent of deepfake artifacts. These features are crucial for enhancing the fairness of the model across diverse demographic groups.



Figure 2: Sample of Facial Images. The real images (top row) and their corresponding fake/synthetic image (bottom row).

## Loss Function

We define the loss function to incorporate both the standard classification loss and a fairness penalty. Specifically, the total loss function  $\mathcal{L}$  is composed of three terms: the classification loss, the demographic loss, and the fairness penalty.

**Classification Loss:** Let  $\mathbf{p} = [p_1, p_2, \dots, p_N]$  be the predicted probabilities and  $\mathbf{y} = [y_1, y_2, \dots, y_N]$  be the true binary labels. The binary cross-entropy loss  $\mathcal{L}_{\text{real}}$  is defined as:

$$\mathcal{L}_{\text{real}} = -\frac{1}{N} \sum_{i=1}^N [y_i \log(p_i) + (1 - y_i) \log(1 - p_i)] \quad (7)$$

where  $N$  is the number of samples,  $y_i$  is the true binary label (0 or 1) for the  $i$ -th sample, and  $p_i$  is the predicted probability for the positive class.

**Demographic Loss:** Let  $\mathbf{p}_{\text{dem},i} = [p_{\text{dem},i,1}, \dots, p_{\text{dem},i,8}]$  represent the predicted probabilities for the  $i$ -th sample across 8 demographic groups, and  $\mathbf{y}_{\text{dem},i} = [y_{\text{dem},i,1}, \dots, y_{\text{dem},i,8}]$  be the one-hot encoded true labels for those groups. The demographic loss  $\mathcal{L}_{\text{dem}}$  is:

$$\mathcal{L}_{\text{dem}} = -\frac{1}{N_{\text{dem}}} \sum_{i=1}^{N_{\text{dem}}} \sum_{c=1}^8 y_{\text{dem},i,c} \log(p_{\text{dem},i,c}) \quad (8)$$

where  $N_{\text{dem}}$  is the number of samples in the demographic component, and the summation runs over the 8 demographic groups.

**Fairness Penalty:** To quantify fairness, we compute the variance of the accuracy rates across different demographic groups. Let  $\mathbf{a}_k$  denote the accuracy for demographic group  $k$ . The variance of accuracies  $\text{Var}_{\text{acc}}$  is:

$$\text{Var}_{\text{acc}} = \frac{1}{K} \sum_{k=1}^K (\text{Accuracy}_k - \text{Accuracy})^2 \quad (9)$$

where  $K$  is the number of demographic groups,  $\text{Accuracy}_k$  is the accuracy for group  $k$ , and  $\text{Accuracy}$  is the mean accuracy across all groups.

**Total Loss:** The total loss function  $\mathcal{L}$  is a combination of the standard loss, the fairness penalty, and the fairness loss. The total loss is defined as:

$$\mathcal{L} = \mathcal{L}_{\text{real}} + \lambda \cdot \text{Var}_{\text{acc}} + \mathcal{L}_{\text{dem}} \quad (10)$$

where  $\lambda$  is a hyperparameter controlling the weight of the fairness penalty.

## Optimization

We employ the Sharpness-Aware Minimization (SAM) optimizer (Foret et al., 2020) as our optimizer. SAM is designed to improve the generalization performance of deep neural networks by encouraging the model to find parameter spaces that are not only good at minimizing the training loss but also robust to small perturbations in the weight space. This results in a flatter and smoother loss landscape, which has been shown to correlate with better generalization.

The SAM optimizer modifies the traditional gradient-based update by introducing a two-step process that seeks to find parameters that minimize both the loss and its sensitivity to perturbations. Let  $\mathcal{L}(\mathbf{w})$  denote the total loss function of the model, where  $\mathbf{w}$  represents the network's parameters. SAM performs the following steps:

**Perturbation Step:** SAM first finds a perturbation  $\boldsymbol{\varepsilon}$  that maximizes the loss function around the current parameters. The perturbation  $\boldsymbol{\varepsilon}$  is chosen to make the loss worse within a neighborhood of the current parameter values:

$$\boldsymbol{\varepsilon}(\mathbf{w}) = \arg \max_{\|\boldsymbol{\varepsilon}\| \leq \rho} \mathcal{L}(\mathbf{w} + \boldsymbol{\varepsilon}) \quad (11)$$

where  $\rho$  controls the size of the neighborhood in which the optimizer searches for the sharpest directions.

**Parameter Update Step:** Once the worst-case perturbation  $\boldsymbol{\varepsilon}$  is found, the SAM optimizer updates the parameters by minimizing the loss function at  $\mathbf{w} + \boldsymbol{\varepsilon}$ :

$$\mathbf{w} \leftarrow \mathbf{w} - \eta \nabla \mathcal{L}(\mathbf{w} + \boldsymbol{\varepsilon}) \quad (12)$$

where  $\eta$  is the learning rate.

This two-step process leads to solutions in flatter regions of the loss surface, improving the model's generalization and robustness.

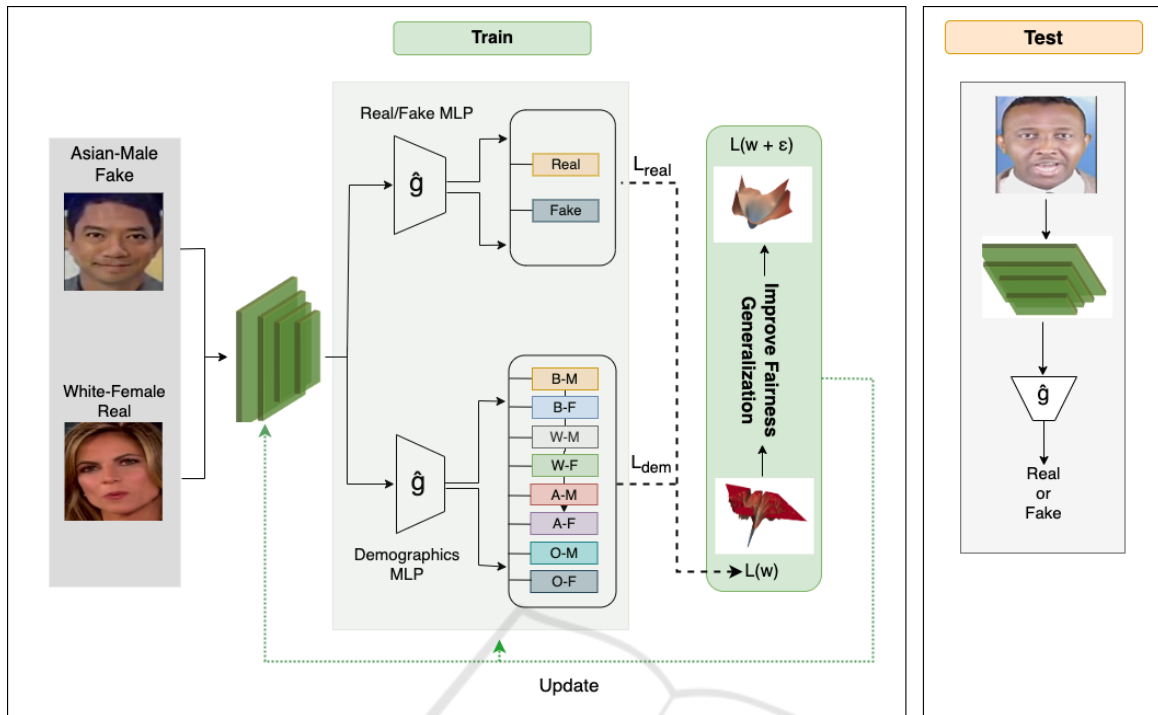


Figure 3: An overview of our proposed method. We utilize EfficientNet (Tan, 2019) to extract deep features from input images for the feature extraction module. For the classification module, two heads are used: the real/fake head predicts whether the input is real or fake, and the demographic head predicts the demographic group. The demographic classification head outputs probabilities for one of eight demographic categories based on both gender and ethnicity. These categories include Black-Male (B-M), Black-Female (B-F), White-Male (W-M), White-Female (W-F), Asian-Male (A-M), Asian-Female (A-F), Other-Male (O-M), and Other-Female (O-F). We use SAM (Sharpness-Aware Minimization) for the optimization module to flatten the loss landscape and enhance fairness generalization across demographic groups.

## 4 EXPERIMENTS

### 4.1 Experimental Settings

#### Dataset

To evaluate the fairness generalization capability of our proposed approach, we train our model on the widely used FaceForensics++ (FF++) dataset (Rossler et al., 2019). For testing, we use FF++, Deepfake Detection Challenge (DFDC) (Dolhansky et al., 2020), and Celeb-DF (Li et al., 2020b). We used the real images from these datasets and created synthetic images as the fake images using the method in Section 3.2.

Since the original datasets do not include demographic information, we follow the approach in (Ju et al., 2024) for data processing and annotation. The datasets are categorized by race and gender, forming the following intersection groups: Black-Male (B-M), Black-Female (B-F), White-Male (W-M), White-Female (W-F), Asian-Male (A-M), Asian-Female (A-F), Other-Male (O-M), and Other-Female (O-F). The

Celeb-DF (Li et al., 2020b) does not contain the Asian sub-group.

#### Evaluation Metrics

For evaluating our approach, we use several performance metrics, including the Area Under the Curve (AUC), Accuracy, and True Positive Rate (TPR). These metrics allow us to benchmark our method against previous works. To assess the degree of fairness of the model, we analyze the accuracy gaps across different races and genders.

#### Baseline Methods

For intra-dataset evaluation, we compare our approach against recent works in deepfake detection fairness, specifically DAW-FDD (2023) (Ju et al., 2024) and Lin et al. (2024) (Lin et al., 2024). For cross-dataset evaluation, our method is compared to the fairness generalization method proposed by Lin et al. (2024) (Lin et al., 2024). To the best of our knowledge, Lin et al. (2024) (Lin et al., 2024) is the prior




Figure 4: Intra-dataset evaluation across different methods. Methods are trained and tested on the same distribution as the training set (FF++ (Rossler et al., 2019)): Performance comparison across gender. See Table 1 for exact numerical values.

work addressing the fairness generalization problem. We also used the Xception (Chollet, 2017) as a baseline for comparison during both intra-dataset evaluation and cross-dataset evaluation. The Xception network (Chollet, 2017) was implemented using transfer learning to perform the deepfake detection task.

### Implementation Details

All experiments are conducted using PyTorch. For our method, we set the batch size to 16 and train the model for 100 epochs. To incorporate fairness, we use a fairness penalty weight  $\lambda$  of 20 (See Eq. 10). For optimization, we employ the SAM optimizer (Foret et al., 2020). The learning rate is set to  $5e-4$ , momentum is configured at 0.9, and weight decay is set to  $5 \times 10^{-3}$ .

## 4.2 Experimental Results

In this section, we discuss the results of our evaluations and present our findings. We start by discussing the results of our intra-dataset evaluations and then, the results of the cross-dataset evaluations.

### Intra-Dataset Evaluation Results

We perform intra-dataset evaluations to measure the model's performance on the same dataset distribution, assessing its ability to fit the training data. As shown in Table 1, our method achieves comparable performance to baseline approaches in deepfake detection. However, it demonstrates a significant improvement over baseline methods in fairness preservation across demographic groups. Figure 4 presents the evaluation results for the gender demographic group, while Figure 5 illustrates the results for the racial demographic group. As evidenced in these figures, the disparity between the Male and Female subgroups is substantially reduced in our approach with a 0.12% difference (in accuracy) between the subgroups, compared




Figure 5: Intra-dataset evaluation across different methods. Methods are trained and tested on the same distribution as the training set (FF++ (Rossler et al., 2019)): Performance comparison across race. See Table 1 for exact numerical values.

to the 1.58%, 2.90% and 3.87% differences between the subgroups for the baseline approaches. Additionally, the disparities among racial groups are significantly lower with our approach with a 0.71% difference (which happens to be between Asian and Black - considering minimum and maximum accuracies) between the subgroups, while the baselines had differences of 5.28%, 1.64% and 1.52%. These results highlight the effectiveness and superiority of our approach in achieving fairness generalization.

### 4.2.1 Cross-Dataset Evaluation Results

We conducted cross-dataset evaluations to assess our model's performance on out-of-distribution data, evaluating its ability to generalize to unseen datasets. As recorded in Table 2, and visualized in Figures 6,7,8 and 9, our approach achieves comparable performance to baseline methods in deepfake detection while significantly outperforming them in fairness generalization across demographic groups.

Figures 6 and 7 visualize the results of evaluations from Table 2 where the model was trained on FF++ and tested on DFDC, while Figures 8 and 9 visualizes the results from Table 2 for training on FF++ and testing on Celeb-DF. These figures demonstrate that our approach exhibits minimal performance disparities across different gender and racial groups compared to the baselines.

As shown in Table 2, on the DFDC dataset, the Xception baseline has the lowest disparity in accuracy across genders. The disparity in accuracy between males and females is 1.85% for Xception, and then 4.5% (Lin et al., 2024) and 8.03% (Ours). Across race on the DFDC dataset, our method has the lowest accuracy disparity across the races at 11.15% (which happens to be between Asian and white) with the baseline recording accuracy disparities of 24.53% and 14.39%.

As shown in Table 2, on the Celeb-DF dataset, our method has a minimal disparity in accuracy across both gender and race. Across gender, we got 15.47%

Table 1: Intra-dataset evaluation compares different deepfake detection approaches by training and testing on the same dataset’s distribution. Results are computed on each dataset’s test set, with the best performance shown in bold.

			FF++ (Rossler et al., 2019)			
			Xception (Chollet, 2017)	DAW-FDD (Ju et al., 2024)	Lin et al. (Lin et al., 2024)	Ours
Gender	Male	TPR (%)	<b>98.62</b>	93.48	92.13	91.91
		AUC (%)	90.23	96.91	97.56	<b>97.71</b>
		ACC (%)	87.03	<b>92.65</b>	92.07	92.39
	Female	TPR (%)	<b>98.85</b>	97.15	98.02	91.89
		AUC (%)	93.48	98.05	<b>98.94</b>	97.72
		ACC (%)	88.61	95.55	<b>95.94</b>	92.27
Race	Black	TPR (%)	<b>99.59</b>	96.36	97.17	91.27
		AUC (%)	95.85	97.92	<b>98.22</b>	97.57
		ACC (%)	89.00	<b>94.67</b>	94.66	92.04
	White	TPR (%)	<b>98.71</b>	95.47	94.81	91.88
		AUC (%)	92.57	97.54	<b>98.21</b>	97.70
		ACC (%)	88.84	<b>94.29</b>	94.09	92.28
	Asian	TPR (%)	<b>98.73</b>	94.40	96.44	92.71
		AUC (%)	89.40	96.39	97.54	<b>97.91</b>
		ACC (%)	83.72	93.04	<b>93.46</b>	92.75
	Other	TPR (%)	<b>98.51</b>	95.77	95.83	91.73
		AUC (%)	89.87	98.23	<b>98.58</b>	97.66
		ACC (%)	84.94	94.68	<b>94.98</b>	92.25




Figure 6: Cross-dataset evaluation across different methods. Methods are trained on (FF++ (Rossler et al., 2019)) and tested on DFDC (Dolhansky et al., 2020). Performance comparison is across gender. See Table 2 for exact numerical values.

disparity in accuracy, while the baselines recorded wider disparities of 40.66% and 25.47% in accuracy. Across race, our method recorded 11.54% as the maximum disparity in accuracy (which happens to be between white and black), the baselines recorded wider disparities of 26.41% and 14.39%.

These results highlight the effectiveness of our approach in achieving fairness generalization in cross-dataset scenarios, demonstrating improved consistency across both racial and gender demographics compared to baseline methods.

## 5 LIMITATIONS

The limitations of our method are:




Figure 7: Cross-dataset evaluation across different methods. Methods are trained on (FF++ (Rossler et al., 2019)) and tested on DFDC (Dolhansky et al., 2020). Performance comparison is across race. See Table 2 for exact numerical values.

- **Dependence on Demographically Annotated Datasets:** The effectiveness of our method relies on the availability of detailed demographic annotations. Such datasets are often difficult to obtain and may not always be comprehensive enough, potentially impacting the fairness assessment.
- **Trade-Off Between Fairness and Detection Performance:** Improving fairness across demographic groups can result in a trade-off with overall detection performance. Enhancements aimed at reducing demographic disparities may lead to a decrease in the model’s overall accuracy in detecting deepfakes.

These limitations highlight the need for further research to address the challenges of dataset dependency and the balance between fairness and detection performance.



Table 2: Cross-dataset evaluation across different methods for deepfake detection. The methods were trained on (FF++ (Rossler et al., 2019)) and tested on DFDC (Dolhansky et al., 2020) and Celeb-DF (Li et al., 2020b).

			DFDC (Dolhansky et al., 2020)			Celeb-DF (Li et al., 2020b)		
			Xception (Chollet, 2017)	Lin et al. (Lin et al., 2024)	Ours	Xception (Chollet, 2017)	Lin et al. (Lin et al., 2024)	Ours
Gender	Male	TPR (%)	<b>85.11</b>	37.81	35.49	<b>98.68</b>	76.89	78.38
		AUC (%)	59.97	60.08	<b>63.30</b>	36.97	66.93	<b>75.44</b>
		ACC (%)	43.75	63.11	<b>66.98</b>	43.29	56.16	<b>66.41</b>
	Female	TPR (%)	<b>77.59</b>	44.36	40.57	<b>99.24</b>	90.19	86.67
		AUC (%)	50.70	<b>60.06</b>	59.12	53.26	75.83	<b>82.09</b>
		ACC (%)	45.60	58.61	<b>58.98</b>	<b>83.95</b>	81.63	81.88
Race	Black	TPR (%)	<b>76.28</b>	45.14	57.58	<b>98.68</b>	76.89	78.38
		AUC (%)	50.40	59.81	<b>68.83</b>	47.63	67.03	<b>72.53</b>
		ACC (%)	39.23	58.66	<b>65.54</b>	63.99	64.10	<b>69.61</b>
	White	TPR (%)	<b>80.82</b>	38.18	32.98	<b>99.24</b>	90.18	86.69
		AUC (%)	56.45	60.14	<b>60.18</b>	53.58	77.02	<b>82.20</b>
		ACC (%)	48.79	59.80	<b>60.24</b>	<b>81.38</b>	80.43	81.15
	Asian	TPR (%)	<b>82.50</b>	42.50	19.16	-	-	-
		AUC (%)	54.75	<b>51.78</b>	43.89	-	-	-
		ACC (%)	22.38	57.04	<b>71.39</b>	-	-	-
	Others	TPR (%)	<b>87.68</b>	55.93	36.18	<b>100</b>	96.66	73.33
		AUC (%)	60.39	<b>72.52</b>	65.38	28.31	64.53	<b>84.97</b>
		ACC (%)	46.91	<b>71.43</b>	65.82	17.22	32.78	<b>70.00</b>




Figure 8: Cross-dataset evaluation across different methods. Methods are trained on (FF++ (Rossler et al., 2019)) and tested on Celeb-DF (Li et al., 2020b). Performance comparison is across gender. See Table 2 for exact numerical values.




Figure 9: Cross-dataset evaluation across different methods. Methods are trained on (FF++ (Rossler et al., 2019)) and tested on Celeb-DF (Li et al., 2020b). Performance comparison is across races. See Table 2 for exact numerical values.

## 6 CONCLUSION

We have presented a novel approach to address the critical challenge of fairness generalization in deepfake detection systems. Our method combines three key innovations: a data-centric strategy using synthetic image generation to balance demographic representation, a multi-task learning architecture that simultaneously optimizes detection accuracy and demographic fairness, and sharpness-aware optimization to enhance generalization capabilities. Our comprehensive evaluations demonstrate the effectiveness of this approach. In intra-dataset testing, our method achieved comparable detection performance to existing approaches while significantly reducing demographic disparities - showing only 0.12% accuracy difference between gender groups compared to up

to 3.87% in baseline methods, and 0.71% difference across racial groups compared to up to 5.28% in baselines. More importantly, in challenging cross-dataset scenarios, our approach demonstrated superior fairness generalization. When tested on the Celeb-DF dataset, our method reduced gender-based accuracy disparities to 15.47% compared to up to 40.66% in baselines, while maintaining strong overall detection performance. These results suggest that our integrated approach of balanced synthetic data, demographic-aware learning, and robust optimization provides a promising direction for developing deepfake detection systems that are both accurate and demographically fair. Future work could explore extending these techniques to other domains where algorithmic fairness and generalization are crucial concerns.

## REFERENCES

- Afchar, D., Nozick, V., Yamagishi, J., and Echizen, I. (2018). Mesonet: a compact facial video forgery detection network. In *2018 IEEE international workshop on information forensics and security (WIFS)*, pages 1–7. IEEE.
- Chen, L., Zhang, Y., Song, Y., Liu, L., and Wang, J. (2022). Self-supervised learning of adversarial example: Towards good generalizations for deepfake detection.
- Chollet, F. (2017). Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1251–1258.
- Das, S., Seferbekov, S., Datta, A., Islam, M. S., and Amin, M. R. (2021). Towards solving the deepfake problem: An analysis on improving deepfake detection using dynamic face augmentation.
- Dolhansky, B., Bitton, J., Pflaum, B., Lu, J., Howes, R., Wang, M., and Ferrer, C. C. (2020). The deepfake detection challenge (dfdc) dataset. *arXiv preprint arXiv:2006.07397*.
- Ezeakunne, U. and Liu, X. (2023). Facial deepfake detection using gaussian processes.
- Foret, P., Kleiner, A., Mobahi, H., and Neyshabur, B. (2020). Sharpness-aware minimization for efficiently improving generalization. *arXiv preprint arXiv:2010.01412*.
- Guan, J., Zhou, H., Gong, M., Ding, E., Wang, J., and Zhao, Y. (2022). Detecting deepfake by creating spatio-temporal regularity disruption.
- Guo, H., Hu, S., Wang, X., Chang, M.-C., and Lyu, S. (2022). Robust attentive deep neural network for detecting gan-generated faces.
- Hazirbas, C., Bitton, J., Dolhansky, B., Pan, J., Gordo, A., and Ferrer, C. C. (2021). Towards measuring fairness in ai: the casual conversations dataset.
- Ju, Y., Hu, S., Jia, S., Chen, G. H., and Lyu, S. (2024). Improving fairness in deepfake detection.
- Li, L., Bao, J., Zhang, T., Yang, H., Chen, D., Wen, F., and Guo, B. (2020a). Face x-ray for more general face forgery detection.
- Li, Y. (2018). Exposing deepfake videos by detecting face warping artif acts.
- Li, Y., Yang, X., Sun, P., Qi, H., and Lyu, S. (2020b). Celeb-df: A large-scale challenging dataset for deepfake forensics. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3207–3216.
- Lin, L., He, X., Ju, Y., Wang, X., Ding, F., and Hu, S. (2024). Preserving fairness generalization in deepfake detection.
- Masood, M., Nawaz, M., Malik, K. M., Javed, A., Irtaza, A., and Malik, H. (2023). Deepfakes generation and detection: State-of-the-art, open challenges, countermeasures, and way forward.
- Nadimpalli, A. V. and Rattani, A. (2022). Gbdf: gender balanced deepfake dataset towards fair deepfake detection.
- Pu, M., Kuan, M. Y., Lim, N. T., Chong, C. Y., and Lim, M. K. (2022a). Fairness evaluation in deepfake detection models using metamorphic testing.
- Pu, W., Hu, J., Wang, X., Li, Y., Hu, S., Zhu, B., Song, R., Song, Q., Wu, X., and Lyu, S. (2022b). Learning a deep dual-level network for robust deepfake detection.
- Rossler, A., Cozzolino, D., Verdoliva, L., Riess, C., Thies, J., and Nießner, M. (2019). Faceforensics++: Learning to detect manipulated facial images. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1–11.
- Shiohara, K. and Yamasaki, T. (2022). Detecting deepfakes with self-blended images.
- Tan, M. (2019). Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*.
- Trinh, L. and Liu, Y. (2021). An examination of fairness of ai models for deepfake detection.
- Wang, J., Wu, Z., Ouyang, W., Han, X., Chen, J., Jiang, Y.-G., and Li, S.-N. (2022). M2tr: Multi-modal multi-scale transformers for deepfake detection.
- Xu, Y., Raja, K., Verdoliva, L., and Pedersen, M. (2023). Learning pairwise interaction for generalizable deepfake detection.
- Xu, Y., Terhörst, P., Raja, K., and Pedersen, M. (2022). A comprehensive analysis of ai biases in deepfake detection with massively annotated databases.
- Zhao, H., Zhou, W., Chen, D., Wei, T., Zhang, W., and Yu, N. (2021a). Multi-attentional deepfake detection.
- Zhao, T., Xu, X., Xu, M., Ding, H., Xiong, Y., and Xia, W. (2021b). Learning self-consistency for deepfake detection.