



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

FedMedICL: Towards Holistic Evaluation of Distribution Shifts in Federated Medical Imaging

Kumail Alhamoud^{1*}, Yasir Ghunaim^{2,3*}, Motasem Alfarra², Thomas Hartvigsen⁴, Philip Torr³, Bernard Ghanem², Adel Bibi³, and Marzyeh Ghassemi¹

¹ MIT, USA

² King Abdullah University of Science and Technology (KAUST), Saudi Arabia

³ University of Oxford, UK

⁴ University of Virginia, USA

Correspondence to: kumail@mit.edu, yasir.ghunaim@kaust.edu.sa

Abstract. For medical imaging AI models to be clinically impactful, they must generalize. However, this goal is hindered by *(i)* diverse types of distribution shifts, such as temporal, demographic, and label shifts, and *(ii)* limited diversity in datasets that are siloed within single medical institutions. While these limitations have spurred interest in federated learning, current evaluation benchmarks fail to evaluate different shifts simultaneously. However, in real healthcare settings, multiple types of shifts co-exist, yet their impact on medical imaging performance remains unstudied. In response, we introduce FedMedICL, a unified framework and benchmark to holistically evaluate federated medical imaging challenges, simultaneously capturing label, demographic, and temporal distribution shifts. We comprehensively evaluate several popular methods on six diverse medical imaging datasets (totaling 550 GPU hours). Furthermore, we use FedMedICL to simulate COVID-19 propagation across hospitals and evaluate whether methods can adapt to pandemic changes in disease prevalence. We find that a simple batch balancing technique surpasses advanced methods in average performance across FedMedICL experiments. This finding questions the applicability of results from previous, narrow benchmarks in real-world medical settings. Code is available at: <https://github.com/m1k2zoo/FedMedICL>.

Keywords: Federated Learning · Continual Learning · Distribution Shifts.

1 Introduction

Machine learning increasingly impacts medical imaging [15,30]. Yet, most FDA-approved models are validated on datasets unrepresentative of real demographic distributions, leading to potential inaccuracies in medical diagnostics [11]. This lack of diversity in data is problematic when models encounter distribution

* Equal contribution. Yasir did part of this work while interning at Oxford.

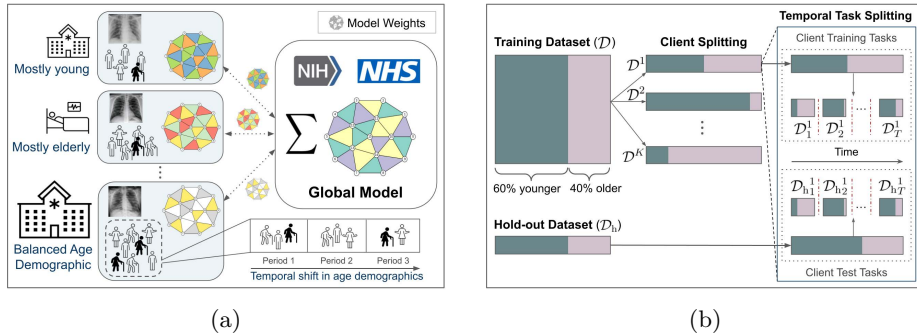


Fig. 1: (a) **Problem Setup**. We model a federated medical imaging scenario, in which siloed hospitals experience demographic imbalances and temporal shifts. (b) **FedMedICL Benchmark Construction**. We construct client datasets (D^1 to D^K), each representing a hospital with unique demographic characteristics and temporal training tasks (D_1^i to D_T^i). We evaluate models on temporally aligned test tasks for testing adaptability to local demographic shifts, and on a hold-out set (D_h) to evaluate generalization to diverse demographics.

shifts [31]. These shifts reflect discrepancies in data characteristics between the training/validation datasets and actual clinical settings. Thus, they can degrade the performance of diagnostic models across diverse patient populations [23,27].

Merging medical datasets from different hospitals, known as data pooling, can enhance training data diversity and address distribution shifts [22]. However, the practical implementation of this approach is often challenged by institutional policies, and slow data access processes [26,29], which severely limit data sharing across healthcare centers. As a result, there is a heavy reliance on smaller, *siloed* datasets, that hinder the deployment of generalizable medical imaging models [11].

As exemplified in Figure 1a, we investigate two key factors impacting the development of robust medical imaging models: the challenges posed by distribution shifts and the siloed nature of medical data. We frame the complications that arise from this complex setup under three types of shifts. These types are:

- ➊ **Label Imbalance**: In individual institutions, some medical conditions are more prevalent, leading to models failing on rarer conditions [1,18]. For instance, a skin disease classification model trained in one region may perform poorly on conditions more common in another geographic region [2,20].
- ➋ **Demographic Imbalance**: Patient demographics vary significantly among healthcare centers. For example, hospitals in American suburbs, which have a higher population of seniors [21], may have fewer data samples from younger patients, leading to models that are less effective for younger demographics [34].
- ➌ **Temporal Distribution Shifts**: Medical data is also affected by temporal changes, including shifts in disease prevalence and patient demographics [28].

An example is the emergence of COVID-19, requiring models to remain effective amidst these shifts. Given these challenges, one cannot help but question:

Are current training and evaluation protocols truly equipped to ensure that deployed medical imaging models perform robustly across different institutions?

Federated learning is one proposed solution, enabling collaborative training across disparate datasets [26], which could improve robustness while maintaining data separation. Yet, current evaluations of federated learning do not explicitly consider label, demographic, and temporal shifts simultaneously. Addressing this gap, we contribute a novel problem setup designed to simulate versatile medical scenarios. We also provide comprehensive benchmarking of existing methods, and we use our setup to emulate federated learning under pandemic conditions.

Problem and Benchmark Contribution. We introduce FedMedICL (Federated Medical Imaging with Continual Learning), a dual-function problem formulation. (i) It offers a unified approach for modeling label, demographic, and temporal shifts in medical imaging. (ii) It acts as a comprehensive testbed for evaluating the efficacy of federated learning in diverse healthcare settings. FedMedICL uses demographic metadata to automatically create federated and continual learning tasks, resulting in a realistic benchmark. We aim to foster reproducible research in this vital setup by releasing our extendable code.

Pandemic Spread Simulation Contribution. We showcase a practical application of FedMedICL: evaluation of model adaptation to a novel disease under pandemic conditions. In our setup, different hospitals may have access to COVID-19 datasets that grow at different rates over time. This evaluation is unexplored in prior federated learning research.

2 Related Work

Distribution Shifts in Medical Imaging. Research in long-tailed recognition [1], fairness [24], and continual learning [6] targets label imbalance [18], demographic [31], and temporal shifts [9] respectively, yet typically operates in isolation. Recent works like SubpopBench [31] have made advances by formalizing subpopulation shifts, yet they do not study them together, and do not model temporal shifts. Similarly, MEDFAIR [34] comprehensively evaluates demographic fairness but leaves label imbalance and temporal shifts unexplored. Furthermore, these studies typically overlook the siloed nature of medical data, a crucial aspect of real-world healthcare settings. In contrast, we offer a holistic approach to evaluate label, demographic, and temporal distribution shifts in medical imaging, while also considering the siloed nature of medical data.

3 FedMedICL: Problem and Benchmark

We study three types of distribution shifts associated with classification for real-world federated medical imaging. Individual hospitals possess small, isolated

datasets, leading to label and demographic imbalances (types ❶ and ❷). Medical data continually evolves, necessitating models that adapt to temporal shifts (type ❸). To replicate these real-world conditions, our benchmark consists of two components that follow the federated learning and continual learning literature [19,9]: Client Splitting, to simulate the distribution of data across medical institutions, and Temporal Task Splitting, to model the evolution of data over time. Next, we formalize the setup for federated and continual learning.

3.1 Background on Federated and Continual Learning

Assume we are given an input image $x \in \mathcal{X}$ and its ground truth diagnosis label $y \in \mathcal{Y}$ with $\mathcal{Y} = \{1, 2, \dots, L\}$. We aim to build a classifier $f_\theta : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Y})$, parameterized by θ , that can correctly classify x , where $\mathcal{P}(\mathcal{Y})$ is the probability simplex over the label space. Furthermore, medical data is often coupled with a set of attributes $a_1, a_2, \dots, a_m \in \mathcal{A}$, *e.g.*, the age or sex of patients.

In federated learning, client splitting divides a dataset \mathcal{D} into K segments $\{\mathcal{D}^1, \dots, \mathcal{D}^K\}$, each corresponding to a hospital. Each hospital may have a different distribution of attributes. Let \mathcal{D}^k be the training distribution at hospital k . \mathcal{D}^k is characterized by the set of probabilities of observing the different attributes at that hospital, denoted by $\{p_i^k\}_{i=1}^m$ where $\mathbb{P}_{(x,y) \sim \mathcal{D}^k}(a = a_i) = p_i^k$. That is, $p_i^k \geq 0$ is the probability that the attribute a_i is observed at hospital k .

In continual learning, temporal task splitting divides a hospital’s dataset \mathcal{D}^i into T sequential segments, $\{\mathcal{D}_1^i, \dots, \mathcal{D}_T^i\}$, which are called temporal tasks.

3.2 Benchmark Construction Methodology

We propose a benchmark construction method that precisely mirrors the complex scenarios in real-world medical federated and continual learning tasks.

Client Splitting. We introduce a scalable method to simulate the distribution of medical data across K clients, each representing a medical institution, by dividing them into two types: *Balanced* and *Skewed*. A Balanced client k maintains a homogeneous demographic distribution, defined as $\forall i, j \in \{1, \dots, m\}$, $p_i^k \approx p_j^k$. Conversely, a Skewed client k is marked by an uneven distribution, with at least one attribute a_i having a significantly higher than average probability, *i.e.*, $p_i^k \gg 1/m$. We create a mix of Balanced and Skewed clients to reflect the diverse demographic imbalances both within and across different hospitals, mimicking the heterogeneity of real-world medical data.

Temporal Task Splitting. We now address the temporal aspect of medical data, necessitating continual learning [9]. We define the *localized split*, which models demographic-based temporal shifts. We also consider novel disease emergence, which is a label-based temporal shift.

1- Localized Split. Hospital admission rates for different demographic groups vary seasonally, potentially due to local factors. To model these localized factors, we let $p_i^k \neq p_i^l \forall l \neq k$. That is, the distribution of attributes at each client evolve

independently and differently from other clients, resulting in cross-institutional data imbalances. Examples of healthcare scenarios that can be modeled by our localized split include demographic variations in emergency rooms [5].

2- Novel Disease Split. We aim to simulate the label imbalance associated with the emergence of a new disease, in which different geographic regions experience different rates of disease spread [25]. Given a novel disease label y_{new} , we define T tasks characterized by $\{\mathcal{D}_1^i, \dots, \mathcal{D}_T^i\}$ for each client’s dataset \mathcal{D}^i . During $T = 1$, we have $\forall i, \mathbb{P}_{(x,y) \sim \mathcal{D}_1^i}(y = y_{\text{new}}) = 0$ for all clients, while in the following tasks some clients encounter a distribution shift with $\exists i, j$ s.t. $\mathbb{P}_{(x,y) \sim \mathcal{D}_{j>1}^i}(y = y_{\text{new}}) > 0$, simulating the introduction of a novel disease into the cross-institutional dataset.

3.3 Evaluation and Datasets

Baselines. We implement baseline methods from various fields: data augmentation (MixUp [33]), domain generalization (SWAD [4]), continual learning (ER [6]), group-imbalanced learning (GB for group-balancing), and class-imbalanced learning (CB for class-balancing, CRT [16]). Each method tackles a single challenge and is unable to handle all scenarios that FedMedICL simulates. Therefore, we adapt these methods by augmenting them with federated averaging (FedAvg) [19]. The modified version of each method is prefixed with “F-” (*e.g.* F-CRT) throughout the paper. For comprehensive comparisons, we implement local training with Empirical Risk Minimization (ERM), which does not use federated learning, serving as a reference point for minimum expected performance. Following MEDFAIR [34], we used a ResNet-18 backbone for all the baselines.

Datasets. We utilize five public medical imaging datasets: CheXpert [14], Fitzpatrick17k [12], HAM10000 [27], OL3I [32], and PAPIA [17]. We also construct a new dataset, CheXCOVID, combining images from CheXpert and the COVID-19 image data collection [8]. Our choice of datasets is based on their scale, coverage of a wide range of imaging tasks, and availability of demographic metadata. A table summarizing these datasets is included in the Appendix.

Reported Metric. These datasets exhibit skewed class distributions, so a naive classifier that overpredicts the majority label achieves a high overall accuracy. Thus, we follow the Long-Tailed Recognition (LTR) literature and measure the LTR Accuracy [1]: $\text{LTR} = \frac{1}{L} \sum_{j=1}^L \frac{1}{N_j} \sum_{\forall i \text{ s.t. } y_i=j} \mathbb{1}\{\arg \max_j f_{\theta}^j(x_i) = j\}$, where N_j denotes the number of instances in the j^{th} class. To quantify label imbalance, we report the Imbalance Factor (IF), $\text{IF} = \frac{\max_i N_i}{\min_i N_i}$ [1], in Figure 2.

Evaluation. As shown in Figure 1b, we evaluate adaptation to local shifts via temporally aligned test tasks, and generalization across diverse demographics via a hold-out set \mathcal{D}_h that preserves the original dataset’s demographic distribution.

4 Benchmark Results

We evaluate the baselines using our localized split framework, detailed in Section 3. This framework models distributed hospitals experiencing independent




Fig. 2: **Benchmarking Demographic Shifts.** We simulate age-based demographic changes over time and benchmark methods in different datasets. The mean LTR accuracy across clients is reported for each method. Except on the PAPILA dataset, no method reliably competes with the simple F-CB baseline.

demographic shifts over time. In this section, demographic groups are defined by skin type for PAPILA and age brackets (*e.g.* 0-40 years) for other datasets, but our framework supports alternative attributes such as sex.

Setup. We set up $K = 10$ clients per dataset, except the largest dataset CheXpert with $K = 50$. This choice ensures similar dataset sizes across clients, providing a fair comparison across datasets. Employing our localized split strategy, we assign $T = 4$ sequential training tasks to each client to simulate seasonal changes in disease patterns and hospital admissions. For example, the prevalence of seasonal flu can alter the demographic composition of hospital patients during flu season. Transitioning to a new task represents evolution in patient age demographics, except for PAPILA, where skin demographics are used. Each client receives $T = 4$ local test sets matching their training data demographics, as depicted in Figure 1b. During each task, training includes multiple communication rounds, each with $M = 5$ iterations and a batch size of 10, followed by federated averaging. At the end of every training task, we assess each client’s model on all previously encountered tasks, following Equation (2) in [10], and report the mean LTR across clients in Figure 2. After the final task, we evaluate each client’s final model on a shared test set \mathcal{D}_h , representing a global demographic distribution. This evaluation is reported in Figure 3.

Results. Figure 2 compares methods across multiple datasets and diverse demographic distributions. Our evaluation under the FedMedICL setup reveals that the straightforward class-balancing approach (F-CB, shown in green) surpasses all other methods in 5/6 datasets. Notably, advanced algorithms such as F-SWAD and F-CRT, despite their strong performance in previous benchmarks like SubpopBench [31] and MEDFAIR [34], fall short in every examined dataset. This discrepancy supports our hypothesis: progress in tackling individual types of shift is not sufficient for deployment in federated settings with simultaneous distribution shifts. Specifically, SWAD’s dense stochastic weight averaging requires many iterations to converge, making it suboptimal for FedMedICL scenarios where quickly adapting to new distributions is crucial. Similarly, CRT’s two-staged paradigm is cumbersome for the dynamic nature of federated and continual learning, which require frequent client communication and swift adap-




Fig. 3: **Performance on New Demographic Distributions.** We report the LTR accuracy on a hold-out test set, as shown in Figure 1b. Results are averaged across clients. No method consistently generalizes better than F-CB.

tation. Results in Figure 3, examining the final model at $T = 4$ against a hold-out set, align with these findings. It suggests that better performance on local client data correlates with better out-of-distribution demographic generalization.

Disparities across Datasets. We explore why performance varies across FedMedICL datasets. The class-balancing method (F-CB, shown in green) is the most effective in most datasets except PAPILA, where the group-balancing method (F-GB, shown in orange) excels. We hypothesize that this variance relates to the different types of imbalance present in each dataset: group imbalance in PAPILA and class imbalance in other datasets. For example, the HAM10000 dataset has 7 labels with an imbalance factor (IF) of 58.3, which indicates a severely skewed label distribution. Similarly, CheXpert has only 2 labels with $IF = 9.6$, meaning that one label encompasses around 90% of the dataset. This is why class-balancing is effective in these datasets. An interesting case is OL3I, having 2 labels with $IF = 22.1$, which is more skewed than CheXpert. Due to its extreme imbalance, methods including F-CB achieve around 50% LTR, which degenerates to the case of simply predicting the majority class.

In PAPILA, we hypothesize that the variation in image characteristics related to age is more significant than class imbalance. We group images by age and compute the contrast feature, which measures the variation in gray-level intensity between neighboring pixels [13]. We compute the Cohen’s D value [7] to compare image contrast features between older adults (60+) and younger adults (below 60). The values are 0.01, 0.14, and 0.40 for CheXpert, HAM10000, and PAPILA respectively. The higher value for PAPILA confirms the significance of differences between the two age groups in this eye fundus dataset. Our analysis aligns with ophthalmological findings that the eye’s structure varies with age [3]. Finally, previous benchmarks found a high accuracy gap between age-based groups in PAPILA [34], further confirming the age variation challenge in its fundus images.

5 Spread of the COVID-19 Pandemic Between Hospitals

FedMedICL can simulate novel disease outbreaks and evaluate algorithms under such conditions. We simulate COVID-19 spreading between hospitals using a




Fig. 4: **Adaptation under Pandemic Conditions.** We simulate two types of hospitals experiencing COVID-19 emergence over four tasks. We report performance on non-COVID labels, in addition to examining how various methods perform in recognizing the novel COVID-19 disease across the four time steps.

subset of our proposed CheXCovid dataset, which augments CheXpert with COVID-19 samples. Following the Novel Disease Split, introduced in Section 3, we consider a scenario with $K = 5$ hospitals. Initially, X-ray screenings across hospitals do not include COVID-19 samples. Over time, the number of samples grows for each hospital, but the rate of increase in COVID-19 admissions varies by geographic region, and hence by hospital [25]. More formally, we assign $T = 4$ sequential training tasks to each hospital, mimicking the evolution of the pandemic. For each task, the number of COVID cases grows. We assume Task 1 is pre-COVID: no hospital has any COVID cases in their training set. Three hospitals experience rapid surges, where approximately 0%, 10%, 50%, and 90% of patients have COVID for each step, respectively. Meanwhile, the remaining two hospitals have slower growth rates, with 0% COVID for the first two tasks, then 10% and 50% for the final two. All models are tested on COVID and non-COVID samples at the end of every task. The pivotal question we address is: *How can our simulated hospitals, experiencing varying rates of COVID-19 admission, utilize federated learning to enhance the detection of the novel disease?*

Results. Figure 4 presents our findings on Non-COVID labels (*i.e.* accuracy on CheXpert labels) and COVID label (recognition of COVID-19). After the pre-COVID training stage (Task 1), all hospital models fail to recognize the novel class. As the tasks progress, ERM exhibits severe forgetting of the Non-COVID conditions, indicated by the severe drop in the ‘Non-COVID Performance’ graph, and demonstrates slow adaptation to COVID-19, as shown by the limited improvement in Task 2 on the ‘COVID Performance’ graph. F-ER maintains stable performance on non-COVID labels, which aligns with expectations, since F-ER is designed to mitigate forgetting. However, it has a relatively flat curve for COVID detection in the first two tasks, suggesting poor performance in early pandemic stages. Conversely, F-CB and F-GB quickly adapt to COVID-19 but at the expense of forgetting previously known diseases. Hence, there is a need for new methods that can balance high plasticity, allowing for the swift detection of emerging pathogens like COVID-19, with sufficient stability, which maintains performance on established conditions.

6 Conclusions and Future Research

We envision FedMedICL as a first step towards holistic evaluation of federated medical imaging. It paves the way for research on distribution shifts in siloed medical datasets, showing that a simple batch-balancing technique outperforms existing methods. Our experiments focus on variations in age and skin type, motivated by evidence of significant performance disparities across these attributes [34]. Yet, we have designed FedMedICL with flexibility in mind, enabling easy integration of additional singular attributes (e.g., device manufacturer) and dual attributes (e.g., intersection of sex and age). FedMedICL can also easily extend to any modality, such as text and tabular data. We encourage researchers to expand our benchmark to include more data attributes and modalities.

Acknowledgments. The authors thank Haoran Zhang for his feedback. Kumail was supported by a Jameel Clinic Fellowship and a Saudi Arabian Cultural Mission Scholarship. Yasir was supported by Saudi Aramco. Marzyeh is supported in part by an NSF CAREER award, a CIFAR AI Chair held at the Vector Institute, and the Gordon and Betty Moore Foundation. This project was also supported by the KAUST Office of Sponsored Research (OSR) under Award No. OSR-CRG2021-4648 and a UKRI grant Turing AI Fellowship (EP/W002981/1). Philip thanks the Royal Academy of Engineering for their support.

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

References

1. Alshammari, S., Wang, Y.X., Ramanan, D., Kong, S.: Long-tailed recognition via weight balancing. In: Conference on Computer Vision and Pattern Recognition (2022)
2. Baghestani, S., Zare, S., Mahboobi, A.A.: Skin disease patterns in hormozgan, iran. *International journal of dermatology* **44**(8), 641–645 (2005)
3. Cavallotti, C., Cerulli, L.: Age-related changes of the human eye. Springer Science & Business Media (2008)
4. Cha, J., Chun, S., Lee, K., Cho, H.C., Park, S., Lee, Y., Park, S.: Swad: Domain generalization by seeking flat minima. *Advances in Neural Information Processing Systems* **34**, 22405–22418 (2021)
5. Chaaban, M.R., Zhang, D., Resto, V., Goodwin, J.S.: Demographic, seasonal, and geographic differences in emergency department visits for epistaxis. *Otolaryngology–Head and Neck Surgery* **156**(1), 81–86 (2017)
6. Chaudhry, A., Rohrbach, M., Elhoseiny, M., Ajanthan, T., Dokania, P.K., Torr, P.H.S., Ranzato, M.: On tiny episodic memories in continual learning (2019)
7. Cohen, J.: Statistical power analysis for the behavioral sciences. Academic press (2013)
8. Cohen, J.P., Morrison, P., Dao, L., Roth, K., Duong, T.Q., Ghassemi, M.: Covid-19 image data collection: Prospective predictions are the future. arXiv preprint arXiv:2006.11988 (2020)

9. Derakhshani, M.M., Najdenkoska, I., van Sonsbeek, T., Zhen, X., Mahapatra, D., Worring, M., Snoek, C.G.: Lifelonger: A benchmark for continual disease classification. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 314–324. Springer (2022)
10. Díaz-Rodríguez, N., Lomonaco, V., Filliat, D., Maltoni, D.: Don’t forget, there is more than forgetting: new metrics for continual learning. In: Workshop on Continual Learning, Neural Information Processing Systems (2018)
11. Ebrahimian, S., Kalra, M.K., Agarwal, S., Bizzo, B.C., Elkholy, M., Wald, C., Allen, B., Dreyer, K.J.: Fda-regulated ai algorithms: trends, strengths, and gaps of validation studies. *Academic radiology* (2022)
12. Groh, M., Harris, C., Soenksen, L., Lau, F., Han, R., Kim, A., Koochek, A., Badri, O.: Evaluating deep neural networks trained on clinical images in dermatology with the fitzpatrick 17k dataset. In: Conference on Computer Vision and Pattern Recognition. pp. 1820–1828 (2021)
13. Haralick, R.M., Shanmugam, K., Dinstein, I.H.: Textural features for image classification. *IEEE Transactions on systems, man, and cybernetics* (6), 610–621 (1973)
14. 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 (2019)
15. Kamnitsas, K., Ledig, C., Newcombe, V.F., Simpson, J.P., Kane, A.D., Menon, D.K., Rueckert, D., Glocker, B.: Efficient multi-scale 3d cnn with fully connected crf for accurate brain lesion segmentation. *Medical image analysis* **36**, 61–78 (2017)
16. Kang, B., Xie, S., Rohrbach, M., Yan, Z., Gordo, A., Feng, J., Kalantidis, Y.: Decoupling representation and classifier for long-tailed recognition. In: International Conference on Learning Representations (2019)
17. Kovalyk, O., Morales-Sánchez, J., Verdú-Monedero, R., Sellés-Navarro, I., Palazón-Cabanes, A., Sancho-Gómez, J.L.: Papila: Dataset with fundus images and clinical data of both eyes of the same patient for glaucoma assessment. *Scientific Data* **9**(1), 291 (2022)
18. Marrakchi, Y., Makansi, O., Brox, T.: Fighting class imbalance with contrastive learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer (2021)
19. McMahan, B., Moore, E., Ramage, D., Hampson, S., y Arcas, B.A.: Communication-efficient learning of deep networks from decentralized data. In: Artificial intelligence and statistics. pp. 1273–1282. PMLR (2017)
20. Paek, S.Y., Koriakos, A., Saxton-Daniels, S., Pandya, A.G.: Skin diseases in rural yucatan, mexico. *International journal of dermatology* **51**(7), 823–828 (2012)
21. Pekmezaris, R., Kozikowski, A., Moise, G., Clement, P.A., Hirsch, J., Kraut, J., Levy, L.C.: Aging in suburbia: An assessment of senior needs. *Educational Gerontology* **39**(5), 355–365 (2013)
22. Petkova, E., Antman, E.M., Troxel, A.B.: Pooling data from individual clinical trials in the covid-19 era. *Jama* **324**(6), 543–545 (2020)
23. Pooch, E.H., Ballester, P., Barros, R.C.: Can we trust deep learning based diagnosis? the impact of domain shift in chest radiograph classification. In: Thoracic Image Analysis: Second International Workshop, TIA 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 8, 2020, Proceedings 2. Springer (2020)
24. Sagawa*, S., Koh*, P.W., Hashimoto, T.B., Liang, P.: Distributionally robust neural networks. In: International Conference on Learning Representations (2020)

25. Sun, Z., Zhang, H., Yang, Y., Wan, H., Wang, Y.: Impacts of geographic factors and population density on the covid-19 spreading under the lockdown policies of china. *Science of The Total Environment* **746**, 141347 (2020)
26. Ogier du Terrail, J., Ayed, S.S., Cyffers, E., Grimberg, F., He, C., Loeb, R., Mangold, P., Marchand, T., Marfoq, O., Mushtaq, E., et al.: Flamby: Datasets and benchmarks for cross-silo federated learning in realistic healthcare settings. *Advances in Neural Information Processing Systems* (2022)
27. Tschandl, P., Rosendahl, C., Kittler, H.: The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific data* (2018)
28. Vokinger, K.N., Feuerriegel, S., Kesselheim, A.S.: Continual learning in medical devices: Fda’s action plan and beyond. *The Lancet Digital Health* (2021)
29. White, T., Blok, E., Calhoun, V.D.: Data sharing and privacy issues in neuroimaging research: Opportunities, obstacles, challenges, and monsters under the bed. *Human Brain Mapping* (2022)
30. Yala, A., Lehman, C., Schuster, T., Portnoi, T., Barzilay, R.: A deep learning mammography-based model for improved breast cancer risk prediction. *Radiology* **292**(1), 60–66 (2019)
31. Yang, Y., Zhang, H., Katabi, D., Ghassemi, M.: Change is hard: A closer look at subpopulation shift. In: *International Conference on Machine Learning* (2023)
32. Zambrano Chaves, J.M., Chaudhari, A.S., Wentland, A.L., Desai, A.D., Banerjee, I., Boutin, R.D., Maron, D.J., Rodriguez, F., Sandhu, A.T., Jeffrey, R.B., et al.: Opportunistic assessment of ischemic heart disease risk using abdominopelvic computed tomography and medical record data: a multimodal explainable artificial intelligence approach. *medRxiv* pp. 2021–01 (2021)
33. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: *International Conference on Learning Representations* (2018)
34. Zong, Y., Yang, Y., Hospedales, T.: Medfair: Benchmarking fairness for medical imaging. In: *International Conference on Learning Representations* (2023)