

# The Secure Future of Medical AI on Privacy-Preserving Federated Learning-A Systematic Review

Fidha Fathima Salim\*, Karthik Unnikrishnan\*\*, Abhin K.S\*\*\*, Vidhula Thomas\*\*\*\*

\*(Integrated MSc Computer Science-Data Science, Nirmala College Muvattupuzha, MG University, Kerala, India  
Email: [fidhafathimasalim73@gmail.com](mailto:fidhafathimasalim73@gmail.com))

\*\* (Integrated MSc Computer Science-Data Science, Nirmala College Muvattupuzha, MG University, Kerala, India  
Email: [karthikunnikrishnan2003@gmail.com](mailto:karthikunnikrishnan2003@gmail.com))

\*\*\* (Integrated MSc Computer Science-Data Science, Nirmala College Muvattupuzha, MG University, Kerala, India  
Email: [abhinks2104@gmail.com](mailto:abhinks2104@gmail.com))

\*\*\*\* (Assistant Professor, Integrated MSc Computer Science-Data Science, Nirmala College Muvattupuzha,  
MG University, Kerala, India  
Email: [vidhulathomas90@gmail.com](mailto:vidhulathomas90@gmail.com))

\*\*\*\*\*

## Abstract:

The study presents a comprehensive review of Federated Learning, outlining a practical method for building strong Artificial Intelligence models in the healthcare industry while maintaining patient privacy. A key issue identified is that privacy laws like GDPR and HIPAA often result in medical data being scattered across different organizations, complicating the development of dependable and broadly applicable AI solutions. FL addresses this by allowing teams to train models together without sharing the actual sensitive data.

Research underscores that traditional Federated Learning faces challenges due to non-IID data distributions and susceptibility to privacy inference attacks. This paper examines advanced frameworks that implement a multi-layered defense system applying techniques like Secure Multi-Party Computation, Homomorphic Encryption, and Differential Privacy to tackle these issues. The key advancements discussed affect hierarchical architectures with edge servers to enhance efficiency, dynamic aggregation strategies to manage data heterogeneity, and adaptive privacy budget allocation to achieve an optimal balance between privacy and utility.

The conceptual framework of this study highlights how these enhanced federated learning approaches offer robust privacy safeguards and achieve remarkable diagnostic accuracy, often matching or even exceeding traditional centralized models. However, despite these achievements, the main challenge—especially when using advanced encryption techniques—remains the substantial processing overhead.

**Keywords** — Federated Learning (FL), Privacy-Preserving AI, Medical AI, Healthcare Analytics, GDPR & HIPAA, Non-IID Data, Differential Privacy (DP), Homomorphic Encryption (HE), Secure Multi-Party Computation (SMPC), Medical Imaging, .

\*\*\*\*\*

## I. INTRODUCTION

Medical data is frequently divided, which hinders AI from achieving its full capabilities. Although this fragmented information can aid in developing tailored treatments and enhancing diagnoses, it still presents difficulties. Regulations such as GDPR and HIPAA mandate that sensitive patient data remain confined to

particular institutions, complicating the creation of robust and broadly applicable AI models.

A viable option is federated learning, allowing various healthcare institutions to collaborate on developing a unified model utilizing their individual data sets. Rather than sharing actual patient data, they share updates regarding the model, which aids in maintaining the privacy of patient information.

Implementing federated learning in the healthcare sector presents certain challenges. In practical medical contexts, data from various hospitals or organizations often differ significantly—they are not independent or identically distributed. This variability can hinder model performance and slow down the learning process. Furthermore, although federated learning prevents the direct sharing of raw data, it still raises privacy concerns. Hackers might employ methods such as deducing whether an individual is included in the dataset or reverse-engineering the model to extract sensitive information. To enhance the safety of the learning process, we must adopt more robust privacy measures, such as Differential Privacy.

This proposal presents a novel federated AI system that emphasizes privacy and is tailored to address the intricate issues within healthcare. To safeguard against inference attacks, our system implements a multi-layered security strategy combined with established privacy methodologies like Differential Privacy and Homomorphic Encryption. Additionally, it features advanced, customized learning techniques, including a dynamic aggregation algorithm and an adaptive node weighting system. These mechanisms assist in managing data distribution disparities and contribute to the creation of an accurate and broadly applicable global model.

## II. LITERATURE REVIEW

In 2022, W. Wang and their team [1] presented a novel framework known as CFPF designed to effectively manage disease outbreaks while safeguarding individuals' privacy. They utilized data from 4G mobile base stations and employed a technique called K-Means clustering to categorize users. Subsequently, they developed a specialized model that integrates CNN and LSTM while implementing a privacy protection method referred to as Local Differential Privacy (LDP). When comparing this model to other robust methods, it demonstrated superior accuracy with a Root Mean Squared Error (RMSE) of 9.89. This indicates that the framework effectively strikes a balance between privacy and practicality during public health emergencies. However, a significant concern is that the framework adopts a risk-based approach to

privacy, which could leave certain high-risk users unprotected, raising serious ethical issues if utilized in the future.

In 2022, T. J. Loftus and their team [2] published findings from using federated learning (FL) in multiple countries to aid in the diagnosis of COVID-19. They trained a DenseNet model using 87,956 chest X-ray images collected from five hospitals around the world. To ensure consistency, these images were standardized through a Common Data Model (CDM), allowing them to evaluate different FL techniques like FedAvg and FedBN. Their research indicated that specialized FL approaches, particularly FedBN, performed better with diverse datasets, achieving an AUROC of 0.78 when applied to new data. This emphasizes the significance of selecting the appropriate algorithm for real-world, heterogeneous data. While the study addressed privacy issues, it did not comprehensively compare or examine advanced privacy safeguarding techniques. This indicates a need for further exploration of these strategies in various contexts.

M. Adnan and his team [3] looked into using differentially private federated learning to study complicated histopathology images in a 2022 study. They used the Cancer Genome Atlas (TCGA) dataset to make a simulated decentralized setup for identifying different types of lung cancer. Their approach included using Differentially Private Stochastic Gradient Descent (DP-SGD) to protect privacy, along with FedAvg for combining results. Because Whole Slide Images (WSIs) are very large, they used Multiple Instance Learning (MIL), where each WSI is shown as a "mosaic" made from smaller, important parts. The results showed that federated learning training worked better than non-collaborative training and matched the performance of centralized models, especially when there were few clients involved. Even with strict privacy settings ( $\epsilon=2.90$ ), the federated learning model kept its strong performance. One big problem was that the models were very sensitive to the DP-SGD settings, and their performance got worse as the data was split among more customers.

In 2023, M. Butt and their research team [4] conducted a study on a fog-based federated learning system aimed at screening for COVID-19. Their findings indicated that a global model outperformed those trained locally. They employed the FedAvg algorithm to train convolutional neural network (CNN) models and incorporated a fog computing layer for facilitating the federated learning process. The research utilized a dataset of chest X-rays related to COVID-19, comprising 21,165 images sourced from simulated hospitals. The outcomes were remarkable: the global federated learning model achieved an F1-score of 94%, significantly surpassing the local model's score of 88%. This highlights how federated learning can produce more robust and generalized models by utilizing data from various sources. Although the model demonstrated high accuracy, the study acknowledged a significant limitation: it did not assess how well the model could safeguard privacy against more sophisticated types of attacks. This indicates the necessity for future research to create frameworks that enhance privacy protection.

In 2023, B. Sindhusaranya and their team [5] presented a novel approach to prevent fraudulent activities in IoMT systems. They integrated federated learning with blockchain technology to develop a system known as FL-BEPP. This solution operates within a fog-cloud environment, utilizing blockchain for secure and immutable record-keeping while employing federated learning to train models at various locations. Their experiments indicated that this method was more energy-efficient and quicker, reducing the processing time from approximately 1.5 minutes to 1.34 minutes. Nevertheless, a significant concern was that they did not evaluate their system using actual healthcare data or assess its accuracy with established metrics. The research also emphasized that privacy is safeguarded due to the inherent security features of both federated learning and blockchain, highlighting the necessity for further testing, particularly with techniques that enhance privacy.

SecureHealth is a proposed model for implementing federated learning in healthcare analytics utilizing various cloud services. It was put

forward by H. Zhang and their colleagues in 2024 [6]. The framework employs three primary privacy techniques—homomorphic encryption, differential privacy, and secure multi-party computation—in a layered approach to safeguard data. This configuration is designed to be compatible with various systems while adhering to legal regulations. The research demonstrated that this layered strategy can effectively yield results comparable to training data consolidated in a single location, although there is a minor decrease in performance, roughly between 3 to 7 percent. While the design is thorough and well-structured, the significant drawback is that homomorphic encryption requires extensive computational resources, estimated to be 10 to 40 times more than what is necessary, posing challenges for practical application.

In 2024, H. Xie and their team [7] introduced a new way to share medical data that improves federated learning with adaptive differential privacy. To reduce communication costs, their approach uses a layered structure with a central server, local edge servers for early data combining, and hospitals or clinics. A major new feature is the adaptive privacy budget method, which changes how much privacy noise is added based on how well the model is learning. This method gives more privacy budget—meaning less noise—in later stages of training to keep the model useful while still protecting privacy. On real-world medical datasets, the framework reduced privacy loss by 85% while achieving 92.5% accuracy, which is only a small 2.3% drop compared to methods that don't protect privacy. It also cut the success rate of membership inference attacks by 87%, showing strong resistance to attacks. The framework's efficiency is mainly supported by both theoretical and experimental evidence in this study, but there's not much discussion about the computational costs or challenges of implementing such a complex, adaptive system in real IT environments.

In 2025, S. M. Orthi and their team [8] introduced a technique for classifying tuberculosis using federated learning, which is effective for data that is not uniformly distributed. They evaluated their method on a dataset consisting of 7,000 chest X-rays

obtained from Kaggle. To ensure data balance, they employed a method known as SMOTE. They trained three different models—ResNet, DenseNet, and SqueezeNet—utilizing a conventional federated learning framework. The ResNet model yielded the highest performance, achieving an accuracy of 96.7% and a F1-score of 97.4%. This indicates that ResNet is highly effective for medical imaging tasks within a federated learning context. Nonetheless, the research had a significant drawback: it relied on the inherent privacy features provided by federated learning. Although SMOTE addressed issues related to uneven data distribution, the study did not investigate alternative methods to safeguard against various attacks, such as differential privacy or homomorphic encryption. This oversight underscores a critical gap in defending models against different kinds of threats.

In 2025, R. Haripriya and their team [9] unveiled a sophisticated FL framework that incorporates a novel adaptive aggregation technique alongside transfer learning. They trained models such as EfficientNetV2 on three medical imaging datasets—TB, Brain Tumor, and Diabetic Retinopathy—that were distributed in a non-IID manner. Their approach utilized an algorithm that alternated between FedAvg and FedSGD based on the circumstances. EfficientNetV2 achieved an impressive accuracy of 97.4% in identifying gliomas, indicating that the adaptive aggregation method could nearly compete with the performance of a centralized system. The findings also indicated that this adaptive strategy outperformed fixed approaches. Nonetheless, a significant drawback of their framework was its substantial computational requirements when implemented with a large number of clients, such as 50 to 100 or more. This points to the necessity for further advancements to enhance scalability and explore techniques such as SMPC for improved security.

In 2025, A. Ali and their team [10] launched a new system named Health-FedNet, which is intended to securely analyze healthcare data. This system was built using information from the MIMIC-III database and features a privacy strategy that incorporates three primary techniques: Homomorphic Encryption, Differential Privacy, and an Adaptive Node

Weighting Mechanism. Health-FedNet achieved an accuracy rate of 92.4%, representing a significant enhancement with a 12% increase in accuracy compared to conventional systems that consolidate all data in one location. This demonstrates the effectiveness of their approach. Nonetheless, the system does face certain challenges. It depends on a reliable central server to collect and process the information, and the implementation of Homomorphic Encryption results in a slowdown, making it approximately 4 to 10 times slower than typical systems. This underscores the necessity for further research to optimize advanced encryption techniques.

Reference	Summary	Dataset Used	Algorithm Used	Key Findings	Limitations and Gaps
Wang et al. (2022)	Using privacy-preserving FL (CFPF) to anticipate crowd flow in epidemic situations. utilizes Local DP and CNN-LSTM	6,000 individuals in Nanjing used 4G mobile data between March 1 and March 22.	Local DP (Planar Laplacian), MFCL (CNN-LSTM), and FL with K-Means clustering.	achieved a high accuracy/privacy balance (RMSE 9.89). Accuracy increased with client clustering.	Mobility data, not clinical, was the main focus. The accuracy vs. privacy trade-off is still open.
Loftus et al. (2022)	Actual world FL implementation for using chest radiographs to diagnose COVID-19. demonstrates how FL can improve generalizability while preserving privacy.	87,956 chest radiographs from five different healthcare systems make up this unique dataset.	DenseNet, CDMs, FL (FedAvg, FedBN, FedAMP, FedProx)	Particularly for limited datasets, FedBN fared better than local models. Performance was better with collaboration than with standalone models.	High data heterogeneity hurts FedAvg. Privacy leaks can still occur as a result of adversarial attacks.
Adnan et al. (2022)	This case study uses TCGA data to classify histopathology images using a differentially	LUAD and LUSC are classified using 2,580 lung cancer WSIs from the Cancer Genome	Federated Averaging (FedAvg), or FL sequestration Rényi DP accountant	The performance of private FL is similar to that of centralized training( $\epsilon$ )	Information is still blurred by FL without DP. Performance deterioration is revealed

	private FL frame. It examines how customer count and distribution(IID vs.non-IID) affect model performance and demonstrates how FL with sequestration can nearly act centralized training delicacy.	Atlas( TCG A). Data that has been divided into different distributions and hospitals.	with DP-SGD Model shop frame for bracket exercising pre-trained DenseNet MEM	= 2.90). FL performs noticeably better than original-only models.As further guests admit data, performance deteriorates	by external confirmation , suggesting that sphere adaption is still an open issue.
Butt et al. (2023)	FL based on fog computing for COVID-19 CXR screening. designed to handle non-IID, unbalanced data.	Database of COVID-19 Radiography (21,165 photos).	Fog computing , CNNs, and FL (FedAvg).	The global FL model scored better on classification measures than the local ones.	The number and quality of local data determine this. Sync delay and coordination overhead in a decentralized system.
Sindhusaranya et al. (2023)	Suggestions Blockchain technology is being used by FL-BEPP to avoid fraud in IoMT. Pay attention to fog-cloud architecture security.	Conceptual framework for ECG, blood pressure monitoring, etc., is not described.	FL, Fog-Cloud architecture, and Blockchain (SHA-256).	Simulations use less power and have less delay than the baseline.	only simulated; actual medical data was not used for testing. The fraud detection model is not thoroughly investigated.
Zhang et al. (2024)	Formulti-cloud healthcare data, SecureHealth FL with DP, Homomorphic Encryption, and SMPC.	Unnamed; assessed using factual medical data.	FL, SMPC, DP( adaptive budget), and HE( RNS-CKKS).	delicacy close to centralized models( difference $\leq 7$ ). robust layers of sequest	10 to 40 x slowness is caused by HE. requires defined data formats, which reduces rigidity.
Xie et al. (2024)	This frame uses allied literacy and adaptive discrimination sequestration to cover medical data. It features a binary-subcaste sequestration system( original and central DP), a hierarchical aggregation structure with edge waiters, and an adaptive sequestration budget to balance sequestration and mileage.	Four real-world medical datasets • CT reviews( 50,000 images) • EHR Records( 10,000 entries) • Clinical Notes( 75,000 textbooks) • Lab Results( 200,000 records)	Framework Hierarchical Federated Learning Privacy • Adaptive DP( Original Central) • sequestration Budget Allocation Model Deep Neural Network( 6 conv layers 3 FC layers) Aggregation Edge-gerçon weighted averaging grade trimming	Achieved 92.5 delicacy( only 2.3 lower thannon-private FedAvg). Reduced sequestration loss by 85 vs. birth DP. eased class conclusion attacks by 87. Adaptive budgeting bettered sequestration-mileage trade-off.	Non-IID data linked as a challenge, but aggregation robustness is only compactly banded. Optimal tuning of sequestration - mileage balance remains a complex issue, especially for different clinical tasks.
Orthi et al. (2025)	Keeping privacy in mind, this study compares SqueezeNet, DenseNet, and ResNet for detecting tuberculosis using chest X-rays with federated learning.	The data used comes from a Kaggle dataset with 7,000 images.	The methods tested include federated learning, SMOTE, ResNet, DenseNet, and SqueezeNet.	ResNet achieved the highest accuracy of 96.7%. When compared to training models on a single set of data, the federated learning setup worked better.	However, there are issues like delays and uneven data distribution across devices. So far, these results haven't been tested in real hospital networks.
HariPriya et al. (2025)	Based on data divergence, adaptive FL aggregation alternates between FedAvg and FedSGD. It	Kaggle datasets: Diabetic Retinopathy, Brain Tumor MRI, Chest X-ray.	Models used: ResNet-RS, FL (FedAvg + FedSGD), GoogLeNet, VGG16,	Compared to static FL, adaptive aggregation enhances accuracy and convergence	Realism is limited by fixed client/round assumptions. VGG16 has a large overhead.

	employs contemporary CNN architectures.		and EfficientNetV2.	e. EfficientNetV2 is the most effective.	
Ali et al. (2025)	Health-FedNet uses adaptive node weighting along with federated learning, differential privacy, and homomorphic encryption.	It works with the clinical database called MIMIC-III.	The system includes differential privacy, homomorphic encryption, federated learning, and adaptive node weighting.	It is 12% more accurate than models that are not decentralized. The weighting helps the model converge faster.	However, the trusted aggregator is a weakness. This setup also slows things down by 4 to 10 times.

### III. METHODOLOGIES

#### A. Federated Learning (FL) Frameworks

Federated Learning (FL) is a decentralized machine learning methodology designed to train a global model collaboratively across multiple guests, similar as hospitals or smart bias, without taking them to partake their sensitive raw data. The process generally begins with a central garçon initializing a global model and distributing it to sharing guests. Each customer also trains this model locally on its own private dataset to produce an streamlined set of model parameters or slants. These original updates, not the underpinning data, are transferred back to the central garçon for aggregation. A crucial challenge in real- world medical operations is that data is frequently miscellaneous and not independent and identically distributed(non-IID) across institutions. For case, different hospitals may use different imaging outfit or have patient populations with varying characteristics, leading to statistical differences in their original datasets. To manage this, some advanced infrastructures use a hierarchical structure with edge waiters that act as intermediate aggregators for near medical institutions. This reduces communication outflow on the central

garçon and can enhance sequestration.

#### a. Federated Averaging (FedAvg)

An aggregation mechanism that is frequently used in FL frameworks is Federated Averaging (FedAvg). This method involves each client receiving the server's current global model, using stochastic gradient descent (SGD) to conduct several local training steps, and then sending back to the server its updated model weights. With weights usually proportional to the size of each client's local dataset, the central server then updates the global model by averaging the incoming model parameters. Mathematically, the objective of FL is to minimize a global loss function

$f(w)$ , which is the average of the local loss functions  $f_i(w)$  for each of the  $n$  clients:

$$\min_w f(w) \quad \text{with} \quad f(w) = \sum_{i=1}^n \frac{n_i}{N} f_i(w)$$

where  $w$  represents the model weights,  $n_i$  is the number of data points on client  $i$ , and  $N$  is the total number of data points across all clients.

#### b. Federated Stochastic Gradient Descent (FedSGD)

Another aggregation approach used in FL is Federated Stochastic Gradient Descent (FedSGD), in which clients create the gradients of the loss function using their local data and send them to the central server instead of the complete model weights. In order to update the parameters of the global model, the server then averages these gradients from each participating client. Compared to FedAvg, this may necessitate more frequent communication, but each round's global model can be updated more finely using gradients.

#### c. Adaptive Aggregation

A revolutionary technique called adaptive aggregation uses the network's current state to dynamically select the best aggregation algorithm, increasing the robustness of federated learning. This method keeps track of the data divergence between clients, which indicates how much the local models deviate from the global model and from one another. The method may employ FedAvg, which is a communication-efficient algorithm, if the divergence is

minimal. A more careful approach, such as FedSGD, might be used if the divergence is substantial, though, to make sure the global model can reliably learn from heterogeneous inputs.

## B. Privacy-Preserving Techniques

### a. Differential Privacy (DP)

Since sensitive information can still be deduced from the shared model updates, FL does not provide a formal privacy guarantee, even though it offers a baseline of privacy by not sharing raw data. By introducing precisely calibrated noise into data, the mathematical framework known as Differential Privacy (DP) offers robust, measurable privacy guarantees. A mechanism is considered

$(\epsilon, \delta)$ -differentially private if, for any two datasets  $D$  and  $D'$  that differ by a single element, the probability of producing any output  $S$  is nearly the same:

$$\Pr[M(D) \in S] \leq \exp(\epsilon) \cdot \Pr[M(D') \in S] + \delta$$

Here,

$\epsilon$  (epsilon) is the **privacy budget**, which controls the trade-off between privacy and utility—a smaller  $\epsilon$  means stronger privacy.

Differentially Private Stochastic Gradient Descent (DP-SGD), which clips gradients to a predetermined norm and then adds Gaussian noise before aggregating them, is a popular machine learning technique for accomplishing this.

### b. Budget Allocation for Privacy and Adaptive Differential Privacy

The privacy-utility trade-off is frequently not optimized during training by static DP techniques. An sophisticated technique called adaptive differential privacy dynamically modifies the noise levels and privacy budget according to training success. Model performance can be maintained while maintaining privacy by allocating a bigger privacy budget (less noise), for instance, at later training stages when gradients are smaller and more susceptible to noise. Throughout the training process, this approach makes it possible to use the total privacy budget

more effectively.

### c. Homomorphic Encryption (HE)

An advanced cryptographic method called homomorphic encryption (HE) enables calculations to be done directly on encrypted data. Before submitting their local model modifications to the central server, clients in FL have the option to encrypt them. To guarantee that the server never sees the plaintext model updates, it can then aggregate these encrypted updates without decrypting them.

### d. Multi-Party Secure Computation (SMPC)

A cryptographic mechanism known as Secure Multi-Party Computation (SMPC) allows several parties to collaboratively compute a function utilizing their private inputs without disclosing those inputs to one another. SMPC can be applied to secure aggregation in FL. Secret "shares" of each client's update are distributed to other participants. No single party can rebuild an individual client's update because the aggregated outcome is calculated from these shares.

## C. Utilization in Imaging Medicine

Because patient privacy is a primary concern, there aren't many large-scale, centralized medical datasets. Since histopathological photos are some of the biggest and most intricate medical images, FL and DP together provide a promising future for collaborative training on sensitive medical data. Using high-resolution Whole Slide Images (WSIs) to classify lung cancer subtypes, such as Lung Adenocarcinoma (LUAD) and Lung Squamous Cell Carcinoma (LUSC), is one particular use.

### a. MIL, or Multiple Instance Learning

The size of WSIs is typically gigapixel, therefore direct processing is not possible. The supervised learning method known as Multiple Instance Learning (MIL) is ideal for this issue. MIL assigns a single label (such as LUAD) to a whole WSI (a "bag"). A model learns to classify the bag based on the collection of its instances, which are little patches ("instances") that make up the WSI. Often, this procedure entails forming the bag from a "mosaic" of sample areas from the WSI.

### b. Classifying Medical Images Using CNN Models

Classification of medical images is a common application for powerful Convolutional Neural Network (CNN) architectures. The DenseNet model was utilized as a feature extractor for patches in the histopathological study. It connects each layer to every other layer in a feed-forward manner. The models ResNet, VGG16, GoogLeNet, and EfficientNetV2 are also often utilized, frequently utilizing transfer learning. After being pre-trained on a sizable dataset such as ImageNet, they are refined on specific medical datasets. This method saves medical data for training and enhances performance by enabling the models to build on a solid foundation of learnt visual elements.

### c. CNN-LSTM with Multiple Factors (MFCL)

One hybrid deep learning model for predicting population flow during epidemics is the Multi-Factors CNN-LSTM (MFCL). Long-term temporal relationships are captured by the LSTM component after the CNN component captures local spatial characteristics from trajectory data and analyzes them as a time series. Additionally, the algorithm takes into account outside variables like holidays and the weather to increase prediction accuracy.

## IV. RESULTS AND DISCUSSIONS

Using the MIMIC- III clinical database, the Health-FedNet frame showed a notable enhancement in performance compared to typical centralized approaches. With a opinion delicacy of 92.4, it outperformed the centralized CNN( 78.2) and LSTM( 80.1) models by 12. It was determined that this enhancement was statistically significant(  $p < 0.01$ ). Its adaptive knot weighting approach, which gives precedence to inputs from advanced- quality data sources, was set up to be effective in reducing vaticination friction through miscellaneous(non-IID) bumps by 35. In addition, the frame demonstrated exceptional effectiveness by lowering quiescence by 18- 54 and bandwidth consumption by 28- 38 when compared to conventional allied literacy

models. Differential sequestration, homomorphic encryption, and adaptive knot weighting all gradationally bettered the model's performance, according to an ablation study.

The main outgrowth of the SecureHealth frame is striking a good balance between model functionality and strong sequestration in amulti-cloud setting. Its model delicacy was within 3- 7 of conventional centralized ways when tested on real-world healthcare datasets, indicating great value while upholding strict sequestration constraints. The frame offers complete protection against a variety of pitfalls, similar as model inversion and class conclusion attacks, thanks to itsmulti-layered sequestration approach, which combines Homomorphic Encryption( HE), Differential sequestration( DP), and SecureMulti-Party Computation( SMPC). It's claimed that the processing cost, which can range from 10 to 40 times depending on the complexity of the model, is a reasonable trade- off for the increased security. also, the frame is made to misbehave with important healthcare laws like GDPR and HIPAA.

This work combined a unique adaptive aggregation system with transfer literacy to gain state- of- the- art delicacy in medical image bracket. The main finding is that dynamic aggregation, which switches between FedAvg and FedSGD according to data divergence, constantly performed better than stationary ways. The frame achieved a peak delicacy of 97.6 on the brain excrescence dataset and 96.2 on the TB dataset by combining this fashion with a contemporary ResNet- RS design. In this allied arrangement, the performance of more recent designs like as ResNet- RS and EfficientNetV2 was constantly better than that of earlier models similar as VGG16 and GoogLeNet. For illustration, the adaptive system using GoogLeNet outperformed FedAvg( 95.9) and FedSGD( 95.5) with an delicacy of 96.3 on the diabetic retinopathy dataset.

The primary discovery of this study is that fog computing-enhanced global Federated Learning (FL) models regularly outperformed localized models trained at individual universities. Across all simulated hospitals, the FL model's accuracy was noticeably higher in studies for COVID-19 screening using chest X-rays. For instance, the accuracy of the FL model was 94%, but the local model at Hospital A was approximately 86%. In a similar vein, Hospital C improved from around 87% to about 96.5%, and Hospital B improved from about 84% to 96%. Additionally, the FL model outperformed the local model in terms of precision, recall, and F1-measure; Hospital C's F1-measure was 94%, while the local model's was 88%. This shows how the collaborative FL architecture can handle non-identically dispersed (Non-IID) data well to provide a more robust and accurate global model.

The findings of this paper are more concerned with vaticination error than bracket delicacy. Compared to all birth models, the suggested Crowd Flow Prediction Framework( CFPF), which employs aMulti-Factors CNN- LSTM model, had a lower error rate. A Mean Absolute Error( MAE) of 7.31 and a Root Mean Square Error( RMSE) of 9.89 were the primary performance pointers. This fared better than cutting- edge models similar as ConvLSTM( RMSE 12.32) and ST- ResNet( RMSE 10.64). One important discovery is the effective balance between sequestration and mileage; as long as the sequestration budget(  $\epsilon$ ) was further than 3, the model was suitable to maintain good delicacy thanks to the frame's threat- grounded Original Differential sequestration( LDP) medium.

Results from a multi-national, real- world COVID-19 opinion study are presented in this composition. The main conclusion is that, particularly when estimated on external data, FL models perform better than models trained at a single institution. The FL models fared far better on an external dataset,

whereas a original model's Area Under the wind( AUROC) fell to as low as 0.56. Chancing the optimal system for miscellaneous(non-IID) data is the most significant outgrowth; FedBN continuously beat other algorithms, including the conventional FedAvg. FedBN's robustness in real- world, varied clinical surrounds was demonstrated by its AUROC values of 0.78 and 0.70 on two external confirmation datasets. The study comes to the establishment conclusion that creating a Common Data Model( CDM) is a necessary step in doing useful, multi-institutional exploration.

The findings of this study place more emphasis on system performance indicators than on diagnostic precision. For Internet of Medical Things (IoMT) applications, it has been demonstrated that the suggested FL-BEPP framework, which combines Federated Learning with Blockchain, greatly lowers system latency and power consumption. The FL-BEPP framework, for example, showed a delay of just 1.34 minutes at a high workload of 750, surpassing baseline techniques that ranged from 1.48 to 1.55 minutes. Under the same load, it also consumed 0.41 Watts of electricity, which was higher than the baseline range of 0.42 to 0.48 Watts. Its ability to detect fraud and secure data while optimizing system resources in a distributed fog-cloud environment is the main finding.

The study's findings demonstrate that a Federated Learning architecture with a ResNet model can efficiently and privately classify TB from chest X-rays. When compared to other models, the suggested ResNet model performed the best, with 96.7% accuracy, 96.8% precision, 98.0% recall, and 97.4% F1-score. This outperformed SqueezeNet (94.18% accuracy) and DenseNet (94% accuracy) in the same federation conditions. There were 495 True Positives and 490 True Negatives in the confusion matrix for the ResNet model, compared to just 8 False Negatives and 10 False Positives. The ROC curve's

high Area Under the Curve (AUC) of 0.97 further demonstrates the model's superior diagnostic performance.

This study empirically demonstrates that differentially private federated learning is a feasible and dependable framework for medical image analysis. In trials simulating a distributed environment with histopathology images, the federated learning (FL) models consistently achieved superior performance compared to models trained in a non-collaborative setting. For illustration, with 4 clients in a non-IID data distribution, the FL model achieved an accuracy of  $0.824 \pm 0.01$ , whereas the average non-collaborative accuracy was only  $0.682 \pm 0.10$ . The results also showed that the FL model's performance was similar to that of a conventional centralized model ( $0.848 \pm 0.02$ ) without taking data sharing. Crucially, when applying Differential Privacy (DP-FL), the framework achieved strong privacy guarantees ( $\epsilon = 2.90$ ) with a test accuracy of  $0.823 \pm 0.01$ , again outperforming non-collaborative models (which ranged from 0.556 to 0.701) and approaching the performance of centralized training ( $0.839 \pm 0.01$ ).

With a dual-layer, adaptive differential privacy approach, Xie et al.'s [7] framework focuses on collaborative medical data modeling. The system demonstrated remarkable accuracy of 92.5% and remarkably robust privacy assurances in experimental assessments on real-world medical datasets. When compared to baseline approaches, it successfully reduced privacy loss by 85%, demonstrating its capacity to optimize the privacy-utility trade-off. By reducing the success rate of membership inference attacks by a noteworthy 87%, the system also showed high resilience against privacy assaults. Its innovative adaptive privacy budget allocation technique, which dynamically modifies noise levels according to training progress,

and its hierarchical architecture with edge servers to boost efficiency are responsible for these outcomes.

The framework by Haripriya et al. [9] is the preferred option for medical image analysis if you want the highest individual accuracy. Its innovative adaptive aggregation system can achieve up to 97.6 accuracy. By employing an adaptive node weighting mechanism, Health-FedNet by Ali et al. [10] achieves an outstanding 92.4 accuracy, a 12% enhancement over standard centralized models, making it the preferred choice for clinical irregular data and managing real-world fluctuations in data quality. SecureHealth by Zhang et al. [6] is the preferred choice if non-supervisory compliance and maximum security are your top priorities. Large-scale, multi-cloud collaborations are robustly defended by its complete multi-layered security result, which combines Homomorphic Encryption (HE), Differential Privacy (DP), and Secure Multi-Party Computation (SMPC). There's a small but manageable accuracy trade-off of roughly 3-7% with this high security. Health-FedNet is the most creative and comprehensive solution overall. Through its innovative adaptive node weighting, it provides real-world resilience, ensures strong privacy, and meets three crucial challenges simultaneously in a unique way. This makes it an appealing option for the majority of cutting-edge, useful healthcare installations.

## V. CONCLUSIONS

Federated Learning (FL) is a very successful, feasible, and trustworthy paradigm for privacy-preserving collaborative healthcare analytics, according to the evaluated research. Research continuously demonstrates that without disclosing private patient information, global federated models can get excellent diagnostic accuracy that is frequently on par with traditional centralized methods. One framework, for example, greatly

reduced privacy loss while achieving 92.5% accuracy.

Standard FL is not impervious to privacy risks, though, as different inference attacks can still expose private data through common model parameters. The literature shows that these difficulties can be successfully overcome by using cutting-edge methods. By incorporating a multi-layered defense, robust frameworks improve security. Differential Privacy (DP), in particular, offers a formal, quantitative foundation for privacy assurances. Innovative approaches such as adaptive privacy budget allocation dynamically modify noise levels according to training progress and data sensitivity in order to maximize the critical trade-off between privacy and usefulness.

In addition, creative results similar as flexible and adaptive aggregation algorithms that take institutional variations into account are used to break the problem of data diversity(non-IID), which is a problem that's current across numerous institutions. It has been demonstrated that FL performs better than non-collaborative training, indeed when considering non-IID data.

Notwithstanding these achievements, the abecedarian downsides of these sequestration-conserving ways continue to be their high computational and transmission expenditure. enforcing hierarchical topologies, in which edge waiters carry out early model aggregation to lower communication outflow and ameliorate sequestration, is a possible volition. In order to make safe, large- scale allied learning a doable option for wide clinical perpetration, unborn exploration should concentrate on both creating further featherlight encryption systems and refining these effective, hierarchical aggregation ways.

## REFERENCES

- [1] Wang, W., Yang, G., Bao, L., Ma, K., & Zhou, H. (2022). A Privacy-Preserving Crowd Flow Prediction Framework Based on Federated Learning during Epidemics. *Security and Communication Networks*, 2022, Article ID 8712597.
- [2] Loftus, T. J., Ruppert, M. M., Shickel, B., Ozrazgat-Baslanti, T., Balch, J. A., Efron, P. A., ... & Bihorac, A. (2022). Federated learning for preserving data privacy in collaborative healthcare research. *Digital Health*, 8.
- [3] Adnan, M., Kalra, S., Cresswell, J. C., Taylor, G. W., & Tizhoosh, H. R. (2022). Federated learning and differential privacy for medical image analysis. *Scientific Reports*, 12, 1953.
- [4] Butt, M., Tariq, N., Ashraf, M., Alsagri, H. S., Moqurrab, S. A., Alhakhani, H. A. A., & Alduraywish, Y. A. (2023). A Fog-Based Privacy-Preserving Federated Learning System for Smart Healthcare Applications. *Electronics*, 12(19), 4074.
- [5] Sindhusaranya, B., Yamini, R., Manimekalai, M. A. P., & Geetha, K. (2023). Federated Learning and Blockchain-Enabled Privacy-Preserving Healthcare 5.0 System: A Comprehensive Approach to Fraud Prevention and Security in IoMT. *Journal of Internet Services and Information Security (JISIS)*, 13(4), 199-209.
- [6] Zhang, H., Feng, E., & Lian, H. (2024). A Privacy-Preserving Federated Learning Framework for Healthcare Big Data Analytics in Multi-Cloud Environments. *Spectrum of Research*, 4(1), 1-18
- [7] Xie, H., Zhang, Y., Zhou, Z., & Zhou, H. (2024). Privacy-Preserving Medical Data Collaborative Modeling: A Differential Privacy Enhanced Federated Learning Framework. *Journal of Knowledge Learning and Science Technology*, 3(4), 340-350.
- [8] Orthi, S. M., Rahman, M. H., Mamun, A. A., Khan, M. N., Siddiq, K. B., Uddin, M., ... & Hossain, M. S. (2025). Federated Learning with Privacy-Preserving Big Data Analytics for Distributed Healthcare Systems. *Journal of Computer Science and Technology Studies*, 7(8), 269-281.
- [9] Haripriya, R., Khare, N., & Pandey, M. (2025). Privacy-preserving federated learning for collaborative medical data mining in multi-institutional settings. *Scientific Reports*, 15, 12482.
- [10] Ali, A., Snášel, V., & Platoš, J. (2025). Health-FedNet: A privacy-preserving federated learning framework for scalable and secure healthcare analytics. *Results in Engineering*, 27, 106484.