

An Examination of Fairness of AI Models for Deepfake Detection

Loc Trinh*, Yan Liu

Department of Computer Science, University of Southern California

{loctrinh, yanliu.cs}@usc.edu

Abstract

Recent studies have demonstrated that deep learning models can discriminate based on protected classes like race and gender. In this work, we evaluate bias present in deepfake datasets and detection models across protected subgroups. Using facial datasets balanced by race and gender, we examine three popular deepfake detectors and find large disparities in predictive performances across races, with up to 10.7% difference in error rate between subgroups. A closer look reveals that the widely used FaceForensics++ dataset is overwhelmingly composed of Caucasian subjects, with the majority being female Caucasians. Our investigation of the racial distribution of deepfakes reveals that the methods used to create deepfakes as positive training signals tend to produce “irregular” faces - when a person’s face is swapped onto another person of a different race or gender. This causes detectors to learn spurious correlations between the foreground faces and fakeness. Moreover, when detectors are trained with the Blended Image (BI) dataset from Face X-Rays, we find that those detectors develop systematic discrimination towards certain racial subgroups, primarily female Asians.

1 Introduction

Synthetic media have become so realistic with the advancement of deep neural networks that they are often indiscernible from authentic content. However, synthetic media designed to deceive poses a dangerous threat to many communities around the world [Cahlan, 2020; Ingram, 2019]. In this context, *Deepfake* videos – which portray human subjects with altered identities or malicious/embarrassing actions - have emerged as a vehicle for misinformation. With the current advancement and growing availability of computing resources, sophisticated deepfakes have become more pervasive, especially to generate revenge pornography [Hao, 2019] and defame celebrities or political targets [Vaccari and Chadwick, 2020]. Hence, there is a critical need for automated systems that can effectively combat misinformation on the internet.

To address this challenge, the vision community has conducted a series of excellent works on detecting deepfakes [Tolosana *et al.*, 2020; Mirsky and Lee, 2021]. Sophisticated facial forgery detection tools [Afchar *et al.*, 2018; Li *et al.*, 2020; Liu *et al.*, 2020] and advanced training sets [Rössler *et al.*, 2019; Jiang *et al.*, 2020] were developed to train detectors capable of identifying deepfakes with high precision. Such results have also seen in real-world impact with Microsoft’s release of Video Authenticator [Burt and Horvitz, 2020], an automated tool trained on the publicly available FaceForensics++ dataset, to analyze a still photo or video to provide a percentage chance that the media is artificially manipulated. It works by detecting the blending boundary of the deepfake and subtle fading or grayscale elements that might not be detectable by the human eye. On the other hand, Facebook has also been pioneering its own system to detect AI-generated profiles and ban hundreds of fake accounts, pages, posts, and social groups¹, along with strengthening its policy on deepfakes and authentic media².

While these works have achieved good progress towards the prediction task, detecting fake videos at a low false-positive rate is still a challenging problem [Li *et al.*, 2020]. Moreover, since most studies focus on the visual artifacts existing within deepfakes, little is discussed about how such systems perform on diverse groups of real people across gender and race, which is the common setting where personal profiles and videos are being audited en masse for authenticity via automated systems. In this context, a small percentage difference in false-positive rates between subgroups would indicate that millions of people of a particular group are more likely to be mistakenly classified as fake.

This draws a connection to fairness in machine learning, where growing concerns about unintended consequences from biased or flawed systems call for a careful and thorough examination of both datasets and models. Gender Shades [Buolamwini and Gebru, 2018] have demonstrated how facial recognition system discriminates across gender and race, showing a large gap in the accuracy of gender classifiers across different intersectional groups: darker-skinned

*Contact Author

¹<https://about.fb.com/news/2019/12/removing-coordinated-inauthentic-behavior-from-georgia-vietnam-and-the-us/>

²<https://about.fb.com/news/2020/01/enforcing-against-manipulated-media/>

females are misclassified in up to 34.7% of cases, while the maximum error rate for lighter-skinned males is only 0.8%. Others have shown that training with biased data has resulted in algorithmic discrimination [Bolukbasi *et al.*, 2016]. Although many works have studied how to create fairer algorithms, and benchmarked discrimination in various contexts [Hardt *et al.*, 2016; Liu *et al.*, 2018], few works have done this analysis for computer vision in the context of synthetic media and deepfakes. Our contributions are as follows:

1. We find that the FaceForensics++ dataset commonly used for training deepfake detectors is overwhelmingly composed of Caucasian subjects, with the majority (36.6%) of videos features female Caucasian subjects.
2. We also find that approaches to generate fake samples as *positive* training signals tend to overwhelmingly produce “irregular” deepfakes - when a person’s face is swapped onto another person of a different race or gender, which leads to detectors learning spurious correlations between foreground faces and *fakeness*.
3. Using facial datasets balanced by gender and race, we find that classifiers designed to detect deepfakes have large predictive disparities across racial groups, with up to 10.7% difference in error rate.
4. Lastly, we observe that when detectors are trained with the Blended Images (BI) from Face X-Rays [Li *et al.*, 2020], we find that detectors develops systematic discrimination towards female Asian subjects.

2 Related Work

2.1 Deepfake Detection

Early deepfake forensic work focused on hand-crafted facial features such as eye colors, light reflections, and 3D head poses / movements. However, these approaches do not scale well to more advanced GAN-based deepfakes. To combat the new generation of deepfakes, researchers leverage deep learning and convolutional networks to automatically extract meaningful features for face forgery detection [Rössler *et al.*, 2019]. Work on shallow networks such as MesoInception4 [Afchar *et al.*, 2018] and Patch-based CNN [Chai *et al.*, 2020] are developed to focus on low and medium level manipulation artifacts. Deep networks, such as Xception [Rössler *et al.*, 2019], also demonstrated success by achieving state-of-the-arts via fine-tuning on ImageNet. Other lines of research examine resolution-inconsistent facial artifacts DSP-FWA [Li and Lyu, 2018] through spatial pyramid pooling modules, blending artifacts via Face X-ray [Li *et al.*, 2020], or temporal artifacts via dynamic prototypes [Trinh *et al.*, 2020]. FakeSpotter [Wang *et al.*, 2020] uses layer-wise neuron behaviors as features in addition to the output of the final-layer.

2.2 Generalizability and Robustness of Detectors

With more advanced deepfake creations, recent works [Cozzolino *et al.*, 2018; Khodabakhsh *et al.*, 2018] have shown that the performance of current detection models *drops* drastically on new types of facial manipulations. Few work call for a closer investigation into the generalizability of deepfake

detectors towards unseen manipulations. In particular, ForensicTransfer [Cozzolino *et al.*, 2018] proposes an autoencoder-based network to transfer knowledge between different but related manipulations via the hidden latent space. Face X-ray [Li *et al.*, 2020] addressed the problem by focusing on the more general blending artifacts as well as creating a blended image dataset to help networks generalize across unseen manipulations. In addition to generalization, recent work have also demonstrated the vulnerability of deepfake detectors to adversarial attacks [Carlini and Farid, 2020], where small tailored perturbations generated via either black-box or white-box attacks can easily fool the networks. This raises a concern about the robustness and commercial readiness of deepfake detectors. In contrast to complex adversarial attacks, our work examines the performances of deepfake detectors on natural images composing of different gender and diverse racial groups, as well as investigating the real-world consequences if deepfake detectors are commercially adopted.

2.3 Algorithmic Fairness and Consequences

Concerns about malicious applications of AI and unintended consequences from flawed or biased systems have propelled many investigations in studying representational and algorithmic bias. Gender Shades [Buolamwini and Gebru, 2018] have demonstrated how facial recognition system discriminates across gender and race, especially for darker-skinned females. [Liu *et al.*, 2018] showed that common fairness criteria may in fact harm underrepresented or disadvantaged groups due to delayed outcomes. [Celis *et al.*, 2019] proposed a framework to combat echo chambers created by highly personalized recommendations on social media that reinforced people’s biases and opinions. In terms of standardized approaches for the field, [Mitchell *et al.*, 2019] and [Gebru *et al.*, 2018] recommend the usage of model cards and datasheets to better document the intended usage of models and data. Although many works have studied how to create fairer algorithms and benchmarked discrimination in various contexts [Hardt *et al.*, 2016; Liu *et al.*, 2018], we conduct a fairness analysis in the context of deepfake detection, which requires bookkeeping of the racial distribution face swaps and providing subgroup-specific deepfakes for audit.

3 Deepfake Detection

We investigate 3 popular deepfake detection models of various sizes, architectures, and loss formulations, all with proven success in detecting deepfake videos. We trained MesoInception4 [Afchar *et al.*, 2018], Xception [Rössler *et al.*, 2019], and Face X-Ray [Li *et al.*, 2020] on the FaceForensics++ dataset, which contains four variants of face swaps. For a fair comparison, we also cross-test the models’ generalizability across datasets with unknown manipulations not seen in FaceForensics++, such as Google’s DeepfakeDetection, Celeb-DF, and DeeperForensics-1.0. To evaluate generalizability, we measure the area under the receiver operating curve (AUC) and the performance of deepfake detectors at a low false-positive rate threshold, using the standardized partial AUC or pAUC (at 10%FPR). We additionally inspect the Equal Error Rate of the models, similarly to [Li *et al.*, 2020].

Model	Test Dataset			Cross-test Datasets								
	FaceForensics++			DFD			DeeperForensics			Celeb-DF		
	AUC	pAUC	ERR	AUC	pAUC	ERR	AUC	pAUC	ERR	AUC	pAUC	ERR
Meso4	0.953	0.858	0.120	0.828	0.766	0.258	0.822	0.742	0.269	0.729	0.575	0.339
Xception	0.935	0.895	0.145	0.913	0.813	0.166	0.941	0.818	0.124	0.778	0.564	0.291
Meso4 + BI	0.935	0.852	0.145	0.883	0.779	0.204	0.881	0.803	0.184	0.713	0.559	0.329
Xception + BI	0.989	0.985	0.029	0.940	0.806	0.143	0.956	0.878	0.105	0.841	0.663	0.235
FaceXRay + BI	0.992	0.989	0.023	0.932	0.839	0.148	0.946	0.849	0.122	0.798	0.679	0.286

Table 1: Test results on FaceForensics++ and cross-testing results on DeepFakeDetection (DFD), DeeperForensics, and Celeb-DF.

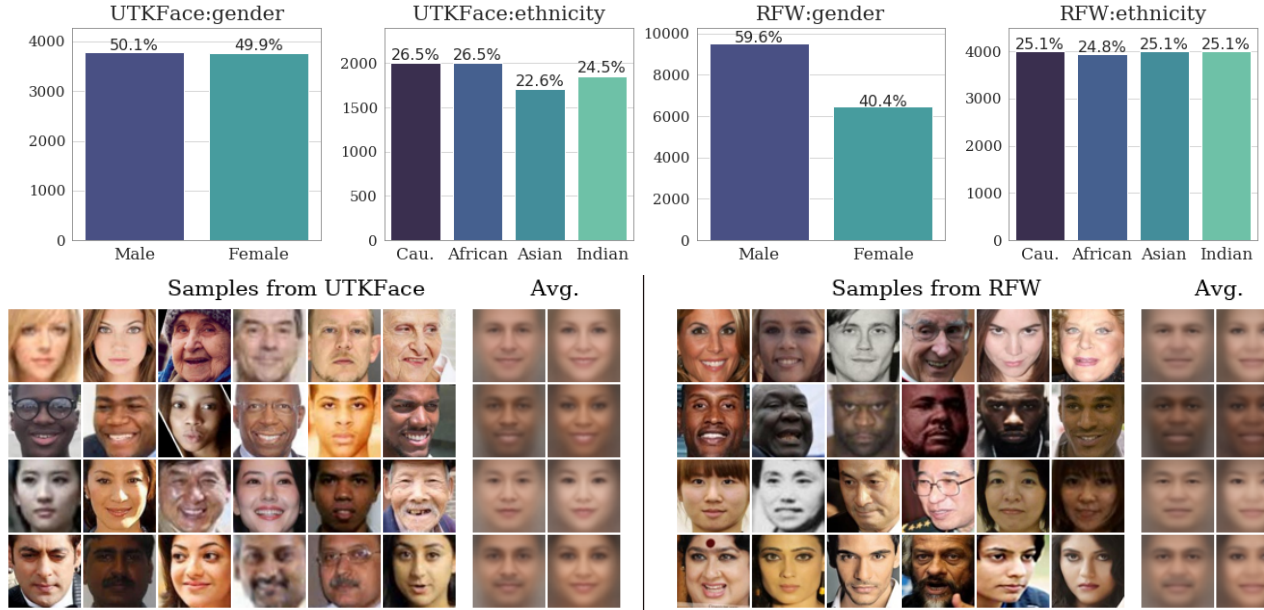


Figure 1: Examples and average faces of both the RFW and UTKFace database, along with their respective gender and racial distributions. In each row from top to bottom: Caucasian, African, Asian, Indian.

Our results match state-of-the-art results, which we then used to audit for fairness. For more detailed information on training, testing, and cross evaluations, see Table 1 and Section A.1 in the Appendix.

4 Deepfake Detection Audit

In this section, we first show the fairness evaluation of multiple popular deepfake detectors across gender groups and intersectional groups. Second, we present an analysis on the representation bias in widely adopted datasets and methodology bias in current state-of-the-art approaches.

We evaluated the deepfake detectors in the previous section, trained using both the FF++ and Blended Image (BI) datasets. Overall, all detectors perform equally on real and deepfake images containing male and female subjects, and all detectors trained with BI perform worst on media with darker African faces. Further analysis of the intersectional subgroups reveals that media with male African faces have the lowest TPR and media with female Asian faces have the highest FPR.

4.1 Key Findings on Evaluated Detectors

- All detectors perform equally on male faces and female faces (0.1 - 0.3% difference in error rate)
- All detectors trained with BI perform worst on darker faces from the African subgroup, especially male African faces (3.5 - 6.7% difference in error rate)
- For detectors trained with BI, faces from the Asian subgroup have the highest FPR, especially female Asian faces (5.2 - 8.1% diff.)
- For detectors trained with BI, faces from the African subgroup have the lowest TPR, especially male African faces (4.7 - 10.7% diff.)
- FaceXRay + BI performs best on Caucasian faces, especially male Caucasian faces (9.8%, 9.5% error rate respectively). Meso4 and Xception detectors (with and without BI) perform best on Indian faces
- The maximum difference in error rate between the best and worst classified subgroups is 10.7%

4.2 Evaluation Methodology

We describe in detail the datasets and metrics utilized in this work to audit deepfake detectors. We adapted racially aware and fair facial recognition datasets labeled with demographic information for our task. For evaluations, we measure AUC and binary classification metrics across different subgroups.

Auditing Datasets

We utilized two face datasets labeled with demographic information: (1) Racial Face-in-the-Wild (RFW) [Wang *et al.*, 2019] and (2) UTKFace [Zhang *et al.*, 2017]. RFW is a dedicated testing dataset manually created for studying racial bias in face recognition. RFW contains four testing subsets, namely Caucasian, Asian, Indian, and African, with images selected from MS-Celeb-1M. Each subset contains about 10K images of 3K individuals for face verification - all with similar distribution with respect to age, gender, yaw pose, and pitch pose. Images in RFW have been carefully and manually cleaned. UTKFace is a large-scale face dataset with a long age span. The dataset consists of over 20K face images with annotations of age, gender, and race. The race labels consist of five groups, namely Caucasian, African, Asian, Indian, and Others (like Hispanic, Latino, Middle Eastern). All images in UTKFace cover large variations in pose, facial expression, illumination, occlusion, and resolution.

For both datasets, we preprocessed images similarly to the deepfake images used for detection training (see Section A.3). Following RFW, we preserve the testing condition of the RFW dataset and do not alter the distribution of the images, which is adapted as the *not-fake* portion of the testing dataset. For UTKFace, despite a large number of available images, it has a quite skewed distribution of racial groups. Hence, we did not use all labeled images but instead down-sampled subgroups to achieve a balanced racial distribution similar to RFW. Figure 1 presents examples and average faces of both the RFW and UTKFace database, along with their respective gender and racial distributions.

To obtain *deepfakes* for the testing dataset along with their demographic labels, we utilized the provided 68 facial landmarks within UTKFace to construct blended images, following the exact methodology as in [Li *et al.*, 2020]. To remain ethnically aware and also maintain demographic information, pairs of faces selected for swapping via the Face X-Rays approach are constrained to be from within the same gender and racial group. We generated 40K blended images per subgroup for a balanced distribution (Figure 2). Our goal is to utilize a deepfake dataset with faithful demographic labels to audit the detectors’ performance on both real and manipulated images.

Evaluation Metrics

We analyze two sets of metrics, binary classification metrics and threshold agnostic Area under the ROC curve (AUC). For classification metrics, similar to Gender Shades [Buolamwini and Gebru, 2018], we follow the gender classification evaluation precedent established by the National Institute for Standards and Technology (NIST) and assess the overall classification accuracy, along with the extension of true positive rate, false positive rate, and error rate (1-PPV) of the intersectional subgroups: $\{\text{male, female}\} \times \{\text{Caucasian, African, Asian, Indian}\}$. Since the FaceForensics++ training dataset is heavily imbalanced, we set the threshold as the value in the range (0.01, 0.99, 0.01) that maximizes the balanced accuracy on the Faceforensics++ validation set. We also evaluated the AUC due to its robustness against class imbalance.




Figure 2: Visualized blended images along with their Face X-Rays for images with low (top row) and high (bottom row) artifacts.

Asian, Indian}. Since the FaceForensics++ training dataset is heavily imbalanced, we set the threshold as the value in the range (0.01, 0.99, 0.01) that maximizes the balanced accuracy on the Faceforensics++ validation set. We also evaluated the AUC due to its robustness against class imbalance.

4.3 Audit Results

Table 2 shows detection performances on gender and racial groups as measured by the AUC, positive predictive value (PPV), error rate, true positive rate (TPR), and false positive rate (FPR) of the 3 deepfake detection models, trained using the FF++ and Blended Image (BI) datasets. We observe disparities in predictive performances between racial groups, which is most apparent in models trained with the BI dataset.

Gender Groups Audit

From Table 2, we observe that all detectors are *equally* accurate in detecting manipulated images containing male and female subjects, with the difference in error rate as low as 0.1 - 0.3%. For four out of five detectors, female subjects have both higher FPR and higher TPR. In the realistic setting where facial profiles on social media are automatically screened via deepfake detectors, FPR indicates that the proportion of real subjects mistakenly identified as fake can be much larger for female subjects than male subjects. This is especially true for the Xception + BI detector, which achieves the best result with error rates of 9.8% on male subjects and 9.9% on female subjects, but nearly twice as large FPR with 7.7% for male subjects and 14.0% for female subjects.

Racial and Intersectional Subgroups Audit

We conduct an intersectional analysis of all detectors on all eight subgroups (M-Cau. F-Cau. M-African, F-African, M-Asian, F-Asian, M-Indian, and F-Indian). As seen in Table 2, we observe large disparities in error rate across race, with the difference in error rate ranging from 3.5 - 7.6% across all detectors. Of note, FaceXRay + BI performs best on Caucasian faces, especially male Caucasians (9.8%, 9.5% error rate respectively). MesoNet and Xception detectors (with and without BI) perform best on Indian faces. Across all detectors, the maximum difference in error rate between the best and worst intersectional subgroups is 10.7%. Figure 3 presents the ratios of FPR and TPR of each subgroup to a reference group, which we have chosen as the “M-Caucasian” group. We notice a stark contrast between FPR and TPR where we observe

Model	Metric	ALL	M	F	Cau.	African	Asian	Indian	M/Cau.	F/Cau.	M/African	F/African	M/Asian	F/Asian	M/Indian	F/Indian
Meso4	AUC	0.614	0.622	0.597	0.589	0.626	0.626	0.614	0.604	0.577	0.588	0.629	0.651	0.601	0.644	0.586
	PPV	0.784	0.764	0.802	0.779	0.794	0.789	0.774	0.795	0.764	0.661	0.940	0.807	0.772	0.797	0.754
	Error Rate	0.435	0.436	0.434	0.462	0.439	0.424	0.415	0.462	0.462	0.458	0.410	0.412	0.436	0.407	0.423
	TPR	0.541	0.517	0.565	0.493	0.524	0.555	0.592	0.476	0.510	0.457	0.591	0.556	0.554	0.578	0.605
	FPR	0.374	0.200	0.225	0.350	0.344	0.371	0.431	0.307	0.394	0.335	0.416	0.333	0.409	0.369	0.494
Xception	AUC	0.810	0.804	0.819	0.793	0.803	0.808	0.841	0.800	0.786	0.788	0.815	0.805	0.812	0.841	0.842
	PPV	0.856	0.827	0.888	0.863	0.838	0.861	0.865	0.864	0.861	0.739	0.957	0.857	0.865	0.863	0.868
	Error Rate	0.267	0.276	0.256	0.301	0.264	0.276	0.225	0.290	0.313	0.301	0.206	0.280	0.273	0.229	0.220
	TPR	0.753	0.749	0.758	0.687	0.783	0.731	0.812	0.704	0.671	0.755	0.812	0.730	0.733	0.808	0.816
	FPR	0.317	0.276	0.205	0.274	0.384	0.295	0.316	0.276	0.272	0.382	0.403	0.304	0.287	0.321	0.310
Meso4 + BI	AUC	0.795	0.811	0.765	0.798	0.79	0.775	0.821	0.818	0.785	0.766	0.78	0.82	0.73	0.845	0.8
	PPV	0.901	0.906	0.897	0.908	0.915	0.878	0.905	0.935	0.888	0.849	0.976	0.913	0.846	0.923	0.889
	Error Rate	0.356	0.355	0.357	0.362	0.385	0.358	0.319	0.384	0.34	0.367	0.413	0.341	0.374	0.324	0.314
	TPR	0.564	0.532	0.596	0.548	0.511	0.58	0.618	0.497	0.6	0.458	0.563	0.577	0.582	0.596	0.64
	FPR	0.155	0.062	0.105	0.138	0.12	0.201	0.162	0.087	0.19	0.116	0.153	0.138	0.264	0.123	0.2
Xception + BI	AUC	0.962	0.964	0.959	0.969	0.951	0.959	0.972	0.972	0.968	0.938	0.958	0.97	0.948	0.977	0.968
	PPV	0.952	0.956	0.949	0.963	0.957	0.933	0.957	0.971	0.955	0.928	0.987	0.956	0.912	0.969	0.945
	Error Rate	0.099	0.098	0.099	0.092	0.119	0.102	0.082	0.092	0.091	0.13	0.102	0.088	0.116	0.076	0.089
	TPR	0.907	0.896	0.919	0.907	0.873	0.923	0.926	0.898	0.916	0.846	0.901	0.918	0.927	0.923	0.93
	FPR	0.114	0.077	0.14	0.088	0.099	0.165	0.104	0.068	0.109	0.094	0.133	0.105	0.224	0.074	0.135
FaceXRay + BI	AUC	0.950	0.95	0.95	0.962	0.936	0.953	0.95	0.963	0.96	0.928	0.937	0.959	0.946	0.946	0.955
	PPV	0.939	0.932	0.947	0.951	0.944	0.931	0.933	0.95	0.952	0.906	0.985	0.938	0.923	0.932	0.933
	Error Rate	0.115	0.115	0.114	0.098	0.133	0.112	0.116	0.095	0.101	0.136	0.13	0.104	0.119	0.122	0.11
	TPR	0.897	0.895	0.899	0.91	0.865	0.912	0.902	0.915	0.905	0.858	0.872	0.914	0.909	0.894	0.911
	FPR	0.145	0.128	0.134	0.118	0.129	0.17	0.163	0.121	0.115	0.127	0.15	0.151	0.189	0.163	0.163

Table 2: Deepfake detection performance on gender and racial groups as measured by the area-under-the-ROC-curve (AUC), positive predictive value (PPV), error rate (1-PPV), true positive rate (TPR), and false positive rate (FPR) of the 3 evaluated deepfake detection models, trained using the standard and Blended Image (BI) approaches.




Figure 3: Ratios of FPR (TPR) for each of the intersectional subgroup to a reference group. In this case, we have chosen the “M-Caucasian” to be the reference group. Purple lines indicates the 20% margins above and below. Red bars indicate the violation of these margins.

that subgroups with Asian or African racial backgrounds have false positive rates as high as three times as that of the reference group. In contrast, the TPRs of all groups are either well-within or around the accepted³ 20% margins of the reference group, indicated by the purple lines.

In addition, there exists a trend in which all detectors trained with BI perform worst on African faces, especially male African faces with 3.5 - 6.7% difference in error rate to the best subgroup. On a closer look, we can see that MesoNet + BI perform worst on female African faces, with 41.3% error rate while Xception + BI and FaceXRay + BI perform worst

on male African faces, with 13.0% and 13.6% error rate respectively. In addition to this trend, we also observe that, for detectors trained with BI, Asian faces have the highest FPR, especially female Asian faces with 5.2 - 8.1% difference in FPR across all BI-trained detectors. Similarly, faces from the African subgroup have the lowest TPR, especially male African faces 4.7 - 10.7% difference in TPR. This trend is uniquely consistent across all three detectors trained with BI (including the state-of-the-art Face X-Ray), even though the detectors all have diverse architectures and training losses.

Analysis of Results

We agree with the findings in [Buolamwini and Gebu, 2018] that using single performance metrics such as AUC or detec-

³Fairness for this metric is in [0.8, 1.2] w.r.t the Four-Fifths rule.

tion accuracy over the entire dataset is not enough to justify massive commercial rollouts of deepfake detectors. Despite high AUC up to 0.962 and detection accuracy up to 90.1% on our deepfake testing dataset, which would allow companies to claim commercial readiness for these detectors on all demographics represented, an intersectional analysis of the detectors shows otherwise.

Our results also show indications of systematic bias in the learning process via generating and using manipulated images for training. Even though training with fake data generated via the BI process helped MesoNet, Xception, and Face X-Ray to improve their overall predictive performances, it also negatively impact predictions on real videos and images. Since fake artifacts are the focus of the detectors given how the training data was prepared, the absence of such artifacts in real and genuine media can lead to unintentional consequences in prediction. Figure 3 plots the ratios of FPR (TPR) for each of the intersectional subgroup to a reference group. We have chosen the “M-Caucasian” to be the reference group. The disparities in FPR suggest that in a real-world scenario, facial profiles of female Asian or female African are 1.5-3 times more likely to be mistakenly labeled as fake than profiles of the male Caucasian. For large scale commercial applications, this would indicate bias against millions of people.

However, we note that the disparities observed are not “intentionally” built into the detectors. Figure 3 (bottom) also demonstrated that the models are indeed focusing on manipulation artifacts as intended, where the ratios of TPR across intersectional subgroups stay well within the 20% margins around the reference group. To the best of our knowledge, the closest work that mentions similar observations about performances between fake and real images is the work by Carlini and Farid [Carlini and Farid, 2020] where the authors use adversarial attacks to change the detectors’ predictions. Noted as surprising by the authors, they found that it is harder to cause real images to be misclassified as fake, requiring up to 7% of image pixels to be flipped, as opposed to causing fake images to be misclassified as real, needing just 1% of pixels required to be flipped. We posit that because of the networks’ focus on the detection of fake artifacts, it is easier to quickly fool the network using its gradient. However, the reverse direction is harder as the network has more trouble coming up with artifacts to “manipulate” a real image.

4.4 Training Distribution and Methodology Bias

To further investigate potential sources for bias in the trained detectors, we analyze both the FaceForensics++ and Blend Image (BI) datasets with respect to their gender and racial distribution. We observe the following key findings:

- Within FF++, 61.7% of all real videos contain a person from the Caucasian group, with 36.0% being female Caucasians.
- For FF++ fake videos, 59.44% are videos of “irregular swaps”. The rest are regular. “Irregular” swaps are when a person’s face is swapped onto another person’s face of a different gender or race.
- For BI blended face images, 65.45% of the images are “irregular swaps”. The rest are regular.




Figure 4: Distribution of intersectional subgroups within FaceForensics++ real videos. 61.7% of all real videos contain a person from the Caucasian group, with 36.0% being female Caucasians.

- For BI images with foreground female Asian faces, 35% are swapped onto female Caucasian faces, 21% onto female Asian faces, and 14% onto female Hispanic faces.

Evaluation Methodology

With the lack of demographic information for videos within the FaceForensics++ dataset, we manually collect ground truth demographic labels. To do so, we annotate each subject into two groups of perceived gender {male, female}, and five groups of perceived race {Caucasian, African, Asian, Indian, Others}. Three graduate annotators are selected for the task, with the assumption that each is of the same skill level in determining gender and racial group. For each subject, the annotators are presented with 5 distinct frames at various times in the video, which displayed the subject at different light angles and poses. We utilized pairwise percent agreement for multiple raters to measure Inter-Rater Reliability (IRR), and the majority label for each subject is selected as the ground truth demographic label. Our annotators achieve 75.93% IRR, which is high for 2 genders and 5 racial groups.

With the demographic labels, we evaluate the percentage of “regular” and “irregular” faceswaps, where “irregular” is defined as a swap where a person’s face is swapped onto another person’s face of a different gender or race. FaceForensics++ provided the IDs for pairs of swaps for all four manipulation methods. Blended Images requires bookkeeping of the target and source faces selected via the BI methodology, which selects a source face from 5000 images where its 68 facial landmarks are closest in Euclidean distance to the target face.

Results

Figure 4 presents the labeled distribution of intersectional subgroups within FaceForensics++ real videos. We observe a strong representation imbalance of gender and racial groups, with the videos containing 58.3% of female subjects and 41.7% of male subjects. The majority of authentic videos are from the Caucasian Group (61.7%), with a major part being female Caucasians (36.0%). Moreover, less than 5% of

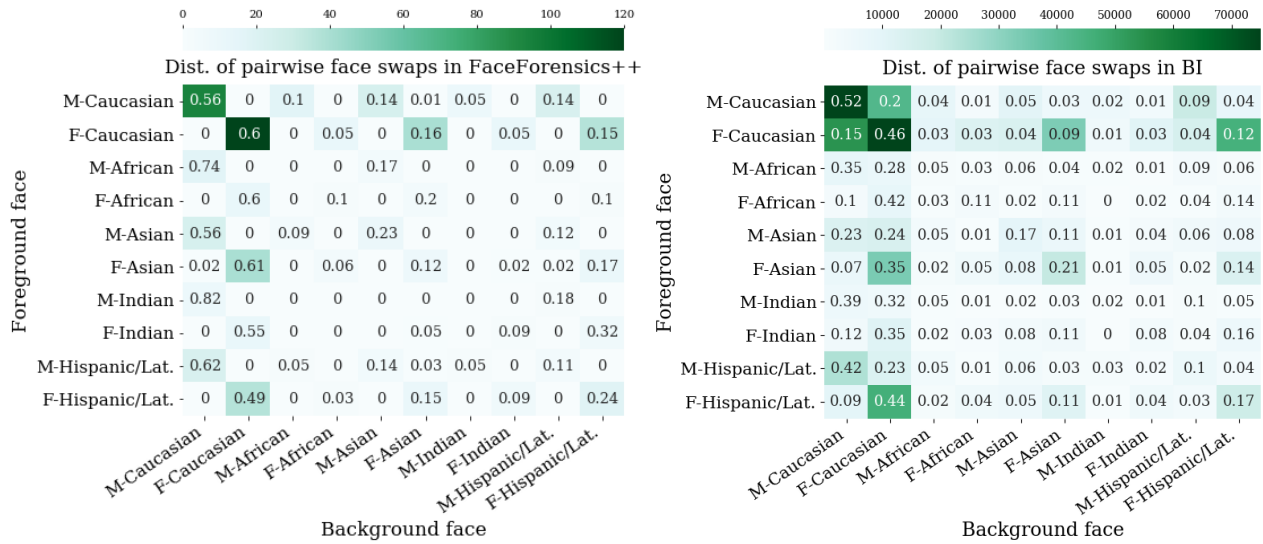


Figure 5: Heatmaps of the distribution of pairwise swaps in FF++ and Blended Image (BI) datasets. Numbers in each row are normalized row-wise to present the percentage of swaps with foreground faces belonging to specific gender and racial group.

real videos contain subjects from the African or Indian group, with male Indian subgroup having the least representation.

Figure 5 plots the heatmap of the distribution of pairwise swaps for manipulated videos/images in FF++ and Blended Image (BI) datasets. Numbers in each square are normalized row-wise to show the percentage of swaps with foreground faces belonging to specific gender and racial group. For FF++ fake videos, 59.44% (428/720) are videos of “irregular swaps”, the rest are “regular” swaps (292/720). 58.75% (423/720) of fakes have female Caucasians and male Caucasians subjects as foreground faces. Zooming in, we can see that the majority (60%) (154/255 videos) of fakes with female Caucasian foreground faces are swapped onto other female Caucasians. On the other hand, the majority (61%) (40/66 videos) of fakes with female Asian foreground faces are swapped onto female Caucasian faces, with only 7% swapping onto female Asian faces. We also observed other types of irregular swaps where female faces are swapped onto male background faces. In the BI dataset, networks are trained with millions of swapped images. Here we sampled 1,000,000 images from the same process to visualize the distribution. Similarly, 65.45% (654,400) of the images are “irregular swaps”, with the scale of BI is much more massive compared to FF++. For BI blended face images with female Asian faces as foreground faces, the majority 35.3% (34,031/96,443 images) are swapped onto female Caucasian faces, 21% onto regular female Asian faces, and 14% onto female Hispanic faces. Hence, given that the networks sees deepfakes with female Asian faces irregularly swapped for most of the time, it is more likely for them to learn a correlation between fakeness and Asian facial features.

Without a specific way to pinpoint the exact source of bias, BI alone may not be fully responsible for misclassification and large disparities in false positive rates. However, we caution against using it for improving deepfake detection perfor-

mances. Perhaps a more racially aware method to generate blended images could be essential for future directions. This would force detection methods to learn to detect meaningful visual artifacts without correlating them to the genders or races of the subjects.

5 Conclusion

As deepfakes become more pervasive, there is a growing reliance on automated systems to combat deepfakes. We argue that practitioners should investigate all societal aspects and consequences of these high impact systems. In this work, we thoroughly measured the predictive performance of popular deepfake detectors on racially aware datasets balanced by gender and race. We found large disparities in predictive performances across races, as well as large representation bias in widely used FaceForensics++. Moreover, a majority of fakes are composed of “irregular” swaps between faces of different gender and races. Our work echoes the importance of benchmark representation and intersectional auditing for increased demographic transparency and accountability in AI systems.

Acknowledgments

This work is supported by the Defense Advanced Research Projects Agency (DARPA) under Agreement No. HR00111990059. The authors would like to express their appreciation to colleagues Caroline Johnston and Nathan Dennler for many useful inputs and valuable comments on this work.

References

- [Afchar *et al.*, 2018] Darius Afchar, Vincent Nozick, Junichi Yamagishi, and Isao Echizen. Mesonet: a compact facial video forgery detection network. *CoRR*, abs/1809.00888, 2018.
- [Bolkubasi *et al.*, 2016] Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *NIPS*, 2016.
- [Buolamwini and Gebru, 2018] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *ACM FAccT*, 2018.
- [Burt and Horvitz, 2020] Tom Burt and Eric Horvitz. New steps to combat disinformation, Sep 2020.
- [Cahlan, 2020] Sarah Cahlan. How misinformation helped spark an attempted coup in gabon. *The Washington Post*, 2020.
- [Carlini and Farid, 2020] Nicholas Carlini and Hany Farid. Evading deepfake-image detectors with white- and black-box attacks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2020.
- [Celis *et al.*, 2019] L Elisa Celis, Sayash Kapoor, Farnood Salehi, and Nisheeth Vishnoi. Controlling polarization in personalization: An algorithmic framework. In *Proceedings of the conference on fairness, accountability, and transparency*, 2019.
- [Chai *et al.*, 2020] Lucy Chai, David Bau, Ser-Nam Lim, and Phillip Isola. What makes fake images detectable? understanding properties that generalize. In *European Conference on Computer Vision*, pages 103–120. Springer, 2020.
- [Cuzzolino *et al.*, 2018] Davide Cuzzolino, Justus Thies, Andreas Rössler, Christian Riess, Matthias Nießner, and Luisa Verdoliva. Forensicttransfer: Weakly-supervised domain adaptation for forgery detection. *arXiv preprint arXiv:1812.02510*, 2018.
- [Gebru *et al.*, 2018] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. Datasheets for datasets. *arXiv preprint arXiv:1803.09010*, 2018.
- [Hao, 2019] Karen Hao. An ai app that “undressed” women shows how deepfakes harm the most vulnerable. *MIT Technology Review*, 2019.
- [Hardt *et al.*, 2016] M. Hardt, E. Price, and Nathan Srebro. Equality of opportunity in supervised learning. In *NIPS*, 2016.
- [Ingram, 2019] David Ingram. A face-swapping app takes off in china, making ai-powered deepfakes for everyone. *NBC*, 2019.
- [Jiang *et al.*, 2020] Liming Jiang, Ren Li, Wayne Wu, Chen Qian, and Chen Change Loy. Deeperforensics-1.0: A large-scale dataset for real-world face forgery detection. In *CVPR*, 2020.
- [Khodabakhsh *et al.*, 2018] A. Khodabakhsh, R. Ramachandra, K. Raja, P. Wasnik, and C. Busch. Fake face detection methods: Can they be generalized? In *2018 International Conference of the Biometrics Special Interest Group (BIOSIG)*, 2018.
- [Li and Lyu, 2018] Yuezun Li and Siwei Lyu. Exposing deepfake videos by detecting face warping artifacts. *CoRR*, abs/1811.00656, 2018.
- [Li *et al.*, 2020] Lingzhi Li, Jianmin Bao, Ting Zhang, Hao Yang, Dong Chen, Fang Wen, and Baining Guo. Face x-ray for more general face forgery detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [Liu *et al.*, 2018] Lydia T Liu, Sarah Dean, Esther Rolf, Max Simchowitz, and Moritz Hardt. Delayed impact of fair machine learning. *arXiv preprint arXiv:1803.04383*, 2018.
- [Liu *et al.*, 2020] Zhengzhe Liu, Xiaojuan Qi, and Philip HS Torr. Global texture enhancement for fake face detection in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8060–8069, 2020.
- [Mirsky and Lee, 2021] Yisroel Mirsky and Wenke Lee. The creation and detection of deepfakes: A survey. *ACM Computing Surveys (CSUR)*, 54(1):1–41, 2021.
- [Mitchell *et al.*, 2019] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*, pages 220–229, 2019.
- [Rössler *et al.*, 2019] Andreas Rössler, Davide Cuzzolino, Luisa Verdoliva, Christian Riess, Justus Thies, and Matthias Nießner. Faceforensics++: Learning to detect manipulated facial images. In *Int. Conf. Comput. Vis.*, 2019.
- [Tolosana *et al.*, 2020] Ruben Tolosana, Ruben Vera-Rodriguez, Julian Fierrez, Aythami Morales, and Javier Ortega-Garcia. Deepfakes and beyond: A survey of face manipulation and fake detection. *arXiv preprint arXiv:2001.00179*, 2020.
- [Trinh *et al.*, 2020] Loc Trinh, Michael Tsang, Sirisha Rambhatla, and Yan Liu. Interpretable and trustworthy deepfake detection via dynamic prototypes. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2020.
- [Vaccari and Chadwick, 2020] Cristian Vaccari and Andrew Chadwick. Deepfakes and disinformation: Exploring the impact of synthetic political video on deception, uncertainty, and trust in news. *Social Media + Society*, 6(1):2056305120903408, 2020.
- [Wang *et al.*, 2019] Mei Wang, Weihong Deng, Jiani Hu, Xunqiang Tao, and Yaohai Huang. Racial faces in the wild: Reducing racial bias by information maximization adaptation network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [Wang *et al.*, 2020] Run Wang, Felix Juefei-Xu, Lei Ma, Xiaofei Xie, Yihao Huang, Jian Wang, and Yang Liu. Fakespotter: A simple yet robust baseline for spotting ai-synthesized fake faces. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2020.
- [Zhang *et al.*, 2017] Zhifei Zhang, Yang Song, and Hairong Qi. Age progression/regression by conditional adversarial autoencoder. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5810–5818, 2017.