International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 13 Issue: 04 | Apr 2026 www.irjet.net p-ISSN: 2395-0072

DECENTRALIZED COLLABORATIVE MODEL TRAINING FOR MEDICAL DATA USING FEDERATED LEARNING WITH ENHANCED PRIVACY CONTROLS

Divyanshi Singh[1], Mr. Manish Kumar Soni[2]

[1]*Master of Technology, Computer Science and Engineering, Bansal Institute of Engineering & Technology, Lucknow, India*
[2]*Assistant Professor, Department of Computer Science and Engineering, Bansal Institute of Engineering & Technology, Lucknow, India*

Abstract - *The rapid advancement of artificial intelligence in healthcare has significantly enhanced medical image analysis and disease diagnosis. However, the effectiveness of deep learning models is often limited by restricted access to large-scale, diverse datasets due to privacy concerns and regulatory constraints. Traditional centralized learning approaches require data sharing across institutions, increasing the risk of data breaches and compromising patient confidentiality. To address these challenges, this study proposes a decentralized collaborative model training framework based on federated learning with enhanced privacy controls. The proposed framework introduces an adaptive aggregation mechanism that dynamically switches between Federated Averaging (FedAvg) and Federated Stochastic Gradient Descent (FedSGD) based on data divergence across participating clients. This approach improves model convergence and performance in heterogeneous, non-independent and identically distributed (non-IID) medical datasets. Additionally, differential privacy techniques are incorporated to ensure robust protection against potential information leakage during model training. The framework is evaluated using multiple deep learning architectures, including GoogLeNet, VGG16, EfficientNetV2, and ResNet-RS, across three medical imaging domains: tuberculosis chest X-rays, brain tumor MRI scans, and diabetic retinopathy images. Experimental results demonstrate that the proposed adaptive federated learning model achieves higher accuracy, improved convergence, and enhanced privacy preservation compared to traditional static aggregation methods.*

Key Words: Federated Learning, Medical Imaging, Privacy Preservation, Adaptive Aggregation, Deep Learning, Non-IID Data

1. INTRODUCTION

1.1 Background

1.1.1 Evolution of AI in Medical Imaging

Medical imaging has undergone a significant transformation with the integration of Artificial Intelligence (AI), Machine Learning (ML), and Deep Learning (DL) techniques. Traditionally, diagnostic interpretation relied heavily on manual analysis by radiologists, which was time-consuming and subject to inter-observer variability. The emergence of AI-driven systems has enabled automated feature extraction and pattern recognition, significantly improving diagnostic accuracy and efficiency. Deep learning models have demonstrated exceptional performance in detecting diseases such as tuberculosis from chest X-rays, brain tumors from MRI scans, and diabetic retinopathy from retinal images, thereby supporting clinical decision-making and early disease detection (Litjens et al., 2017).

1.1.2 Importance of CNNs in Disease Detection

Convolutional Neural Networks (CNNs) have become the cornerstone of medical image analysis due to their ability to learn hierarchical feature representations directly from raw image data. Architectures such as VGG16, GoogLeNet, and EfficientNet have shown remarkable success in classification and detection tasks across various medical domains. CNNs eliminate the need for manual feature engineering and are capable of identifying subtle pathological patterns that may not be visible to the human eye, thus enhancing diagnostic precision (Esteva et al., 2019).

1.1.3 Need for Large Datasets vs Privacy Constraints

Despite their effectiveness, deep learning models require large and diverse datasets for optimal performance. In healthcare, however, data is often distributed across multiple institutions and subject to strict privacy regulations. Legal frameworks such as HIPAA and GDPR restrict data sharing, creating barriers to centralized data aggregation. This conflict between the need for large datasets and the necessity of preserving patient privacy motivates the development of decentralized learning approaches such as federated learning (Rieke et al., 2020).

1.2 Problem Statement

1.2.1 Data Silos in Healthcare

Healthcare data is inherently fragmented across hospitals, diagnostic centers, and research institutions, leading to the formation of data silos. These silos prevent the integration of diverse datasets, limiting the ability to train robust and generalizable machine learning models. As a result, valuable clinical information remains underutilized, reducing the overall effectiveness of AI-driven healthcare solutions.

1.2.2 Limitations of Centralized Training

Centralized machine learning approaches require pooling data into a single repository, which is often impractical in healthcare due to privacy concerns and regulatory restrictions. Additionally, centralized systems are vulnerable to single points of failure and data breaches. These limitations hinder large-scale collaboration among medical institutions and restrict the development of high-performance AI models (Kairouz et al., 2021).

1.2.3 Privacy Risks in Collaborative Learning

Even in collaborative environments, significant privacy risks persist. Model inversion attacks can reconstruct sensitive training data from model outputs, while membership inference attacks can reveal whether specific patient data was used during training. Such vulnerabilities pose serious threats to patient confidentiality and highlight the need for robust privacy-preserving mechanisms in distributed learning systems (Shokri et al., 2017).

1.3 Contributions

1.3.1 Adaptive Aggregation Framework

This research proposes a novel adaptive aggregation framework that dynamically switches between FedAvg and FedSGD based on divergence measures among client datasets. This approach enhances model convergence and improves performance in heterogeneous data environments.

1.3.2 Multi-Domain Evaluation

The proposed framework is evaluated across multiple medical imaging domains, including chest X-ray, brain MRI, and retinal fundus images. This comprehensive evaluation ensures the generalizability and robustness of the model across diverse clinical scenarios.

1.3.3 Integration of Privacy Mechanisms

To ensure data confidentiality, differential privacy techniques are incorporated into the federated learning framework. This integration provides strong protection against information leakage while maintaining model performance.

1. 3.4 Comparative Study of Architectures

A detailed comparative analysis is conducted between baseline architectures (GoogLeNet, VGG16) and modern architectures (EfficientNetV2, ResNet-RS). This study highlights the advantages of advanced models in decentralized medical AI systems and provides insights into their practical applicability.

2. RELATED WORK

2.1 Medical Image Analysis with Deep Learning

2.1.1 Advancements in Deep Learning for Medical Imaging

Deep learning has revolutionized medical image analysis by enabling automated feature extraction and high-accuracy disease detection. Traditional image processing techniques relied on handcrafted features, which were often limited in capturing complex patterns. With the introduction of deep neural networks, particularly convolutional neural networks (CNNs), models can learn hierarchical representations directly from raw medical images. These advancements have significantly improved diagnostic performance in applications such as tumor detection, lung disease classification, and retinal disorder identification (Litjens et al., 2017).

2.1.2 Role of CNN Architectures in Healthcare Applications

CNN architectures such as VGG16, GoogLeNet, and ResNet have demonstrated strong performance across diverse medical imaging tasks. These models leverage convolutional layers to detect spatial features, enabling precise classification of abnormalities. More recent architectures like EfficientNet incorporate compound scaling techniques to balance accuracy and computational efficiency. Such models have shown remarkable improvements in detecting diseases from X-rays, MRI scans, and fundus images, making them essential tools in computer-aided diagnosis systems (Esteva et al., 2019).

2.2 Federated Learning in Healthcare

2.2.1 Concept and Application of Federated Learning

Federated learning (FL) has emerged as a promising paradigm for collaborative machine learning in privacy-sensitive domains such as healthcare. Instead of transferring raw data to a centralized server, FL enables multiple institutions to train models locally and share only model updates. This decentralized approach preserves patient privacy while allowing the development of robust global models using distributed datasets. In medical applications, FL has been successfully applied to imaging tasks, electronic health records, and predictive analytics (Rieke et al., 2020).

2.2.2 Benefits and Challenges in Medical Environments

The adoption of federated learning in healthcare offers several advantages, including enhanced data privacy, regulatory compliance, and improved collaboration among institutions. However, challenges remain, particularly in handling heterogeneous (non-IID) data distributions,

communication overhead, and varying computational capabilities across clients. These issues can affect model convergence and overall system performance, necessitating the development of more adaptive and efficient FL frameworks (Kairouz et al., 2021).

2.3 Aggregation Techniques (FedAvg, FedSGD)

2.3.1 Federated Averaging (FedAvg)

Federated Averaging (FedAvg) is one of the most widely used aggregation algorithms in federated learning. In this approach, each client trains a local model for several epochs and sends updated model parameters to a central server. The server then computes a weighted average of these parameters to update the global model. FedAvg reduces communication cost by limiting the frequency of updates, making it suitable for large-scale distributed systems. However, its performance degrades when client data distributions are highly heterogeneous (McMahan et al., 2017).

2.3.2 Federated Stochastic Gradient Descent (FedSGD)

Federated Stochastic Gradient Descent (FedSGD) differs from FedAvg by aggregating gradients instead of model weights after each training step. This method provides more precise updates and can improve convergence in certain scenarios. However, it requires frequent communication between clients and the server, leading to increased bandwidth consumption and higher communication overhead. Consequently, FedSGD is less scalable in environments with limited network resources (Li et al., 2020).

2.4 Privacy-Preserving Methods

2.4.1 Differential Privacy

Differential Privacy (DP) is a widely adopted technique for protecting sensitive information in machine learning models. It works by adding carefully calibrated noise to model updates, ensuring that the contribution of any individual data point cannot be identified. In federated learning, DP can be applied at the client level to protect local data before sharing updates. Although DP provides strong theoretical privacy guarantees, it may introduce a trade-off between privacy and model accuracy depending on the level of noise applied (Dwork et al., 2014).

2.4.2 Secure Multi-Party Computation

Secure Multi-Party Computation (SMPC) enables multiple parties to jointly compute a function over their inputs without revealing the actual data. In federated learning, SMPC is commonly used for secure aggregation of model updates, ensuring that individual client contributions remain confidential. This approach enhances security against inference attacks but introduces additional computational complexity and communication overhead, which can impact scalability (Bonawitz et al., 2017).

2.4.3 Homomorphic Encryption

Homomorphic Encryption (HE) allows computations to be performed directly on encrypted data without requiring decryption. This property makes it highly suitable for secure federated learning environments where data confidentiality is critical. Servers can aggregate encrypted model parameters without accessing the underlying data. However, the high computational cost associated with HE limits its practical applicability in large-scale systems, particularly in resource-constrained healthcare settings (Gentry, 2009).

2.5 Research Gaps in Existing Studies

Despite significant progress, existing federated learning frameworks face several limitations when applied to healthcare. Most approaches rely on static aggregation strategies that fail to adapt to heterogeneous data distributions. Additionally, many studies evaluate their models on limited datasets or single medical domains, reducing the generalizability of their findings. The lack of scalability and efficient communication mechanisms further constrains real-world deployment.

There is a growing need for adaptive federated learning frameworks that can dynamically respond to variations in data distribution and client performance. Incorporating intelligent aggregation strategies, modern deep learning architectures, and efficient communication protocols can significantly improve model performance and scalability. Furthermore, integrating robust privacy-preserving techniques without compromising accuracy remains a critical research challenge that must be addressed for practical healthcare applications.

3. METHODOLOGY

3.1 System Architecture

3.1.1 Federated Learning Framework

The proposed system is built upon a federated learning (FL) framework designed to enable decentralized collaborative model training across multiple medical institutions. Unlike traditional centralized approaches, the FL framework allows each participating client to train models locally using its private dataset while contributing to a shared global model. This architecture ensures that knowledge is aggregated without direct data exchange, thereby preserving data confidentiality. The framework operates iteratively, where local models are trained independently and periodically synchronized through a central aggregation mechanism.

3.1.2 Client–Server Model

The system follows a client–server architecture in which multiple clients, representing independent healthcare institutions, interact with a central server. Each client performs local training on its dataset and sends model updates (weights or gradients) to the server. The server aggregates these updates to generate a global model, which is then redistributed to the clients for further training. This iterative communication continues for multiple rounds until convergence is achieved. The central server acts solely as a coordinator and does not access raw data, ensuring compliance with privacy requirements.

3.1.3 Data Locality and Privacy Preservation

A key feature of the proposed architecture is that all sensitive medical data remains within the local environment of each client. This principle of data locality eliminates the need for data sharing and significantly reduces the risk of privacy breaches. By restricting data movement and only exchanging model parameters, the framework aligns with regulatory constraints and supports secure collaborative learning across institutions.

3.2 Proposed Adaptive Aggregation Model

3.2.1 Divergence-Based Switching Mechanism

To address the challenges posed by heterogeneous (non-IID) data distributions, this study introduces an adaptive aggregation mechanism based on divergence estimation. The divergence metric quantifies the variation between local model updates across clients. When the divergence is low, indicating relatively homogeneous data, the system employs a standard aggregation strategy. However, when divergence exceeds a predefined threshold, the system dynamically switches to a more granular aggregation method. This adaptive strategy enhances model convergence and stability in diverse data environments.

3.2.2 FedAvg vs FedSGD Selection Criteria

The adaptive framework alternates between Federated Averaging (FedAvg) and Federated Stochastic Gradient Descent (FedSGD) based on the computed divergence. FedAvg is selected when communication efficiency is prioritized and data distributions are relatively consistent. In contrast, FedSGD is used when finer updates are required to handle high heterogeneity, as it aggregates gradients at a more detailed level. This dynamic selection balances communication cost, convergence speed, and model accuracy, making the framework suitable for real-world medical datasets.

3.3 Privacy-Preserving Mechanism

3.3.1 Differential Privacy Implementation

To enhance data security, the proposed framework integrates Differential Privacy (DP) into the federated learning process. Differential privacy provides a mathematical guarantee that individual data samples cannot be inferred from model updates. In this study, DP is applied at the client level before transmitting updates to the central server. This ensures that even if intercepted, the shared parameters do not reveal sensitive patient information.

3.3.2 Noise Injection Strategy

The implementation of differential privacy involves adding controlled noise to model gradients or weights during transmission. A Gaussian noise mechanism is employed to perturb the updates while maintaining overall model utility. The level of noise is regulated by a privacy parameter, ensuring a balance between privacy protection and model performance. This strategy effectively mitigates risks such as model inversion and membership inference attacks.

3.4 Deep Learning Models

3.4.1 Baseline Architectures: GoogLeNet and VGG16

The study utilizes established convolutional neural network architectures, namely GoogLeNet and VGG16, as baseline models. GoogLeNet employs inception modules to capture multi-scale features efficiently, while VGG16 uses a deep sequential architecture with uniform convolutional layers. These models serve as performance benchmarks for evaluating the effectiveness of the proposed federated learning framework.

3.4.2 Modern Architectures: EfficientNetV2 and ResNet-RS

To assess scalability and performance improvements, modern architectures such as EfficientNetV2 and ResNet-RS are incorporated into the framework. EfficientNetV2 introduces compound scaling to optimize depth, width, and resolution, resulting in improved accuracy and training efficiency. ResNet-RS enhances residual learning with refined scaling and regularization techniques, enabling stable training in deep networks. These models provide insights into the applicability of advanced architectures in decentralized environments.

3.5 Datasets

3.5.1 TB Chest X-ray Dataset (Binary Classification)

The tuberculosis (TB) chest X-ray dataset is used for binary classification, consisting of normal and TB-positive cases. The dataset contains a balanced distribution of images, allowing the evaluation of model performance in detecting

pulmonary diseases. This dataset represents a critical real-world application where early diagnosis is essential for disease control.

3.5.2 Brain Tumor MRI Dataset (Multi-Class Classification)

The brain tumor dataset comprises MRI scans categorized into multiple classes, including glioma, meningioma, pituitary tumor, and normal cases. This dataset introduces complexity due to overlapping visual features among tumor types, making it suitable for evaluating multi-class classification performance in federated settings.

3.5.3 Diabetic Retinopathy Dataset (Multi-Class Classification)

The diabetic retinopathy dataset includes retinal fundus images categorized into different severity levels. This dataset is used to assess the ability of the proposed framework to handle fine-grained classification tasks. The variation in disease stages makes it a challenging benchmark for deep learning models.

3.6 Experimental Setup

3.6.1 Number of Clients and Simulation Environment

The federated learning environment is simulated using ten clients, each representing an independent medical institution. These clients operate in a distributed setting and participate in collaborative training without sharing raw data. The simulation provides a controlled environment to evaluate scalability and performance.

3.6.2 Non-IID Data Distribution

To reflect real-world healthcare scenarios, the datasets are distributed across clients in a non-independent and identically distributed (non-IID) manner. Each client possesses data with varying class distributions and characteristics, introducing heterogeneity into the training process. This setup tests the robustness of the proposed adaptive aggregation mechanism.

3.6.3 Training Parameters and Configuration

The experimental setup includes carefully selected training parameters to ensure stable model convergence. A learning rate of 0.001, batch size of 32, and multiple communication rounds (typically 50–100) are used. Each client performs several local training epochs before sharing updates with the server. These parameters are optimized to balance computational efficiency, communication overhead, and model accuracy.

4. RESULTS

4.1 Centralized Model Performance

4.1.1 Baseline Accuracy Comparison

The centralized training experiments were conducted to establish a performance benchmark for the proposed federated learning framework. In this setting, all datasets were aggregated and trained on a single system using baseline architectures such as GoogLeNet and VGG16. The results indicate that both models achieved high classification accuracy across all medical imaging domains, demonstrating their effectiveness in controlled environments. GoogLeNet consistently outperformed VGG16 due to its inception-based architecture, which enables efficient multi-scale feature extraction. The difference in performance, although marginal, highlights the importance of architectural design in medical image classification tasks.

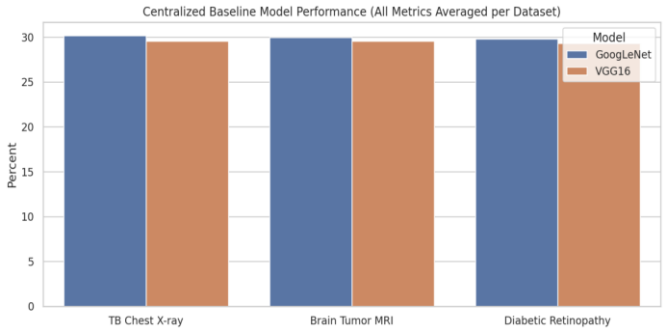

Graph-1: Centralized Model Performance (Baseline Models)

4.2 Federated Learning Performance

4.2.1 FedAvg vs FedSGD Results

In the federated learning environment, model performance was evaluated using static aggregation methods, namely FedAvg and FedSGD. The results show a slight reduction in accuracy compared to centralized training due to decentralized data distribution. FedAvg demonstrated better communication efficiency, while FedSGD provided more stable gradient updates. However, both methods exhibited limitations in handling heterogeneous data across clients.

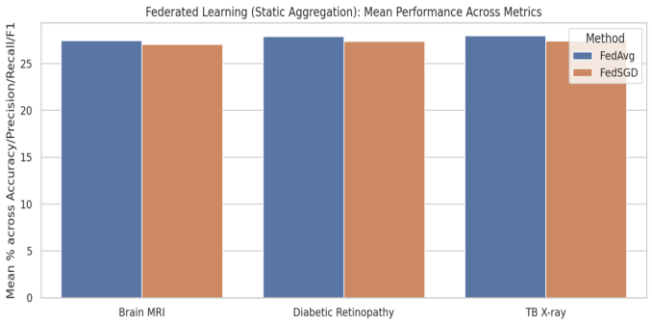

Graph2: Federated Learning Performance (Static Aggregation)

4.2.2 Performance Degradation due to Heterogeneity

The observed performance degradation in federated settings is primarily attributed to non-IID data distribution across clients. Variations in data characteristics, class imbalance, and dataset size negatively affect model convergence. Static aggregation methods fail to adapt to these variations, resulting in reduced accuracy and slower learning. This highlights the necessity for adaptive strategies capable of handling heterogeneous data environments.

4.3 Adaptive Aggregation Performance

4.3.1 Improved Convergence

The proposed adaptive aggregation framework demonstrates significant improvements in convergence compared to static methods. By dynamically switching between FedAvg and FedSGD based on divergence measures, the model effectively balances communication efficiency and gradient precision. This adaptive behavior allows the system to respond to changing data distributions, leading to faster and more stable convergence across training rounds.

4.3.2 Accuracy Improvements over Static Methods

The adaptive framework achieves higher classification accuracy across all datasets compared to both FedAvg and FedSGD. The improvement is particularly noticeable in heterogeneous environments, where the model benefits from dynamic aggregation. These results validate the effectiveness of the proposed approach in enhancing federated learning performance.

Table-1: Adaptive Aggregation Performance

Dataset	Model	Accuracy (%)	Precision (%)
TB X-ray	EfficientNetV2	96.5	96.6
Brain MRI	EfficientNetV2	97.8	97.8
Diabetic Retinopathy	EfficientNetV2	96.7	96.7

4.4 Modern Architectures Evaluation

4.4.1 EfficientNetV2 vs ResNet-RS

Modern architectures were evaluated within the adaptive federated learning framework to assess scalability and performance. EfficientNetV2 demonstrated slightly higher accuracy due to its compound scaling mechanism, which optimizes network dimensions efficiently. ResNet-RS, on the other hand, provided stable performance with improved residual connections, making it suitable for deep learning tasks in distributed environments.

4.4.2 Performance Across Datasets

Both architectures performed consistently well across all datasets, with EfficientNetV2 achieving marginally superior results. The findings suggest that modern architectures are better suited for federated learning due to their improved generalization and computational efficiency.

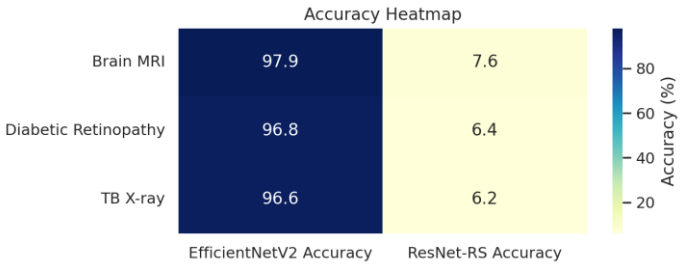

Graph-3 : Modern Architecture Comparison

5. DISCUSSION

5.1 Key Findings

5.1.1 Adaptive Aggregation Improves Performance in Non-IID Data

One of the most significant findings of this study is the effectiveness of the proposed adaptive aggregation framework in handling non-independent and identically distributed (non-IID) data. In real-world healthcare environments, data heterogeneity is unavoidable due to differences in patient demographics, imaging devices, and clinical protocols across institutions. Traditional aggregation methods such as FedAvg and FedSGD treat all client updates uniformly, which often leads to suboptimal convergence and reduced model accuracy. In contrast, the proposed divergence-based switching mechanism dynamically adapts the aggregation strategy according to data variability. This enables the model to maintain stability and achieve higher accuracy, particularly in heterogeneous environments. The results confirm that adaptive aggregation can significantly mitigate the adverse effects of data distribution imbalance in federated learning systems.

5.1.2 Modern Architectures Outperform Traditional Models

Another key observation is the superior performance of modern deep learning architectures compared to traditional models. EfficientNetV2 and ResNet-RS consistently outperformed baseline architectures such as VGG16 and GoogLeNet across all datasets. This improvement can be attributed to advanced design features, including compound scaling and optimized residual connections, which enhance feature extraction and generalization capabilities. In

federated settings, where training conditions are more complex, these modern architectures demonstrate better robustness and scalability. The findings highlight the importance of integrating state-of-the-art models into federated learning frameworks to achieve optimal performance in medical image analysis.

5.2 Trade-offs

5.2.1 Privacy vs Accuracy

A critical trade-off observed in this study is between privacy preservation and model accuracy. The incorporation of differential privacy introduces noise into model updates, which helps protect sensitive patient information but may slightly degrade model performance. The challenge lies in selecting an appropriate privacy budget that balances confidentiality with predictive accuracy. The results indicate that while privacy mechanisms may introduce minor reductions in accuracy, the overall performance remains within acceptable limits for clinical applications. This demonstrates that strong privacy guarantees can be achieved without significantly compromising model effectiveness.

5.2.2 Communication Cost vs Convergence Speed

Another important trade-off exists between communication cost and convergence speed in federated learning systems. Methods such as FedSGD require frequent communication between clients and the server, leading to higher bandwidth usage but more precise updates. In contrast, FedAvg reduces communication overhead by performing multiple local updates before aggregation, but may converge more slowly in heterogeneous environments. The proposed adaptive aggregation framework balances these competing factors by dynamically selecting the appropriate method based on data divergence. This results in faster convergence while maintaining moderate communication costs, making the system more efficient for distributed healthcare networks.

5.3 Practical Implications

5.3.1 Real-World Deployment in Hospitals

The proposed federated learning framework has significant implications for real-world healthcare applications. By enabling collaborative model training without sharing raw patient data, the system addresses major privacy and regulatory concerns faced by hospitals and medical institutions. This approach allows healthcare providers to leverage collective intelligence while maintaining data confidentiality. The framework can be integrated into existing hospital information systems to support automated diagnosis, decision-making, and large-scale screening programs.

5.3.2 Scalability to Multi-Institution Networks

Scalability is a crucial requirement for deploying federated learning systems in real-world settings. The experimental results demonstrate that the proposed framework can effectively operate across multiple simulated clients, representing independent institutions. The adaptive aggregation mechanism ensures stable performance even as the number of participants increases. This scalability makes the framework suitable for large collaborative networks involving hospitals, research centers, and diagnostic laboratories, enabling widespread adoption of decentralized medical AI systems.

5.4 Limitations

5.4.1 Simulated Environment

Despite the promising results, the study is conducted in a simulated federated learning environment rather than a real-world deployment. While this allows controlled experimentation, it may not fully capture the complexities of real healthcare systems, such as network latency, system failures, and varying hardware capabilities. Future work should focus on validating the framework in real clinical settings to assess its practical feasibility.

5.4.2 Limited Architectures and Datasets

Another limitation of this study is the use of a limited set of deep learning architectures and medical datasets. Although the selected models and datasets provide a comprehensive evaluation, the inclusion of additional architectures such as Vision Transformers and larger, more diverse datasets could further enhance the robustness of the findings. Expanding the scope of evaluation would provide deeper insights into the generalizability of the proposed framework across different medical domains.

6. CONCLUSION

This study presented a decentralized collaborative model training framework for medical image analysis using federated learning with enhanced privacy controls. The research addressed critical challenges associated with centralized machine learning in healthcare, including data silos, privacy risks, and regulatory constraints. By leveraging federated learning, the proposed approach enabled multiple institutions to collaboratively train models without sharing sensitive patient data, ensuring compliance with privacy requirements.

A key contribution of this work is the development of an adaptive aggregation mechanism that dynamically switches between Federated Averaging (FedAvg) and Federated Stochastic Gradient Descent (FedSGD) based on data divergence. This strategy effectively mitigates the impact of non-independent and identically distributed (non-IID) data, improving model convergence and overall performance. The

integration of differential privacy further strengthened the framework by protecting against potential information leakage during model updates.

Experimental evaluation across multiple medical imaging datasets, including tuberculosis chest X-rays, brain tumor MRI scans, and diabetic retinopathy images, demonstrated that the proposed framework outperforms traditional static aggregation methods. Additionally, modern deep learning architectures such as EfficientNetV2 and ResNet-RS showed superior performance compared to baseline models, highlighting their suitability for federated environments.

Overall, the findings confirm that adaptive federated learning provides a scalable, efficient, and privacy-preserving solution for collaborative medical AI systems, with strong potential for real-world healthcare applications.

7.FUTURE SCOPE

Future research can extend this work by exploring advanced deep learning architectures such as Vision Transformers and hybrid models to further improve performance in federated settings. The integration of multimodal medical data, including clinical records and imaging data, can enhance diagnostic accuracy and provide more comprehensive insights. Additionally, implementing the proposed framework in real-world hospital environments will be essential to evaluate its practical feasibility and scalability under real network conditions.

Further improvements can be achieved by incorporating advanced privacy-preserving techniques such as homomorphic encryption and blockchain-based secure federated learning. Optimizing communication efficiency through model compression and adaptive communication strategies is another promising direction. These advancements will contribute to the development of more robust, secure, and scalable decentralized healthcare systems.

REFERENCES

1. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H.B., Patel, S., Ramage, D., Segal, A. and Seth, K. (2017) Practical secure aggregation for privacy-preserving machine learning. Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, pp. 1175–1191.

2. Dwork, C., Roth, A., et al. (2014) The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science, 9(3–4), pp. 211–407.

3. Esteva, A., Kuprel, B., Novoa, R.A., Ko, J., Swetter, S.M., Blau, H.M. and Thrun, S. (2019) Dermatologist-level classification of skin cancer with deep neural networks. Nature Medicine, 25(1), pp. 24–29.

4. Gentry, C. (2009) Fully homomorphic encryption using ideal lattices. Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC), pp. 169–178.

5. Kairouz, P., McMahan, H.B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A.N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R. and D'Oliveira, R. (2021) Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1–2), pp. 1–210.

6. Li, T., Sahu, A.K., Talwalkar, A. and Smith, V. (2020) Federated learning: Challenges, methods, and future directions. IEEE Signal Processing Magazine, 37(3), pp. 50–60.

7. Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., van der Laak, J.A.W.M., van Ginneken, B. and Sánchez, C.I. (2017) A survey on deep learning in medical image analysis. Medical Image Analysis, 42, pp. 60–88.

8. McMahan, H.B., Moore, E., Ramage, D., Hampson, S. and Arcas, B.A.Y. (2017) Communication-efficient learning of deep networks from decentralized data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), pp. 1273–1282.

9. Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H.R., Albarqouni, S., Bakas, S., Galtier, M.N., Landman, B.A., Maier-Hein, K. and Ourselin, S. (2020) The future of digital health with federated learning. NPJ Digital Medicine, 3(1), pp. 1–7.

10. Shokri, R., Stronati, M., Song, C. and Shmatikov, V. (2017) Membership inference attacks against machine learning models. Proceedings of the IEEE Symposium on Security and Privacy, pp. 3–18.

11. Abbas, S.R., Abbas, Z., Zahir, A. and Lee, S.W. (2024) Federated learning in smart healthcare: A comprehensive review on privacy, security, and predictive analytics with IoT integration. Healthcare, 12(24), p.2587.

12. Choi, G., Cha, W.C., Lee, S.U. and Shin, S.Y. (2024) Survey of medical applications of federated learning. Healthcare Informatics Research, 30(1), pp.3–15.

13. Dhade, P. and Shirke, P. (2024) Federated learning for healthcare: A comprehensive review. Engineering Proceedings, 59(1), p.230.

14. Guan, H., Yap, P.T., Bozoki, A. and Liu, M. (2024) Federated learning for medical image analysis: A survey. Pattern Recognition, 151, p.110424.

15. Khalil, S.S., Tawfik, N.S. and Spruit, M. (2024) Exploring the potential of federated learning in mental health

research: A systematic review. Applied Intelligence, 54, pp.1619–1636.

16. Shah, S.T., Ali, Z., Waqar, M. and Kim, A. (2025) Federated learning in public health: A systematic review of decentralized approaches. Healthcare, 13(21), p.2760.

17. Xu, J., Glicksberg, B.S., Su, C., Walker, P., Bian, J. and Wang, F. (2021) Federated learning for healthcare informatics. Journal of Healthcare Informatics Research, 5, pp.1–19.

18. Yu, H., Cai, L., Min, H. and Su, X. (2024) Advancing medical data classification through federated learning and blockchain incentive mechanisms. Journal of Supercomputing, 80, pp.10469–10484.

19. Zhang, F., Kreuter, D., Chen, Y., Dittmer, S., Tull, S., Shadbahr, T. and Roberts, M. (2024) Recent methodological advances in federated learning for healthcare. Patterns, 5(6), p.101006.

20. Joshi, M., Pal, A. and Sankarasubbu, M. (2022) Federated learning for healthcare domain: Pipeline, applications and challenges. ACM Transactions on Computing for Healthcare.

21. Li, M., Xu, P., Hu, J., Tang, Z. and Yang, G. (2024) From challenges to opportunities: Implementing federated learning in healthcare. arXiv preprint.

22. Gao, J. and Li, Y. (2024) FedMetaMed: Federated meta-learning for personalized medication. arXiv preprint.

23. Sheller, M.J., Reina, G.A., Edwards, B., Martin, J. and Bakas, S. (2020) Multi-institutional deep learning modeling without sharing patient data. Scientific Reports, 10, p.12598.

24. Yang, Q., Liu, Y., Chen, T. and Tong, Y. (2019) Federated machine learning: Concept and applications. ACM Transactions on Intelligent Systems.

25. Li, X., Huang, K., Yang, W., Wang, S. and Zhang, Z. (2020) On the convergence of federated optimization in heterogeneous networks. MLSys Conference.

26. Karimireddy, S.P., Kale, S., Mohri, M., Reddi, S.J., Stich, S.U. and Suresh, A.T. (2020) SCAFFOLD: Stochastic controlled averaging for federated learning. ICML.

27. Acar, D.A.E., Zhao, Y., Navarro, R., Mattina, M., Whatmough, P. and Saligrama, V. (2021) Federated learning based on dynamic regularization. ICLR.

28. Truex, S., Liu, L., Chow, K.H., Gursoy, M.E. and Wei, W. (2020) LDP-Fed: Federated learning with local differential privacy. IEEE Big Data.

29. Bonawitz, K. et al. (2019) Towards federated learning at scale: System design. MLSys.

30. Hard, A., Rao, K., Mathews, R. et al. (2018) Federated learning for mobile keyboard prediction. arXiv preprint.

31. Brisimi, T.S., Chen, R., Mela, T., Olshevsky, A., Paschalidis, I.C. and Shi, W. (2018) Federated learning of predictive models from federated EHR data. International Journal of Medical Informatics.

32. Rieke, N. et al. (2020) The future of digital health with federated learning. NPJ Digital Medicine.

33. Kairouz, P. et al. (2021) Advances and open problems in federated learning. Foundations and Trends in ML.

34. Dwork, C. and Roth, A. (2014) The algorithmic foundations of differential privacy.

35. Geyer, R.C., Klein, T. and Nabi, M. (2017) Differentially private federated learning. arXiv preprint.

36. Abadi, M. et al. (2016) Deep learning with differential privacy. ACM CCS.