

# Demographic Bias in Presentation Attack Detection of Iris Recognition Systems

Meiling Fang<sup>1,2</sup>, Naser Damer<sup>1,2</sup>, Florian Kirchbuchner<sup>1,2</sup>, Arjan Kuijper<sup>1,2</sup>

<sup>1</sup>Fraunhofer Institute for Computer Graphics Research IGD, Darmstadt, Germany

<sup>2</sup>Mathematical and Applied Visual Computing, TU Darmstadt, Darmstadt, Germany

Email: meiling.fang@igd.fraunhofer.de

**Abstract**—With the widespread use of biometric systems, the demographic bias problem raises more attention. Although many studies addressed bias issues in biometric verification, there are no works that analyze the bias in presentation attack detection (PAD) decisions. Hence, we investigate and analyze the demographic bias in iris PAD algorithms in this paper. To enable a clear discussion, we adapt the notions of differential performance and differential outcome to the PAD problem. We study the bias in iris PAD using three baselines (hand-crafted, transfer-learning, and training from scratch) using the NDCLD-2013 [18] database. The experimental results point out that female users will be significantly less protected by the PAD, in comparison to males.

**Index Terms**—PAD bias, iris PAD, differential performance.

## I. INTRODUCTION

Biometric recognition systems like face, iris, or fingerprint have gained wide deployment within various application fields. This has raised two significant questions. One is concerned with the security of recognition systems against presentation attacks. The other is related to fairness, i.e., if the system performs the same for every different demographic group.

General presentation attacks include printed images, replaying videos, synthetic images. Iris presentation attack also contains textured contact lenses. Several researchers have already addressed these PAD problem, e.g., face PAD [7], [16], [23] and iris PAD [2].

Demographic bias in face recognition was discussed by [11], [19], [25]. Prior works on this issue indicated that face recognition systems obtained inherently a lower performance for certain demographic groups [11], [19]. For example, females are more challenging to recognize compared to males [19]. However, no previous works address demographic bias in PAD decisions. As iris recognition has become increasingly popular, e.g., iris recognition technology is deployed on consumer smartphones, we address the demographic bias issues in iris PAD systems in this work.

Howard *et al.* [11] defined four recognition related terms to facilitate demographic bias research. We adapt and reuse two of these terms to investigate the demographic bias in PAD systems: 1) **Differential Performance** represents a difference in the bona fide or attack decision distribution between specific

demographic groups independent of any decision threshold, 2) **Differential Outcome** describes a difference in Attack Presentation Classification Error Rate (APCER) or Bonafide Presentation Classification Error Rate (BPCER) rates between different demographic groups relative to a decision threshold. We utilize these terms to demonstrate our experiments and examine the bias in Sec.IV and Sec.V.

This work presents the first analysis of demographic bias in PAD systems (of any modality). Specifically, we study gender demographic bias in both hand-crafted and deep learning-based iris PAD algorithms. This is done under the variation of the training data bias scenarios. We point out that the PAD decisions tend to be significantly less accurate for female subjects in most of the experimental setups.

## II. RELATED WORK

*a) Bias in Biometric Systems:* Danks *et al.* [3] defined the five types of algorithmic biases: 1) training data bias, 2) algorithmic focus bias, 3) algorithmic processing bias, 4) transfer context bias, 5) interpretation bias. The study of algorithmic bias in autonomous systems has caused researchers to think about bias in biometric systems. Despite the improved performance of biometric recognition systems such as fingerprint, iris or face, the fairness of the recognition systems is still questionable. The current processes in biometric systems include (among others) recognition, classification and presentation attack detection. Biometric recognition is a technology that can identify or verify a person by using biometric information [14], [15]. For example, facial or iris recognition systems compute a similarity score between subjects from images or a sequence of video frames. The classification here refers to assigning demographics or other labels to biometric samples [4]. Presentation attack refers to using a presentation to fool the system in an effort to induce an incorrect decision such as using printed images or display videos. In contrast, presentation attack detection systems are expected to automatically determine such attacks [12], [22].

Drozowski *et al.* [5] summarized an overview of the topic of algorithmic bias in the context of biometrics and also contributed a comprehensive survey of the existing literature on biometric bias estimation and mitigation. According to recent studies, biometric recognition and classification systems do not perform equally well in different demographic groups. Several face recognition relevant studies indicated that the

This research work has been funded by the German Federal Ministry of Education and Research and the Hessen State Ministry for Higher Education, Research and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE.

female, Black, and younger (18–30y.o.) cohorts are inherently more difficult to recognize [11], [19]. Such phenomena that recognition systems have a lower performance for specific demographic groups is denoted as demographic bias. The study of demographic bias provides a chance to find the recognition system problems and mitigate the biases.

*b) Bias in PAD systems:* Many studies have already pointed demographic bias in face or fingerprint recognition systems [11], [19], [20], [27] and obtained similar conclusions. Moreover, there is one study that experimentally related the bias to biometric attacks, namely the face morphing attack [26]. Garcia *et al.* [26] suggested that morphing attacks are much more successful for Asian females. However, the authors did not study bias in attack detection performance. Since iris biometric features are deployed in many security scenarios, the decision bias in the processes involved must be analyzed and minimized. Unfortunately, no previous works experimentally addressed the demographic bias in PAD systems in general, or the iris PAD specifically. The primary difficulty is limited data resources. Czajka *et al.* [2] summarized iris PAD databases. Parts of the databases cannot be shared due to information protection, others were collected many years ago and are unrealistically simple. The rest of the databases contained almost no demographic information. In Sec.IV, we find that only one database, NDCLD-2013 [18], can be used to investigate the demographic bias. Even this database also has an imbalanced number of attack samples of males and females. Despite the limited database availability, the demographic bias in iris PAD is an open issue and is of great interest. With the widespread use of iris recognition technology, we believe the bias problem needs to be analyzed to enable future mitigation efforts.

### III. IRIS PRESENTATION ATTACK DETECTION

Iris recognition systems are susceptible to the presentation attacks, like printed iris images, replaying videos, or contact lenses. As a result, numerous iris PAD solutions have been proposed to target iris vulnerability issues [6], [9]. Czajka *et al.* [2] summarized the state-of-the-art iris PAD algorithms in detecting different categories of attack. Various iris PAD frameworks are grouped into hand-crafted and deep-learning-based methods. Hand-crafted features such as Local Binary Pattern (LBP) [8], [21], [28], Binarized statistical image features (BSIF) as well as their variations have been the research focal points until 2015 and made a remarkable contribution to iris PAD problems. After that, with the rapid development and application of deep learning in multiple domains, especially, computer vision field, neural-network-based PAD methods showed up out of nowhere. However, the deeper the network structure e.g., VGG16 [24], the higher the computational needs. Such algorithms with high computational requirements are hard to deploy in mobile devices. Therefore, Howard *et al.* [10] released two new MobileNet models: MobileNetV3-Large and MobileNetV3-Small to target high and low resource use cases. We utilize the MobileNetV3-Small structure to be one of our PAD baselines.

To investigate the effects of demographic bias on different types of iris PAD systems. These systems are based on 1) hand-crafted features, 2) transfer learning of a general-purpose pre-trained network, 3) a trained from scratch network with computationally-efficient. We use the following baselines to examine the differential performance and differential outcome:

- 1) **LBP + SVM:** LBP [8] has been proposed to address the texture-based classification tasks. We adapt the LBP feature in conjunction with a Support Vector Machine (SVM) [1] using a linear kernel to detect the presentation attacks. The threshold, which determines the label of the iris image, is chosen based on the development subset. Finally, the trained SVM model with a specific threshold is used to make a decision in the testing set, whether the iris image is bona fide or attack.
- 2) **VGG-16 + PCA + SVM:** VGG-16 [24] is a convolutional neural network for image classification. Considering that VGG-16 is one of the most widely used feature extractors, we utilize the pre-trained VGG-16 to encode iris features. After extraction of VGG-16 features, a linear dimensionality reduction technique, Principal Component Analysis (PCA), is used to project the extracted feature from each iris image with  $7 \times 7 \times 512$  dimension into a 128 dimensional sub-space. Similar to LBP, those lower-dimensional features are fed to SVM [1] with a linear kernel and make a prediction (bona fide or attack) based on the threshold choosing from development set.
- 3) **MobileNetV3-Small:** We use MobileNetV3-Small [10] as a network structure to train from scratch. The contrast-enhanced full iris images in the training set are fed to the MobileNetV3-Small model. Again, the threshold is picked up to the development set and used to predict the label of iris image (bona fide and attack). In this stage, we use the *early stopping criteria* to stop training once there is no improvement in 5 epochs. The batch size is set to 16, limited by the GPU memory.

### IV. EXPERIMENTAL SETUP

In this experiment, we use the NDCLD-2013 database [18] including textured contact lens attack collected by the University of Notre Dame. The detailed iris image capture information has been published in [18]. Briefly, this database contained 5100 images of 330 subjects, which acquired with a LG4000 and an IrisGuard AD100 iris capture device. At the time of the acquisition, Doyle *et al.* [18] have divided the database into training and testing subsets, which was guaranteed to have an even distribution of genders and ethnicity. Nonetheless, we demonstrate experiments only in one database and only on the gender cohort. The reasons are:

- 1) The number of the public iris PAD database is much less than the face and iris recognition databases. Due to the awareness of privacy security, demographic information is more difficult to collect. To our knowledge, only NDCLD-2013 [18] and NDCLD-2015 [17] database had demographic information. While the NDCLD-2013

included gender, race and age categories, the NDCLD-2015 had gender and eye color information.

- 2) There are only two female subjects with attack samples in the NDCLD-2015 database. Similarly, the eye color was largely unbalanced in this database. Therefore, we have to disregard the NDCLD-2015 database, though we believe that the impact of eye color bias in iris PAD systems is very interesting and worth the investigation.
- 3) As shown in Fig. 1, blue (BF-NL) denotes the number of subjects wearing no contact lens, green (BF-SL) is the number of subjects wearing the soft lenses and green (Attack) is the number of subjects wearing the textured lenses. It should be noticed that race/ethnicity and age categories are unbalanced: 1) no attack samples in Black and Asian categories, 2) unbalance age groups (18–30y.o.), and (30–70y.o.). Therefore, for statistical significance, we only analyze the demographic bias over gender.

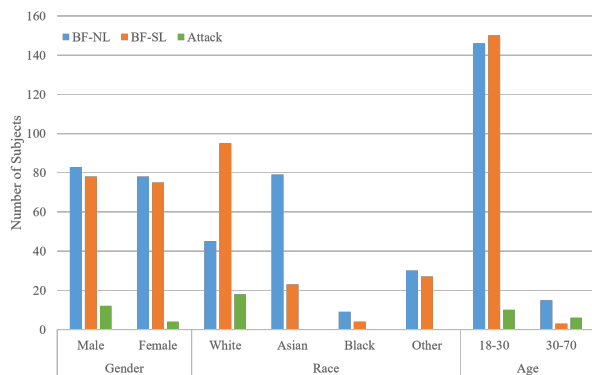


Fig. 1: Subject-specific distribution of the demographic variables by different types of representation in database NDCLD-2013 [18]. BF-NL is the bona fide iris image, BF-SL is the iris image with the soft lenses, which is also reported as bona fide, Attack is the iris image with textured lenses.

Given the limited data resources, we are only able to study gender demographics in iris PAD systems. Two sub-database LG4000 and AD100 are combined into one database. Train-test-split is defined in [18]. A summary of the training and testing samples for each type of representation can be found in Tab.I. Additionally, we split the training subset into identity-disjoint training (80%) and development (20%) sets. In this study, soft lens iris images are used as bona fide samples.

Gender	Male			Female		
	BF-NL	BF-SL	Attack	BF-NL	BF-SL	Attack
# Training	600	600	920	600	600	280
# Testing	250	250	370	250	250	130

TABLE I: Number of images used for training and testing per gender. Training and test sets are subject-disjoint. A total of 5100 iris images are used in this study.

For gender demographics, two experiments are conducted. Because subjects cannot be balanced in 5-fold, we repeat each experiment 5 times to replace the 5-fold cross-validation

and both experiments are subject-disjoint. The first experiment measures the iris PAD performance on males versus females. Specifically, the training and development set contains iris images from both male and female subjects. Then, we separately evaluate males and females set. The result from the first experiment can be found in Tab.II. In addition, we report the attack and bona fide decision distribution in the gender cohort from three iris PAD baselines to analyze the differential performance mentioned in Sec.I (see Fig.2, Fig.3, Fig.4). The second experiment investigates the influence of the training subset on iris PAD performance, which means that we use iris images from one gender as a training and development set. For example, we train a model by using only male data from the training set and compute a threshold by fixing APCER at 0.01% in the male development set. Then, this model and threshold are applied separately to males and females testing set. Subsequently, we are able to analyze the outcome differences in the scenarios where one gender is exclusively used for training and threshold assignment. The detailed results of three PAD methods from this experiment can be found in Tab.III, Tab.IV and Tab.V.

The following metrics are used to measure the iris PAD algorithm performance:

- **Attack Presentation Classification Error Rate (APCER):** The proportion of attack images incorrectly classified as bona fide samples.
- **Bonafide Presentation Classification Error Rate (BPCER):** The proportion of bonafide images incorrectly classified as attack samples.
- **Half Total Error Rate (HTER):** corresponds to the average of BPCER and APCER.

The APCER, BPCER, and HTER that we report are calculated on the same threshold, which chosen by fixing APCER at 0.01% in the development subset. The APCER and BPCER follows the standard definition presented in the ISO/IEC 30107-3 [13].

## V. RESULTS AND ANALYSIS

In this section, we analyze the experimental results described in Sec.IV. First, the difference in bona fide and attack PAD decision distribution between males and females, independent of any threshold, is demonstrated to measure the differential performance. Second, a difference in BPCER and APCER between males and females regarding a specific threshold is discussed to examine the differential outcome.

a) *Differential Performance in the PAD decision distribution:* The ideal distribution of PAD decision is that the decision of the attack samples is close to zero and the decision of the bona fide samples is close to one. In this subsection, we represent the PAD decision distributions between the gender cohort, which are generated by three algorithms: LBP, MobileNetV3-Small, and VGG-16 method. First, Fig.2 shows that most bona fide decisions of males are distributed between 0.7 to 1.0 by LBP method. In contrast, the attack and the bona fide decisions of females are relatively difficult to separate. Second, deep-learning-based algorithms obtain a much better

performance compared to the LBP method. As shown in Fig.3 and Fig.4, the attack and bona fide decisions generated by MobileNetV3-Small and VGG-16 baselines are relatively separated. Although both gender cases achieved good results, we find out that there is still a small difference in the PAD decision distribution. By observing the zoomed region (in Fig. 3) in MobileNetV3-Small distribution, we notice that male decisions are denser. Moreover, the female cohort performs slightly worse than males, which is noted by the smaller gap between the attack and the bona fide distribution from VGG-16 (see Fig.4). Together, those evidences point out that female presentation subjects are potentially more difficult to detect by iris PAD systems.




Fig. 2: PAD Decision Distribution by LBP method between bona fide (green) and attack (red) on gender demographic. The female’s PAD decision distribution is less separable.




Fig. 3: PAD Decision Distribution from MobileNetV3-Small method between bona fide (green) and attack (red) on gender demographic.

b) *Differential Outcome in APCER and BPCER:* As shown in Tab.II, both hand-crafted-based and neural network-based algorithms performed worse on the females than males. This finding is consistent with the results of gender cohort researches in face recognition systems [11], [19]. LBP method performs significantly worse on females. The BPCER of the LBP method on females (42.60%) is much higher than on males (21.85%) with the same number of bona fide images in the testing set. MobileNetV3-Small and VGG-16 obtain higher APCER on females, but, the same BPCER on both cohorts. Excluding the gender bias effect, one of the possible reasons is due to a small number of female attack samples in the training data as described in [3].




Fig. 4: PAD Decision Distribution from VGG-16 method between bona fide (green) and attack (red) on gender demographic. The decision gap of female samples between bona fides and attacks are smaller than that of males.

The results of the second experiment possess more information about the nature of the discrepancy. In most cases, better testing performances are achieved on the trained gender cohort when training exclusively on males or females. Also, the LBP method’s performance indicates the significant error increase when training only on mixed data (males + females) than training on males/females in Tab.III. One possible reason is that fewer data cause the underfit or overfit of the classifier. As shown in Tab. III, even if the female samples have been learned, the female group performs worse than males. Hence, the LBP method shows that it is more difficult to decide if a female subject is bona fide or attack. Besides, the VGG-16 method demonstrates a much worse generalizability on the female cohort (Tab.V). The model, which is trained only by female samples, marks male attack samples as bona fide images. MobileNetV3-Small (Tab.IV) provide the consistent generalizability of the females with the VGG-16 method. There is still a possibility that iris PAD error on females can be reduced by merely changing the distribution of the female samples. In general, we notice that different PAD solutions perform significantly better on male samples, even in some cases where females samples are exclusively used for training.

Metric	LBP			MobileNetV3-Small			VGG-16		
	Male	Female	Both	Male	Female	Both	Male	Female	Both
APCER	12.70	8.46	11.6	8.11	9.23	8.40	12.43	16.15	13.40
BPCER	31.00	42.6	36.8	0.20	0.20	0.20	0	0	0
HTER	21.85	<b>25.53</b>	24.2	4.16	<b>4.72</b>	4.30	6.22	<b>8.08</b>	6.7

TABLE II: Results of three iris PAD methods. The training and development set comprise of both males and females samples. The threshold is chosen by fixing APCER at 0.01% in development set. Then, the evaluation is applied on males, females and mixed testing set separately. The bold number denotes the worst results. All HTER results indicate that the female samples are harder to classify.

## VI. CONCLUSION

In this paper, we address the demographic bias problem in iris PAD systems. Since no previous work discussed bias in PAD, we adapt two terms used for verification bias (differential performance and outcome) to analyze the demographic bias on

Train	LBP					
	Male			Female		
	Male	Female	Both	Male	Female	Both
APCER	14.59	16.92	15.2	14.32	16.15	14.78
BPCER	32.00	43.80	39.70	60.60	67.00	63.80
HTER	23.29	<b>30.36</b>	27.45	37.46	<b>41.58</b>	38.29

TABLE III: Results of the LBP methods. The training set and development set contain iris images of only one gender. The evaluation is performed on males, females and mixed gender, separately. Training on female cohort achieves much worse HTER than training on males. One possible reason is that Female cohort has less subjects.

Train	MobileNetV3-Small					
	Male			Female		
	Male	Female	Both	Male	Female	Both
APCER	2.16	7.69	3.60	44.32	10.76	35.6
BPCER	2.00	0	1.00	6.00	7.60	6.80
HTER	2.08	<b>3.85</b>	2.30	<b>25.16</b>	9.18	21.20

TABLE IV: Results of the MobileNetV3-Small methods. The training set and development set contain iris images of only one gender. The evaluation is performed on males, females and mixed gender, separately.

Train	VGG-16					
	Male			Female		
	Male	Female	Both	Male	Female	Both
APCER	18.38	70.77	29.80	100	13.08	75.80
BPCER	0	0	0	0	0	0
HTER	9.19	<b>35.39</b>	14.90	<b>50.00</b>	6.54	37.90

TABLE V: Results of the VGG-16 methods. The training set and development set contain iris images of only one gender. The evaluation is performed on males, females and mixed gender, separately.

PAD performance. Two kinds of experiments are designed to investigate demographic bias in iris PAD accuracy, as well as the impact of the training data bias. The results achieved by three PAD baselines showed a significant difference in the performance and outcome of the algorithms between male and female samples in most experimental setups.

## REFERENCES

- [1] C. Cortes and V. Vapnik, "Support-vector networks," in *Machine Learning*, 1995, pp. 273–297.
- [2] A. Czajka and K. W. Bowyer, "Presentation attack detection for iris recognition: An assessment of the state-of-the-art," *ACM Comput. Surv.*, vol. 51, no. 4, pp. 86:1–86:35, 2018.
- [3] D. Danks and A. J. London, "Algorithmic bias in autonomous systems," in *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, C. Sierra, Ed. ijcai.org, 2017.
- [4] A. Dantcheva, P. Elia, and A. Ross, "What else does your biometric data reveal? A survey on soft biometrics," *IEEE Trans. Information Forensics and Security*, vol. 11, no. 3, pp. 441–467, 2016.
- [5] P. Drozdowski, C. Rathgeb, A. Dantcheva, N. Damer, and C. Busch, "Demographic bias in biometrics: A survey on an emerging challenge," *CoRR*, vol. abs/2003.02488, 2020.
- [6] M. Fang, N. Damer, F. Boutros, F. Kirchbuchner, and A. Kuijper, "Deep learning multi-layer fusion for an accurate iris presentation attack detection," in *23th FUSION, 2020, July 6-9, 2020*. IEEE, 2020.
- [7] A. George, Z. Mostaani, D. Geissenbuhler, O. Nikisins, A. Anjos, and S. Marcel, "Biometric face presentation attack detection with multi-channel convolutional neural network," *IEEE Trans. Information Forensics and Security*, vol. 15, pp. 42–55, 2020.
- [8] P. Gupta, S. Behera, M. Vatsa, and R. Singh, "On iris spoofing using print attack," in *ICPR 2014, Stockholm, Sweden, August 24-28, 2014*. IEEE Computer Society, 2014, pp. 1681–1686.
- [9] S. Hoffman, R. Sharma, and A. Ross, "Convolutional neural networks for iris presentation attack detection: Toward cross-dataset and cross-sensor generalization," in *IEEE CVPR Workshops 2018, Salt Lake City, UT, USA, June 18-22, 2018*. IEEE Computer Society, pp. 1620–1628.
- [10] A. Howard, M. Sandler, G. Chu, L. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan, Q. V. Le, and H. Adam, "Searching for mobilenetv3," *CoRR*, 2019.
- [11] J. Howard, Y. Sirotin, and A. Vemury, "The effect of broad and specific demographic homogeneity on the imposter distributions and false match rates in face recognition algorithm performances," *IEEE BTAS.*, 2019.
- [12] International Organization for Standardization, "ISO/IEC DIS 30107-1:2016: Information Technology – Biometric presentation attack detection – P. 1: Framework," 2017.
- [13] —, "ISO/IEC DIS 30107-3:2016: Information Technology – Biometric presentation attack detection – P. 3: Testing and reporting," 2017.
- [14] International Organization for Standardization and International Electrotechnical Committee, "ISO/IEC 19795-1:2006: Information Technology – Biometric Performance Testing and Reporting - Part 1: Principles and Framework," April 2006.
- [15] —, "ISO/IEC 2382-37:2017: Information Technology – Vocabulary - Part 37: Biometrics, 2nd ed." February 2017.
- [16] S. Jia, G. Guo, Z. Xu, and Q. Wang, "Face presentation attack detection in mobile scenarios: A comprehensive evaluation," *Image Vis. Comput.*, vol. 93, p. 103826, 2020.
- [17] J. S. D. Jr. and K. W. Bowyer, "Robust detection of textured contact lenses in iris recognition using BSIF," *IEEE Access*, vol. 3, pp. 1672–1683, 2015.
- [18] J. S. D. Jr., K. W. Bowyer, and P. J. Flynn, "Variation in accuracy of textured contact lens detection based on sensor and lens pattern," in *IEEE BTAS 2013, Arlington, VA, USA, September 29 - October 2, 2013*. IEEE, 2013, pp. 1–7.
- [19] B. Klare, M. J. Burge, J. C. Klontz, R. W. V. Bruegge, and A. K. Jain, "Face recognition performance: Role of demographic information," *IEEE Trans. Information Forensics and Security*, vol. 7, no. 6, pp. 1789–1801, 2012.
- [20] B. Lu, J. Chen, C. D. Castillo, and R. Chellappa, "An experimental evaluation of covariates effects on unconstrained face verification," *CoRR*, 2018.
- [21] H. Mandalapu, R. Ramachandra, and C. Busch, "Empirical evaluation of texture-based print and contact lens iris presentation attack detection methods," in *Proceedings of the 3rd ICBEA 2019, Stockholm, Sweden, May 29-31, 2019*. ACM, 2019.
- [22] S. Marcel, M. S. Nixon, J. Fierrez, and N. W. D. Evans, Eds., *Handbook of Biometric Anti-Spoofing - Presentation Attack Detection, Second Edition*, ser. Advances in Computer Vision and Pattern Recognition. Springer, 2019.
- [23] R. Shao, X. Lan, J. Li, and P. C. Yuen, "Multi-adversarial discriminative deep domain generalization for face presentation attack detection," in *IEEE CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*. Computer Vision Foundation / IEEE, 2019, pp. 10023–10031.
- [24] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *ICLR 2015, San Diego, CA, USA, May 7-9, 2015*, Y. Bengio and Y. LeCun, Eds., 2015.
- [25] P. Terhörst, J. N. Kolb, N. Damer, F. Kirchbuchner, and A. Kuijper, "Post-comparison mitigation of demographic bias in face recognition using fair score normalization," *CoRR*, 2020.
- [26] R. Vicente-Garcia, L. Wandzik, L. Grabner, and J. Krüger, "The harms of demographic bias in deep face recognition research," in *2019 International Conference on Biometrics, ICB 2019, Crete, Greece, June 4-7, 2019*. IEEE, 2019, pp. 1–6.
- [27] M. Wang, W. Deng, J. Hu, J. Peng, X. Tao, and Y. Huang, "Racial faces in-the-wild: Reducing racial bias by deep unsupervised domain adaptation," *CoRR*, 2018.
- [28] H. Zhang, Z. Sun, and T. Tan, "Contact lens detection based on weighted LBP," in *20th ICPR ICPR 2010, Istanbul, Turkey, 23-26 August 2010*. IEEE Computer Society, 2010, pp. 4279–4282.