From Competition to Collaboration: Making Toy Datasets on Kaggle Clinically Useful for Chest X-Ray Diagnosis Using Federated Learning

Pranav Kulkarni¹ Adway Kanhere^{1,2} Paul H. Yi¹ Vishwa S. Parekh¹ ¹University of Maryland Medical Intelligent Imaging (UM2ii) Center Department of Diagnostic Radiology and Nuclear Medicine University of Maryland School of Medicine Baltimore, MD 21201 {pkulkarni, pyi, vparekh}@som.umaryland.edu ²Department of Biomedical Engineering Johns Hopkins University Baltimore, MD 21218 akanher1@jhu.edu

Abstract

Chest X-ray (CXR) datasets hosted on Kaggle, though useful from a data science competition standpoint, have limited utility in clinical use because of their narrow focus on diagnosing one specific disease. In real-world clinical use, multiple diseases need to be considered since they can co-exist in the same patient. In this work, we demonstrate how federated learning (FL) can be used to make these toy CXR datasets from Kaggle clinically useful. Specifically, we train a single FL classification model ('global') using two separate CXR datasets - one annotated for presence of pneumonia and the other for presence of pneumothorax (two common and life-threatening conditions) - capable of diagnosing both. We compare the performance of the global FL model with models trained separately on both datasets ('baseline') for two different model architectures. On a standard, naive 3-layer CNN architecture, the global FL model achieved AUROC of 0.84 and 0.81 for pneumonia and pneumothorax, respectively, compared to 0.85 and 0.82, respectively, for both baseline models (p>0.05). Similarly, on a pretrained DenseNet121 architecture, the global FL model achieved AUROC of 0.88 and 0.91 for pneumonia and pneumothorax, respectively, compared to 0.89 and 0.91, respectively, for both baseline models (p>0.05). Our results suggest that FL can be used to create global 'meta' models to make toy datasets from Kaggle clinically useful, a step forward towards bridging the gap from bench to bedside.

1 Introduction

Chest X-Ray (CXR) is the most commonly ordered medical imaging study globally and is critical for screening many life threatening conditions (e.g., pneumonia). Accordingly, many large-scale public CXR datasets have been released through curation of [8, 2, 3, 5]. These, in turn, have resulted in numerous data science competitions hosted on platforms like Kaggle (e.g., RSNA pneumonia detection challenge [7]), resulting in expert-level performance for disease diagnoses.

Although useful from a data science competition standpoint, these Kaggle-hosted CXR datasets have limited clinical utility because of their narrow focus on one single diagnostic task. For example, the two Kaggle CXR competitions hosted by RSNA and SIIM have focused on diagnosis of a single

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

disease, like pneumonia and pneumothorax [9]. Although impressive results have resulted from these competitions, their utility is limited given the dozens of diagnoses that could present in real-world clinical practice. Therefore, a method to harmonize these toy datasets to train a clinically-useful model could revolutionize how small, narrowly-focused datasets can be leveraged in aggregate for development of clinically-relevant deep learning models.

We propose CheXViz, a federated learning (FL) framework for training a single model on spatially distributed datasets with different disease annotations into a 'global' meta-deep learning model. Briefly, FL is a machine learning technique that approaches the problem from a multi-domain and multi-task perspective. By using a decentralized and distributed approach, consisting of a central server and nodes, a global 'meta' model can be trained to generalize distributed tasks with noniid labels. During each training step ('FL round'), every node trains a local model. Then, the weights across all nodes are aggregated by the central server and redistributed back to the nodes. In medical imaging, FL has enabled training of large-scale 'global' deep learning models using datasets spread across multiple institutions without sharing sensitive patient data. In this preliminary work, we demonstrate the utility of CheXViz for training a single model to diagnose pneumonia and pneumothorax using two toy datasets from Kaggle can be made clinically useful using FL.



Figure 1: Illustration of the CheXViz framework

2 Methods

2.1 CheXViz

We developed CheXViz using a multi-task FL setup. The CheXViz model is initialized as a deep neural network consisting of two distinct blocks - a representation block and a task block. The CheXViz model is distributed across all the participating nodes to train their tasks. During training, only the weights corresponding to the representation block are aggregated and redistributed by the central server back to the nodes, thereby preserving task-related information for each node in their task block. Figure 1 illustrates the CheXViz framework.

2.2 Experimental Setup

We evaluated the CheXViz framework for the task of training a generalized global model that can classify cases of pneumonia and pneumothorax using distributed and non-iid CXR datasets from the RSNA Pneumonia Detection and the SIIM-ACR Pneumothorax Segmentation competitions on Kaggle. Our experiments involved two different deep network architectures: A naive convolutional neural network (CNN) consisting of 3 convolutional layers ('standard' model) and a pre-trained DenseNet121 architecture using transfer learning (TL) [1]. We implemented Federated Averaging (FedAvg) for model weight aggregation as described in [4].

The standard baseline models for pneumonia and pneumothorax classification were trained for 300 epochs using a learning rate scheduler with an initial learning rate of 1e-3. During FL, we trained the standard global model for 300 FL rounds where for each round, models were locally trained for 1 epoch before aggregation. The DenseNet121 model was initialized with pre-trained ImageNet weights ('base model') and a new classification head was trained for 30 epochs with a learning rate of 1e-3 while the rest of the base model weights were frozen and then fine tuned with a slower learning rate of 1e-5 to prevent overfitting. For the baseline models, the models were fine tuned for 150 epochs. For FL, we utilized FedAvg during the fine tuning step of TL. The global model was fine tuned for 150 FL rounds with 1 epoch before aggregation with a learning rate of 1e-5. All the models were trained and evaluated on a Google Cloud VM with four NVIDIA T4 GPUs.

We computed sensitivity, specificity, area under precision-recall (AUPR) curve, and the area under receiver operating characteristic (AUROC) curve for each model. The AUROC values between the FL global model and the task's baseline model were compared using bootstrapping and a paired t-test. Statistical significance was defined as p < 0.05.



3 Results

Figure 2: ROC curves obtained from baseline and CheXViz models evaluated across both the datasets.

The CheXViz framework trained 'meta' models demonstrated excellent performance compared to the baseline models for the diagnostic classification of pneumonia and pneumothorax abnormalities, as shown in Figure 2. For pneumonia classification with the standard model architecture, the baseline and FL models achieved a validation AUROC of 0.85 and 0.84 respectively. We observed a similar trend with the DenseNet121 architecture, where the baseline and FL models achieved a validation AUROC of 0.89 and 0.88 respectively. For pneumothorax classification with the standard model architecture, the baseline and FL models achieved a validation AUROC of 0.82 and 0.81 respectively.

Again, a similar result was achieved with the DenseNet121 architecture, where the baseline and FL models both achieved a validation AUROC of 0.91. The results are detailed in Appendix Table 1.

We further visualized the Grad-CAM outputs for evaluating the explainability and generalizability of the models [6]. Our preliminary analysis suggests that the heatmaps from CheXViz models demonstrate higher and focused activations within the lungs, compared to the baseline models as shown in Appendix Figure 3. For future work, we intend to outline a methodology to quantify the generalizability of models for CXR classification using Grad-CAM heatmaps.

4 Discussion

Although Kaggle CXR datasets and data science competitions have made an indelible impact on data science and AI for healthcare, they are still a far cry from being clinically useful datasets. This is understandable, given the challenges in curating expert-level annotations for diseases, and ostensibly why these Kaggle-hosted competitions have focused largely on single diseases[7, 9]. Nevertheless, there is a gap between these toy datasets and clinical utility. Put another way, it is unclear how to use these datasets to train clinically-useful models capable of detecting multiple diseases. Our findings demonstrate that our FL framework (CheXViz) can be used to create global 'meta' models to make toy datasets from Kaggle clinically useful, a large step forward towards bridging the gap from bench to bedside. Although preliminary in nature and focusing on only two datasets, our framework and results are extensible to any number of datasets and disease labels, as well as tasks beyond classification (e.g., segmentation and object detection). It is our hope that our work can be a first step towards moving Kaggle CXR datasets from competition to collaboration and transform these toy datasets into clinically useful models.

References

- G. Huang et al. Densely connected convolutional networks. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017. doi: https://doi.org/10.48550/arXiv.1608. 06993.
- [2] J. Irvin, P. Rajpurkar, M. Ko, Y. Yu, S. Ciurea-Ilcus, C. Chute, H. Marklund, B. Haghgoo, R. Ball, K. Shpanskaya, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 590–597, 2019.
- [3] A. E. Johnson, T. J. Pollard, N. R. Greenbaum, M. P. Lungren, C.-y. Deng, Y. Peng, Z. Lu, R. G. Mark, S. J. Berkowitz, and S. Horng. Mimic-cxr-jpg, a large publicly available database of labeled chest radiographs. arXiv preprint arXiv:1901.07042, 2019.
- [4] B. McMahan and more. Communication-Efficient Learning of Deep Networks from Decentralized Data. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, volume 54 of Proceedings of Machine Learning Research, pages 1273–1282. PMLR, 20–22 Apr 2017.
- [5] H. Q. Nguyen, K. Lam, L. T. Le, H. H. Pham, D. Q. Tran, D. B. Nguyen, D. D. Le, C. M. Pham, H. T. Tong, D. H. Dinh, et al. Vindr-cxr: An open dataset of chest x-rays with radiologist's annotations. *Scientific Data*, 9(1):1–7, 2022.
- [6] R. R. Selvaraju et al. Grad-cam: Visual explanations from deep networks via gradient-based localization. *Proceedings of the IEEE international conference on computer vision*, 2017. doi: https://doi.org/10.1007/s11263-019-01228-7.
- [7] G. Shih, C. C. Wu, S. S. Halabi, M. D. Kohli, L. M. Prevedello, T. S. Cook, A. Sharma, J. K. Amorosa, V. Arteaga, M. Galperin-Aizenberg, et al. Augmenting the national institutes of health chest radiograph dataset with expert annotations of possible pneumonia. *Radiology. Artificial intelligence*, 1(1), 2019.
- [8] 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.

[9] P. H. Yi, T. K. Kim, E. Siegel, and Y.-F.-A. N. Demographic reporting in publicly available chest radiograph data sets: Opportunities for mitigating sex and racial disparities in deep learning models. *Journal of the American College of Radiology*, 19(1 Pt B):192–200, 2021.

Societal Impact

Our work has a potential positive societal impact by taking an important step towards translational research. Our work can be a first step towards moving Kaggle CXR datasets from competition to collaboration and transform these toy datasets into clinically useful models. However, the use of FL can potentially lead to privacy and security risks, making the system vulnerable to client and server attacks, thereby leading to potentially negative societal impacts. We are actively working to address these vulnerabilities in FL systems.

A Appendix

Table 1: Model Metrics							
Task	Model	Loss	Sensitivity	Specificity	AUPR	AUROC	p-value
Pneumonia	Standard Standard w/ FL	0.38 0.39	82.57 78.01	71.97 74.41	0.63 0.61	0.85 0.84	0.10
	DenseNet121 DenseNet121 w/ FL	0.34 0.35	84.90 80.08	76.76 79.79	0.71 0.70	0.89 0.88	0.19
Pneumothorax	Standard Standard w/ FL	0.41 0.42	74.13 80.22	75.89 69.87	0.54 0.52	0.82 0.81	- 0.71
	DenseNet121 DenseNet121 w/ FL	0.31 0.31	91.30 84.57	76.31 83.10	0.73 0.73	0.91 0.91	- 0.76

Sandard Model
Sandard Model<

(a) Pneumonia

(b) Pneumothorax

