
“Ensemble Knowledge Distillation for Federated Semi-Supervised Medical Image Classification”

Shivani S, Suma L

Dept. of CSE, Bangalore Institute of Technology, Bangalore 560002

shivani7632@gmail.com, sumal@bit-bangalore.edu.in

Abstract.

Federated Learning (FL) allows clients (e.g., hospitals) to train shared models while preserving privacy and not sharing raw data. Most FL methods assume all the data is labeled, which is not the case for medical imaging because of the high cost and the specialized skills needed to perform data annotation. Medical FL clients possess a large volume of unlabeled data but only a small amount of labeled data. In this work, we present EKDFSSL (Ensemble Knowledge Distillation with Federated Semi-Supervised Learning) that extends FSSL to the case of detecting pneumonia in X-ray images. Client models in FSSL gain the ability to learn from unlabeled data due to the presence of the global model as a teacher and the application of consistency regularization to the client models. This enhances the stability and reliability of the model predictions in the presence of input data variations. At the server side, knowledge distillation is performed to reduce overfitting and to improve model performance by aggregating the client models. This technique of Federated Learning extends the client models ability of aligning their predictions with the global model, improving generalization. Our method is privacy preserving and, in the case of Federal Learning, makes the best use of unlabeled data. Our method provides a scalable, privacy preserving, and practical solution to medical imaging that suffers from a lack of annotated data.

Keywords. EKDFSSL, Federated Learning (FL), Semi-Supervised Learning (SSL), Federated Semi-Supervised Learning (FSSL), Knowledge Distillation (KD), Pneumonia Detection, Chest X-ray, Privacy-Preserving AI, Consistency Regularization (CR), Ensemble Learning (EL)

1. INTRODUCTION

The rapid growth of artificial intelligence in healthcare has made medical image analysis a valuable asset for supporting timely and accurate diagnoses of various medical conditions. However, when considered in conjunction with the existing limitations of imaging modalities, the analysis of medical images continues is hampered by the limited availability of datasets and continual concerns for patient privacy.

Pneumonia diagnosis from chest X-ray images presents multiple challenges. As well as having to pay expert radiologists to review the images, the provision of labeled examples is extremely costly. Moreover, both the Health Insurance Portability and Accountability Act and the General Data Protection Regulation present constraints on sharing data between different medical institutions. Additionally, traditional deep learning systems face serious obstacles in privacy, communication costs, and practicality in operating in real world healthcare systems, due to a reliance on centralized data collection. Because of the issues outlined above, there is a compelling case for the development of learning systems that are privacy preserving and adaptable to both labeled and unlabeled data.

In the last few years, the ability to collaboratively develop machine learning models without the need for data sharing has resulted in federated learning becoming one of the most popular privacy preserving frameworks. However, despite the benefits of this novel approach, most federated learning techniques are designed with the assumption that all involved parties have access to some form of labeled data. However, this is most definitely not the case as most healthcare institutions will tend to have predominantly unlabeled data, like chest X-ray images, with only a small volume of data that is labeled. This means a considerable challenge for the machine learning community is to develop systems which are privacy preserving and address the above imbalance while providing high quality outcomes.

To tackle the above challenge, we propose a framework called EKDFSSL (Ensemble Knowledge Distillation for Federated Semi-Supervised Learning) that detects pneumonia from chest X-ray images. Our approach combines consistency regularization and ensemble knowledge distillation in a federated learning environment. Herein, the global model provides guidance to client models, which helps promote the internal models to predict consistently for different augmentations of the same unlabeled data. This stabilizes the learning process and protects against the global model’s performance worsening. On the server side, a knowledge distillation ensemble is formed to combine the client model predictions and facilitate learning with limited labeled data while combating overfitting.

There are several aspects of this work that advance the state of the art. The first is the design of a framework for federated semi-supervised learning that operates in a realistic environment where all labeled data exists solely on the server, while the clients are equipped solely with unlabeled data. The second is augmentations-based consistency regularization technique in a global model to elevate the learning process of clients. The third is the incorporation of ensemble knowledge distillation to combine the models of clients to improve the overall model. The fourth is the enhancement of privacy due to minimal data transfer. Finally, we developed a web-based (Flask) interface to conduct secure and easy user interactive pneumonia screening.

This work introduces a novel pneumonia detection system that is both scalable and considers user privacy, and that has the potential to be used in real world healthcare settings.

2. LITERATURE SURVEY

2.1. *Ensemble Knowledge Distillation for Federated Semi-Supervised Image Classification*

This work presents a federated semi-supervised learning framework for image classification where the labels exist at the server. Here, the global model acts as a teacher in a consistency

regularization framework to enhance learning from unlabeled data at the client side. To counter the overfitting problem due to the scarce labeled data, an ensemble knowledge distillation loss is used for server-side retraining. The performance of the framework is evaluated on several federated datasets such as Fashion-MNIST, SVHN, CIFAR-10, and CIFAR-100 and the results outperformed existing frameworks such as FedMatch, FedRGD, FedMix, and GDST. The framework is designed to achieve high accuracy for both IID and non-IID data. The current work is also limited by the assumption that all clients possess unlabeled data whereas all the server has is labeled data, an assumption that may not hold in many real world cases [1].

2.2. *Clients Help Clients: Alternating Collaboration for Semi-Supervised Federated Learning*

One way to mitigate communication costs while dealing with data heterogeneity in semi-supervised federated learning involves alternating collaboration. Similarity-based clustering helps to form groups of clients that deal with comparable data distributions to generate robust pseudo-labels. To counteract bias from non-IID data, dissimilarity-based clustering is employed. To optimize the labeling and support balancing, a reinforcement learning strategy is incorporated. Compared to baseline methods, the proposed method demonstrated an accuracy increase, along with a reduction in communication costs. Although clustering at each training round does not account for the variability in client computational resources, it does introduce additional computation costs [2].

2.3. *Federated Semi-Supervised Learning for Attack Detection in IIoT*

A novel federated semi-supervised scheme for intrusion detection has been proposed for the Industrial Internet of Things (IIoT). Using federated averaging, local autoencoders are first trained on the available unlabeled data, and then the local models are aggregated at the server. The server then provides labeled data and trains the added classification layers. Detection accuracy is achieved while maintaining privacy through the scheme. The sparse labeled datasets and communication costs used in the experiments illustrate the high detection accuracy and performance of the proposed scheme. The lack of convolutional layers in the proposed design may prevent the learning and detection of complicated spatial intrusion patterns. In this design, the thresholds used for pseudo-labeling are fixed and may not be optimal throughout the training process [3].

2.4. *Knowledge-Driven Federated Learning for AIS Signal Identification*

The proposed approach integrates knowledge-based and semi-supervised Federated Learning (FL) methods to address classification of AIS signals under the constraint of unbalanced labeling. The approach uses an HSIC-based similarity consistency to align feature representation across clients and a voting-based method for the generation of pseudo-labels to enhance the quality of the labels. The approach is able to attain a high level of classification accuracy across disparate domains, and thus is able to handle imbalanced data. One of the drawbacks is that the computation of HSIC requires parameter tuning and the voting method increases the communication burden due to additional model exchanges [4].

2.5. *Heterogeneity-Aware Semi-Supervised Federated Learning*

The proposed approach deals with the different levels of computation and communication capability of clients by incorporating them into the clustering process to minimize the effect

of slow clients. Using a Multi-Armed Bandit approach, different levels of training iterations are implemented for the different clusters, which improves the overall process. The results showed a significant improvement in convergence and a reduction in the communication burden. One of the drawbacks of the approach is that the framework is based on the assumption that heads of clusters are trustworthy, which is not always the case, and that adaptive optimization increases the burden on computation at the server [5].

2.6. *Privacy-Preserving Generative Network for Smart Grids*

A framework for anomaly detection in smart grids has been developed that combines generative networks with semi-supervised federated learning. This framework is built on the combination of temporal convolutional blocks and the geometric median for aggregation under the non-IID assumption. The framework is capable of having a high degree of accuracy in classifying both binary and multi-class examples. One of the challenges with this approach is that the use of generative adversarial networks (GANs) is known to produce unstable solutions and can cause mode collapse. Moreover, more layers in the architecture increase the computational burden [6].

2.7. *Robust Semi-Supervised Federated Learning for UAV Image Recognition*

An image recognition framework, built with a robust semi-supervised federated learning approach, is designed for systems that utilize drones. This framework is flexible to the models of the labels, either at the client or at the server, and incorporates a parameter mixing strategy to integrate supervised, unsupervised, and global models. This framework also utilizes a frequency-based aggregation approach to remedy the non-IID assumption and employs a hyperparameter self-adjustment strategy. Despite the framework having stable performance across multiple conditions, the main drawback is the heavy computational burden and memory usage caused by the large number of models [7].

2.8. *Random Sampling Consensus Federated Semi-Supervised Learning*

To improve model reliability in non-IID situations, a federated semi-supervised learning method that uses random sampling consensus has been developed. By randomly selecting clients to generate sub-models and aggregating the sub-models using a distance-based weighting method, the method is able to focus on the reliable models and optimize the performance to a greater extent. This method has been experimentally found useful with a wide variety of datasets and numbers of clients. The random sampling approach may ignore valuable client inputs. Additionally, the distance-based weighting method has a high computational complexity that will increase with the number of clients [8].

2.9. *Advanced Model-Contrastive Federated Learning via Knowledge Distillation*

MOON-KD is the name of an advanced federated learning method that combines model-contrastive learning and knowledge distillation in the aim of attaining improved performance given non-IID data. The global model, in this context, functions as a teacher, instructing local models on the appropriate changes to be made via the distillation loss, which is integrated with the supervised contrastive loss, in a given objective function. The alignment of local models with the global model is, thus, retained, and relative model performance is improved. In particular, it has been found, in experimental work carried out using the CIFAR-10 dataset, that MOON-KD outperforms the baseline methods of MOON and FedAvg, and that the distillation and contrastive losses, in particular, define the

boundaries of the overall model performance. There is no personalization for the client within this method, and there may be a large communication burden as there is an exchange of large model parameters. There are proposed method improvements that will fall within the confines of model compression, privacy, and more dynamic aggregation methods [9].

2.10. Federated Learning Algorithm Based on Knowledge Distillation

FedDistill is an approach using Knowledge Distillation to tackle issues of disparate model architectures and data distributions at client sites in Federated Learning (FL). Here, clients store both a Personal and a Global model. The Personal model acts as a teaching model that distills knowledge to the Global model, enabling learning across diverse model architectures. Evaluation of the approach on several datasets, like MNIST and CIFAR-10, shows both local and global model accuracy gains. The approach also brings a communication cost advantage due to the many local updates it allows prior to model aggregation. However, this approach still suffers from challenges of the classical averaging frameworks and is sensitive in performance to the local training iteration ceiling. It also requires tuning to understand the ideal training parameters for local iterations [10]. The model cannot effectively be used for real time detection due to high computational demands. The model's complexity is due to additional layers, high parameter counts, and intricate MRS-YOLO feature extraction. This all results in a significant increase in memory utilization [10].

3. METHODOLOGY

This study aims to develop the framework for building a privacy-preserving pneumonia detection model from CXR images using the Federated Semi-Supervised Learning (FSSL) approach. The foundation of the proposed framework is the EKDFSSL, which combines client-side unsupervised learning and server-side supervised learning.

In the first step, the chest X-ray images are distributed across client nodes, which in this case are assumed to be hospitals and diagnostic centers. Each of these client nodes focuses on a large amount of unlabeled data, while the central server holds a small amount of labeled data. This situation is typical in practice, where there is a considerable amount of unlabeled data and a limited amount of sensitive labeled medical data.


A global model is first created at the server and communicated to a subset of selected clients. Each client is then able to perform a training step on their unlabeled data. One approach to learning from unlabeled data is through the use of consistency regularization. In this framework, two different augmented versions of the same data sample are used. One is a weak version and the other is a strong version, and the model is trained to produce consistent output for both. The global model acts as a teacher to the client model, which learns to develop stable, meaningful representations.

Once local training is complete, each client communicates the updated model parameters to the server. The server then uses the Federated Averaging (FedAvg) algorithm to develop an enhanced global model. Knowledge is combined in the aggregation process, which supports privacy, as it does not involve the communication of any raw data.

The server then conducts a supervised re-training phase on its collected labeled dataset. Here, the model is optimized using a combination of the cross-entropy loss and the ensemble

knowledge distillation loss. In the distillation process, predictions from several client models serve as an ensemble teacher and offer a more informative supervision to the global model.

The international model is then assessed using a test data set. The complete system is implemented via a web interface built with Flask. This enables health care workers to use the model for pneumonia screenings with the assurance of data safety and privacy.



3.1. Data Distribution

For the system we propose, gross examination of chest X rays data is shared among a central server and several clients (hospitals). The server stores a small labeled data set of normal and pneumonia cases, while each client stores a large data set of unlabeled chest X ray images. This distributed system imitates most real world medical situations where labeled

data is of a small volume, sensitive, and contains private medical information. It contains the unrefined patient data and creates a variety of situations within the environment. The system is designed to better the performance of the system in the real world.

3.2. *Client-Side Training using Consistency Regularization*

For the client based training, we use the semi-supervised learning based on the consistency regularization (FixMatch) on the unlabeled data. For every input image, we generate two images, one with weak augmentation and one with strong augmentation. The weakly augmented image is sent to the global model, which in this case acts as a teacher to generate soft pseudo-labels. The local model (the student) is then tasked with the strongly augmented image and generates a prediction.

The consistency loss along with a Kullback-Leibler (KL) divergence helps bring the teacher and student outputs closer. The student model is then fine-tuned with an optimizer, such as stochastic gradient descent (SGD). This approach helps the model learn from unlabeled data and keeps the predictions consistent for the different transformations applied to the data.

3.3. *Server Aggregation using Federated Averaging (FedAvg)*

Once the local training is complete, the updated model weights from a client are communicated to the central server. The server then applies the FedAvg algorithm to aggregate the models. In this method, the contribution made by a client is dependent on the client's dataset size. The resulting global model is enhanced with the knowledge from all the clients, and the privacy of the clients is maintained as the raw data from the clients is not accessed.

3.4. *Server-Side Retraining using Ensemble Knowledge Distillation (EKD)*

The server focuses on model performance enhancement by retraining with its labeled data. This stage utilizes Ensemble Knowledge Distillation (EKD). This method generates predictions for the same input from the global and several client models, which are then combined through an ensemble averaging method to create a more powerful supervisor signal.

Training is performed using a combination of cross-entropy loss (for labeled data) and knowledge distillation loss (based on ensemble predictions). During this training, control for each loss is introduced. This method significantly improves model performance when labeled data is insufficient.

3.5. *Output and Visualization*

The server focuses on model performance enhancement by retraining with its labeled data. This stage utilizes Ensemble Knowledge Distillation (EKD). This method generates predictions for the same input from the global and several client models, which are then combined through an ensemble averaging method to create a more powerful supervisor signal.

Training is performed using a combination of cross-entropy loss (for labeled data) and knowledge distillation loss (based on ensemble predictions). During this training, control for each loss is introduced. This method significantly improves model performance when labeled data is insufficient.

4. ARCHITECTURE




Fig 1: Architecture Diagram

The server enhances model performance by retraining with the labeled dataset. The process uses Ensemble Knowledge Distillation (EKD). Given the same input, the global model and various client models generate predictions, which, after being synthesized with an ensemble averaging technique, provide a more robust supervisory signal.

The training process uses a combination of cross-entropy loss (for labeled data) and knowledge distillation loss (derived from ensemble predictions). A progressive weighting factor assigns balance to the losses during the training process. The model's generalization and accuracy are improved with the limited labeled data.

4.1. Client Data Module

The system targets a distributed architecture with various client nodes (hospitals) as data sources. Each client has a collection of locally stored chest X-ray images, most being unlabeled, for data privacy. Furthermore, the central server stores a small labeled dataset with normal and pneumonia cases.

We intentionally split the data to mimic the challenge in the medical field, where there is limited labeled and sensitive data, but there exists a considerable volume of unlabeled data across various medical institutions. The module ensures that data remain in their respective clinics and are only utilized for local computation.

4.2. Client-Side Training Module

This module manages the local model training for each client. Relevant features from chest X-ray images are extracted using a convolutional neural network (CNN). To enhance model performance with available unlabeled data, the system employs a consistency-based learning approach due to the limitation of labeled data at the client side.

The global model is sent from the server to the clients, while the clients make revisions to the global model using their local datasets. The training process is designed to make the model produce robust and consistent predictions, even when the input data vary

significantly. When the training process is complete, the model is packaged, and only the revised model parameters are sent to the server.

4.2.1. Feature Extraction Module

The central server coordinates the learning process. When the clients update their models, they send the model updates to the central server. The server updates the model without needing to see the raw data. The server uses a technique called Federated Averaging to combine model parameters. This technique allows the server to learn from the clients and keeps the data private. After the model is updated, the server sends the global model to the clients to continue the training process. This framework supports a learning process that is both iterative and collaborative.

4.2.2. Server-Side EKD Module

To improve the global model, the server adds an EKD Module. In this case, the EKD Module combines the predictions of multiple client models and the aggregated global model.

The model is retrained with the labeled data the server has to improve the model's accuracy and its ability to generalize. The system is able to optimally combine both labeled and unlabeled datasets. The system also shows an increased robustness to unstandardized medical image data.

4.2.3. Output and Evaluation Module

The final global model is integrated into a web-based system where users can submit chest X-ray images. The system processes the images and classifies the cases as normal or pneumonia with confidence score. The system is evaluated with traditional classification metrics of accuracy, precision, recall and F1-score. These metrics are benchmarks of the system's effectiveness and reliability in real world clinical applications. Also, the performance is evaluated based on the number of iterations of the training rounds to check if the model is improving.

5. DATASET DESCRIPTION

The dataset is composed of chest X-ray images (normal cases and pneumonia cases) collected from open sourced medical imaging. The dataset is purposefully arranged to comply with the federated learning framework. There is a small amount of labeled data at the central server and large amounts of unlabelled data at the various client locations (i.e. hospitals). This data arrangement imitates real world clinical situations where there is small amount of labeled data and large amount of unlabelled data readily available.

5.1. Data Sources

The chest X-ray data was obtained from sites like Kaggle that host open source medical data. These data repositories contain X-ray data that are specifically collected for pneumonia detection.

The server's dataset contains labeled data that are clearly classified to facilitate supervised learning and model assessment. These data repositories are highly regarded in academic research and therefore can be used to evaluate and benchmark model's performance.

5.2. *Dataset Composition*

There are two primary classes in this dataset.

Normal class contains Chest X-ray images which are devoid of infectious signs.

The Pneumonia class contains images that depict signs of lung infection.

In the described system, the labeled dataset is centralized, and the client devices hold unlabeled images. This allows for combining the supervised learning at the server and semi-supervised learning at the client, respectively.

5.2.1. *Data Annotation*

The labeled dataset contains annotations that are at the image level and classify the image to either the Normal or Pneumonia class. These annotations are utilized in the training and evaluation at the server.

5.2.2. *Data Preprocessing*

Before the images are utilized for training, they undergo some basic preprocessing to ensure uniformity. The images are resized, the pixel values are normalized and converted to the appropriate image input for the model.

These help to improve the consistency of the training and ensure the model is able to learn the important and the correct patterns from the data.

5.2.3. *Data Augmentation*

For the purpose of creating a robust model, various data augmentation techniques are employed in the training, which may include flipping, rotation, scaling, and brightness.

These changes allow the model to become more flexible and adaptive to the real-world challenges posed by medical imaging, such as variable image quality, fluctuating orientation, and poor or imperfect lighting.

5.2.4. *Dataset Split*

For evaluation purposes, the labeled dataset is partitioned into training, validation, and testing sets to promote efficient learning and allow for a more accurate assessment of performance. The model training employs roughly 80% of the overall dataset, where the remaining 20% is allocated for validation and testing. The training dataset allows the model to learn critical distinguishing attributes of the chest X-ray images. The validation dataset is used to control the learning process to prevent overfitting.

After training, the model is tested with the testing dataset to assess the performance of the model beyond the constraints of the training dataset. This is the most accurate indication of the model's ability to generalize in the context of a real-world medical diagnosis. The dataset splitting strategy promotes a fair assessment and enhances the reliability of the results from the experiments.

6. RESULTS

We assess the federated semi-supervised medical image classification framework for chest X-ray images dispersed across multiple hospital clients. Each hospital develops its own local CNN model on the local chest X-ray data. To avoid sharing local chest X-ray data, hospitals send the learned model weights to the central server. This maintains data privacy and security during the training process.

The central server is essential to enhance the overall model. After receiving model updates from all the clients, the central server uses the Federated Averaging (FedAvg) algorithm to construct a global model. Furthermore, Ensemble Knowledge Distillation (EKD) is applied to improve the classification performance by bridging the knowledge gap between client models, thereby achieving efficient distributed learning.

We demonstrate that our DenseNet121-based model has a superior performance on medical image classification when compared to the ResNet model. DenseNet121 uses a principal design that allows the model to learn which features are necessary to classify images, thus providing a better mechanism to identify and learn the important features of chest X-Ray images. Our model has an accuracy of almost 83%, which demonstrates that our model is reliable in distinguishing between normal and abnormal chest X-Ray images.

The training process demonstrates stability with multiple rounds of communication between the client and server. Training loss is decreasing over time, and so the model is learning. The validation loss, similarly, shows stability, indicating that the model is able to generalize to unseen data, with no significant overfitting.

The usability of the system is further enhanced with the addition of the Flask-based web interface, which allows users to upload chest X-ray images and receive predictions. The system also maintains a database of prediction logs, model info, and prediction-related reports. The system provides an efficient, safe, and privacy-preserving way to implement federated medical image classification with deep learning.




Fig 2: ResNet Accuracy

The performance of the ResNet model during training was inferior and it was not able to accomplish chest X-ray classification results on par with DenseNet121. There was some improvement to training accuracy, but it was slow and inconsistent. The validation accuracy was also low, but with more frequent fluctuations, which shows poor generalization to chest X-ray images.




Fig 3: DenseNet Accuracy

DenseNet121 reached an impressive nearly 87% validation accuracy, while ResNet struggled to surpass 70% accuracy for most of the training rounds. It was fairly evident in the training and validation results that ResNet was not as suitable for the medical image classification task, while DenseNet121 offered better learning, feature extraction, and consistency for effective pneumonia detection.




Fig 4: Cofusion Matrix of ResNet




Fig:5 Cofusion Matrix of DenseNet.

The confusion matrix illustrates that DenseNet121 was superior to ResNet in classifying pneumonia in chest X-ray images. DenseNet121 classified the majority of pneumonia instances and had a low number of false-negative instances, making it an excellent choice for medical detection tasks. In comparison, ResNet had an increased number of false-positive instances and was unable to distinguish pneumonia images from normal images. The overall image classification results depict that DenseNet121 is the most accurate and reliable of the two for the task of classifying pneumonia in chest X-rays.

7. CONCLUSION

This paper proposes a system for automated pneumonia detection in chest X-ray images using deep learning. The system utilizes DenseNet121 for feature extraction and classification as a result of the benefits that the architecture and transfer learning provide. The system also implements knowledge distillation to simplify the model while retaining the diagnostic strength. The results of the experimental evaluations suggest that the model is capable of maintaining a validation accuracy of 87% and a pneumonia detect rate of 96.7%, indicating that the model is capable of accurately identifying pneumonia cases with a limited number of missed cases.

Training and validation results show good convergence during training, with decreasing loss values and successful feature learning. Training and validation accuracy values show good generalization, with limited overfitting. The analysis of the confusion matrix shows the model is efficient, with only a small number of pneumonia cases being misclassified. While normal case classification accuracy is moderate, the model's priority is on the classification of pneumonia cases, which is important in medical diagnosis, to reduce the chances of missing a pneumonia infection.

The method proposed shows a high degree of applicability in the medical field in that it aids the radiologist in the early detection of pneumonia and in the lessening of the overall workload associated with diagnosis. The framework will be useful in future research to be extended to the classification of other respiratory diseases, COVID-19 and tuberculosis, and the presence of lung opacity in X-ray images. In addition, lightweight distilled models may enable portable edge devices to be used for diagnosis in real time, and combine with explainable AI methods to improve clinical readability by visually annotating areas of pneumonia on chest X-ray images.

8. FUTURE WORK

The proposed system shows strong performance in detecting pneumonia from chest X-ray images. There are, however, additional features that can be added to the system to increase its capability and applicability in a clinical setting. The system currently classifies images as Normal and Pneumonia. However, the system can be developed to classify additional respiratory illnesses such as COVID-19, tuberculosis, lung opacity, pneumothorax, and pulmonary fibrosis. This would develop a more comprehensive system for the classification of medical images in the screening of respiratory illnesses.

The classification system can be further improved through the development of advanced teacher-student frameworks in knowledge distillation. More advanced heterogeneous distillation frameworks, where the teacher and student are of different architectures, can help decrease the computational burden of the classification system, while still maintaining classification performance. The quality of knowledge distillation can also be improved by the use of adaptive temperature scaling and attention transfers, among other techniques.

The proposed system is also adaptable for low-resource healthcare settings, such as rural hospitals, emergency care systems, and mobile healthcare units. Lightweight distilled student models can be deployed to the edge for real-time diagnosis with a lower computational burden. Future work can also incorporate dynamic chest imaging and

fluoroscopy with a focus on improving temporal features with more consistency in the diagnosis.

9. REFERENCE

- [1] E. Shang, H. Liu, J. Zhang, R. Zhao, and J. Du, "Ensemble knowledge distillation for federated semi-supervised image classification," *Tsinghua Science and Technology*, vol. 30, no. 1, pp. 112-123, Feb. 2025.
- [2] Z. Jiang, Y. Xu, H. Xu, Z. Wang, and C. Qiao, "Clients help clients: Alternating collaboration for semi-supervised federated learning," in *Proc. IEEE 40th Int. Conf. Data Engineering (ICDE)*, 2024, pp. 1847-1860.
- [3] O. Aouedi, K. Piamrat, G. Muller, and K. Singh, "Federated semisupervised learning for attack detection in industrial internet of things," *IEEE Trans. Ind. Informat.*, vol. 19, no. 1, pp. 995-1005, Jan. 2023.
- [4] Q. Tian, S. Wang, and M. Wang, "FSLBSC: Knowledge-driven federated learning for unbalanced labeling AIS signal identification," *IEEE Trans. Intell. Transp. Syst.*, 2025.
- [5] Y. Xu, Z. Yao, H. Xu, Y. Liao, Y. Zhu, and X. Fu, "Heterogeneity-aware semi-supervised federated learning with client-to-client collaboration," *IEEE Trans. Mobile Comput.*, 2025.
- [6] M. Abdel-Basset, N. Moustafa, and H. Hawash, "Privacy-preserved generative network for trustworthy anomaly detection in smart grids: A federated semisupervised approach," *IEEE Trans. Ind. Informat.*, vol. 19, no. 1, pp. 995-1005, Jan. 2023.
- [7] Z. Zhang, S. Ma, Z. Yang, Z. Xiong, J. Kang, Y. Wu, K. Zhang, and D. Niyato, "Robust semisupervised federated learning for images automatic recognition in internet of drones," *IEEE Internet Things J.*, vol. 10, no. 7, pp. 5733-5746, Apr. 2023.
- [8] X. Liang, Y. Lin, H. Fu, L. Zhu, and X. Li, "RSCFed: Random sampling consensus federated semi-supervised learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2022, pp. 1-10.
- [9] Y. Chen, "Advanced model-contrastive federated learning via knowledge distillation," in *Proc. IEEE 3rd Int. Conf. Control, Electron. Comput. Technol. (ICCECT)*, 2025.
- [10] D. Jiang, C. Shan, and Z. Zhang, "Federated learning algorithm based on knowledge distillation," in *Proc. Int. Conf. Artif. Intell. Comput. Eng. (ICAICE)*, 2020.
- [11] S. Caldas, J. Konecny, H. B. McMahan, and A. Talwalkar, "Expanding the reach of federated learning by reducing client resource requirements," *arXiv preprint arXiv:1812.07210*, 2018.
- [12] C. Xie, S. Koyejo, and I. Gupta, "Asynchronous federated optimization," *arXiv preprint arXiv:1903.03934*, 2019.

- [13] T. Nishio and R. Yonetani, "Client selection for federated learning with heterogeneous resources in mobile edge," in Proc. IEEE Int. Conf. Communications (ICC), Shanghai, China, 2019, pp. 1-7.
- [14] D. Yang, Z. Xu, W. Li, A. Myronenko, H. R. Roth, S. Harmon, S. Xu, B. Turkbey, E. Turkbey, X. Wang, et al., "Federated semi-supervised learning for COVID region segmentation in chest CT using multi-national data from China, Italy, Japan," *Med. Image Anal.*, vol. 70, p. 101992, 2021..
- [15] A. Y. Zhu, Y. Liu, J. J. Q. Yu, and X. Yuan, "Semi-supervised federated learning for travel mode identification from GPS trajectories," *IEEE Trans. Intell. Transport. Syst.*, vol. 23, no. 3, pp. 2380-2391, 2022.
- [16] Y. Zhao, H. Liu, H. Li, P. Barnaghi, and H. Haddadi, "Semi-supervised federated learning for activity recognition," *arXiv preprint arXiv:2011.00851*, 2020.
- [17] W. Jeong, J. Yoon, E. Yang, and S. J. Hwang, "Federated semi-supervised learning with inter-client consistency & disjoint learning," *arXiv preprint arXiv:2006.12097v3*, 2020.
- [18] Z. Zhang, Y. Yang, Z. Yao, Y. Yan, J. E. Gonzalez, K. Ramchandran, and M. W. Mahoney, "Improving semi-supervised federated learning by reducing the gradient diversity of models," in Proc. IEEE Int. Conf. Big Data (Big Data), Orlando, FL, USA, 2021, pp. 1214-1225.
- [19] Z. Zhang, S. Ma, J. Nie, Y. Wu, Q. Yan, X. Xu, and D. Niyato, "Semi-supervised federated learning with non-IID data: Algorithm and system design," in *Proc. IEEE 23rd Int. Conf. High Performance Computing & Communications; 7th Int. Conf. Data Science & Systems; 19th Int. Conf. Smart City; 7th Int.
- [20] Z. Zhang, S. Ma, Z. Yang, Z. Xiong, J. Kang, Y. Wu, K. Zhang, and D. Niyato, "Robust semi-supervised federated learning for images automatic recognition in Internet of drones," *IEEE Internet Things J.*, vol. 10, no. 7, pp. 5733-5746, 2023.
- [21] X. Liu, L. Zhu, S. T. Xia, Y. Jiang, and X. Yang, "GDST: Global distillation self-training for semi-supervised federated learning," in Proc. IEEE Global Communications Conf. (GLOBECOM), Madrid, Spain, 2021, pp. 1-6.