

# Are Sex-Based Physiological Differences the Cause of Gender Bias for Chest X-Ray Diagnosis?

Weng, Nina; Bigdeli, Siavash; Petersen, Eike; Feragen, Aasa

Published in: Clinical Image-Based Procedures, Fairness of AI in Medical Imaging, and Ethical and Philosophical Issues in Medical Imaging

Link to article, DOI: 10.1007/978-3-031-45249-9\_14

Publication date: 2023

Document Version Early version, also known as pre-print

# Link back to DTU Orbit

Citation (APA):

Weng, N., Bigdeli, S., Petersen, E., & Feragen, A. (2023). Are Sex-Based Physiological Differences the Cause of Gender Bias for Chest X-Ray Diagnosis? In *Clinical Image-Based Procedures, Fairness of AI in Medical Imaging, and Ethical and Philosophical Issues in Medical Imaging* (pp. 142-152). Springer. https://doi.org/10.1007/978-3-031-45249-9\_14

#### **General rights**

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

• Users may download and print one copy of any publication from the public 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 public portal

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

# Are Sex-based Physiological Differences the Cause of Gender Bias for Chest X-ray Diagnosis?

Nina Weng<sup>1</sup>, Siavash Bigdeli<sup>1</sup>, Eike Petersen<sup>1</sup>, and Aasa Feragen<sup>1</sup>

Technical University of Denmark, Denmark {ninwe,sarbi,ewipe,afhar}@dtu.dk

Abstract. While many studies have assessed the fairness of AI algorithms in the medical field, the causes of differences in prediction performance are often unknown. This lack of knowledge about the causes of bias hampers the efficacy of bias mitigation, as evidenced by the fact that simple dataset balancing still often performs best in reducing performance gaps but is unable to resolve all performance differences. In this work, we investigate the causes of gender bias in machine learning-based chest X-ray diagnosis. In particular, we explore the hypothesis that breast tissue leads to underexposure of the lungs and causes lower model performance. Methodologically, we propose a new sampling method which addresses the highly skewed distribution of recordings per patient in two widely used public datasets, while at the same time reducing the impact of label errors. Our comprehensive analysis of gender differences across diseases, datasets, and gender representations in the training set shows that dataset imbalance is not the sole cause of performance differences. Moreover, relative group performance differs strongly between datasets, indicating important dataset-specific factors influencing male/female group performance. Finally, we investigate the effect of breast tissue more specifically, by cropping out the breasts from recordings, finding that this does not resolve the observed performance gaps. In conclusion, our results indicate that dataset-specific factors, not fundamental physiological differences, are the main drivers of male-female performance gaps in chest X-ray analyses on widely used NIH and CheXpert Dataset.

Keywords: Algorithmic fairness · Gender biases · Chest X-ray

# 1 Introduction

AI fairness receives increased attention with the escalating demand for examining the validity and responsibility of AI methods. This is particularly crucial in the medical field, where automatic and intelligent decision-making algorithms could easily lead to unfair treatment without the awareness of fairness.

A series of studies have assessed fairness in various medical imaging settings, including chest x-rays [14,20], retinal imaging [5], brain MRI [16,24], and cardiac MRI [18]. While different types of bias mitigation techniques have been applied [6, 15, 26, 27], there is currently very limited work that seeks to diagnose the causes of bias [17], enabling the targeted selection of bias mitigation methods. Supporting the urgency of this type of bias reasoning, [27] shows, by comparing the performance of different bias mitigation strategies, that simply balancing datasets is still one of the best strategies for mitigating gender bias for chest X-ray diagnosis. At the same time, studies have shown that while the level of group representation affects group-wise performance, balancing datasets alone does not guarantee equal performance across groups [14], see fig. 1 for an example. This implies that without further investigation of the causes of bias, our at-



AUROC

0.7



Male

Female

tempts to mitigate bias might be very limited in success, as also evidenced by the well-known 'leveling down' phenomenon [28].

In this paper, we investigate causes of gender bias in machine learning-based chest X-ray diagnosis, where significant performance disparities between genders are observed. Previous works [1, 3, 7, 12] have suggested that breast tissue might lead to impaired image quality, and hence lower performance, of chest X-ray-based diagnostic classifiers. Here, we perform a series of experiments to analyze whether the observed gender bias is indeed a result of physiological sex differences. In short, our contributions include:

- 1. A new way of sampling training and test sets from publicly available chest X-ray datasets that reduces the influence of potential confounders, such as a highly skewed distribution of the number of recordings per patient and missing disease labels, on training and analysis. In particular, we propose to sample just a single recording per patient, preferring samples with a disease-positive label. A comparison of different sampling strategies provides further strong evidence of label errors in publicly available datasets.
- 2. A combination of our proposed sampling method with training sets of varying gender ratios to provide **a comprehensive re-analysis of gender differences** in model performance across multiple diseases in two wellknown datasets (CheXpert [11] and NIH ChestX-ray8 [25]). Our results indicate that imbalanced datasets are not the only cause of performance differences, and, crucially, that gender-based performance differences differ between datasets even for the same disease.

3. Further experiments designed to study whether female breasts cause diagnostic classifier bias. We assess how cropping out the breasts from recordings affects model performance, and we consider inter-dataset transfer of models for more evidence of inherent gender bias from the dataset. We find that female breasts do not appear to be a strong cause of model performance disparities and that there appear to be other, presently unknown dataset-dependent factors influencing model performance.

# 2 Related Work

Following claims of radiologist-level performance in machine learning-based chest X-ray disease classification [19], performance disparities of such disease classifiers between patient groups have come under increased scrutiny [14, 20, 21]. Larrazabal et al. [14] analyzed the effect of gender imbalance in training sets on performance disparities in the trained classifier, finding a strong link between the two. Moreover, their results show how performance in some groups remains poor even if models are trained *only* on subjects from that group; refer to fig. 1 for an example. Seyved-Kalantari et al. [20,22] showed that state-of-the-art classifiers consistently and selective underdiagnosed historically underserved patient populations, such as non-white and female patients. Later, Zhang et al. [27] investigated a range of possible bias mitigation techniques, finding that simple group balancing still appears to be the most successful mitigation technique; this result has also been confirmed in other contexts [10]. Taken together, these results emphasize the importance of well-representative datasets. However, as Larrazabal et al. [14] had shown, group balancing alone cannot alleviate all performance differences, thus emphasizing the urgent need for more nuanced investigations into the *sources* of bias to enable successful bias mitigation.

In the wake of these important studies, there have been several investigations into potential sources of bias in chest X-ray datasets. In response to the study of Seyyed-Kalantari et al. [22], Bernhardt et al. [4] and Glocker et al. [8] pointed out the importance of properly accounting for confounding factors in bias analyses, such as age and disease distributions between patient groups. Bernhardt et al. [4] moreover underlined the challenge of properly evaluating performance differences if *label biases* affect both the training and test sets. There is reason for concern in this regard, since multiple studies have reported high error rates in (NLPderived [11]) chest X-ray disease labels [23,27]. Our results in this study provide further independent evidence of widespread label errors in these databases.

Separately from these methodological issues, the performance differences between male and female patients led Ganz et al. [7] to speculate that an important cause might lie in female breasts occluding the recordings of important lung regions. Indeed, the confounding effect of female breasts on clinical chest x-ray interpretation is well-known [1, 3, 12] and physiologically plausible: additional breast tissue results in the underexposure of lung tissue. In this work, we take a step to systematically assess the effect of female breast tissue on machine learning-based chest X-ray diagnosis.

# 3 Methods

#### 3.1 Datasets

We consider two datasets: ChestX-ray8 (NIH) [25] and CheXpert [11]. As the NIH dataset only contains frontal images, we also only use those views from the CheXpert to enable a fair comparison, resulting in 112,120 recordings from 30,850 patients in the NIH dataset and 190,299 recordings from 64,224 patients in the CheXpert dataset. Both datasets slightly over-represent male subjects (54% males vs. 46% females in NIH, 56% males vs. 45% females in CheXpert); refer to table 1 in the supplementary material for further details.

#### 3.2 Sampling Strategy

**Motivation.** We observe that in both datasets, the number of recordings per patient is very uneven, ranging from 1 to 89 (CheXpert) and 1 to 184 (NIH). In particular, less than 25% of controls and over 50% of patients have more than 5 scans. This results in few patients with many recordings strongly influencing the training process and the final model, as well as strong distribution shifts between different data splits.

In addition, like outlined above, it has been shown that *disease labels au*tomatically derived from patient records are unreliable [27]. Particularly, it has been observed [2] that the commonly used text mining method worked poorly with the hospital record of "no change from previous", which would be wrongly marked as "no finding". Labels might be especially unreliable in patients with many recordings.

For these reasons outlined above, we select only one sample from each patient to avoid over-representation with a preference for positive labelled samples.

**Principles underlying the proposed sampling strategy.** We designed our sampling strategy to conform with the following principles:

- Utilize one sample per patient: reducing risk of distribution shift between splits and over-reliance on individual subjects.
- Prioritize diseased samples when selecting the single sample per patient: reducing risk of label bias towards "no finding" like outlined above.
- Keep disease prevalence constant across splits: reducing risk of distribution shift between splits.
- Allow disease prevalence to vary between protected groups: ensuring that our assessment is realistic, by utilizing realistic prevalences across groups.
- For training and validation, draw a fixed-size sample with a predefined percentage of female subjects (0%, 50% or 100%): enabling an assessment of the influence of training set composition on model performance. (For testing, we always draw a fixed-size sample with an equal number of males/females.)
- Use an identical test set when evaluating the same training split at different gender ratios: enabling more reliable performance comparisons.

For further details on our sampling scheme, refer to Algorithm 1 in the supplementary material.



**Fig. 2.** Model performance (AUROC) in male (blue) and female (yellow) test subjects for training datasets with differing gender representation. Results are only shown for the six disease labels that are present in both datasets; for results on the remaining labels refer to the supplementary material. Icons in the bottom right of each plot indicate the observed performance trends across gender ratios; refer to Section 4.1.

#### 3.3 Experimental Settings

All experiments are carried out using the PyTorch framework with a pretrained ResNet50 [9] and the Adam optimizer [13] (learning rate  $10^{-6}$ , 20 epochs, batch size 64). We split the datasets into 60% / 10% / 30% train, validation, and test sets per gender ten times based on disease prevalence, and then train separate single-label classifiers for all disease labels in both datasets. During training, data augmentation is applied by random horizontal flipping, rotation (degree  $\leq 15^{\circ}$ ) and scaling (from 0.9 to 1.1) with a probability of 0.5 each. Following [14], experiments are run on three different gender ratios, i.e. 0%, 50% and 100% females. Performances are evaluated by the area under the receiver-operating characteristic curve (AUROC).

#### 4 Results

#### 4.1 Model performance across diseases, gender ratios, and datasets

Figure 2 displays male and female test subject performance under varied intraining gender ratios for the common 6 disease labels in both datasets. Moreover, a stylized summary of the observed trends across gender ratios (' $\times$ ', '>', '<', or '=') is marked in the figure. We highlight three key observations.

Dataset imbalance is not the sole cause of performance differences. In some dataset–disease combinations, such as Pneumothorax–NIH, males keep outperforming females regardless of the proportion of women in the training set. Similar and opposite trends (males having worse performance than females across gender ratios) could be observed in other disease labels, marked by '>' and '<' in fig. 2. If performance differences were caused solely by training set imbalance, then the majority group should consistently outperform the minority group (resulting in an '×' trend shape), which is not observed in most of the cases.



Fig. 3. Comparison of results from 10 varied splits when training and testing using three different sampling strategies for three disease labels. Three rows present one sample per patient (npp=1) with/without prioritizing diseased samples, and the original setting from [14] which uses all the samples with multi-label learning. Blue box represents male performance and yellow refers female.

Performance trends sometimes differ strongly between datasets for the same diseases. Consider again the case of Pneumothorax classification, and compare the results between the two datasets: On NIH, males demonstrate significantly better performance regardless of training gender distribution, while on CheXpert, the trend has reversed completely. This suggests that there are dataset-specific factors unrelated to fundamental biological sex differences that strongly influence performance disparities.

Higher performance and higher variance on NIH compared to CheXpert. Across all six shared disease labels, test-set prediction performance is considerably better on NIH compared to CheXpert, despite the smaller NIH training set size. At the same time, the NIH performance also shows a larger variance across diseases, which appears likely to be related to the smaller test set size.

#### 4.2 Comparison of different sampling strategies

To analyze the effect of our proposed sampling strategy, we first compare the results when training and evaluating using three different sampling strategies: 1) one sample per patient, preferring disease-positive samples (our proposed strategy), 2) one sample per patient, drawing the sample randomly and not preferring disease-positive samples, and 3) using all samples per patient with the same split and multi-label training as [14]. Note that in the latter case, we use larger training set sizes. As shown in fig. 3, in both datasets, test set performance is



Fig. 4. Comparison of results when drawing the training set according to one sampling strategy and the test set according to another, for the disease label Pneumonia. Error bars computed from a single training run using test set bootstrapping.

consistently and considerably improved when using the disease-preferring sampling strategy over a random sampling of one recording per patient; we interpret this as evidence of strong label noise in the "no finding" label like discussed in section 3.2 and like reported previously in other datasets automatically labeled using the CheXpert labeller [27]. The difference between the two setups is especially strong in the NIH dataset, leading to the hypothesis that NIH dataset might suffer from more widespread label errors compared to CheXpert.

Since label noise confounds not only training but also evaluation, we also present results for one disease when the training on one strategy and the test set according to another (fig. 4). Due to the potential data leakage through train and test set, instead of using the splits from [14], we implemented the experiments with all samples (npp=None) using modified proposed strategy without sampling. The results indicate label errors are more likely occur in NIH than CheXpert, as the results in CheXpert are more stable across sampling strategies while the performance drops when testing on all samples in NIH.

#### 4.3 Breast cropping does not mitigate gender biases

To assess specifically whether differences in male and female breast physiology account for the observer performance differences, we perform an additional experiment. We simply crop the lower two fifths of each recording to ensure that the images contain only the parts above the breast for both genders. An illustration and the results of this experiment are provided in fig. 5. Compared to fig. 2, while overall performance drops slightly for both gen-



**Fig. 5.** The illustration of cropped chest x-rays and the results.

ders, this intervention does not close the performance gap between both genders.

# 4.4 Dataset Bias v.s. Model Bias

To further investigate the effect of dataset-specific factors, we also evaluated classifiers trained for one disease on CheXpert using NIH test sets, and vice versa (fig. 6). We observe that large male-female performance gaps only arise when models are tested on NIH, regardless of which dataset the model is trained on.



Fig. 6. Inter-dataset inference (50% female in training, *Pneumothorax*) with booststrapping.

# 5 Discussion & Conclusions

We interpret the results of this work and address the key findings as follows.

Supporting evidence of frequent label errors in the "no finding" label. In agreement with previous reports on other datasets [27], our analyses provide supporting evidence of the existing hypothesis on frequent mislabeling of records as "no finding" using the proposed simple and easily applicable sampling method.

Male-female performance gaps are influenced, but not fully explained, by training set representation. In line with the previous results of Larrazabal et al. [14], we find that training set representation does (in some diseases strongly) influence male and female model performance. It does not, however, fully explain the performance gaps: in some cases, even when trained on a fully female dataset, models still perform worse on women (and vice versa).

Male-female performance gaps are not primarily caused by breast shadows. It had been previously hypothesized [7] that breast shadows might play an essential role in gender bias in ML-based chest X-ray diagnosis, which is not supported by our findings: breast cropping does not mitigate the performance gaps. Additionally, the varied bias trends between datasets also contradict this hypothesis.

Biological differences may not be the main driver of male-female performance gaps. The performance gaps should be expected to be consistent across datasets (or at least within the same disease), if biological sex differences were the main driver of performance gaps, which has not been observed in this work.

Dataset-specific factors strongly influence male-female performance gaps. As the previous hypotheses on the origins of male-female performance gaps have been tentatively rejected considering the results, further research should continue investigating other potential sources of bias. Those could be the distribution of various confounders, the prevalence of further label errors, or differing recording quality [8, 17, 27]. We conclude that there must be further, at present unknown dataset-specific factors driving the observed performance gaps. Acknowledgements. Work on this project was partially funded by the Independent Research Fund Denmark (DFF, grant number 9131-00097B), Denmark's Pioneer Centre for AI (DNRF grant number P1), a Google Award for inclusion research, the Novo Nordisk Foundation through the Center for Basic Machine Learning Research in Life Science (MLLS, grant number NNF20OC0062606).

# References

- 1. Chest. https://radiologykey.com/chest-11/, accessed: 2023-6-20
- Half a million x-rays! First impressions of the Stanford and MIT chest x-ray datasets. https://laurenoakdenrayner.com/2019/02/25/ half-a-million-x-rays-first-impressions-of-the-stanford-and-mit-chest-x-ray-datasets/, accessed: 2023-6-20
- 3. Alexander, C.: The elimination of confusing breast shadows in chest radiography. Australasian Radiology **2**(2), 107–108 (1958)
- Bernhardt, M., Jones, C., Glocker, B.: Potential sources of dataset bias complicate investigation of underdiagnosis by machine learning algorithms. Nature Medicine 28(6), 1157–1158 (jun 2022). https://doi.org/10.1038/s41591-022-01846-8
- Burlina, P., Joshi, N., Paul, W., Pacheco, K.D., Bressler, N.M.: Addressing artificial intelligence bias in retinal diagnostics. Translational Vision Science & Technology 10(2), 13–13 (2021)
- Daneshjou, R., Vodrahalli, K., Novoa, R.A., Jenkins, M., Liang, W., Rotemberg, V., Ko, J., Swetter, S.M., Bailey, E.E., Gevaert, O., Mukherjee, P., Phung, M., Yekrang, K., Fong, B., Sahasrabudhe, R., Allerup, J.A.C., Okata-Karigane, U., Zou, J., Chiou, A.S.: Disparities in dermatology AI performance on a diverse, curated clinical image set. Science Advances 8(32) (aug 2022). https://doi.org/10.1126/sciadv.abq6147
- Ganz, M., Holm, S.H., Feragen, A.: Assessing bias in medical ai. In: Workshop on Interpretable ML in Healthcare at International Connference on Machine Learning (ICML) (2021)
- Glocker, B., Jones, C., Bernhardt, M., Winzeck, S.: Algorithmic encoding of protected characteristics in chest x-ray disease detection models. eBioMedicine 89, 104467 (mar 2023). https://doi.org/10.1016/j.ebiom.2023.104467
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
- Idrissi, B.Y., Arjovsky, M., Pezeshki, M., Lopez-Paz, D.: Simple data balancing achieves competitive worst-group-accuracy. In: Schölkopf, B., Uhler, C., Zhang, K. (eds.) Proceedings of the First Conference on Causal Learning and Reasoning. Proceedings of Machine Learning Research, vol. 177, pp. 336-351. PMLR, https: //proceedings.mlr.press/v177/idrissi22a.html
- Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciurea-Ilcus, S., Chute, C., Marklund, H., Haghgoo, B., Ball, R., Shpanskaya, K., et al.: Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In: Proceedings of the AAAI conference on artificial intelligence. vol. 33, pp. 590–597 (2019)
- 12. Jenkins, P.: Making sense of the chest x-ray: a hands-on guide. CRC Press (2013)
- Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)

- 10 Weng et al.
- Larrazabal, A.J., Nieto, N., Peterson, V., Milone, D.H., Ferrante, E.: Gender imbalance in medical imaging datasets produces biased classifiers for computer-aided diagnosis. Proceedings of the National Academy of Sciences 117(23), 12592–12594 (2020)
- Pakzad, A., Abhishek, K., Hamarneh, G.: CIRCLe: Color invariant representation learning for unbiased classification of skin lesions. In: Lecture Notes in Computer Science, pp. 203–219. Springer Nature Switzerland (2023). https://doi.org/10.1007/978-3-031-25069-9\_14
- Petersen, E., Feragen, A., da Costa Zemsch, M.L., Henriksen, A., Wiese Christensen, O.E., Ganz, M., Initiative, A.D.N.: Feature robustness and sex differences in medical imaging: A case study in mri-based alzheimer's disease detection. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 88–98. Springer (2022)
- Petersen, E., Holm, S., Ganz, M., Feragen, A.: The path toward equal performance in medical machine learning. Patterns 4(7) (2023). https://doi.org/10.1016/j.patter.2023.100790
- Puyol-Antón, E., Ruijsink, B., Mariscal Harana, J., Piechnik, S.K., Neubauer, S., Petersen, S.E., Razavi, R., Chowienczyk, P., King, A.P.: Fairness in cardiac magnetic resonance imaging: assessing sex and racial bias in deep learning-based segmentation. Frontiers in cardiovascular medicine 9, 859310 (2022)
- Rajpurkar, P., Irvin, J., Zhu, K., Yang, B., Mehta, H., Duan, T., Ding, D., Bagul, A., Langlotz, C., Shpanskaya, K., Lungren, M.P., Ng, A.Y.: Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning (Nov 2017). https://doi.org/10.48550/ARXIV.1711.05225
- Seyyed-Kalantari, L., Liu, G., McDermott, M., Chen, I.Y., Ghassemi, M.: Chexclusion: Fairness gaps in deep chest x-ray classifiers. In: BIOCOMPUTING 2021: proceedings of the Pacific symposium. pp. 232–243. World Scientific (2020)
- Seyyed-Kalantari, L., Zhang, H., McDermott, M.B.A., Chen, I.Y., Ghassemi, M.: Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. Nature Medicine 27(12), 2176–2182 (dec 2021). https://doi.org/10.1038/s41591-021-01595-0
- Seyyed-Kalantari, L., Zhang, H., McDermott, M.B., Chen, I.Y., Ghassemi, M.: Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. Nature medicine 27(12), 2176–2182 (2021)
- 23. Smit, A., Jain, S., Rajpurkar, P., Pareek, A., Ng, A., Lungren, M.: Combining automatic labelers and expert annotations for accurate radiology report labeling using BERT. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP). Association for Computational Linguistics (2020). https://doi.org/10.18653/v1/2020.emnlp-main.117
- Stanley, E.A., Wilms, M., Mouches, P., Forkert, N.D.: Fairness-related performance and explainability effects in deep learning models for brain image analysis. Journal of Medical Imaging 9(6), 061102–061102 (2022)
- 25. Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M.: Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2097–2106 (2017)
- Wu, Y., Zeng, D., Xu, X., Shi, Y., Hu, J.: FairPrune: Achieving fairness through pruning for dermatological disease diagnosis. In: Lecture Notes in Computer Science, pp. 743–753. Springer Nature Switzerland (2022). https://doi.org/10.1007/978-3-031-16431-6\_70

- Zhang, H., Dullerud, N., Roth, K., Oakden-Rayner, L., Pfohl, S., Ghassemi, M.: Improving the fairness of chest x-ray classifiers. In: Conference on Health, Inference, and Learning. pp. 204–233. PMLR (2022)
- Zietlow, D., Lohaus, M., Balakrishnan, G., Kleindessner, M., Locatello, F., Schölkopf, B., Russell, C.: Leveling down in computer vision: Pareto inefficiencies in fair deep classifiers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10410–10421