



King's Research Portal

DOI:

[10.1007/978-3-030-87199-4_39](https://doi.org/10.1007/978-3-030-87199-4_39)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Puyol-Antón, E., Ruijsink, B., Piechnik, S. K., Neubauer, S., Petersen, S. E., Razavi, R., & King, A. P. (2021). Fairness in Cardiac MR Image Analysis: An Investigation of Bias Due to Data Imbalance in Deep Learning Based Segmentation. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 12903, 413-423. https://doi.org/10.1007/978-3-030-87199-4_39

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Fairness in Cardiac MR Image Analysis: An Investigation of Bias Due to Data Imbalance in Deep Learning Based Segmentation

Esther Puyol-Antón¹, Bram Ruijsink^{1,2}, Stefan K. Piechnik⁷, Stefan Neubauer⁷, Steffen E. Petersen^{3,4,5,6}, Reza Razavi^{1,2}, and Andrew P. King¹

¹ School of Biomedical Engineering & Imaging Sciences, King's College London, UK

² Guy's and St Thomas' Hospital, London, UK.

³ William Harvey Research Institute, NIHR Barts Biomedical Research Centre, Queen Mary University London, Charterhouse Square, London, EC1M 6BQ, UK

⁴ Barts Heart Centre, St Bartholomew's Hospital, Barts Health NHS Trust, West Smithfield, EC1A 7BE, London, UK

⁵ Health Data Research UK, London, UK

⁶ Alan Turing Institute, London, UK

⁷ Division of Cardiovascular Medicine, Radcliffe Department of Medicine, University of Oxford, UK.

Abstract. The subject of 'fairness' in artificial intelligence (AI) refers to assessing AI algorithms for potential bias based on demographic characteristics such as race and gender, and the development of algorithms to address this bias. Most applications to date have been in computer vision, although some work in healthcare has started to emerge. The use of deep learning (DL) in cardiac MR segmentation has led to impressive results in recent years, and such techniques are starting to be translated into clinical practice. However, no work has yet investigated the fairness of such models. In this work, we perform such an analysis for racial/gender groups, focusing on the problem of training data imbalance, using a nnU-Net model trained and evaluated on cine short axis cardiac MR data from the UK Biobank dataset, consisting of 5,903 subjects from 6 different racial groups. We find statistically significant differences in Dice performance between different racial groups. To reduce the racial bias, we investigated three strategies: (1) stratified batch sampling, in which batch sampling is stratified to ensure balance between racial groups; (2) fair meta-learning for segmentation, in which a DL classifier is trained to classify race and jointly optimized with the segmentation model; and (3) protected group models, in which a different segmentation model is trained for each racial group. We also compared the results to the scenario where we have a perfectly balanced database. To assess fairness we used the standard deviation (SD) and skewed error ratio (SER) of the average Dice values. Our results demonstrate that the racial bias results from the use of imbalanced training data, and that all proposed bias mitigation strategies improved fairness, with the best SD and SER resulting from the use of protected group models.

Keywords: Fair AI · Segmentation · Cardiac MRI · Inequality

1 Introduction

Fairness in artificial intelligence (AI) is a relatively new but fast-growing research field which deals with assessing and addressing potential bias in AI models. For example, an early landmark paper [2] found differences in performance of a video-based gender classification model for different racial groups. With AI models starting to be deployed in the real world it is essential that the benefits of AI are shared equitably according to race, gender and other demographic characteristics, and so efforts to ensure the fairness of deployed models have generated much interest. Most work so far has focused on computer vision problems but some applications in healthcare are starting to emerge [15, 22].

Recently, deep learning (DL) models have shown remarkable success in automating many medical image segmentation tasks. In cardiology, human-level performance in segmenting the main structures of the heart has been reported [1], and researchers have proposed to use these models for tasks such as automating cardiac functional quantification [19]. These methods are now starting to move towards wider clinical translation.

It has long been well understood that cardiac structure and function, as well as the mechanisms leading to cardiovascular disease, vary according to demographic characteristics such as race and gender [11]. For example, the Multi-Ethnic Study of Atherosclerosis (MESA) [24] showed that there are racial differences in regional left ventricular (LV) systolic function in a large cohort study of adults. In addition, there are profound race-associated disparities among those who are affected by and die from cardiovascular disease [13]. Inequalities in detection of disease likely play a role in these differences. Therefore, it is surprising that no work to date has investigated potential bias in AI models for cardiac image analysis. To the best of our knowledge, the closest related work was on assessing differences in radiomics features by gender [18].

In this paper, we perform the first analysis of the fairness of DL-based cardiac MR segmentation models, focusing in particular on the impact of gender or race imbalance in the training data. To the best of our knowledge, this is also the first analysis of the fairness of segmentation models in general. We discover significant bias in performance between racial groups when trained using one of the largest and widely used public databases of cardiac magnetic resonance (MR) data (the UK Biobank). We propose algorithms to mitigate this bias, resulting in a fairer segmentation model that ensures that no racial group will be disadvantaged when segmentations of their cardiac MR data are used to inform clinical management.

2 Background

In recent years, addressing fairness concerns in AI models has been an active research area. To date, most previous works have focused on classification tasks such as image recognition, although Hwang *et al.* [7] investigated bias in image-to-image translation. In an extensive search, we did not find any previous work on bias mitigation for image segmentation tasks.

At a high level, fairness techniques can be grouped into two categories: (i) ‘fairness through awareness’, also known as bias mitigation strategies or discrimination-aware classification, that aim to make the AI algorithms more aware of the protected attributes⁸ by making predictions independently for each protected group [4]; and (ii) ‘fairness through unawareness’, which assumes that if the model is unaware of the protected attributes while making decisions, the decisions will be fair. This second approach has been the most common one used to train AI models. However, it has been shown to be unsuccessful in many cases [21, 25] due to the correlation of protected attributes with other variables in the data.

Focusing now on bias mitigation algorithms (a.k.a. ‘fairness through awareness’), DL pipelines contain three possible points of intervention to mitigate unwanted bias: the training data, the learning procedure, and the output predictions, and these are associated with three corresponding classes of bias mitigation strategy: pre-processing, in-processing, and post-processing:

1) Pre-processing approaches modify the training dataset to remove the discrimination before training an AI model. Common strategies used are under-sampling, over-sampling [21] or sample weighting [9] to neutralize discriminatory effects; data generation [14] or data augmentation [12] using generative adversarial networks to balance the training dataset; or training using a balanced dataset [20].

2) In-processing approaches try to modify state-of-the-art learning algorithms in order to remove discrimination during the model training process via model regularization. There are two main approaches: implicit regularization that adds implicit constraints which disentangle the association between model prediction and fairness sensitive attributes [3]; and explicit regularization that adds explicit constraints through updating the model’s loss function to minimize the performance difference between different protected groups [14, 23].

3) Post-processing approaches correct the output of an existing algorithm to satisfy the fairness requirements. Common strategies used are equalized odds post-processing [5] that solves a linear program to find probabilities with which to change output labels to optimize equalized odds; calibrated equalized odds post-processing [17] that optimizes over calibrated classifier score outputs to find probabilities with which to change output labels with an equalized odds objective; and reject option classification [10] that gives favorable outcomes to unprivileged protected groups and unfavorable outcomes to privileged protected groups in a confidence band around the decision boundary with the highest uncertainty.

⁸ In fair AI, the *protected attribute(s)* are the ones for which fairness needs to be ensured, e.g. gender or race. A set of samples with the same value(s) for the protected attribute(s) are known as a *protected group*.

3 Methods

To investigate whether training data imbalance can lead to bias between racial and/or gender groups in automated cardiac MR segmentation, we conducted a comparative study using five different approaches. The baseline approach was based on fairness through unawareness. To enable us to test the hypothesis that any observed bias was due to data imbalance, we also trained the segmentation network using a (smaller) race and gender balanced database. We subsequently investigated three approaches based on fairness through awareness (two pre-processing approaches: stratified batch sampling and protected group models and one in-processing method that we call a ‘fair meta-learning for segmentation’). These approaches are illustrated in detail in Fig. 1.

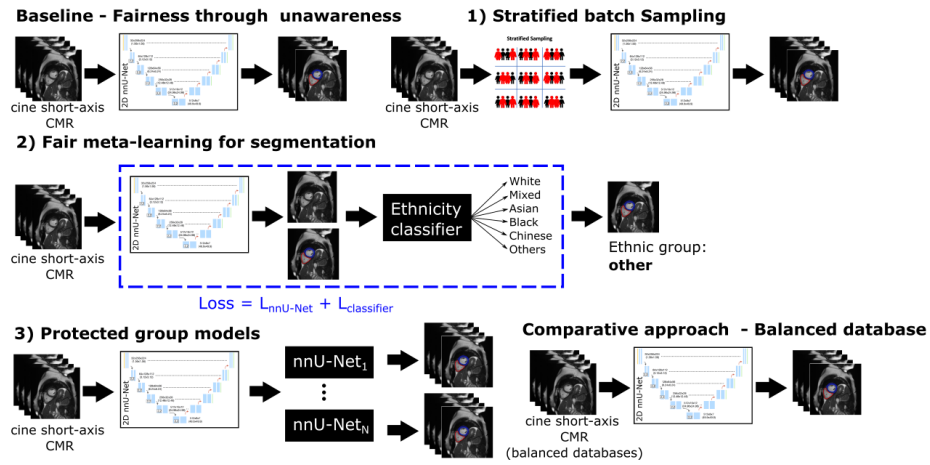


Fig. 1. Diagram showing the different strategies used for bias mitigation. Baseline is based on fairness through unawareness, and acts as the baseline approach. (1)-(3) are three approaches based on fairness through awareness: (1) stratified batch sampling, (2) ‘fair meta-learning for segmentation’ and (3) protected group models. The comparative approach of using a smaller balanced database allows us to test the hypothesis that observed bias is due to data imbalance.

Segmentation network: For all of the approaches we used the nnU-Net network [8] for automatic segmentation of the left ventricle blood pool (LVBP), left ventricular myocardium (LVM) and right ventricle blood pool (RVBP) from cine short-axis cardiac MR slices at end diastole (ED) and end systole (ES). Details of the data used and experimental set up are provided in Section 4.

Baseline - Fairness through unawareness: We trained the nnU-net model blinded to the protected attributes (i.e. race and gender). This will be considered the baseline approach in the rest of the paper and its performance can be directly compared with other published cardiac MR segmentation techniques [19, 1].

Comparative approach - Balanced database: This strategy aims to evaluate the scenario where the database is perfectly balanced. To this end, we randomly selected for each racial group the same number of subjects ($N=87$), which corresponds to the number of subjects in the smallest racial group. In total, 522 subjects were used for training this model.

Approach 1 - Stratified batch sampling: This strategy aims to modify the training sampling strategy to remove the discrimination before training. For each training batch, the data are stratified by the protected attribute(s) and samples are selected to ensure that each protected group is equally represented. This approach has previously been used to train fair classifiers [9].

Approach 2 - Fair meta-learning for segmentation: This strategy aims to add a meta-fair classifier to the segmentation network to train a model that not only segments cardiac MR images but also performs classification of the protected attribute(s). As the classifier we used a DenseNet network [6] and the input of the classifier was the cardiac MR image as one channel and the output of the nnU-Net segmentation network as a second channel. We formulate the problem as a multi-task learning problem where both networks are jointly optimized. The idea is that the classification network prevents the learning of a dominant group from negatively impacting the learning of another one. Both networks were first trained for 500 epochs independently, and then jointly trained for another 500 epochs without deep supervision. This is a novel approach for segmentation fairness but is based on the work proposed by Xu *et al.*, [23] for classification fairness, which combined a ResNet-18 network for facial expression recognition with a fully connected network for protected attribute classification.

Approach 3 - Protected group models: In contrast to Approaches 1 and 2, this strategy assumes that the protected attributes are available at inference time (as well as training time) and aims to train a different segmentation model for each protected group. As the number of subjects for each protected group can vary, which could impact performance, we initially trained the nnU-Net model for 500 epochs using the full training database, and then fine-tuned independent nnU-Net models for each of the protected groups for another 500 epochs. This approach is similar in concept to [20] in which classifiers were first trained using an unbalanced database and then fine-tuned using balanced ones. However, we choose to fine-tune using only individual protected group data and we are not aware of previous work that has taken this approach.

Evaluation metrics: In the literature, several fairness metrics have been proposed but because fairness research has so far focused on classification tasks all of them are tailored for such tasks. Because we deal with a segmentation task, we propose to evaluate our approaches using a segmentation overlap metric and two new segmentation fairness metrics. To measure segmentation overlap we use the average Dice similarity coefficient (DSC) over the three classes, i.e. LVBP, LVM and RVBP. As fairness metrics, we utilize the standard deviation (SD) and skewed error ratio (SER) of the average DSC values. The standard deviation reflects the amount of dispersion of the average DSC values between different protected groups. The SER is computed by the ratio of the highest

error rate to the lowest error rate among different protected groups and it can be formulated as $SER = \frac{\max_g(1-DSC_g)}{\min_g(1-DSC_g)}$, where g are the protected groups. The SD and SER fairness metrics were adapted from [20], which used classification accuracy and classification error rate for SD and SER respectively.

4 Materials and Experiments

We demonstrated our approach for fairness in DL-based cardiac MR segmentation using data from the UK Biobank [16]. The dataset used consisted of ED and ES short-axis cine cardiac MR images of 5,903 subjects (61.5 ± 7.1 years). Data on race and gender were obtained from the UK Biobank database and their distribution is summarized in Figure 2. For all subjects, the LV endocardial and epicardial borders and the RV endocardial border were manually traced at ED and ES frames using the cvi42 software (version 5.1.1, Circle Cardiovascular Imaging Inc., Calgary, Alberta, Canada). Each segmentation was generated by 1 of 10 experts who followed the same guidelines and were blinded to race and gender. Each expert contoured a random sample of images containing different races and genders. For all approaches (apart from the balanced database comparative approach), we used the same random split into training/validation/test sets of 4,723/590/590 subjects respectively. For the balanced database approach the training/validation/test set sizes were 417, 105 and 590. All models were trained on a NVIDIA GeForce GTX TITAN X. All networks were trained for 1,000 epochs and optimized using stochastic gradient descent with ‘poly’ learning rate policy (initial learning rate of 0.01 and Nesterov momentum of 0.99).

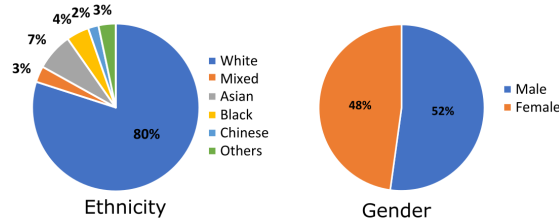


Fig. 2. The percentage of gender and different races in the UK Biobank dataset used in this paper.

Bias assessment: The first experiment aims to assess gender and racial bias for the baseline approach (i.e. ‘fairness through unawareness’). Table 1 shows the DSC values for LVBP, LVM and RVBP at ED and ES as well as the overall average DSC. Unpaired Student’s t -tests (significant difference reported for p -value < 0.01 with Scheffe post hoc test) were used for comparison between the full test set and each of the protected groups for the average DSC. From the results, we can make several important observations. First, there is a clear racial

bias. For example, when the racial distribution is 80% white and 20% others, the nnU-net model reaches 93.51% average DSC for white subjects, but this dramatically decreases to less than 86% for black subjects and less than 85% for mixed race subjects. Second, the results show that there is no significant gender bias. This suggests that the accuracy of the baseline segmentation model for each group is correlated with its representation in the training set. The results of the comparative approach (Balanced database, see bottom row of Table 2) support this conclusion, achieving the lowest SD and SER. As expected, the smaller training database size led to a significant reduction in accuracy for all racial groups.

Bias mitigation strategies: Since Table 1 showed that the baseline segmentation model had a racial bias but no gender bias, in this second experiment we focus only on the use of bias mitigation techniques to reduce racial bias. Table 2 shows the comparison between the baseline approach and the three approaches for bias mitigation. The results show that all mitigation strategies reduce the racial bias but that Approach 3 (Protected group models) achieved the lowest SD and SER. Note, however, that Approach 3 requires knowledge of the protected attributes at inference time. The best-performing approach that does not require such knowledge is Approach 1 (Stratified batch sampling). For Approach 2 (Fair meta-learning for segmentation), the average accuracy of the protected attribute classifier was 0.77 and the precision and recall values per group varied between 0.69 and 0.89.

Table 1. Dice similarity coefficient (DSC) for the LV blood pool (LVBP), LV myocardium (LVM) and RV blood pool (RVBP) at end diastole (ED) and end systole (ES) for Baseline - fairness through unawareness. The ‘Avg’ column that represents the average DSC values across LVBP, LVM and RVBP at both ED and ES has been boldfaced for easier comparison with Table 2. The first row reports the DSC for the full database, the second and third rows report DSC by gender and the remaining rows report DSC by racial group. Asterisks indicate statistically significant differences between the full population and each protected group for the average DSC.

DSC (%) for Baseline —Fairness through unawareness							
	ED			ES			Avg
	LVBP	LVM	RVBP	LVBP	LVM	RVBP	
Total	93.48	83.12	89.37	89.37	86.31	80.61	87.05
Male	93.58	83.51	88.82	90.68	85.31	81.00	87.02
Female	93.39	82.71	89.90	89.59	86.60	80.21	87.07
White	97.33	93.08	94.09	95.06	90.58	90.88	93.51*
Mixed	92.70	78.94	86.91	86.70	82.54	79.32	84.52*
Asian	94.53	87.33	90.51	90.13	88.94	81.94	88.90*
Black	92.77	85.93	89.49	89.42	85.74	71.91	85.88*
Chinese	91.81	74.51	85.74	86.39	85.12	79.34	83.82*
Others	91.74	78.94	89.50	88.53	84.96	80.27	85.66*

Figure 3 shows some sample segmentation results for different racial groups with both high and low DSC.




Fig. 3. Illustration of the segmentation results at ED for each racial group with high (top row) and low (bottom row) DSC.

Table 2. Comparison between baseline model (i.e. ‘fairness through unawareness’) and three approaches for bias mitigation. The values shown for each race represent the average DSC values across LVBP, LVM and RVBP at both ED and ES. The ‘Avg’ column is the average such value across all races (without weighting for the number of subjects in each group). We boldface the SD (lower is better) and SER (1 is the best) metrics since these are the important fairness criteria. The bottom row shows the results for using a smaller gender and race balanced training database.

Approach	Segmentation							Fairness	
	White	Mixed	Asian	Black	Chinese	Others	Avg	SD	SER
Baseline - Fairness through unawareness	93.51	84.52	88.90	85.88	87.63	85.66	87.68	3.25	2.38
1. Stratified batch sampling	90.88	93.84	93.65	93.07	94.35	93.50	93.22	1.22	1.62
2. Fair meta-learning for segmentation	92.75	88.03	90.64	89.60	88.18	88.27	89.58	1.86	1.65
3. Protected group models	91.03	93.17	93.34	92.15	93.04	93.08	92.64	0.89	1.35
Comparative approach - Balanced database	79.32	80.98	80.37	79.78	80.82	80.72	80.33	0.65	1.09

5 Discussion

To the best of our knowledge, this paper has presented the first study of fairness in AI-based image segmentation, focusing on the problem of data imbalance. We have shown, for the first time, that racial bias exists in DL-based cardiac MR segmentation models. Our hypothesis is that this bias is a result of the unbalanced nature of the training data, and this is supported by the results

which show that there is racial bias but not gender bias when trained using the UK Biobank database, which is gender-balanced but not race-balanced. We also found very low bias when training using a balanced database. One reason why bias has yet to be investigated in DL-based cardiac MR segmentation is that in current clinical practice the outputs of such models are typically manually modified by a clinician. However, this modification process is time-consuming and prone to error, so does not eliminate the bias completely. Furthermore, in the future, as AI models become more clinically accepted, it may be that the role of clinician oversight decreases or is removed completely, in which case lack of bias in DL-based segmentation models will become a critical feature.

We have proposed three ‘fairness through awareness’ bias mitigation strategies. The first strategy (Stratified batch sampling) ensures that in each batch all protected groups are equally represented. The second strategy (Fair meta-learning for segmentation) aims to make the model more aware of the protected attribute(s) by simultaneously training a protected attribute classifier. The last strategy (Protected group models) aims to train independent segmentation models for each protected group. Our results show that all three strategies improve fairness, with lower SD and SER values compared to the baseline model. These results are in keeping with previous works where the strategy of ‘fairness through unawareness’ has shown to be insufficient in preventing bias in several problems [21, 25].

It is important to note that the best-performing bias mitigation strategy (Approach 3) requires knowledge of the protected attribute at inference time. In general in medical image segmentation, such patient information is available via electronic health records. However, there are likely to be scenarios in which it is not available. In the case where the protected attribute(s) is not available at inference time it would be possible to use either Approach 1 or Approach 2. Approach 3 could be seen as a ‘best achievable’ model for such situations or where the protected attributes have been inadequately registered.

In conclusion, in this work we have highlighted the concerning issue of bias in DL-based image segmentation models. We have proposed three bias mitigation techniques, which were inspired by work from the literature on fairness in classification but which are all novel in the context of image segmentation. In the future, we aim to extend the current work to other kinds of existing fairness interventions, such as combinations of pre-processing, in-processing and post-processing strategies, as well as extending the analysis to other imaging modalities.

Acknowledgements

This work was supported by the EPSRC (EP/R005516/1 and EP/P001009/1), the Wellcome EPSRC Centre for Medical Engineering at the School of Biomedical Engineering and Imaging Sciences, King's College London (WT 203148/Z/16/Z) and has been conducted using the UK Biobank Resource under application numbers 17806 and 2964. SEP, SN and SKP acknowledge the BHF for fund-

ing the manual analysis to create a cardiac MR imaging reference standard for the UK Biobank imaging resource in 5000 CMR scans (PG/14/89/31194). We also acknowledge the following funding sources: the NIHR Biomedical Research Centre at Barts, EU’s Horizon 2020 (grant no. 825903, euCanShare project), the CAP-AI programme funded by the ERDF and Barts Charity, HDR UK, the Oxford NIHR Biomedical Research Centre, the Oxford BHF Centre of Research Excellence and the MRC eMedLab Medical Bioinformatics infrastructure (MR/L016311/1).

References

1. Bernard, O., et al.: Deep learning techniques for automatic MRI cardiac multi-structures segmentation and diagnosis: is the problem solved? *IEEE transactions on medical imaging* **37**(11), 2514–2525 (2018)
2. Buolamwini, J., Gebru, T.: Gender shades: Intersectional accuracy disparities in commercial gender classification. In: *Conference on fairness, accountability and transparency*. pp. 77–91. PMLR (2018)
3. Das, A., Dantcheva, A., Bremond, F.: Mitigating bias in gender, age and ethnicity classification: a multi-task convolution neural network approach. In: *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*. pp. 0–0 (2018)
4. Dwork, C., et al.: Fairness through awareness. In: *Proceedings of the 3rd innovations in theoretical computer science conference*. pp. 214–226 (2012)
5. Hardt, M., Price, E., Srebro, N.: Equality of opportunity in supervised learning. *arXiv preprint arXiv:1610.02413* (2016)
6. Huang, G., et al.: Densely connected convolutional networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 4700–4708 (2017)
7. Hwang, S., et al.: Fairfacegan: Fairness-aware facial image-to-image translation. *arXiv preprint arXiv:2012.00282* (2020)
8. Isensee, F., et al.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods* **18**(2), 203–211 (2021)
9. Kamiran, F., Calders, T.: Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems* **33**(1), 1–33 (2012)
10. Kamiran, F., Karim, A., Zhang, X.: Decision theory for discrimination-aware classification. In: *2012 IEEE 12th International Conference on Data Mining*. pp. 924–929. IEEE (2012)
11. Kishi, S., et al.: Race–ethnic and sex differences in left ventricular structure and function: The coronary artery risk development in young adults (CARDIA) study. *Journal of the American Heart Association* **4**(3), e001264 (2015)
12. Lu, K., et al.: Gender bias in neural natural language processing. In: *Logic, Language, and Security*, pp. 189–202. Springer (2020)
13. Mody, P., et al.: Most important articles on cardiovascular disease among racial and ethnic minorities. *Circulation: Cardiovascular Quality and Outcomes* **5**(4), e33–e41 (2012)
14. Ngxande, M., Tapamo, J.R., Burke, M.: Bias remediation in driver drowsiness detection systems using generative adversarial networks. *IEEE Access* **8**, 55592–55601 (2020)
15. Obermeyer, Z., et al.: Dissecting racial bias in an algorithm used to manage the health of populations. *Science* **366**(6464), 447–453 (2019)

16. Petersen, S.E., et al.: UK Biobank’s cardiovascular magnetic resonance protocol. *Journal of cardiovascular magnetic resonance* **18**(1), 1–7 (2015)
17. Pleiss, G., et al.: On fairness and calibration. arXiv preprint arXiv:1709.02012 (2017)
18. Raisi-Estabragh, Z., et al.: Variation of cardiac magnetic resonance radiomics features by age and sex in healthy participants from the UK Biobank. *European Heart Journal* **41**(Supplement_2), ehaa946–0197 (2020)
19. Ruijsink, B., Puyol-Antón, et al.: Fully automated, quality-controlled cardiac analysis from CMR: validation and large-scale application to characterize cardiac function. *Cardiovascular Imaging* **13**(3), 684–695 (2020)
20. Wang, M., Deng, W.: Mitigating bias in face recognition using skewness-aware reinforcement learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9322–9331 (2020)
21. Wang, Z., Qinami, K., et al.: Towards fairness in visual recognition: Effective strategies for bias mitigation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8919–8928 (2020)
22. Wilder, B., et al.: Clinical trial of an AI-augmented intervention for HIV prevention in youth experiencing homelessness. arXiv preprint arXiv:2009.09559 (2020)
23. Xu, T., et al.: Investigating bias and fairness in facial expression recognition. In: *European Conference on Computer Vision*. pp. 506–523. Springer (2020)
24. Yoneyama, K., et al.: Cardiovascular magnetic resonance in an adult human population: serial observations from the multi-ethnic study of atherosclerosis. *Journal of Cardiovascular Magnetic Resonance* **19**(1), 1–11 (2017)
25. Zhang, B.H., Lemoine, B., Mitchell, M.: Mitigating unwanted biases with adversarial learning. In: *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*. pp. 335–340 (2018)