Federated Learning Platform for Privacy-Preserving Medical Predictions

Saranya Devi. T. V[1], Ms. E. Uva Shakthi[2]
[1]*M. E student, Jaya Engineering college*
[2]*Assistant Professor, Jaya Engineering college*

***Abstract*—Privacy concerns and healthcare regulations often prevent hospitals from sharing patient data for developing accurate machine learning models. Federated Learning (FL) provides a solution by enabling multiple healthcare institutions to collaboratively train a global model without exchanging sensitive patient information. This project presents a Federated Learning platform for privacy-preserving medical prediction, where each hospital trains a local model using its own data and shares only the model parameters with a central server. The server aggregates these updates using the Federated Averaging (FedAvg) algorithm to generate an improved global model. To further enhance privacy, Differential Privacy is applied to protect the transmitted model updates from potential information leakage. The proposed system supports secure collaboration among healthcare organizations while ensuring compliance with data privacy regulations. Experimental results demonstrate that the federated model achieves improved prediction performance compared to isolated local models, without compromising patient confidentiality. This approach enables the development of reliable medical prediction systems and promotes secure, collaborative healthcare research.**

***Index Terms*—Federated Learning, Medical Prediction, Differential Privacy, Privacy Preservation, Healthcare, Federated Averaging (FedAvg).**

I. INTRODUCTION

The quick development of Artificial Intelligence (AI) and Machine Learning (ML) has greatly changed the healthcare industry by allowing more accurate disease prediction, better medical diagnosis, and tailored treatment plans. These technologies depend on a lot of patient data to create dependable predictive models. However, healthcare data is spread out across many hospitals and medical centers, making it hard to collect all this data in one place. This is because of privacy laws, legal rules, and ethical issues. Sharing personal health information might break these regulations and make data breaches more likely.

Federated Learning (FL) has become a useful way for machine learning that lets different healthcare organizations work together to train a common model without sharing actual patient data. Instead of sending personal medical records, each organization uses its own data to train a local model and sends only the model's parameters to a central server. The server then combines these parameters using the Federated Averaging (FedAvg) method and sends the updated global model back to all the participating organizations. To protect privacy even more, Differential Privacy is used by adding carefully managed noise to the model updates before they are sent, which helps prevent any information from being leaked. The proposed Federated Learning platform offers a safe and privacy-focused system for predicting healthcare insurance outcomes.

It allows for team-based model training while keeping patient information confidential, improves the accuracy of predictions, and follows healthcare privacy laws. This method shows how distributed AI can support trustworthy medical decisions without exposing sensitive health data.

II. RELATED WORK

Recent studies have shown that Federated Learning (FL) is effective for enabling collaborative machine learning while protecting data privacy. Mc Mahan and others introduced the Federated Averaging (FedAvg) algorithm, which allows multiple clients to train a common global model without sharing raw data. Bonawitz and others suggested secure aggregation methods to safeguard client updates during the model

aggregation process. To further improve privacy, Differential Privacy has been widely used by adding controlled random noise to model updates, which reduces the chance of information leaks. Researchers have also looked into Homomorphic Encryption and Secure Multi-Party Computation to boost the security of federated systems. In the healthcare field, Sheller and others, as well as Rieke and others, showed that Federated Learning can achieve prediction results that are similar to those from centralized methods, all while keeping patient information private. These findings indicate that FL has great potential for medical diagnosis and healthcare analysis without the need to share sensitive patient data.

III. PROPOSED SYSTEM

The proposed system introduces a Federated Learning platform for privacy-preserving medical prediction across multiple healthcare institutions. Instead of sharing sensitive patient data, each hospital trains a local machine learning model using its own dataset. The trained model parameters are securely transmitted to a central server, where the Federated Averaging (FedAvg) algorithm combines them to generate an improved global model. Differential Privacy is applied to the model updates before transmission to protect sensitive information from potential attacks. The updated global model is then distributed back to all participating hospitals for the next training round. This collaborative learning approach improves prediction accuracy while ensuring data privacy, security, and compliance with healthcare regulations.

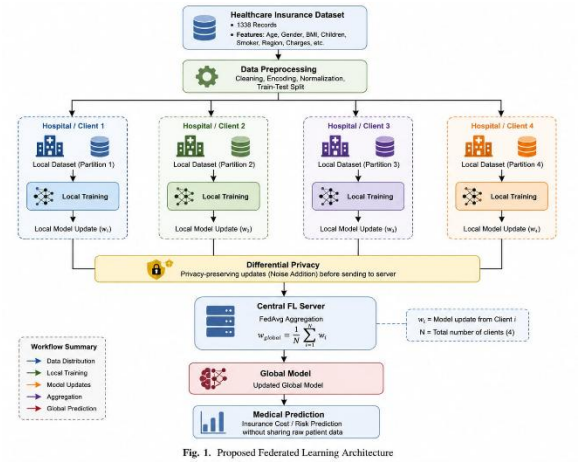

Fig. 1. Proposed Federated Learning Architecture

Fig. 1. System Architecture of the Proposed Federated Learning Platform

Figure 1 shows the overall structure of the proposed Federated Learning-based healthcare insurance prediction system. The system includes several healthcare institutions that each train their own machine learning models using their own patient data. Instead of sharing personal medical records, only the model parameters from each institution are sent to a central federated server. The server combines these parameters using the Federated Averaging (FedAvg) algorithm to create a better global model, which is then sent back to all the participating hospitals. This design allows for collaborative model training while keeping patient information private and following healthcare data security laws.

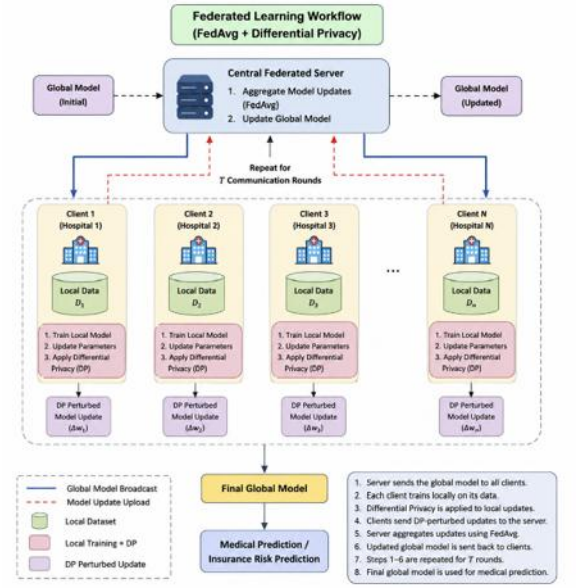

Fig. 2. Federated Learning Workflow with FedAvg and Differential Privacy

Figure 2 shows the workflow of the proposed Federated Learning framework designed for privacy-protected medical prediction. The process starts with the central federated server sending the global model to all participating healthcare institutions. Each hospital then uses its own patient data to train the model locally, without sharing personal medical information. Before sending the updated model parameters back to the server, Differential Privacy is applied to prevent any possible data leaks. The central server aggregates the updates received from all participating models through the Federated Averaging (FedAvg) technique, generating a more robust global model. This refined model is then utilized to provide reliable and privacy-preserving medical predictions.

IV. METHODOLOGY

The system proposed uses a Federated Learning framework to predict healthcare insurance outcomes securely. Any missing data in the dataset is addressed, categorical variables are converted into numerical codes using suitable methods, and numerical data is scaled to improve the efficiency of the model training process. The cleaned dataset is divided into four separate groups, each representing a different healthcare organization. Each organization trains a machine learning model using only its own data, making sure that personal patient details are never shared outside the local system. After completing each round of training, only the updated model parameters are sent to a central server. Before sending these updates, a Differential Privacy technique is applied to introduce carefully controlled random noise, which protects sensitive information and lowers the chance of recovering private patient details.

The central server combines all the updates from the different organizations using the Federated Averaging (FedAvg) method to create a better overall model. Once this combination is done, the improved global model is shared back with all the participating organizations. Each organization then uses this updated model as the base for the next round of local training. This process continues for ten rounds of communication until the model reaches a stable state. Finally, the trained global model is tested using the Mean Squared Error (MSE) to assess how well it makes predictions, all while keeping patient data confidential throughout the entire training process.

Table 1. Healthcare Insurance Dataset Description

Attribute	Data Type	Description
Age	Integer	Age of the insured individual in years.
Sex	Categorical	Gender of the individual (Male/Female).
BMI	Float	Body Mass Index (BMI) representing body fat based on height and weight.
Children	Integer	Number of dependent children covered under the insurance plan.
Smoker	Categorical	Smoking status of the individual (Yes/No).
Region	Categorical	Residential region (Northeast, Northwest, Southeast, Southwest).
Charges	Float	Annual medical insurance charges (Target variable).

Table 1 outlines the characteristics included in the healthcare insurance dataset. These features include demographic and lifestyle information, which are processed before being used in federated learning. The goal is to predict the medical insurance charges.

V. EXPERIMENTAL RESULTS

To assess its performance, the proposed federated learning system was applied to a healthcare insurance dataset comprising records of 1,338 patients. Before the training process began, the data was prepared by converting categorical variables into numerical codes, scaling numerical values, and correcting any inconsistencies. The dataset was split into 1,070 samples for training and 268 samples for testing. The training data was evenly distributed among four simulated healthcare organizations, and each organization trained its own local model independently. The Flower framework was employed to set up the federated learning environment. Model updates from each organization were combined using the Federated Averaging (FedAvg) method over ten rounds of communication to create a strong global model. During the training process, all patient data remained on the respective organizations' systems, showcasing how federated learning preserves privacy. The final global model was assessed using the Mean Squared Error (MSE) metric, resulting in a test MSE of 323,418,400. 0000. The results show that working together through federated learning allows for accurate predictions while ensuring patient privacy and lowering the risks linked to sharing data in a central location. These findings highlight the potential of federated learning for securely predicting healthcare insurance costs in dispersed medical settings.

Table 2: Experimental Results

Parameter	Value
Dataset Size	1338 records
Training Samples	1070
Testing Samples	268
Number of Clients	4
Communication Rounds	10
Aggregation Algorithm	FedAvg
Evaluation Metric	Mean Squared Error (MSE)
Final Test MSE Loss	323418400. 0000

Table 2 presents the experimental results of the proposed Federated Learning-based healthcare insurance prediction system. The healthcare insurance

dataset consists of 1,338 records, which were divided into 1,070 training samples and 268 testing samples. The training process was carried out using four client nodes in a simulated federated environment, and the local model updates were aggregated using the Federated Averaging (FedAvg) algorithm over 10 communication rounds. The final global model was evaluated using the Mean Squared Error (MSE) metric, achieving a test MSE loss of 323,418,400. 0000. These results demonstrate that the proposed framework enables collaborative model training while preserving patient privacy, making it suitable for secure healthcare insurance prediction applications.

VI. CONCLUSION

This work presented a federated learning framework for healthcare insurance prediction that enables multiple healthcare institutions to collaboratively train a machine learning model without exchanging patient records. By combining the Federated Averaging (FedAvg) algorithm with Differential Privacy, the framework improves data protection while allowing effective model training across distributed datasets. The experimental results indicate that the proposed approach achieves reliable prediction performance while maintaining patient privacy and supporting healthcare data security requirements. Although the current implementation produced encouraging results, there is still scope for improving both accuracy and system efficiency in future developments.

VII. FUTURE WORK

Several enhancements can be considered to improve the proposed system. One possible direction is the use of more advanced aggregation techniques that can reduce the influence of unreliable or malicious client updates during federated training. Privacy protection can also be strengthened by integrating mechanisms such as Secure Multi-Party Computation (SMPC) and Homomorphic Encryption during model aggregation. In addition, evaluating the framework on larger and more diverse datasets collected from multiple healthcare organizations would provide a better assessment of its scalability and predictive capability. Future work may also investigate deep learning models and deploy the framework in cloud or edge computing environments to support real-time healthcare applications.

REFERENCES

[1] McMahan, H. B., et al., “Communication-Efficient Learning of Deep Networks from Decentralized Data,” in Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), 2017. Available: https://proceedings.mlr.press/v54/mcmahan17a.html

[2] Konečný, J., et al., “Federated Learning: Strategies for Improving Communication Efficiency,” 2016. Available: https://arxiv.org/abs/1610.05492

[3] Bonawitz, K., et al., “Practical Secure Aggregation for Privacy-Preserving Machine Learning,” in Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS), 2017. doi: 10.1145/3133956.3133982.

[4] Dwork, C., and Roth, A., The Algorithmic Foundations of Differential Privacy. 2014. doi: 10.1561/0400000042.

[5] Shokri, R., and Shmatikov, V., “Privacy-Preserving Deep Learning,” in Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS), 2015, pp. 1310–1321. doi: 10.1145/2810103.2813687.

[6] Acar, A., et al., “A Survey on Homomorphic Encryption,” ACM Computing Surveys, 2018. doi: 10.1145/3214303.

[7] Mohassel, P., and Zhang, Y., “SecureML: A System for Scalable Privacy-Preserving Machine Learning,” in IEEE Symposium on Security and Privacy, 2017. doi: 10.1109/SP.2017.12.

[8] Sheller, M. J., et al., “Federated Learning in Medicine: Multi-Institutional Collaboration,” Scientific Reports, vol. 10, 2020. doi: 10.1038/s41598-020-69250-1.

[9] Rieke, N., et al., “The Future of Digital Health with Federated Learning,” NPJ Digital Medicine, vol. 3, 2020. doi: 10.1038/s41746-020-00323-1.

[10] Xu, J., et al., “Federated Learning for Healthcare Informatics,” Journal of Healthcare Informatics

Research, 2021. doi: 10.1007/s41666-020-00082-4.

[11] Li, T., et al., "Federated Learning: Challenges, Methods, and Future Directions," IEEE Signal Processing Magazine, vol. 37, no. 3, pp. 50–60, 2020. doi: 10.1109/MSP.2020.2975749.

[12] Kaissis, M., et al., "Secure, Privacy-Preserving and Federated Machine Learning in Medical Imaging," Nature Machine Intelligence, vol. 2, pp. 305–311, 2020. doi: 10.1038/s42256-019-0137-9.

[13] Yang, Q., et al., "Federated Machine Learning: Concept and Applications," ACM Transactions on Intelligent Systems and Technology, vol. 10, no. 2, 2019. doi: 10.1145/3298981.

[14] Kairouz, S., et al., "Advances and Open Problems in Federated Learning," Foundations and Trends in Machine Learning, vol. 14, nos. 1–2, pp. 1–210, 2021. doi: 10.1561/2200000083.

[15] Zhao, Y., et al., "Federated Learning with Non-IID Data," 2018. Available: https://arxiv.org/abs/1806.00582

[16] Zhu, L., et al., "Deep Leakage from Gradients," 2019. Available: https://arxiv.org/abs/1906.08935

[17] Smith, V., et al., "Federated Multi-Task Learning," 2017. Available: https://arxiv.org/abs/1705.10467

[18] Yang, Q., et al., "Federated Transfer Learning," IEEE Intelligent Systems, vol. 34, no. 4, pp. 70–74, 2019. doi: 10.1109/MIS.2019.2908142.

[19] Lu, Y., et al., "Blockchain and Federated Learning for Secure Data Sharing," IEEE Transactions on Industrial Informatics, vol. 16, no. 6, pp. 4177–4186, 2020. doi: 10.1109/TII.2019.2941905.

[20] Nguyen, D., et al., "Federated Learning for IoT: A Comprehensive Survey," IEEE Communications Surveys & Tutorials, vol. 23, no. 3, pp. 1622–1658, 2021. doi: 10.1109/COMST.2021.3053824.

[21] Li, A., et al., "Privacy-Preserving Federated Brain Tumor Segmentation," 2019. Available: https://arxiv.org/abs/1910.00962

[22] Yang, B., et al., "Federated Learning Systems: Vision and Challenges," IEEE Internet of Things Journal, vol. 6, no. 5, pp. 8763–8776, 2019. doi: 10.1109/JIOT.2019.2943119.

[23] Abadi, M., et al., "Deep Learning with Differential Privacy," in Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS), 2016. doi: 10.1145/2976749.2978318.

[24] Geyer, K., et al., "Differentially Private Federated Learning," 2017. Available: https://arxiv.org/abs/1712.07557

[25] Truex, S., et al., "A Hybrid Approach to Privacy-Preserving Federated Learning," 2019. Available: https://arxiv.org/abs/1812.03224

[26] Hard, R., et al., "Federated Learning for Mobile Keyboard Prediction," 2018. Available: https://arxiv.org/abs/1811.03604

[27] Chen, M., et al., "Machine Learning for Healthcare," IEEE Access, vol. 5, pp. 26825–26839, 2017. doi: 10.1109/ACCESS.2017.2787665.

[28] Liu, X., et al., "Deep Learning in Medical Ultrasound Analysis," Engineering, vol. 5, no. 2, pp. 261–275, 2019. doi: 10.1016/j.eng.2019.01.019.

[29] Chen, T., and Guestrin, C., "XGBoost: A Scalable Tree Boosting System," in Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016. doi: 10.1145/2939672.2939785.

[30] Goodfellow, I., Bengio, Y., and Courville, A., Deep Learning. Cambridge, MA, USA: MIT Press, 2016. Available: https://www.deeplearningbook.org

[31] Krizhevsky, A., Sutskever, I., and Hinton, G. E., "ImageNet Classification with Deep Convolutional Neural Networks," in Advances in Neural Information Processing Systems, 2012.

[32] He, K., Zhang, X., Ren, S., and Sun, J., "Deep Residual Learning for Image Recognition," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016. doi: 10.1109/CVPR.2016.90.

[33] Brown, T., et al., "Language Models are Few-Shot Learners," in Advances in Neural Information Processing Systems (NeurIPS), 2020. Available: https://arxiv.org/abs/2005.14165

[34] Hochreiter, S., and Schmidhuber, J., "Long Short-Term Memory," Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997. doi: 10.1162/neco.1997.9.8.1735.

[35] LeCun, Y., Bengio, Y., and Hinton, G., "Deep Learning," Nature, vol. 521, pp. 436–444, 2015. doi: 10.1038/nature14539.

[36] Recht, B., et al., "Hogwild: A Lock-Free Approach to Parallelizing Stochastic Gradient Descent," in Advances in Neural Information Processing Systems, 2011. Available: https://papers.nips.cc/paper/2011

[37] Zaharia, M., et al., "Apache Spark: A Unified Engine for Big Data Processing," Communications of the ACM, vol. 59, no. 11, pp. 56–65, 2016. doi: 10.1145/2934664.

[38] Nia, N. G., Kaplanoglu, E., and Nasab, A., "Evaluation of Artificial Intelligence Techniques in Disease Diagnosis and Prediction," Computers in Biology and Medicine, 2023.

[39] Kumar, Y., Koul, A., Singla, R., and Ijaz, M. F., "Artificial Intelligence in Disease Diagnosis: A Systematic Literature Review, Synthesizing Framework and Future Research Agenda," Journal of Ambient Intelligence and Humanized Computing, vol. 14, no. 7, pp. 8459–8486, 2023.

[40] Rimal, Y., et al., "Comparative Analysis of Heart Disease Prediction Using Logistic Regression, SVM, KNN, and Random Forest with Cross-Validation for Improved Accuracy," Scientific Reports, vol. 15, Art. no. 13444, 2025.

[41] Mohamed, A., Abdelrehim, M., and Al-Barazie, R., "Context Matters in Machine Learning Based Disease Prediction with Insights from Diverse Clinical and Symptom Data," Scientific Reports, vol. 15, Art. no. 42669, 2025.