

# Gender Imbalance in Medical Imaging Datasets Produces Biased Classifiers for Computer-aided Diagnosis

Agostina J. Larrazabal    Nicols Nieto    Victoria Peterson  
Diego H. Milone        Enzo Ferrante

May 26, 2020

## Abstract

Artificial intelligence systems for computer assisted diagnosis and image-based screening are being adopted worldwide by medical institutions. In such a context, generating fair and unbiased classifiers becomes of paramount importance. The research community of medical image computing is making great efforts in developing more accurate algorithms to assist medical doctors in the difficult task of disease diagnosis. However, little attention is paid to the way databases are collected and how this may influence the performance of AI systems. Our study sheds light on the importance of gender balance in medical imaging datasets used to train AI systems for computer assisted diagnosis. We provide empirical evidence supported by a large scale study, based on three deep neural network architectures and two well-known publicly available X-ray image datasets used to diagnose various thoracic diseases under different gender imbalance conditions. We found a consistent decrease in performance for under-represented genders when a minimum balance is not fulfilled. This raises the alarm for national agencies in charge of regulating and approving computer assisted diagnosis systems, which should include explicit gender balance and diversity recommendations. We also establish a new open problem for the academic medical image computing community which needs to be addressed by novel algorithms endowed with robustness to gender imbalance.

Artificial intelligence (AI) influences almost every aspect of our daily life. The media articles we read, the movies we watch, even the driving roadmap we take are somehow influenced by these systems. In particular, the rise of AI in healthcare during the last few years is changing the way medical doctors diagnose, especially when dealing with medical images. AI systems can not only augment the information provided by such images with useful annotations [20, 19], but they are also starting to take autonomous decisions by performing computer assisted diagnosis (CAD) [8, 5].

Although the interest in performing fair and unbiased evaluations of AI medical systems exists since the 80's [3], the ethical aspects of AI have gained relevance in the last few years. It has been shown that human bias, such as gender and racial bias, may not only be inherited but also amplified by AI systems in multiple contexts [34, 14, 1, 26]. For example, face recognition systems have shown to exhibit accuracy disparities depending on gender and ethnicity, being darker-skinned females the most misclassified group [2]. This tendency of AI systems to learn biased models, which reproduce social stereotypes and underperform in minority groups, is specially dangerous in the context of healthcare [33, 4].

In the last years, the research community of gendered innovations has largely contributed to create awareness and integrate sex and gender analyses into all phases of basic and applied research [25]. However, such assessment in the context of medical imaging and CAD remains largely unexplored. In this work, we perform a large-scale study that quantifies the influence of gender imbalance in medical imaging datasets used to train AI-based CAD systems. It is worth mentioning that most of the existing work dealing with imbalanced data in the context of deep learning focuses on cases where it is related to the target classes [10, 17]. In our study, this would translate to an imbalance in terms of number of patients per pathology. However, note that in this case the imbalance is given by a demographic variable different from the target class: the gender, which is generally neglected. Our results show that using gender imbalanced datasets to train deep learning based CAD systems may affect the performance in pathology classification for minority groups.

## Results and discussion

A model based on deep neural networks, which achieves state-of-the-art results when diagnosing 14 common thoracic diseases using X-ray images [24], was implemented to perform computer assisted diagnosis. We employed the area under the receiver operating characteristic curve (AUC) [9] to quantify its performance. Figure 1 shows the experimental results obtained when training the classifier under different gender imbalance ratios. In 1.A, the boxplots aggregate the results for 20 experiments using fully imbalanced datasets. The blue boxes represent the performance for models trained only with male images, while orange boxes indicate training with female-only images. Both models are evaluated over male-only (A, top) and female-only (A, bottom) test images. A consistent decrease in performance is observed when using male patients for

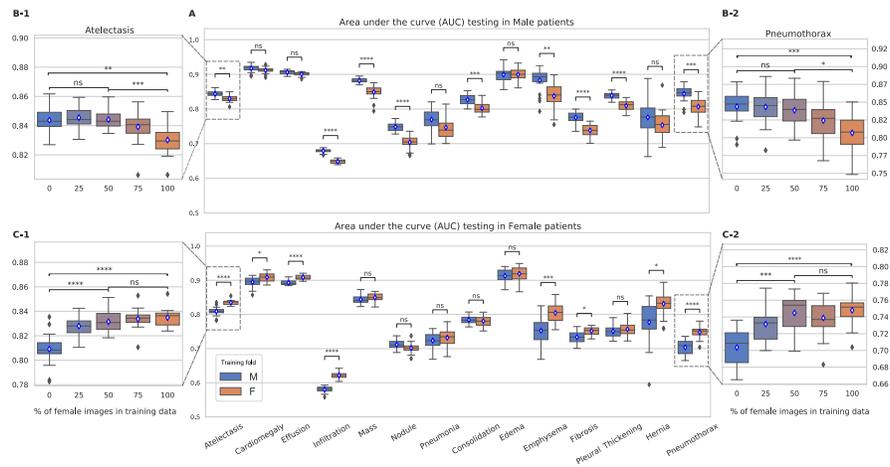


Figure 1: Experimental results for a DenseNet-121 [13] classifier trained with images from the NIH dataset [32, 23] for 14 thoracic diseases under different gender imbalance ratios. **(A)** The boxplots aggregate the results for 20 folds, training with male-only (blue) and female-only (orange) patients. Both models are evaluated given male and female test folds. A consistent decrease in performance is observed when using male patients for training and female for testing (and viceversa). **(B-C)** Area under the receiver operating characteristic curve (AUC) achieved for two exemplar diseases under a gradient of imbalance gender ratios, from 0% of female images in training data to 100%, with increments of 25%. **(B-1,2)** show the results when testing on male patients, while **(C-1,2)** presents the results when testing on female patients. Statistical significance according to Mann-Whitney U test is denoted by \*\*\*\* ( $p \leq 0.00001$ ), \*\*\* ( $0.00001 < p \leq 0.0001$ ), \*\* ( $0.0001 < p \leq 0.001$ ), \* ( $0.001 < p \leq 0.01$ ) and ns ( $p > 0.01$ ).

training and female for testing (and viceversa). The same tendency was confirmed when evaluating three different deep learning architectures in two X-ray datasets with different pathologies.

We also explored intermediate imbalance scenarios, where both female and male patients were present in the training dataset but considering different proportions (0/100%, 25/75% and 50/50%). Figures 1.B and 1.C show the average classification performance for two exemplar diseases, Pneumothorax and Atelectasis, under such gradient of gender imbalance ratios (indicated with the percentage of female patients used for training). We found that even with a 25%/75% imbalance ratio, the average performance across all diseases in the minority class is significantly lower than a model trained with a perfectly balanced dataset. Moreover, we did not find significant differences in performance between models trained with a gender balanced dataset (50% male and 50% female) and an extremely imbalanced dataset from the same gender. In other words, a CAD system trained with a diverse (and balanced) dataset achieved the best performance for both genders. Altogether, our results indicate that diversity provides additional information and increases the generalization capability of AI systems. Thereafter, it also suggests that diversity should be prioritized when designing databases used to train machine learning-based CAD systems.

Our study shows that gender imbalance in medical imaging datasets produces biased classifiers for computer-aided diagnosis based on CNNs, with significantly lower performance in under-represented groups. We provide experimental evidence in the context of X-ray image classification for such potential bias, aiming to raise the alarm not only within the medical image computing community but also for national agencies in charge of regulating and approving medical systems. As an example, let us take the US Food and Drug Administration (FDA) agency. Even though they have released several documents related to the importance of gender/sex issues in the design and evaluation of clinical trials and medical devices [30], when looking at the specific guidelines to obtain the certification to market medical computer assisted systems [29, 28], there is no explicit mention to gender/sex as one of the relevant demographic variables that should describe the sampled population. Similar issues are observed in the medical imaging community. Albeit a few datasets provide this information at the subject level, most public datasets of similar characteristics do not contain gender/sex information at the patient level to date (e.g. the recent MIMIC-CXR [16] X-ray dataset or the REFUGE database of opthalmological images [22], just to name a few). The same tendency is observed in many of the datasets included in a recent analysis of 150 databases from grand challenges on biomedical image analysis [21], which provides recommendations for database and challenge design, where there is no explicit mention to the importance of sex/gender demographic information.

In general, it is well known that CNNs tend to learn representations useful to solve the task they are being trained for. When we go from male to female images (or viceversa), structural changes in the images appear, leading to a change in data distribution which explains the decrease in performance. Algorithmic solutions to such *domain adaptation* problem [31] should be engineered, specially

in cases when it is difficult to obtain gender-balanced datasets (e.g. ABIDE I [7]).

## Datasets

We use the NIH Chest-XRay14 dataset [32, 23] which includes 112,120 chest X-ray images from 30,805 patients, labeled with 14 common thorax diseases (including hernia, pneumonia, fibrosis, emphysema, edema, cardiomegaly, pleural thickening, consolidation, mass, pneumothorax, nodule, atelectasis, effusion and infiltration). Labeling was performed according to an automatic natural language processing (NLP) analysis of the radiology reports. The dataset provides demographic information including patient’s gender: 63,340 (56.5%) images for male and 48,780 (43.5%) for female patients. Following the demographic variables reported in the original dataset publication [32], we used the term *gender* to characterize our imbalance study. However, given that some anatomical attributes are reflected in X-ray images, the term *sex* could be more accurate according the Sex and Gender Equity in Research (SAGER) guidelines [12]. The CheXpert database [15] was also used to confirm that our observations generalize for different datasets. It contains 224,316 chest radiographs of 65,240 patients with diagnostic information ( $\sim 60\%$  male and  $\sim 40\%$ ). The uncertainty labels included in CheXpert were interpreted as negative following the U-Zeros approach discussed in the original paper [15].

## Deep learning model

Deep neural networks are machine learning methods with multiple abstraction levels, which compose simple but non-linear modules transforming representations at one level into a representation at a higher, slightly more abstract level [18]. A special type of deep neural network, known as convolutional neural networks (CNNs), was used to implement the CAD system [24, 32]. Results shown in Figure 1 correspond to a Densely Connected CNN (DenseNet) architecture with 14 outputs, one for each disease [13].

We adopted a Keras implementation of the DenseNet-121 which has shown to achieve state-of-the-art results in X-ray image classification [24]. The network has 121 convolutional layers and a final fully connected layer producing a 14-dimensional output, after which we apply an element-wise sigmoid non-linearity. A model pre-trained on ImageNet [6] was used to initialize the network weights. We trained it end-to-end using Adam optimizer with standard parameters ( $\beta_1=0.9$  and  $\beta_2=0.999$ ), a batch size of 32 and an initial learning rate of 0.001 that was decayed by a factor of 10 each time the validation loss plateaus after an epoch. Additionally, we evaluated two other CNN architectures, the ResNet [11] and the Inception-v3 [27], confirming that our observations generalize for different neural models.

## Methodology

Since images can be labeled with multiple diseases, we implemented an automatic method to construct random splits which guarantees that male and female folds will have the same number of images per pathology. Given a frontal X-ray image, the CAD system predicts the presence or absence of the 14 thoracic diseases. Two models were trained in each experiment, one considering male-only while the other one female-only training datasets. Intermediate imbalance scenarios were also analyzed, in which female and male images were presented in the training dataset at different proportions (0/100%, 25/75% and 50/50%). To avoid other sources of bias, care was taken to guarantee by training data construction that male and female folds include the same number of pathological cases per class. For the NIH Chest-XRay14, every split included 48,568 images. For the CheXpert dataset every split included 27,147 images. The same experiment was performed 20 times, using different random splits. In the testing phase, both models were evaluated in male and female patients separately. The classification performance was measured by the well-known area under the receiving operating curve (AUC) [9].

## Data Availability

The NIH Chest-XRay14 dataset is publicly available at <https://nihcc.app.box.com/v/ChestXray-NIHCC>.

The CheXpert dataset is publicly available at <https://stanfordmlgroup.github.io/competitions/chexpert/>.

The source code of the original CNNs is publicly available at <https://github.com/brucechou1983/CheXNet-Keras>.

The modified version of this code with our auxiliary scripts, the data splits used in our experiments and the additional results for all the CNN architectures in both datasets can be accessed at [https://github.com/N-Nieto/GenderBias\\_CheXNet](https://github.com/N-Nieto/GenderBias_CheXNet).

## References

- [1] T. Bolukbasi, K.-W. Chang, J. Y. Zou, V. Saligrama, and A. T. Kalai. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Advances in neural information processing systems*, pages 4349–4357, 2016.
- [2] J. Buolamwini and T. Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *ACM FAccT*, pages 77–91, 2018.
- [3] B. Chandrasekaran. On evaluating artificial intelligence systems for medical diagnosis. *AI magazine*, 4(2):34–34, 1983.

- [4] D. S. Char, N. H. Shah, and D. Magnus. Implementing machine learning in health care addressing ethical challenges. *The New England journal of medicine*, 378(11):981, 2018.
- [5] J. De Fauw, J. R. Ledsam, B. Romera-Paredes, S. Nikolov, N. Tomasev, S. Blackwell, H. Askham, X. Glorot, B. ODonoghue, D. Visentin, et al. Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nature medicine*, 24(9):1342, 2018.
- [6] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [7] A. Di Martino, D. Oconnor, B. Chen, K. Alaerts, J. S. Anderson, M. Assaf, J. H. Balsters, L. Baxter, A. Beggiato, S. Bernaerts, et al. Enhancing studies of the connectome in autism using the autism brain imaging data exchange II. *Scientific data*, 4:170010, 2017.
- [8] A. Esteva, B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau, and S. Thrun. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639):115, 2017.
- [9] T. Fawcett. An introduction to roc analysis. *Pattern recognition letters*, 27(8):861–874, 2006.
- [10] G. Haixiang, L. Yijing, J. Shang, G. Mingyun, H. Yuanyue, and G. Bing. Learning from class-imbalanced data: Review of methods and applications. *Expert Systems with Applications*, 73:220–239, 2017.
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [12] S. Heidari, T. F. Babor, P. De Castro, S. Tort, and M. Curno. Sex and gender equity in research: rationale for the sager guidelines and recommended use. *Research Integrity and Peer Review*, 1(1):2, 2016.
- [13] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [14] M. Hutson et al. Even artificial intelligence can acquire biases against race and gender. *Science Magazine*, 10, 2017.
- [15] J. Irvin, P. Rajpurkar, M. Ko, Y. Yu, S. Ciurea-Ilcus, C. Chute, H. Marklund, B. Haghighi, R. Ball, K. Shpanskaya, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI Conference*, volume 33, pages 590–597, 2019.

- [16] A. E. Johnson, T. J. Pollard, S. J. Berkowitz, N. R. Greenbaum, M. P. Lungren, C.-y. Deng, R. G. Mark, and S. Horng. Mimic-cxr, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific Data*, 6, 2019.
- [17] J. M. Johnson and T. M. Khoshgoftaar. Survey on deep learning with class imbalance. *Journal of Big Data*, 6(1):27, 2019.
- [18] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436, 2015.
- [19] R. Lindsey, A. Daluiski, S. Chopra, A. Lachapelle, M. Mozer, S. Sicular, D. Hanel, M. Gardner, A. Gupta, R. Hotchkiss, et al. Deep neural network improves fracture detection by clinicians. *Proceedings of the National Academy of Sciences*, 115(45):11591–11596, 2018.
- [20] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. Van Der Laak, B. Van Ginneken, and C. I. Sánchez. A survey on deep learning in medical image analysis. *Medical image analysis*, 42:60–88, 2017.
- [21] L. Maier-Hein, M. Eisenmann, A. Reinke, S. Onogur, M. Stankovic, P. Scholz, T. Arbel, H. Bogunovic, A. P. Bradley, A. Carass, et al. Why rankings of biomedical image analysis competitions should be interpreted with care. *Nature communications*, 9(1):1–13, 2018.
- [22] J. I. Orlando, H. Fu, J. B. Breda, K. van Keer, D. R. Bathula, A. Diaz-Pinto, R. Fang, P.-A. Heng, J. Kim, J. Lee, et al. Refuge challenge: A unified framework for evaluating automated methods for glaucoma assessment from fundus photographs. *Medical image analysis*, 59:101570, 2020.
- [23] C. Qin, D. Yao, Y. Shi, and Z. Song. Computer-aided detection in chest radiography based on artificial intelligence: a survey. *Biomedical engineering online*, 17(1):1–23, 2018.
- [24] P. Rajpurkar, J. Irvin, K. Zhu, B. Yang, H. Mehta, T. Duan, D. Ding, A. Bagul, C. Langlotz, K. Shpanskaya, et al. Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning. *arXiv preprint arXiv:1711.05225*, 2017.
- [25] L. Schiebinger and M. Schraudner. Interdisciplinary approaches to achieving gendered innovations in science, medicine, and engineering1. *Interdisciplinary Science Reviews*, 36(2):154–167, 2011.
- [26] G. Stanovsky, N. A. Smith, and L. Zettlemoyer. Evaluating gender bias in machine translation. *arXiv preprint arXiv:1906.00591*, 2019.
- [27] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of IEEE CVPR*, pages 2818–2826, 2016.

- [28] U.S. Food and Drug Administration. Computer-Assisted Detection Devices Applied to Radiology Images and Radiology Device Data - Premarket Notification [510(k)] Submissions, 2012. Online; accessed 23 March 2020.
- [29] U.S. Food and Drug Administration. Clinical Performance Assessment: Considerations for Computer-Assisted Detection Devices Applied to Radiology Images and Radiology Device Data - Premarket Approval (PMA) and Premarket Notification [510(k)] Submissions , 2020. Online; accessed 23 March 2020.
- [30] U.S. Food and Drug Administration. Understanding Sex Differences at FDA, 2020. Online; accessed 23 March 2020.
- [31] M. Wang and W. Deng. Deep visual domain adaptation: A survey. *Neurocomputing*, 312:135–153, 2018.
- [32] X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri, and R. M. Summers. 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*, pages 2097–2106, 2017.
- [33] J. Wiens, S. Saria, M. Sendak, M. Ghassemi, V. X. Liu, F. Doshi-Velez, K. Jung, K. Heller, D. Kale, M. Saeed, et al. Do no harm: a roadmap for responsible machine learning for health care. *Nature medicine*, pages 1–4, 2019.
- [34] J. Zou and L. Schiebinger. AI can be sexist and racistits time to make it fair. *Nature*, 559:324–326, 2018.