Federated Learning for Accurate Labeling of Chest X-Ray Scans

Shanker Ram
Department of Research, Nysha.ai
West Valley College
Saratoga, USA
sram@mywvm.wvm.edu

Yajat Nagaraj Kiran
Department of Research, Nysha.ai
West Valley College
Saratoga, USA
ynagaraj@mywvm.wvm.edu

Arnav Bhute
Department of Research, Nysha.ai
Prosper High School
Texas, USA
arnav.bhute.2026@gmail.com

Tashvi Khare
Department of Research, Nysha.ai
Westview High School
Oregon, USA
khare.tashvi@gmail.com

***Abstract*—Federated learning is an increasingly common technique used within machine learning that allows multiple devices to collectively train a model without necessitating the centralization of data. This approach is highly valuable within medical tasks, where privacy concerns within patient datasets can be mitigated through the decentralization of machine learning training. Within past literature, there have remained difficulties in constructing well annotated, large chest X-ray datasets due to these patient privacy concerns. In this paper, we seek to demonstrate the validity of federated learning by training a deep learning model on decentralized Chest X-ray imaging data. We utilize the publicly available NIH Chest X-ray dataset to train our model. Five clients were trained over 10 rounds, and a ResNet-34 global model was initialized and moved to a GPU, where clients iterated over each round to update the model. We initialized a new parameter accumulation dictionary for each round that was outfitted with Secure Aggregation algorithm with in-built additive homomorphic encryption of local parameters towards parameter averaging. The model achieved a validation loss of 0.09 and an accuracy of 0.83. These results indicate that the outlined federated learning approach was able to approach benchmark clinical grade accuracy, demonstrating the effectiveness of federated learning in advanced medical imaging analysis with the preservation of patient privacy.**

***Index Terms*—Federated Learning, Chest X-ray Imaging, Medical Imaging, Deep Learning, ResNet-34, Patient Privacy, Decentralized Machine Learning, Healthcare Data, Distributed Training, Validation Accuracy**

I. Introduction

Chest radiography or the chest X-ray is among one of the most common medical imaging exams in the world, with industrialized nations reporting 23.8 percent of their population acquiring erect-view chest X-ray images [1]. Chest radiography imaging exams provide a benefit over traditional clinical examinations because they remain accessible, affordable, portable in hospitals, and generally non-invasive with less radiation exposure, and can be used for advanced prognostic indications. Currently, chest radiography is typically used as a first-line imaging tool to diagnose a wide range of thoracic diseases, particularly those affecting the lungs, heart, and other regions of the chest cavity, such as pneumonia and tuberculosis, and lung cancers. However, even with the widespread multipurpose image applications of X-ray, accurate patient diagnosis remains a widespread challenge because patient abnormalities must be interpreted diligently. Further, diagnoses are prone to inter-observer variability due to the process of visual diagnosis itself being intrinsically subjective.

Recently, developments within surrogate and deep learning based imaging applications have simplified the automated analysis across a wide range of medical imaging, such as MRI and X–ray [2]. For example, image annotation tools using deep learning algorithms have achieved clinical grade accuracy in the end-to-end detection, classification, and eventual prognosis of diseases such as breast cancer, lung cancer, and diabetic retinopathy. However, because these models involve millions of parameters, deep learning models must train on large, well-annotated, and high-quality image data to mitigate overfitting and achieve generalizing performance and training stability. [5] Obtaining datasets with requirements such as the ones outlined above remains a challenge due to logistical hurdles, patient privacy obligations, and the associated high costs of manual or automated medical image annotations [3,4]. To address such data acquisition challenges, federated learning may be an emerging solution, allowing multiple institutions to collaboratively train a large scale model without centralizing datasets within a server. This occurs when a global model is initialized and sent to different clients, who can train the global model on private institutional datasets. After client-side training is complete, the central server can aggregate parameters through a range of techniques before updating the global model. Finally, the updated global model is sent to clients in the next round for training, creating an iterative training process until the global model converges. Machine learning training via Federative learning approaches have demonstrated performance comparable to centrally hosted machine learning models at hospitals. [6]

In this paper, we showcase the effective use of a federated learning approach to train a ResNet-34 Convolutional Neural Network on Chest X-ray scans towards accurate thoracic disease classification. Our approach not only preserves patient privacy but additionally achieves benchmark accuracy comparable to centrally hosted models, showcasing the potential for federated learning to improve medical image analysis on multiple private datasets while protecting patient confidentiality and privacy.

A. ML in Medical Analysis

Recent advancements in deep learning have made a significant impact on the field of medical imaging, providing tools for automatic analysis of complex medical image data. For example, an application of Convolution Neural Networks (CNNs) trained on chest X-ray data has been shown to exceed average radiologist performance on the F1 Metric [7]. Such results underscore the potential of deep learning models to enhance diagnostic precision in radiology. Similarly, within an independent breast cancer detection study by six radiologists, the CNN performed better than all the human readers, with the area under the receiver operating characteristic curve (AUROC) for the CNN improving over human performance by an absolute margin of 11.5 percent[8]. Another recent study used the Archimedes-assisted Henry Gas Optimization Algorithm + EC classification method, and were able to receive an accuracy of 0.95 with tuning percentage of 70, for using chest x-ray images to diagnose pneumonia, which is higher than previous techniques used to detect this condition[26].

Further applications of deep learning are shown in the field of lung cancer detection, where researchers have built a custom CNN architecture enhanced with the Channel Attention and Spatial Attention (SA) mechanisms. Channel Attention mechanisms enhance the importance of channels by re-weighting feature maps, obtained from the convolutional layer, by channel importance. These channels can range from edges and texture patterns to color values and higher-level features. Spatial Attention mechanisms enhance the spatial importance of certain regions within spatial context through aggregation of feature maps into spatial attention maps and then re-weighting original feature maps in terms of spatial importance. Both of these attention mechanisms were used within the lung cancer detection algorithm, which improved feature extraction capabilities and provided a more effective detection scheme.[9]

B. Federated Learning (FL) in Medical Applications

As mentioned above, despite the numerous advancements of deep learning within medical image analysis, privacy restrictions and the large amount of data necessary for generalizable performance makes deep learning algorithms challenging to train. Federated Learning has shown to be a promising approach to solving this problem, as the technique facilitates the training of a global model on decentralized, private datasets, without having to exchange much of the data itself, which also helps during potential collaborations between multiple organizations. The parameters obtained from training on each of these private, institutional datasets can then be aggregated and used to update the final parameters of the global model. This technique can even lead to improvements over central server models because the global model can be trained and validated on multiple private datasets, mitigating overfitting and improving generalized performance. For example, in a federated learning approach to brain tumor segmentation, the segmentation results indicated that federated learning based models outperformed centralized and local models on numerous data distributions [10]. Federated approaches particularly excelled in small distribution scenarios and outperformed the centrally run models significantly. [10]

Another example of Federated Learning can be seen with an Alzheimer's disease detection model where researchers implemented two types of aggregation algorithms, Federated Averaging (FedAvg) and Secure Aggregation (SegAgg). [11] Both FedAvg and SecAgg demonstrated robust performance, with minimal degradation in model performance. However, losses did increase slightly in relative performance degradation (RPD) with a greater number of clients and imbalance training distributions. Between the aggregation models, SecAgg was much more immune to membership inference attacks, enhancing patient security. Further, FedAvg suffered a noticeable decline in performance with non-IID data, leading to performance differences in local and global data distributions. [11, 19] This study highlights the importance of selecting appropriate aggregation algorithms based on data distributions and privacy requirements. [17, 19]

II. Dataset

The National Institutes of Health (NIH) Chest X-ray dataset is a collection of over 100,000 X-ray images and corresponding thoracic disease labels made publicly available by the National Institutes of Health. The dataset was constructed from the collection of X-ray scans from over 30,000 patients, many of whom had advanced lung diseases. The labels of the fourteen common thoracic diseases in the dataset were: atelectasis, consolidation, infiltration, pneumothorax, edema, emphysema, fibrosis, effusion, pneumonia, pleural thickening, cardiomegaly, nodule mass, and hernia. The fifteenth label was "No Findings" for patients with no recorded thoracic disease visible from X-ray.

The dataset was split into ten files, each with thousands of scans. These files were used for attempting federated learning. To be compatible with Kaggle, the original TAR archive files were converted to zip files.

For the label of every X-ray scan, NIH used a weakly supervised multi-label image classification and disease localization formulation. Additionally, natural language processing (NLP) data mining techniques were used to aggregate information from radiological reports. The labeling accuracy is estimated to be greater than 90%.

III. Preprocessing

Using the PyTorch library, the same series of unique image augmentations-resizing, normalizing, augmenting data, and

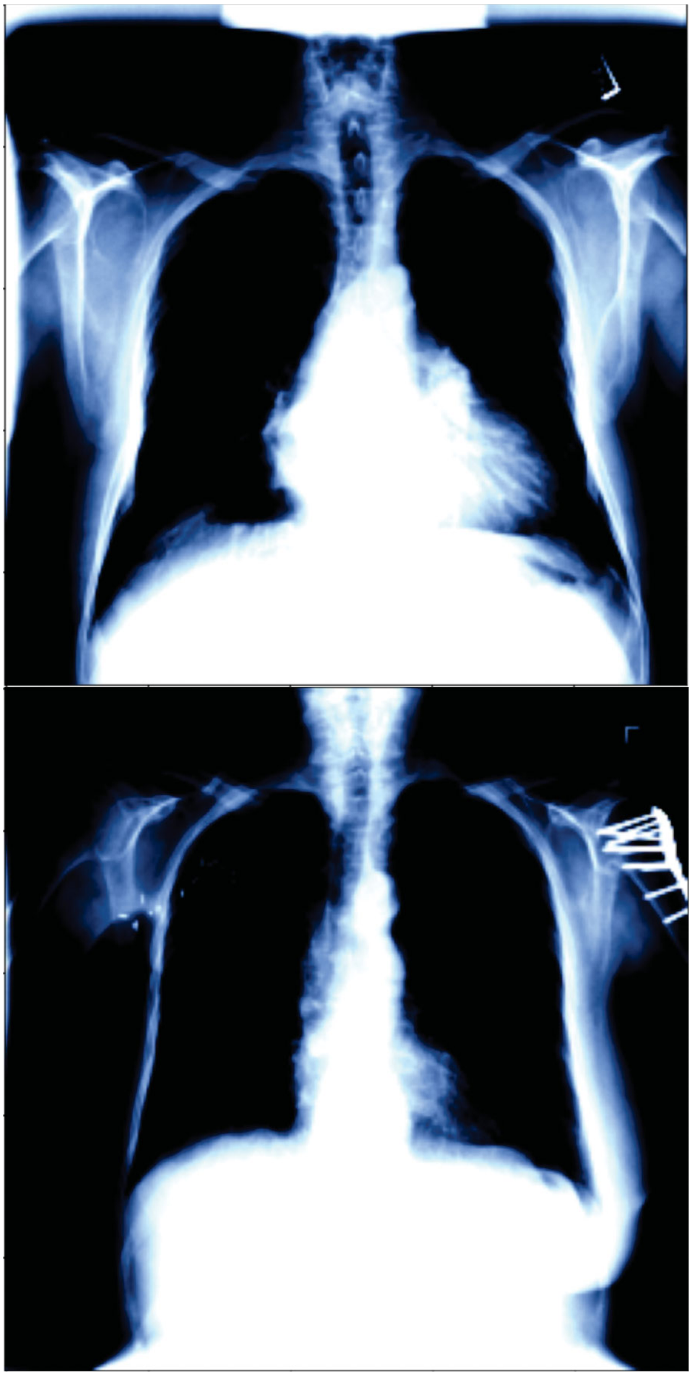

Fig. 1. Examples of Chest X-Ray Scans

converting images to tensors-was done for each X-ray scan. Uniformly resizing the images to a pixel dimension of 240 x 240 balanced computational efficiency with high enough resolution for good disease classification. [12] All pixel values were normalized, therefore, with respect to precomputed means of RGB for standardizing the intensity levels across all images. This enables the model to learn better by reducing variability across the dataset. [12]

In the dataset management for the federated learning framework, each client's dataset is split into three subsets: 70% of the data for training, 15% for validation, and 15% for testing. The split in data ensures every client has substantial data to train their local ResNet-34 model with sufficient data for validation and testing in the performance evaluation. [12] For the global model, a similar split was considered for the whole dataset by which we can evaluate how the global model performed versus the individual local models each being trained on smaller subsets.

Comparing the local and global models, one could observe that the latter aggregated knowledge across all the clients to yield more generalizable results

IV. Architecture

After testing multiple deep learning architectures with federated settings, the final deep learning architecture selected for the global model was the ResNet-34 deep CNN, a variant of the Microsoft ResNet architecture with 34 layers (shown in Figure 2a.) The ResNet is trained using residual learning, where "shortcut" paths are created to retain features from previous layers and combat gradient degradation. [12] Specifically to federated learning, ResNet can maintain gradient flow across decentralized training nodes, making the model robust for federated settings. [14]

The federated settings were achieved by applying a federated learning framework composed of Secure Aggregation and Homomorphic Encryption. These algorithms would encrypt all parameters obtained from local models that were to be transferred to the central server, from which the parameter accumulation dictionary was used to aggregate global parameters for the centrally-run ResNet. Each local device would initialize its own instance of a ResNet and train the model locally. The central server would then periodically update the global model from the aggregated parameters. We iteratively repeated training through rounds until convergence was achieved. This process is depicted in Figure 2b.

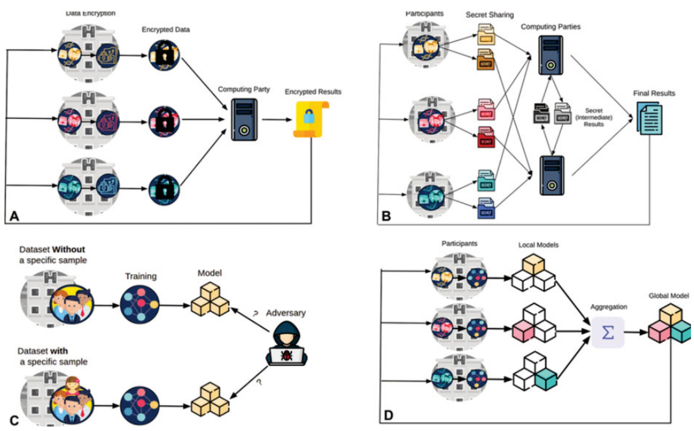

Fig. 2. Privacy-Preserving Techniques Visualization

The other deep learning architectures that we tested for the role of global model were VGGNets, [13] InceptionNets, [15] and MobileNets.[16] However, these architectures had unique challenges that prevented their inclusion as the global model within the federated settings.

First, VGGNets were slower in training as compared to the ResNet global model due in part to the significant number of parameters (138 million parameters) because there are no residual connections and the large number of fully connected layers at the end of the network. [13] InceptionNet was also tested, but the inception modules and the memory required to run parallel computations for performing multiple computations had a prolonged runtime and memory costs for clients. [15] Lastly, MobileNets were not used because of depth wise separable convolutions that reduced the number of parameters too drastically, leading to a tradeoff with lower accuracy scores. [16]

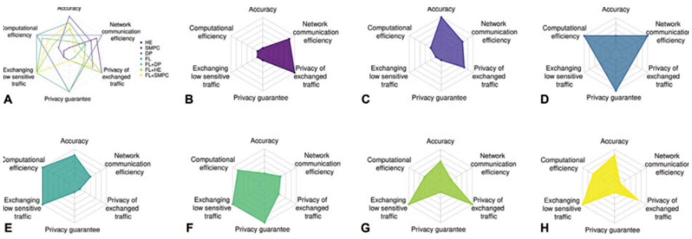

Fig. 3. Privacy-Preserving Techniques Visualization Pt. 2

For transferring parameters from local models into global models, a variety of approaches exist to update global parameters from local training. The options we favored most heavily for the transfer of parameters were between 1.) parameter accumulation dictionaries [20, 21] 2.) gradient-based updates [21] and 3.) model update deltas. [27] Parameter accumulation dictionaries are structures that hold the local parameters and typically transformed with some form of aggregation algorithm (e.g. Seg Agg, FedAvg) [19] to result in composite parameters that are used to update the model. [20, 21] Gradient-based updates aggregate gradients of local clients and apply them to the model with gradient-descent like updates to find optimal parameters. [21, 22] Model update deltas is the process in which local models calculate the difference between local parameters and the initialized parameters of the global model sent to the clients. [27] This occurs per round, and the differences are aggregated, which are then applied to the global model. We chose the parameter accumulation dictionary because it does not require extra computations to calculate deltas and gradients such as those found in the other methods, aiding us in avoiding burdening client computations. Further, the parameter accumulation dictionary would allow us to use Secure Aggregation and additive Homomorphic Encryption for patient privacy, as detailed extensively below. [17, 18] The choice of parameter accumulation dictionary can be changed in the future to find optimal transfer methods for both clients and the global model.

Lastly, we had to make the decision in choosing the Secure Aggregation algorithm [17] over the Federated Averaging algorithm. [24] This was because the Secure Aggregation algorithm most aligned with the eventual goal of protecting patient privacy within medical databases. While FedAvg has proven to be faster and typically the predominant algorithm applied within federated learning algorithms, we did not use it because 1.) there remain no intrinsic privacy mechanisms within FedAvg algorithm. [24] If a foreign entity is to enter the centrally-run server, all local parameters used to train the global model can be openly witnesses. [17] By comparison, the Secure Aggregation algorithm can be outfitted with advanced security schemes such as the Additively Homomorphic Encryption [18], that encrypts all individual parameters and only provides the model with the average aggregated parameter. With the goal of protecting healthcare privacy, the chances of a honest-but-curious type of server (server that follows protocol but may use private information from the data) or external attacks is too great, and so, the Secure Aggregation Algorithm with Homomorphic Encryption was the federated algorithm we chose for our training process. [17, 18]

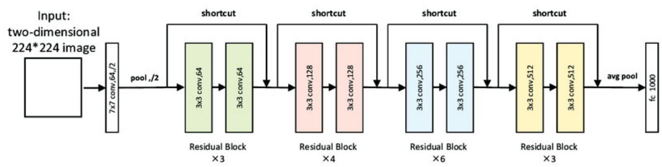

Fig. 4. Examples of ResNet-34 Architecture

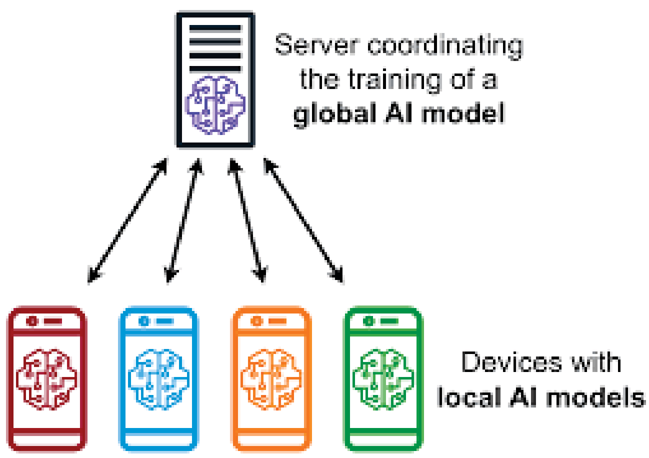

Fig. 5. Federated Learning Example

V. Training and Results

We used a total of five clients for model training, each of which locally trained a ResNet-34 model by using Secure Aggregation on divided sections of the dataset. The number of clients was determined based on the availability of devices that had enough memory to locally train the ResNet. We ran federated learning for 10 rounds, dividing the dataset into batches of size 64 for training. Each client ran for 8 epochs with an optimized learning rate of 0.001 through the Learning Rate Range Test that ensured effective convergence.

Finally, initialize the global model and transfer it to the GPU for training, then iterate through each of the rounds of each client before updating the global model. A new parameter accumulation dictionary was started at the beginning of each round in order to simplify the aggregation of parameters. It is efficient to update the global model after each client has finished its training in every round. The best validation loss encountered during the training process was 0.09, where the best accuracy reached was 0.83, hence really strong performance and generalization in the context of federated learning.

While the focus was on Secure Aggregation for privacy preservation, we suspect that the use of Federated Averaging could improve performance to the model. FedAvg allows clients to make multiple updates of the local gradient descent before aggregating those parameters into the global model, thus enabling the combination of the best gradients out of each local model. [24, 25] This may make the global model generalize even better, especially when the data across clients is non-IID. Given the promising results we have gotten, FedAvg could be an avenue worth pursuing in further experiments.

VI. Conclusion

This work shows how federated learning can be used to enhance medical image classification with consideration for some critical privacy and data security concerns in healthcare. We investigated a new concept in the use of federated learning for the classification of thoracic diseases using decentralized X-ray image scans on the NIH Chest X-ray dataset. Considering that medical data are highly sensitive and often the process for dataset acquisition is very laborious, federated learning will provide an emerging solution for multiple institutions to collaboratively train a shared model without compromising patient privacy and decentralized datasets.

Comparing local and global models, one could observe that the latter aggregated knowledge across all the clients to yield more generalizable results. While the local models were optimized to perform better on their respective subsets of the overall data, the global model enriched its knowledge regarding the whole data distribution from the collective specialized parameters aggregate from various subsections of the dataframe. By keeping simply the parameters as opposed to sharing the exact same proportion of the database, the ResNet is able to update accuracy.

We do a literature review of related works on federated learning applications to medical imaging and decide upon the NIH Chest X-ray dataset. The dataset consists of more than 100,000 images across 14 different thoracic diseases with a ”No Findings” category. First, data preprocessing was performed to transform all image scans into the standardized format for the ResNet-34. Next, we implemented a federated learning framework by training and testing multiple deep learning architectures in a decentralized setting. The experiments we conducted showed that the ResNet-34 model, a deep convolutional neural network famous for residual learning, performed better compared to other architectures such as VGGNet, InceptionNet, and MobileNet. The federative learning framework was outfitted with an parameter accumulation dictionary and Secure Aggregation with in-built additive homomorphic encryption. With these measures, we build an optimal architecture to protect patient privacy by encrypting the local parameters sent to the server, before finally using the parameter dictionary for a smooth convergence scheme. The peak performance of the ResNet-34 global model was at 0.83 accuracy and 0.09 validation loss. In this case, the use of federated learning pushed very close to current benchmark performance of 0.87 and validation loss of 0.06, showcasing the applicability of using a federated system to protect patient privacy while retaining accuracy, towards clinical integration.

Results have proven the feasibility of federated learning for real-world medical data to reach generalizable performance. The framework not only sustains high accuracy but also protects patient privacy by decentralized training. Since data privacy has become one of the most important barriers in medical research, federated learning can potentially open a lot more collaborations across healthcare institutions with no need for extensive data sharing. Our approach showed the potential for federated learning to provide accuracy equal to or greater than a centralized model while eliminating many logistical headaches and most of the privacy risks evident in traditional machine learning.

A. Future Works

Further, in our future research work, we want to extend this work by comparing ResNet-34 with other state-of-the-art models for image classification. We have seen that ResNet-34 performed well in the federated learning setting; however, more alternatives may be considered, such as EfficientNet, DenseNet, or even Transformer-based architectures. These models might give a large spectrum of trade-offs regarding training speed, memory consumption, and smooth gradient descent, especially in decentralized settings where computational resources differ across devices. This will provide a broader insight into which architectures will work best for medical imaging use cases in federated learning.

Another essential direction to work towards is efficient hyperparameter tuning. In the present work, we used fixed hyperparameters for training from manual tests such as the learning rate range test, while in future works, wider ranges will be explored, such as learning rate, batch size, and number of local epochs. Other aggregation algorithms will be tried, too, including FedAvg, FedProx, and personalized federated learning approaches. Gaining insight into how different hyperparameter settings affect decentralized learning and their ability to preserve patient privacy will be the key to fine-tuning training on a variety of tasks and datasets.

Scalability is another critical aspect for which we have designed our investigative approach. The federated learning system considered in this work was limited to five clients. The next step will be the expansion of the number of participating clients in order to investigate whether the proposed model scales well and can generalize over a greater variety of devices with diverse computation power. We also plan to apply the proposed federated learning framework on larger and more diverse medical imaging datasets, such as those from CT scans or MRI scans. In that way, the generalization ability of the proposed approach will be able to be assessed for various types of medical data beyond chest X-rays, potentially in the context of zero-shot or few-shot models.

Another promising direction for future work is to incorporate enhanced measures for privacy and security. Although federated learning by its nature provides several advantages in terms of privacy, we aim to incorporate techniques related to differential privacy and secure multiparty computation to enhance data security even further. This would ensure that even while the data of patients is shared across decentralized networks, no breach in data privacy occurs, and model accuracy is effectively retained. By incorporating multiparty computations, extra security layers can be built into the framework which would contribute to the same shared goal of mitigating data breaches to protect patient privacy.

Eventually, the deployment of a federated learning system within a clinical workflow would remain our critical objective. We envision collaboration with hospitals and other healthcare institutions to test our model on non-public datasets in order to understand the practical challenges of implementing federated learning in a decentralized health network. This would provide an opportunity to share important lessons learned from the operational process, handling diverse data formats, and managing infrastructure limitations. We will also look into how advances in Cross-silo, Hierarchical Federated Learning-one can allow better cross-site performances in large-scale distributed healthcare networks.

These will serve to further fine-tune the models toward clinical-grade performance for medical image analysis while ensuring patient privacy.

REFERENCES

[1] "SOURCES AND EFFECTS OF IONIZING RADIATION United Nations Scientific Committee on the Effects of Atomic Radiation UNSCEAR 2008 Report to the General Assembly with Scientific Annexes VOLUME I UNITED NATIONS," 2010. Accessed: Sep. 09, 2024. [Online]. Available: https://www.unscear.org/docs/publications/2008/UNSCEAR_2008_Annex-A-CORR.pdf

[2] R. Najjar, "Redefining Radiology: A Review of Artificial Intelligence Integration in Medical Imaging," Diagnostics, vol. 13, no. 17, p. 2760, Aug. 2023, doi: https://doi.org/10.3390/diagnostics13172760.

[3] B. Sistaninejhad, H. Rasi, and P. Nayeri, "A Review Paper about Deep Learning for Medical Image Analysis," Computational and Mathematical Methods in Medicine, vol. 2023, p. e7091301, May 2023, doi: https://doi.org/10.1155/2023/7091301.

[4] W. G. Panhuis et al., "A systematic review of barriers to data sharing in public health," BMC Public Health, vol. 14, no. 1, Nov. 2014, doi: https://doi.org/10.1186/1471-2458-14-1144.

[5] N. Rieke et al., "The future of digital health with federated learning," npj Digital Medicine, vol. 3, no. 1, pp. 1–7, Sep. 2020, doi: https://doi.org/10.1038/s41746-020-00323-1.

[6] W. Li et al., "Privacy-preserving Federated Brain Tumour Segmentation," arXiv:1910.00962 [cs], Oct. 2019, Available: https://arxiv.org/abs/1910.00962

[7] P. Rajpurkar et al., "CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning," arXiv.org, 2017. https://arxiv.org/abs/1711.05225

[8] S. M. McKinney et al., "International evaluation of an AI system for breast cancer screening," Nature, vol. 577, no. 7788, pp. 89–94, Jan. 2020, doi: https://doi.org/10.1038/s41586-019-1799-6.

[9] Zia UrRehman et al., "Effective lung nodule detection using deep CNN with dual attention mechanisms," Scientific Reports, vol. 14, no. 1, Feb. 2024, doi: https://doi.org/10.1038/s41598-024-51833-x.

[10] B. Evaldsson, "Federated Learning for Brain Tu- mor Segmentation." Accessed: Sep. 09, 2024. [Online]. Available: https://liu.diva-portal.org/smash/get/diva2:1847651/FULLTEXT01.pdf

[11] A. Mitrovska, P. Safari, K. Ritter, Behnam Shariati, and Johannes Karl Fischer, "Secure federated learning for Alzheimer's disease detection," Frontiers in aging neuroscience, vol. 16, Mar. 2024, doi: https://doi.org/10.3389/fnagi.2024.1324032.

[12] He, F., Liu, T., and Tao, D. (2020). Why resnet works? residuals generalize. IEEE transactions on neural networks and learning systems, 31(12), 5349-5362.

[13] Koloskova, A., Lin, T., Stich, S. U., and Jaggi, M. (2019). Decentralized deep learning with arbitrary communication compression. arXiv preprint arXiv:1907.09356.

[14] Alippi, C., Disabato, S., and Roveri, M. (2018, April). Moving convolutional neural networks to embedded systems: the alexnet and VGG-16 case. In 2018 17th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN) (pp. 212-223). IEEE.

[15] Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. (2016). Rethinking the inception architecture for computer vision. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 2818-2826).

[16] Sood, S., Singh, H., Malarvel, M., and Ahuja, R. (2021, October). Significance and limitations of deep neural networks for image classification and object detection. In 2021 2nd International Conference on Smart Electronics and Communication (ICOSEC) (pp. 1453-1460). IEEE.

[17] Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., ... and Seth, K. (2016). Practical secure aggregation for federated learning on user-held data. arXiv preprint arXiv:1611.04482.

[18] Aono, Y., Hayashi, T., Wang, L., and Moriai, S. (2017). Privacy-preserving deep learning via additively homomorphic encryption. IEEE transactions on information forensics and security, 13(5), 1333-1345.

[19] Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... and Zhao, S. (2021). Advances and open problems in federated learning. Foundations and trends® in machine learning, 14(1–2), 1-210.

[20] Yaghoobi, M., Daudet, L., and Davies, M. E. (2009). Parametric dictionary design for sparse coding. IEEE Transactions on Signal Processing, 57(12), 4800-4810.

[21] Aharon, M., and Elad, M. (2008). Sparse and redundant modeling of image content using an image-signature-dictionary. SIAM Journal on Imaging Sciences, 1(3), 228-247.

[22] Li, T., Sahu, A. K., Zaheer, M., Sanjabi, M., Talwalkar, A., and Smith, V. (2020). Federated optimization in heterogeneous networks. Proceedings of Machine learning and systems, 2, 429-450.

[23] Jayaram, K. R., Muthusamy, V., Thomas, G., Verma, A., and Purcell, M. (2022, December). Adaptive aggregation for federated learning. In 2022 IEEE International Conference on Big Data (Big Data) (pp. 180-185). IEEE.

[24] McMahan, B., Moore, E., Ramage, D., Hampson, S., and y Arcas, B. A. (2017, April). Communication-efficient learning of deep networks from decentralized data. In Artificial intelligence and statistics (pp. 1273-1282). PMLR.

[25] Chen, M., Shlezinger, N., Poor, H. V., Eldar, Y. C., and Cui, S. (2021). Communication-efficient federated learning. Proceedings of the National Academy of Sciences, 118(17), e2024789118.

[26] Nalluri, S., Sasikala, R. (2024). Pneumonia screening on chest X-rays with optimized ensemble model. Expert Systems with Applications, 242, 122705.

[27] R. Torkzadehmahani et al., "Privacy-Preserving Artificial Intelligence Techniques in Biomedicine," Methods of Information in Medicine, Jan. 2022, doi: https://doi.org/10.1055/s-0041-1740630.

[28] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated Learning: Challenges, Methods, and Future Directions," IEEE Signal Processing Magazine, vol. 37, no. 3, pp. 50–60, May 2020, doi: https://doi.org/10.1109/msp.2020.2975749.

[29] P. M. Mammen, "Federated Learning: Opportunities and Challenges," arXiv:2101.05428 [cs], Jan. 2021, Available: https://arxiv.org/abs/2101.05428