Privacy-Preserving Federated Learning Framework for Cardiovascular Disease Risk Prediction under Non-IID Data

Jiazhide Liu

Aberdeen School of Artificial Intelligence and Data Science, South China Normal University, Foshan, China

20233803015@m.scnu.edu.cn

Abstract. With the rapid growth of digital healthcare data and increasing concerns over data security and regulatory compliance, the need for privacy-preserving collaborative learning has become more urgent than ever. Nowadays, Cardiovascular Disease (CVD) has become the leading cause of global mortality, while traditional centralized medical model training is facing problems like severe data privacy barriers and data island. This work designs a privacy-compliant federated learning architecture to realize clinical heart disease risk prediction, FedAvg algorithm will be used to enable cross-institutional collaborative training without sharing raw patient data. At the same time, the research will adopt K-Means based non-IID to simulate real-world medical data heterogeneity. Experimental results show that the proposed framework achieves competitive performance compared with centralized training, the optimal test accuracy is 0.8704, exceeding the result of conventional centralized training. Therefore, this framework can provide a feasible privacy-preserving solution for cross-hospital clinical collaboration and offer a practical approach for future distributed medical risk prediction.

Keywords: Federated Learning (FL), Cardiovascular Disease (CVD), K-Means-based, non-IID.

1. Introduction

Nearly 32% of worldwide deaths were caused by Cardiovascular Disease (CVD) based on the World Health Organization [1]. Therefore, it is important and urgent to realize early and accurate diagnosis of heart disease to reduce mortality rates, as timely intervention can significantly improve patient outcomes. Meanwhile, there is a problem that traditional clinical diagnosis relies heavily on experienced physicians, and data-driven predictive models that adopted to assist clinical decision-making have shown great potential. However, most of the healthcare data involves sensitive personal information and is restricted by strict privacy regulations, which makes it extremely difficult to aggregate scattered patient data from multiple medical institutions for centralized model training. Strict privacy regulations such as HIPAA and GDPR are also prohibiting the direct sharing of patient health records, leading to a "data island" problem that hinders the development of high-performance predictive models.

At the same time, Federated Learning (FL) has been proposed to address this dilemma, enabling collaborative model training without exposing private data. Besides, extensive research has validated the feasibility and great potential of FL in healthcare. For instance: As highlighted by Rieke et al. [2], FL can realize cross-institutional collaboration under privacy regulations while complying with strict privacy regulations. And FL can achieve performance comparable to centralized training was proved by Sheller et al. [3], without requiring any patient data sharing. As for the cardiovascular disease prediction, through a multisite study, Chen et al. validating that FL-based models outperform local models trained on single-site data [4]. Besides, Li et al. further confirmed that FL can effectively preserve data privacy while maintaining high prediction accuracy [5]. For recent applications specifically, Zhang et al. applied FL to COVID-19 prognosis [6], showing that cross-institutional FL can significantly improve model generalization. In addition, Dou et al. demonstrated that FL can effectively address the data heterogeneity problem in medical datasets [7]. The core Federated Averaging (FedAvg) algorithm that this paper adopted was first proposed by McMahan et al. [8], which laid the foundation of modern federated learning frameworks. But most existing studies are

assuming IID data across clients, which is inconsisting with real-world heterogeneous patient distributions across hospitals.

In this paper, it will present a FedAvg-based federated framework for heart disease prediction, the dataset will be public Heart Disease UCI due to its applicability and authority. And then, a K-Means-based non-IID partitioning strategy will be adopted to simulate cross-hospital data heterogeneity. After that it will conduct comprehensive parameter sensitivity analysis, covering client number, network structure, and optimizer, with a centralized baseline will be established to verify the framework's performance. This study will illustrate the results that the proposed framework can reach a similar performance to centralized training while preserving data privacy, which is a practical privacy-preserving solution for collaborative medical prediction.

2. Method

The technical implementation of the privacy-protected heart disease prediction framework will be elucidated in this section, encompassing dataset preparation, design of the federated learning architecture, and details of experimental configuration.

2.1. Dataset Preparation

This study utilizes the publicly available UCI dataset of heart disease, which can be accessed through the OpenML platform (data ID=53) [9]. As a widely validated benchmark dataset in the field of clinical risk prediction, this dataset has been extensively employed in recent studies on federal heart disease prediction to verify the effectiveness of cross-institutional collaboration frameworks [9]. This dataset comprises a total of 303 clinical samples, with each sample possessing 13 diagnostic features, including age, gender, resting blood pressure, type of chest pain, resting electrocardiogram results, serum cholesterol, fasting blood glucose, exercise-induced ST segment depression, maximum heart rate, exercise-induced angina, exercise peak ST segment slope, number of major vessels displayed under thalassemia, and fluoroscopy. The target variable is a binary label which is used to show whether the heart disease is presence or absence, it contains the original multi-class diagnostic results into a standard binary classification task in order to facilitate clinical risk prediction. In terms of data preprocessing, a stratified training-testing set split with a ratio of 8:2 was first performed to ensure consistent distribution of positive and negative samples in both the training and testing sets, ultimately resulting in 242 training samples and 61 testing samples. Subsequently, to eliminate the impact of different feature scales on model training, standard normalization was applied to the feature space. Besides, aims at simulating the real-world heterogeneity of data across multiple medical institutions, this study adopted a non-independent and identically distributed (non-IID) splitting strategy based on K-Means. Previous research has already confirmed that due to differences in patient demographics and collection protocols, medical data from different institutions naturally exhibit non-IID characteristics [7]. Specifically, the training data was clustered into multiple groups by using the K-Means algorithm, where the K clusters are set to be consistent with the number of clients. To ensure the reproducibility of the partitioning process, the K-Means algorithm is configured with a fixed random state of 42 and 10 initializations, which effectively avoids the randomness of clustering results. Each cluster corresponds to a local client dataset. In this way the study can effectively simulate the non-IID data distribution that exists between different hospitals in real-world scenarios.

2.2. Federated Learning-based Prediction

Federated Learning (FL) is a distributed machine learning paradigm that allows numerous participants to collaboratively train a shared model without their private raw data exposure. This is a perfect way to address the data island problem in the medical field. As a comprehensive survey has summarized, FL has become the standard solution for privacy-preserving collaborative medical model training, which can effectively break the data barrier while complying with strict privacy regulations [10]. In this work, the FedAvg algorithm, the core framework of modern federated

learning proposed by McMahan et al. [8], will be implemented to construct the cross-institutional prediction system.

This is a framework employing a typical client-server architecture. The central server is responsible for initializing the global model and coordinating the training process, and each client still maintains its own local medical data. In each global communication round, the server will first broadcast the current global model parameters to all participating clients. Then, each client will initialize its local model using the received global parameters and conduct several rounds of local training by its private dataset. After local training is completed, the client will only upload the updated model parameters to the server, the original data will not be leaked. The server then aggregates these local parameters through a weighted approach based on the size of each client's dataset to update the global model. This process is iteratively repeated until the global model converges. As for the local prediction model, it will adopt a flexible Multi-layer Perceptron (MLP) structure, which equips with adjustable hidden layers, in order to map 13-dimensional clinical features to binary predictions of heart disease risk.

2.3. Experimental Hyperparameter Configuration

In this article, all experiments will be implemented based on the PyTorch deep learning framework. Existing research on federated cardiovascular risk assessment will be consistent with this[11]. To ensure the full reproducibility of the experiments, a fixed random seed of 42 will be chosen for all random processes. The random processes include data splitting, model initialization, and training shuffling. In addition, the initial learning rate will be set to 0.001 so that the model can converge stably. The SGD optimizer, it should be noted, additional hyperparameter tuning attempts will be conducted to explore its potential performance. For example, the experiment will test different learning rates of 0.01, 0.005, 0.001, and 0.0005. However, even after comprehensive hyperparameter tuning, the performance of SGD still failed to reach the level of adaptive optimizers in this non-IID federated scenario. Furthermore, three different optimizers, ranging from SGD, Adam, and RMSprop, were evaluated to study their impact on the performance of federated training.

In the federated training process, the global communication rounds will be set to 20 rounds, with each client performing 5 local training epochs per round and a fixed batch size of 8. This is based on the optimization considerations of efficiency and computing power. To verify the effectiveness of the proposed framework, a centralized training baseline with the same model structure and hyperparameters will also be established. This baseline uses all training data for direct model training, so that it has an ability to provide a reference for comparing the performance of FL.

3. Results and discussion

This section will present the experimental results of the proposed federated heart disease prediction framework and the analysis of the impact of different hyperparameters on model performance. Additionally, the discussion of the implications of the findings will also be shown.

Table 1. Summary of experimental results across different groups

Experiment Group	parameter	Optimizer	Final Accuracy	Final Loss
Centralized	-	-	0.8519	0.6304
Client Number	2	-	0.8333	0.8985
Client Number	4	-	0.8148	0.7318
Client Number	8	-	0.8704	0.4852
Client Number	16	-	0.8148	0.3761
Network Layers	1	-	0.8519	0.4168
Network Layers	2	-	0.8519	0.4819
Network Layers	3	-	0.8333	0.6797
Optimizer	-	SGD	0.3519	0.7015
Optimizer	-	Adam	0.8519	0.4788
Optimizer	-	RMSprop	0.8333	0.9468

According to Table 1, the proposed federated framework achieves competitive performance compared with the centralized baseline. Notably, when the client number is set to 8, the model reaches the highest test accuracy of 0.8704, which even outperforms the centralized training baseline (0.8519). There are some previous studies proving that FL can leverage the distributed heterogeneous data to improve the generalization ability of the model, rather than suffering from performance degradation, which is very consistent with the finding in this study [12].

Regarding the impact of client numbers, it can be said from Table 1 that the model performance increases at first and then decreases when the number of clients grows. As it is illustrated, the optimal performance is finally achieved with 8 clients, the reason why 8 clients are chosen is that too few clients (2 or 4) will lead to insufficient data diversity, while too many clients (16) will introduce excessive data heterogeneity which will hinder the convergence of the global model. It aligns with the findings of previous hyperparameter sensitivity studies, which pointed out that moderate data heterogeneity across institutions can enhance the model's generalization. Meanwhile excessive heterogeneity will lead to performance degradation in federated training [13].

When it comes to the optimizer selection, the results indicate that the Adam optimizer achieves the best performance, with a final accuracy of 0.8519. At the same time, the SGD optimizer only achieves 0.3519, which is far below the acceptable level. This is due to the adaptive learning rate of Adam can better adapt to the non-IID data distribution in FL. On the contrary, the fixed learning rate of SGD fails to handle the heterogeneous gradient updates across different clients [14]. Besides, the RMSprop optimizer also achieves a reasonable performance of 0.8333, which is slightly lower than Adam.

In terms of network structure, the model with 1 or 2 hidden layers had achieved a similar performance. However, the 3-layer network leads to a drop in accuracy. This suggests that for the small-scale clinical dataset, a shallow network is more sufficient to capture the feature patterns, so a deeper network will cause overfitting, which is consistent with the observation of previous medical federated learning studies [15].

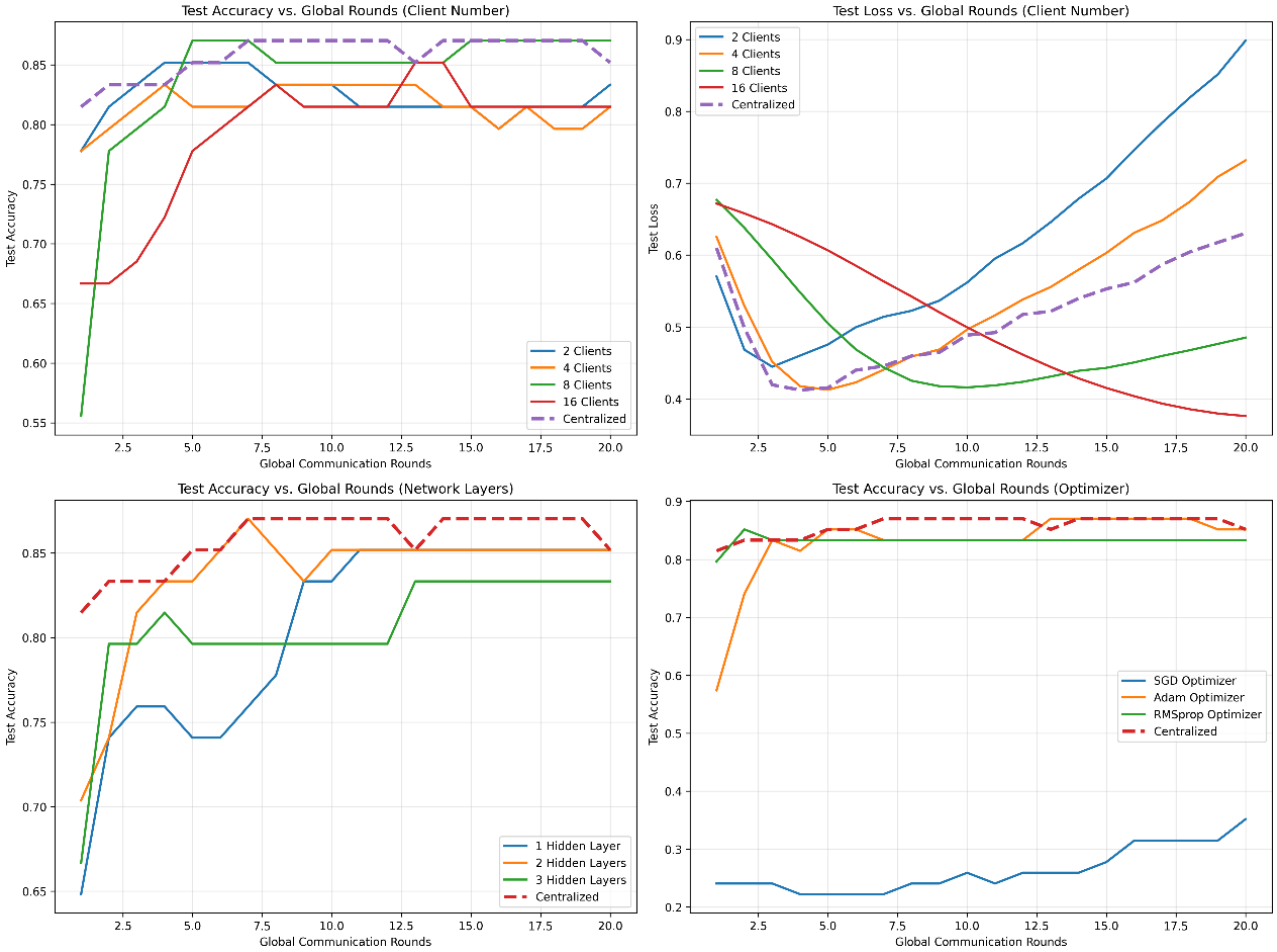

Fig. 1 Convergence curves of test accuracy and loss across different experimental groups (Picture credit: Original)

As illustrated in Fig. 1, the convergence curves are further confirming these findings. The 8-client group converges the fastest and achieves the lowest final loss, meanwhile, other client groups show slower convergence or higher fluctuation. Moreover, the SGD group fails to converge effectively during the 20 communication rounds, which further verifies the inadaptability of SGD in this scenario.

Overall, the results indicate that the federated framework envisions can effectively achieve privacy-preserving prediction for heart disease. With a proper hyperparameter configuration, it can even outperform the centralized baseline, and finally provide a feasible solution for cross-institutional clinical collaboration.

4. Conclusion

This study addressed data privacy and data island problems in medical machine learning. A federated learning framework for heart disease prediction is built, enabling collaborative training without raw data sharing. The experiment also verified that the framework could achieve comparable performance to centralized training, with the highest optimal accuracy of 0.8704. In the future, more work will focus on optimizing the framework so that more complex heterogeneous data can be handled and integrating differential privacy to further enhance privacy protection.

References

[1] Roth, G. A., Mensah, G. A., Johnson, C. O., et al. Global Burden of Cardiovascular Diseases and Risk Factors, 1990-2019: Update From the GBD 2019 Study. Journal of the American College of Cardiology, 2020, 76(25): 2982-3021.

[2] Rieke, N., Hancox, J., Li, W., et al. The future of digital health with federated learning. npj Digital Medicine, 2020, 3: 119.

[3] Sheller, M. J., Edwards, B., Reina, G. A., et al. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. Scientific Reports, 2020, 10(1): 12598.

[4] Chen, Y., Shi, Y., Zhang, L., et al. Federated learning for cardiovascular disease risk prediction: A multisite study. Journal of the American Medical Informatics Association, 2023, 30(5): 879-888.

[5] Li, Y., Jiang, X., Wang, H., et al. Federated learning for privacy-preserving healthcare predictive modeling. IEEE Journal of Biomedical and Health Informatics, 2020, 24(11): 3219-3229.

[6] Zhang, L., Wang, B., Li, X., et al. Federated learning for COVID-19 prognosis from chest CT scans across multiple institutions. Nature Communications, 2021, 12(1): 6015.

[7] Dou, Q., So, C. W., Jiang, M., et al. Privacy-preserving federated learning for medical imaging. Nature Machine Intelligence, 2022, 4(3): 239-251.

[8] McMahan, H. B., Moore, E., Ramage, D., et al. Communication-efficient learning of deep networks from decentralized data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, 2017: 1273-1282.

[9] Dua, D. and Graff, C. UCI Machine Learning Repository: Heart Disease Dataset. Irvine, CA: University of California, School of Information and Computer Sciences, 2019. [Online]. Available: https://archive.ics.uci.edu/ml/datasets/heart+disease

[10] Linardos, A., Kushibar, K., Walsh, S., et al. Federated learning for multi-center imaging diagnostics: a simulation study in cardiovascular disease. Scientific Reports, 2022, 12(1): 4743.

[11] Dubey, M., Tembhurne, J. V., Makhijani, R. Enhancing federated learning through differential privacy: Introducing FedHybrid for multicenter diverse heart disease datasets. IEEE Transactions on Emerging Topics in Computational Intelligence, 2025.

[12] Xu, J., Glicksberg, B. S., Su, C., et al. Federated learning for healthcare informatics. Journal of Healthcare Informatics Research, 2021, 5(1): 1-19.

[13] Alom, M. S., Akhi, S. S., Borsha, S. N., et al. Federated machine learning for cardiovascular risk assessment: A decentralized XGBoost approach. 2025 International Conference on Quantum Photonics, Artificial Intelligence, and Networking (QPAIN), 2025.

[14] Zhou, Y., Ram, P., Salonidis, T., et al. Single-shot Hyper-parameter Optimization for Federated Learning: A General Algorithm & Analysis. arXiv preprint arXiv:2202.08338, 2022.

[15] Wang, Z., Kuang, W., Zhang, C., et al. FEDHPO-BENCH: A Benchmark Suite for Federated Hyperparameter Optimization. Proceedings of the 40th International Conference on Machine Learning, 2023: 23451-23468.