Analyzing Cross-Population Domain Shift in Chest X-Ray Image Classification and Mitigating the Gap with Deep Supervised Domain Adaptation

Aminu Musa[1], Mariya Ibrahim Adamu[2], Habeebah Adamu Kakudi[2], Monica Hernandez[4], and Yusuf Lawal[3]

[1] Federal University Dutse, Nigeria
musa.aminu@fud.edu.ng
[2] Bayero University Kano, Nigeria
hakakudi.cs@buk.edu.ng
[3] Aminu Kano Teaching Hospital, Nigeria
[4] Aragon Institute on Engineering Research, Spain, PID2022-138703OB-I00

Abstract. Medical image analysis, empowered by artificial intelligence (AI), plays a crucial role in modern healthcare diagnostics. However, the effectiveness of machine learning models hinges on their ability to generalize to diverse patient populations, presenting domain shift challenges. This study investigates the domain shift problem within chest X-ray classification, with a particular emphasis on cross-population variations, especially within underrepresented groups. We examine the domain shift of a supervised version of Adversarial Domain Adaptation (ADA) across three distinct population datasets (sources), using a Nigerian chest X-ray dataset as the target dataset. By evaluating model performance, we quantify the disparities between the source and target populations. Our experiments revealed varying model performance when trained on the source domain and evaluated on the target domain. To address this variability, we propose a supervised domain adaptation technique that leverages labeled data from both domains for fine-tuning. The results demonstrate significant enhancements in model accuracy for chest X-ray classification in the Nigerian dataset. This research underscores the importance of domain-aware model development in AI-driven healthcare, contributing to addressing cross-population domain-shift challenges in medical imaging.

Keywords: Chest X-ray, deep-learning, domain shift, domain adaptation

1 Introduction

The discipline of medical image analysis has experienced a profound transformation through the integration of artificial intelligence (AI) technologies [1]. AI models have demonstrated promising capabilities in analyzing medical images, thereby assisting medical professionals in achieving accurate diagnoses and devising effective treatment plans [2]. Notably, recent advancements in computer vision have facilitated the utilization of deep learning in medical image analysis, fostering the development of accurate and efficient models for disease diagnosis [3]. Nevertheless, the application of AI techniques to medical imaging in the context of diverse populations and datasets presents various challenges, including issues related to generalization, biases, and safety concerns [4].

Chest X-ray analysis plays a crucial role in the diagnosis and management of a wide range of respiratory and cardiovascular conditions [5]. Nowadays, AI models are increasingly used with this imaging modality for a wide range of tasks such as image segmentation, image registration, analysis of respiration motion, detection of anatomical features, disease diagnosis, and prognosis [6]. The performance of deep learning models is being widely explored for these tasks [7].

Traditionally, the interpretation and analysis of medical images is carried out by expert radiologists or physicians [3]. Due to the lack of professional radiologists in the developing world and the enormous stress faced by physicians as a result of heavy workload, human interpretation of medical images is prone to errors. This results in longer time for initial diagnosis, delays in the follow up, and chance of misdiagnosis. AI models are known to be accurate in classification and object detection tasks [8]. Therefore, integrating AI models in chest X-ray analysis and interpretation may greatly improve the process, and offer accurate and timely analysis to support decision-making. However, variations in patient demographics, imaging protocols, and equipment across different populations can result in significant domain shifts, potentially hindering the performance of AI models trained on one population when applied to another, especially if they come from different races [4].

Many research works focus on investigating domain shifts as a result of different equipment, different hospitals, or cross-modality [7, 9–11]. These models tend to be biased when tested on samples from underrepresented groups in the dataset. One of the primary causes of model bias has been attributed to the lack of diversity of the training set where minority populations are underrepresented [4]. The challenges become further difficult to address, as there are scarce publicly available datasets from these underrepresented groups.

In this paper, we study the critical issue of domain shift in chest X-ray image classification, particularly concerning diverse populations. We recognize that conventional machine learning and deep learning models, even when pretrained on extensive datasets, may encounter challenges in adapting to new populations due to inherent domain differences. Our focus extends to quantifying population-based domain shifts within African chest X-ray datasets from Nigeria, where unique demographic and clinical characteristics can introduce domain-specific complexities in accurately classifying X-ray images. By scrutinizing and addressing this domain shift, our objective is to address the generalizability and clinical applicability of AI-driven chest X-ray image classification across diverse populations.

To achieve this goal, we propose and evaluate a domain adaptation approach aimed at mitigating the adverse effects of domain shift. We address domain adaptation employing Domain Adversarial Neural Networks (DANNs), which prioritizes the adjustment of invariant features within the target domain to narrow the disparity between the pre-trained source domain and the Nigerian target domain. Our analysis entails a comprehensive examination of domain shift in chest X-ray image classification, utilizing a meticulously curated dataset representative of the Nigerian population. We delve into the intricacies of feature-level domain adaptation techniques, elucidating their application in enhancing the performance of AI models across disparate domains. Through extensive experimentation, we demonstrate the efficacy of our proposed approach in addressing domain shift and enhancing the accuracy of chest X-ray image classification. Consequently, our contributions advance the realm of AI-driven healthcare solutions by fostering adaptability to diverse populations.

2 Literature Review

In recent years, the fusion of artificial intelligence (AI) with medical image analysis has paved the way for transformative advancements in healthcare diagnostics [12]. Among these applications, chest X-ray image classification holds particular significance for its role in diagnosing a wide range of respiratory and cardiovascular conditions [13]. Several X-ray diagnosis systems based on machine learning have been proposed, achieving promising performance [5, 14, 15]. However, the successful application of AI models in this domain often hinges on their adaptability to account for the distinct characteristics of diverse patient populations [16]

The term cross-population domain shift describes variations in data distributions across different populations. This is a widely recognized challenge in medical image analysis, especially with the advent of deep-learning [17]. Extensive studies underscore the impact of domain shifts on the performance of AI models. For example, Rajpurkar et al. highlighted that models trained on data from one population may not generalize effectively to others, leading to wide disparities in performance [18].

Domain adaptation (DA) techniques have emerged as a promising avenue to tackle the issue of domain shift [19]. Feature-level adaptation techniques have garnered significant attention, demonstrating their ability to align feature distributions [20]. Ganin et al. proposed domain-adversarial training and unsupervised domain adaptation [21, 22], which have demonstrated success in mitigating domain discrepancies. In addition, He et al. introduced a classification-aware semi-supervised domain adaptation technique, which shows promise for leveraging limited labeled data effectively [23]. Techniques such as data augmentation through latent space interpolation [24] have also emerged as valuable tools for enhancing model generalization.

While the concept of domain adaptation has been well-explored in computer vision, a real value can be seen in the context of medical imaging applications in healthcare domains [25]. Given the unique challenges inherent in chest X-ray image classification across diverse datasets [26], it becomes evident that domain shift originating from variations across different patient populations could introduce patient-specific discrepancies. Addressing this pressing issue is crucial for ensuring the generalizability and effectiveness of AI-based diagnostic systems in clinical practice. Although DA techniques adapt to variations in data from different population are highly pertinent, they are relatively limited. Some preliminary studies have begun to explore DA techniques on different medical imaging datasets. For instance, Feng et al. applied domain adaptation methods to different chest X-ray datasets, highlighting the potential of these techniques to adapt models to new populations and domains [27]. Seyyed-Kalantari et al. investigated the issue of bias in machine learning assisted diagnosis system, in under-served patient population [28]. Their work confirmed the presence of cross-population variability in medical imaging analysis. The comprehensive survey conducted by Yu et al. provides valuable insights into the landscape of domain adaptation techniques in medical image analysis, offering a roadmap for future research directions in this rapidly evolving field [25].

Domain shift is one of the major issues that undermines the generalizability potential of deep learning based diagnosis systems [29]. DA can be categorized into to two broad categories: 1) Unsupervised DA, which tries to align data distributions of the features in the feature space with label source data, mapped to unlabelled target dataset [30], [31]. 2) Supervised methods, which align the distribution gap between labelled source and target domain [32], [33]. Other works emphasize the importance of measuring domain shift for deep learning models in medical images [34]. In the work of Stacke et al., the authors proved the existence of domain shift in medical images data as result of different data acquisition pipeline, different medical facility, or over time [9]. Results from their experiments demonstrate how the proposed measure outperforms existing methods for measuring domain shift and uncertainty, by having a strong association with performance decline when testing a model across a wide range of different types of domain changes [9].

3 Datasets and Methods

3.1 Nigerian X-ray dataset

The target domain dataset used in this research is a locally collected dataset of chest X-rays from the Radiology Centre of Aminu Kano Teaching Hospital (AKTH), Nigeria. The dataset consists of 6 345 X-ray images. The images were annotated by three different physicians into three different categories: pneumonia, tuberculosis, and normal X-rays. The distribution of the images per category is: 2 340 X-ray images diagnosed with tuberculosis by physicians, 1 445 diagnosed with pneumonia, and 2 560 termed as normal X-rays. The image dimension was 299×299 pixels with 72 pixels/inch DPI. The dataset will be released for public use under Apache license 2 on Kaggle. Figure 1 shows representative sample images from this dataset.

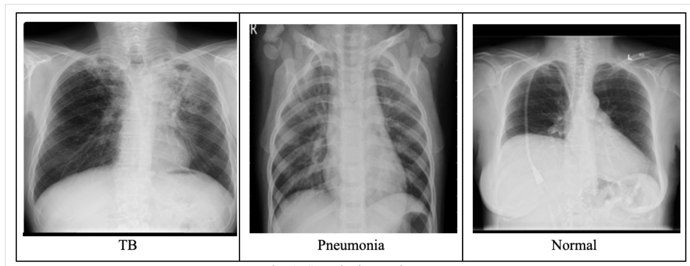

Fig. 1. Representative images from the Nigerian dataset. Left, image from the tuberculosis (TB) group. Centre, image from the pneumonia (P) group. Right, image from the healthy normal (HC) group.

3.2 Baseline X-ray datasets

Three different chest X-ray datasets were selected as baseline datasets for our study. The first dataset is the chest X-ray classification dataset, collected at Guangzhou Medical Center, China [5]. The dataset is made of 5 863 images. The images were annotated by expert physicians into pneumonia and normal X-rays. The second dataset is VinDr-CXR, which consist of 18 000 images annotated by 17 different physicians into 6 different classes of diagnosis including pneumonia and tuberculosis [35]. The dataset was collected in Hanoi Medical University Hospital, Vietnam. The third dataset is COVID-19 radiography database, organized by a team of researchers from Qatar University, Doha [36]. The dataset has 1 345, pneumonia X-rays, 1 341 normal X-rays, and 1 200 COVID-19 X-rays. In the following, we will refer to the datasets as China, Vietnam, and Doha.

3.3 Baseline models and domain shift analysis

We employ a DenseNet201 CNN model to obtain the baseline models [37]. The architecture is selected to be one of the best known models in classifying X-ray images. Three different models were trained with the three baseline datasets selected for this study, respectively, and the best performance in our Nigerian test set was recorded. The models were trained with ImageNet weights using fine tuning, by freezing the initial layers, to ensure that only weights of the unfrozen layers are updated during training. These allows the model to preserve the generic features while learning domain-specific representations through the unfrozen layer. We replaced the last layer with a dense layer and applied the softmax activation function with three labels. All the models utilized a similar hyper-parameter setup: categorical cross-entropy as loss function and restricting the computation of the loss (for training and validation) to the labels from the target domain. Stochastic gradient descent was used as the optimizer, with the initial learning rate of 0.001, momentum of 0.9, and the mini-batch size to 32. After each epoch, we reduced the learning rate by a factor of 10 if the validation loss did not improve. The batch images underwent data augmentation

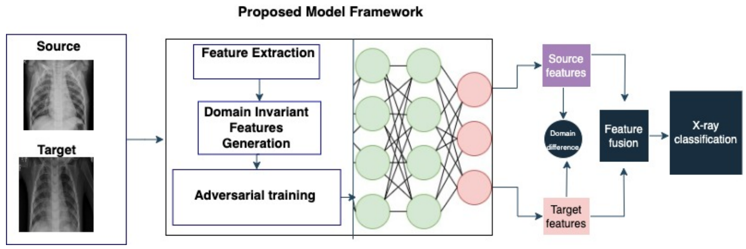

Fig. 2. Architecture of our Adversarial Neural Network (ADA-NN), where the input data comes from two distinct datasets (a labeled source and a target), and fed into separate feature extraction blocks. These features are then merged through a feature fusion block before being passed to a classifier, while domain adversarial training encourages the feature extractor to learn domain-invariant representations, enhancing the model's ability to generalize across domains.

using common techniques such as scaling, rotation around the image center, translation relative to the image extent, and zooming in. Finally, all the networks were evaluated on the target domain data. Five fold cross-validation was used, utilizing the 10% of the samples from the source data as test set. The Nigerian dataset is used separately as the target test dataset to evaluate the domain adaptation performance.

3.4 Adaptation with Domain Adversarial Neural Networks (DANNs)

Adversarial Domain Adaptation (ADA) is one of the most powerful domain adaptation techniques employed to mitigate the challenges posed by domain shifts in machine learning, specifically within the context of computer vision tasks [38]. It offers a principled framework for adapting models trained on a source domain to perform effectively on a target domain, even if there are significant distribution differences between the two domains. The idea of leveraging adversarial learning is to create domain-invariant feature representations. This is achieved through a two-step process: 1) Feature Extractor and Label Predictor. This step works similarly to a conventional deep learning architecture. The feature extractor transforms the input data (images) into a feature space, while the label predictor performs the primary task of image classification. 2) Domain Discriminator. A neural network performs the task of determining the domain (source or target) of the feature representations extracted by the feature extractor.

Domain adaptation is achieved through adversarial learning. Specifically, the feature distribution between source and target domains are effectively aligned through adversarial adaptation in the shared feature space. The feature extractor learns to produce domain-invariant features for the task, making the model adaptable to new domains different from the source population datasets, without requiring extensive retraining of the model on the target datasets.

Dataset	**Accuracy (%)**	**Precision (%)**	**F1-score (%)**	**AUC**
China	79.58 ± 0.72	74.0 ± 3.74	78.20 ± 2.74	0.88 ± 0.14
Vietnam	96.32 ± 0.12	93.0 ± 0.44	96.12 ± 0.74	0.97 ± 0.04
Doha	94.89 ± 1.74	92.22± 7.74	95.07 ± 1.16	0.95 ± 0.03

Table 1. Performance of DenseNet models trained and tested on China, Vietnam, and Doha datasets, respectively.

Dataset	**Accuracy (%)**	**Precision (%)**	**F1-score (%)**	**AUC**
China	62.45 ± 3.31	54.34 ± 3.33	60.45 ± 1.08	0.65 ± 0.02
Vietnam	66.02 ± 2.77	45.67 ± 2.08	65.43 ± 3.22	0.66 ± 0.03
Doha	71.70 ± 2.54	68.41 ± 8.21	69.00 ± 5.05	0.73 ± 0.04

Table 2. Performance of DenseNet models trained on China, Vietnam, and Doha datasets and tested on the Nigerian dataset.

Dataset	**Accuracy (%)**	**Precision (%)**	**F1-score (%)**	**AUC**
China	81.85 ± 8.59	79.07 ± 3.47	83.72 ± 2.28	0.93 ± 0.03
Vietnam	89.40 ± 0.44	83.02 ± 9.89	90.01 ± 0.61	0.92 ± 0.08
Doha	88.70 ± 7.46	78.44 ± 3.17	70.70 ± 5.46	0.94 ± 0.16
Nigeria	90.08 ± 2.25	87.29 ± 0.34	89.75± 0.93	0.96 ± 0.01

Table 3. Performance of ADA models trained on China, Vietnam, Doha, and Nigerian dataset and tested on the Nigerian dataset.

While ADA is popularly used in unsupervised settings where the datasets are unlabelled, we propose a modification for supervised ADA. The domain discriminator loss function is updated based on its ability to distinguish between features from the source and target domains. This adversarial objective ensures that the feature extractor learns domain-invariant features. For the discrimination step, the objective is to minimize the combined classification loss from both domains. The approach depicted in Figure 2 aims at illustrating the general workflow of our proposed ADA technique.

4 Experimental Results

To evaluate the effectiveness of our domain adaptation approach, we used the Nigerian chest X-ray dataset for both feature alignment and testing. Specifically, the 80% of the Nigerian dataset was employed as unlabeled target data during the feature alignment process. This subset was used to train the domain discriminator and align the feature distributions between the source and target domains. The remaining 20% of the Nigerian dataset was reserved for testing purposes only, ensuring that the evaluation of the model's performance on the target domain was conducted on data not seen during training or feature alignment. This approach guarantees a fair assessment of the generalization capabilities of the model

4.1 Domain shift analysis

In this section, we assess the performance of the DenseNet models trained in the baseline datasets and tested in our Nigerian dataset. Table 1 and Table 2 gather the performance results of the three models tested on both the source and target domains, respectively. Our analysis reveals a significant decline in performance, with accuracies ranging from 79.70% to 96.70% in the source domains, contrasting sharply with the 62.45% to 71.70% accuracy observed in the target domain. This notable decrease underscores the presence of a cross-population domain shift among the datasets, with the Chinese model demonstrating the most pronounced decline. Consequently, it becomes evident that DenseNet struggles to generalize effectively to our target domain when solely trained on baseline source domains.

4.2 ADA results

Table 3 presents the results of our adversarial models over the Nigerian test set. The metrics show a considerable performance improvement accross all the datasets, with an accuracy now ranging from 81.85 to 89.4% when tested in the target domain. The accuracy of the ADA model trained over the Nigerian test set reached the 90.08%, which is close to Vietnam and Doha performance. These results indicate the effectiveness of ADA on handling the cross-population domain shifts.

5 Conclusions

In this paper, we investigated the potential cross-population bias of deep-learning models in a relevant medical domain such as X-ray classification. Our research introduces a new chest X-ray data for the Nigerian population and uncovers cross-population domain shifts in deep-learning-based X-ray classification models. We investigate the extent of domain shift between different sources and target populations and propose the use of domain adversarial networks as a domain adaptation strategy. Through empirical evaluation and analysis, we have quantified the extent of the downgrade in performance due to the domain shifts.

Our results revealed notable disparities in classification performance across the different adaptation scenarios. Particularly, our scenarios, where the source and target populations exhibited substantial demographic and clinical diversity, resulted into high domain shift challenges. These findings underscore the necessity of domain adaptation techniques for enhancing model generalizability across diverse populations. By thoroughly characterizing domain shift, we illuminate the path toward mitigating its effects and improving the robustness of chest X-ray image classification models. This analysis informs our exploration of domain adaptation strategies, including supervised fine-tuning and domain-specific adaptations, as effective means of addressing these domain discrepancies and achieving better adaptation outcomes.

The paper proposed a solution by leveraging the adversarial adaptation framework ADA to address the challenge of cross-population domain shift in chest X-ray image classification. We showed the effectiveness of ADA in mitigating domain shift and greatly improving classification performance for diverse patient populations, contributing to the advancement of AI-driven healthcare solutions.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Bajwa, J., Munir, U., Nori, A., Williams, B.: Artificial intelligence in healthcare: transforming the practice of medicine. Future Healthc J **8** (2021) e188–e194
2. Singh, V.K., Rashwan, H.A., Abdel-Nasser, M., Akram, F., Haffar, R., Pandey, N., Arenas, M., Romani, S., Puig, D.: A computer-aided diagnosis system for breast cancer molecular subtype prediction in mammographic images. Elsevier eBooks (2021) 153–178
3. Reardon, S.: Rise of robot radiologists. Nature **576** (2019) S54–S58
4. Ricci, A., Echeveste, R., Ferrante, E.: Addressing fairness in artificial intelligence for medical imaging. Nature Communications **13** (2022)
5. : Chest X-Ray, https://www.hopkinsmedicine.org/health/treatment-tests-and-therapies/chest-xray (2019)
6. Shen, D., Wu, G., Suk, H.I.: Deep learning in medical image analysis. Annual Review of Biomedical Engineering **19** (2017) 221–248
7. Luo, Y., Zheng, L., Guan, T., Yu, J., Yang, Y.: Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. ArXiv (2018)
8. : Imagenet, https://www.image-net.org/challenges/lsvrc/2017/index.php (2017)
9. Stacke, K., Eilertsen, G., Unger, J., Lundström, C.: Measuring domain shift for deep learning in histopathology. IEEE Journal of Biomedical and Health Informatics **25** (2021) 325–336
10. Hong, J., Yu, S.C.H., Chen, W.: Unsupervised domain adaptation for cross-modality liver segmentation via joint adversarial learning and self-learning. Applied Soft Computing **121** (2022) 108729
11. Xing, F., Bennett, T.D., Ghosh, D.: Adversarial domain adaptation and pseudo-labeling for cross-modality microscopy image quantification. Lecture Notes in Computer Science **11840** (2019) 740–749
12. Kermany, D.S., Goldbaum, M., Cai, W., Valentim, C.C., Liang, H., Baxter, S.L., McKeown, A., Yang, G., Wu, X., Yan, F., Dong, J., Prasadha, M.K., Pei, J., Ting, M.Y., Zhu, J., Li, C., Hewett, S., Dong, J., Ziyar, I., Shi, A., Zhang, R., Zheng, L., Hou, R., Shi, W., Fu, X., Duan, Y., Huu, V.A., Wen, C., Zhang, E.D., Zhang, C.L., Li, O., Wang, X., Singer, M.A., Sun, X., Xu, J., Tafreshi, A., Lewis, M.A., Xia, H., Zhang, K.: Identifying medical diagnoses and treatable diseases by image-based deep learning. Cell **172** (2018) 1122–1131.e9
13. Musa, A., Adam, F.M., Ibrahim, U., Zandam, A.Y.: Learning from small datasets: An efficient deep learning model for covid-19 detection from chest X-Ray using dataset distillation technique. In: 2022 IEEE Nigeria 4th International Conference

on Disruptive Technologies for Sustainable Development (NIGERCON). (2022) 1–6
14. Won Jo, S., Seok, J.: A study on deep learning-based classification for pneumonia detection. In: 2022 13th International Conference on Information and Communication Technology Convergence (ICTC). (2022) 1496–1498
15. Vinoth, R., Subalakshmi, S., Thamaraichandra, S.: Pneumonia detection from chest X-Ray using alexnet image classification technique. In: 2023 7th International Conference on Intelligent Computing and Control Systems (ICICCS). (2023) 1307–1312
16. Ranjan, V., Harit, G., Jawahar, C.: Domain adaptation by aligning locality preserving subspaces. ICAPR 2015 - 2015 8th International Conference on Advances in Pattern Recognition (2015)
17. Cohen, J.P., Hashir, M., Brooks, R., Bertrand, H.: On the limits of cross-domain generalization in automated X-Ray prediction. ArXiv (2020)
18. Rajpurkar, P., Irvin, J., Zhu, K., Yang, B., Mehta, H., Duan, T., Ding, D.Y., Bagul, A., Langlotz, C., Shpanskaya, K., Lungren, M.P., Ng, A.Y.: Chexnet: Radiologist-level pneumonia detection on chest X-Rays with deep learning. ArXiv (2017)
19. Lenga, M., Schulz, H., Saalbach, A.: Continual learning for domain adaptation in chest X-Ray classification. Proceedings of Machine Learning Research (2020)
20. Zhao, T.: Seismic facies classification using different deep convolutional neural networks. Seg Technical Program Expanded Abstracts (2018)
21. Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. In: International conference on machine learning, PMLR (2015) 1180–1189
22. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., March, M., Lempitsky, V.: Domain-adversarial training of neural networks. Journal of machine learning research **17** (2016) 1–35
23. He, G., Liu, X., Fan, F., You, J.: Classification-aware semi-supervised domain adaptation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. (2020) 964–965
24. Liu, X., Zou, Y., Kong, L., Diao, Z., Yan, J., Wang, J., Li, S., Jia, P., You, J.: Data augmentation via latent space interpolation for image classification. In: 2018 24th International Conference on Pattern Recognition (ICPR), IEEE (2018) 728–733
25. Yu, M., Guan, H., Fang, Y., Yue, L., Liu, M.: Domain-prior-induced structural mri adaptation for clinical progression prediction of subjective cognitive decline. In: International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI), Springer (2022) 24–33
26. Shelke, A., Inamdar, M., Shah, V., Tiwari, A., Hussain, A., Chafekar, T., Mehendale, N.: Chest X-Ray classification using deep learning for automated covid-19 screening. Sn Computer Science **2** (2021) 300
27. Feng, Y., Xu, X., Wang, Y., Lei, X., Teo, S.K., Zheng, J., Ting, Y., Zhen, L., Zhou, J.T., Liu, Y., Tan, C.H.: Deep supervised domain adaptation for pneumonia diagnosis from chest X-Ray images. IEEE Journal of Biomedical and Health Informatics **26** (2022) 1080–1090
28. 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** (2021) 2176–2182
29. Pooch, E.H.P., Ballester, P., Barros, R.C.: Can we trust deep learning based diagnosis? the impact of domain shift in chest radiograph classification. In Petersen, J., San José Estépar, R., Schmidt-Richberg, A., Gerard, S., Lassen-Schmidt, B., Jacobs, C., Beichel, R., Mori, K., eds.: Thoracic Image Analysis, Cham, Springer International Publishing (2020) 74–83

30. Long, M., Cao, Y., Wang, J., Jordan, M., Edu, J.: Learning transferable features with deep adaptation networks. Proceedings of Machine Learning Research (2015)
31. Thiam, P., Lausser, L., Kloth, C., Blaich, D., Liebold, A., Beer, M., Kestler, H.A.: Unsupervised domain adaptation for the detection of cardiomegaly in cross-domain chest X-Ray images. Frontiers in Artificial Intelligence **6** (2023) 1056422
32. He, B., Chen, Y., Zhu, D., Xu, Z.: Domain adaptation via wasserstein distance and discrepancy metric for chest X-Ray image classification. Research Square (2023)
33. He, B., Chen, Y., Zhu, D., Xu, Z.: Domain adaptation via wasserstein distance and discrepancy metric for chest X-Ray image classification. Scientific Reports **14** (2024) 2690
34. Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
35. : VinDr-CXR: An open dataset and benchmarks for disease classification and abnormality localization on chest radiographs — vindr (2020)
36. Rahman, T.: Covid-19 radiography database (2020)
37. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017) 2261–2269
38. HassanPour Zonoozi, M., Seydi, V.: A survey on adversarial domain adaptation. Neural Process Lett **55** (2023) 2429–2469